# Why Order Flow Explains Exchange Rates

Martin D. D. Evans[*]

Richard K. Lyons

9 November 2001

**Abstract**

Evans and Lyons (1999) find that order flow accounts for about two-thirds of variation in the DM/$ rate. Though never tested, the underlying cause of order flow in their model is portfolio shifts unrelated to macroeconomic information (e.g., shifts in risk preferences or shifts in hedging demands). This paper tests whether order flow is caused (in part) by macroeconomic information. Using a two-equation system—one for price and one for order flow—we examine the links between order flow and macro announcements. We find empirical support for the macro-information channel: about one-third of variation in order flow is due to macro announcements. For prices, this implies that about 20 percent of price variation is due to announcement-induced order flow. Another 10 percent of price variation is due to the direct effect of announcements on price (i.e., not involving order flow).

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

---

[*] Respective affiliations are Georgetown University and NBER, and UC Berkeley and NBER. We thank the following for helpful comments: Jon Faust, Antonio Mello, Andrew Rose and seminar participants at U.C. Berkeley and the Empirical Exchange Rate Models conference (Sept. 2001). Both authors thank the National Science Foundation for financial assistance.

# Why Order Flow Explains Exchange Rates

> Thus, economists today still have very limited information about the relationship between equilibrium exchange rates and macroeconomic fundamentals. Accordingly, it is hardly conceivable that rational market participants with complete information about macroeconomic fundamentals could use that information to form precise expectations about the future market-clearing levels of exchange rates.
>
> Isard (1995, p. 182)

Traditional models of exchange rates are based wholly on publicly available macroeconomic information. In specifications of this type, incremental information cannot be conveyed by trading because public information is already impounded in price. As an empirical matter, however, models within this framework perform poorly (Meese and Rogoff 1983a, Frankel and Rose 1995), leading some researchers to conclude that the main determinants of exchange rates are not macroeconomic (Flood and Rose 1995).[1]

A recent paper by Evans and Lyons (1999) adopts a different approach. They relax the strong assumption that all price-relevant information is available publicly. This makes room for trading to convey incremental information. Their model predicts that a variable called "order flow" acts as a proximate cause of exchange-rate movements.[2] This prediction is borne out empirically: they find that order flow accounts for about two-thirds of daily variation in the DM/$ rate. In their model, the underlying cause of this variation (i.e., the cause of the order flow) is private portfolio shifts—unrelated to macroeconomic information—that induce portfolio-balance effects on price. (These shifts can be due to changing risk tolerance, changing hedging demand, changing liquidity demand, etc.) Though

[1] These models do perform better over longer horizons, however (Meese and Rogoff 1983b, Chinn 1991, Mark 1995).

[2] Order flow is not synonymous with trading volume. Order flow—a concept from microstructure finance—refers to *signed* volume. Trades can be signed in microstructure models depending on whether the "aggressor" is buying or selling. (The dealer posting the quote is the passive side of the trade.) For example, a sale of 10 units by a trader acting on a dealer's quotes is order flow of –10.

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

consistent with their results, this underlying cause is never tested directly, leaving the nagging question of what is really driving the order flow.

This paper addresses whether macroeconomic information is an underlying cause of order flow. We do so by examining the links between order flow and macro announcements. Our model distinguishes three sources of exchange-rate variation. The first source mirrors traditional models—public announcement information that is impounded in price immediately and directly (i.e., with no role for order flow). The second source is an indirect effect of public announcement information that operates via induced order flow. The third source of exchange rate variation mirrors Evans and Lyons (1999), namely order flow that is not related to public announcement information.

When brought to the data, our model speaks rather clearly: all three sources of price variation are significant. The indirect effect of announcements (via order flow) accounts for twice as much price variation as the direct effect of announcements (not involving order flow). Thus, even when one would expect order flow's role to be muted, it is still the key driver of price variation. Together, these two announcement effects account for about 30 percent of price variance. The third source—order flow unrelated to announcements—accounts for another (roughly) 30 percent of price variance.

That announcement-related order flow is central to exchange rate determination suggests a possible explanation for the Meese and Rogoff (1983) findings. To understand why, write the price of foreign exchange, $P_t$, in the standard way as a function of current and expected future macro fundamentals: $P_t = g(f_t, f_{t+1}^e)$. In our model, price-setters learn about changes in $f_{t+1}^e$ by observing order flow, and they set prices rationally based on this information. Thus, while macro (macro expectations) is the underlying cause of price changes, order flow here serves as a proximate cause. Now, if the macro variables that order flow is forecasting are largely beyond the one-year horizon, then the empirical link between exchange rates and macro variables will be loose (even if "forecasts" are based on realized future $f_{t+1}$'s out to one year). That macro empirical results are more positive at horizons beyond one year is consistent with this “anticipation” hypothesis.

Though the literature linking exchange rates and announcements is well developed, none of it uses order flow to sort out the relationship. The existing literature has two branches: the first addresses the direction of exchange-rate changes (first moments) and the second, later branch addresses exchange-rate volatility (second moments). A common finding of the first branch is that directional effects of announcements are difficult to detect over horizons of even a couple days because the stream of other factors affecting price swamps them. At higher frequencies, effects are often statistically significant (particularly for the employment report and money-supply announcements), but goodness-of-fit statistics are generally disappointing.[3] The second, later branch of this literature—which focuses on announcement effects on volatility—is partly a response to the difficulty in accounting for exchange-rate first moments.[4] This work finds that the largest exchange-rate moves are linked to the arrival of macroeconomic announcements. On the other hand, though major announcements dominate the volatility picture at release, their ability to account for changing volatility is less than that of systematic features such as ARCH and time-of-day effects (Andersen and Bollerslev 1998).

Finally, our empirical results have an interesting auxiliary implication: market participants appear to be using different models. It has long been theorized that this is the case (i.e., that traders believe in different mappings from fundamentals to price; see, e.g., Frankel and Froot 1990 and Isard 1995). Yet, firm empirical evidence for the hypothesis is lacking. The testable implication we examine here is based on the following (rather general) result from models of risky-asset trade: when (1) information is publicly observed and (2) all market participants agree on the mapping from that information to price, then price

[3] See, for example, Cornell (1982), Engel and Frankel (1984), Hakkio and Pearce (1985), Ito and Roley (1987), Hardouvelis (1988), and Ederington and Lee (1995). For a recent paper that is more successful at finding systematic first-moment effects of announcements see Andersen et al. (2001).

[4] See, for example, Goodhart et al. (1993), DeGennaro and Shrieves (1997), and Andersen and Bollerslev (1998). See also the work on bond prices and announcements, e.g., Fleming and Remolona (1999) and Balduzzi et al. (2001).

adjustment occurs independently of order flow. By using data on macro announcements we insure that the first of these conditions is met. Our finding that the adjustment of price to announcements depends on order flow suggests a violation of the second condition—that all participants agree on the mapping.[5] Order flow appears to convey information about differing individual assessments of announcements' relevance.[6]

The remainder of the paper is in five sections. The next section presents a model for understanding how announcements can induce order flow that subsequently drives price (not a property of traditional exchange-rate models). Section 2 presents our empirical model for disentangling the three sources of price adjustment noted above. Section 3 describes the data. Section 4 presents our empirical results. Section 5 concludes.

## 1. Modeling the Link Between Order Flow and Announcements

Our model of trading serves several important purposes. First, it is designed to accommodate data at the daily frequency (unlike existing transaction-frequency models). To do so, it breaks the trading day into three basic rounds that capture the essentials of how major currency market operate. Second, the model provides a clear null under which causation runs from order flow to price, with interdealer flow serving as the means by which dispersed information is learned. The importance of interdealer flow in the model not only corresponds well to reality, but also produces equations that are estimable: interdealer flow data are available because this type of trading occurs electronically (unlike data on trades between dealers and the public). Third, the model shows why order flow's price impact should persist; this clarification is important for those who believe that

[5] Strictly speaking, our statement regarding the link between order flow and price adjustment bears on the sufficiency of the two conditions for the independence result. One would need necessity as well to conclude definitively based our evidence that agents disagree on the mapping.

[6] This specification is most sensible in an environment where the data-generating process is time varying. Note that differing assessments need not be irrational: in a world where the true model is not obvious, model formulation is likely to be costly, leading to rational disagreement in equilibrium, despite having observed the same macroeconomic information. To formalize this, consider a setting in which agents are able to process a common signal more precisely if they pay a fixed cost to observe a better model (in lieu of paying a fixed cost to observe the signal itself, as in Grossman and Stiglitz 1980).

trades can have only fleeting "indigestion" effects on price.

There are two basic types of information that order flow can convey. The first is information about the stream of future cash flows (i.e., numerators in a security-valuation model). In foreign exchange, this stream includes future interest differentials. The second is information about market-clearing risk premia (i.e., valuation denominators). The trading model we develop here includes both. It adopts a simultaneous-trade approach (see Lyons 2001). Where applicable, we compress our presentation of the model below by using results established for other simultaneous-trade models. At the center of the model is macro announcements that change the information content of order flow, and thereby change the signal processing performed by dealers upon that order flow.

Consider an infinitely lived, pure-exchange economy with two assets, one riskless (with gross return equal to one) and one risky. The periodic (daily) payoff on the risky asset – foreign exchange – is denoted $R_t$, where $R_t$ is composed of a series of increments:

$$R_t = \sum_{i=1}^{t} \Delta R_i \tag{1}$$

The $DR_t$ increments are i.i.d. Normal$(0, s_R^2)$ and are observed publicly each day before trading. These increments represent changes in public information whose impact on the value of foreign exchange is clear to all participants (e.g., changes in interest rates).

The foreign exchange market has two participant types, customers and dealers. There is a continuum of customers, indexed by $z \in [0,1]$ and $N$ dealers, indexed by $i$. The mass of customers on [0,1] is large (in a convergence sense) relative to the $N$ dealers. (This will insure that dealers have a comparative disadvantage in holding overnight positions.) Customers and dealers have constant absolute risk aversion (CARA) and maximize utility of the following form:

$$U_t = E_t\left\{-\sum_{s=0}^{\infty} d^s \exp(-q\, c_{t+s})\right\}$$

where $E_t$ is the expectations operator conditional on agent's information at time t, and $c_{t+s}$ is consumption in period t+s. We assume that all agents have the same time discount factor $d$ and risk aversion parameter $q$. The specifics of the trading environment are described below (with formal setup of the dealer's problem presented in the appendix).

The timing of the model is summarized in Figure 1. Within each 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).

**Figure 1: Daily Timing**

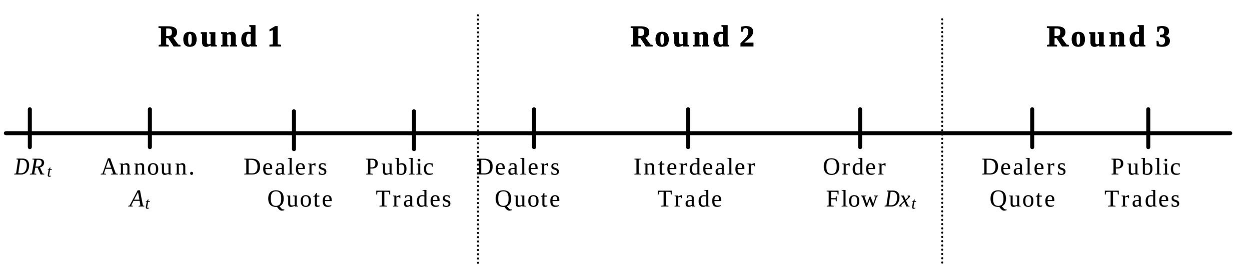


Notes: The daily payoff $R_t$ is paid at the beginning of each day, which includes that day's payoff increment $DR_t$ (publicly observed). The variable $A_t$ is an indicator variable that equals 1 if a macro announcement takes place, 0 otherwise (announcements are publicly observed). The variable $Dx_t$ denotes signed interdealer order flow (publicly observed).

Each day begins with payment of $R_t$, which includes the new payoff increment $\Delta R_t$. Following the payoff realization each day, nature chooses whether to produce a macro announcement. Let the indicator variable $A_t$ equal 1 if there is an announcement on day t, 0 if not. The information in any announcement comes in two parts. The first part is a common-knowledge (or "mean") part: all agents agree that this first part's impact on the exchange rate should be $B_t$. This exchange rate effect $B_t$ arises because this part of the announcement's information is correlated with the next period's payoff, $\Delta R_{t+1}$. The second part of the information in any announcement is dispersed information: this is independent, incremental information about $\Delta R_{t+1}$ that is dispersed among the non-dealer customers (e.g., mutual funds, hedge funds, currency overlay managers).[7] This second part of the announcement's information must be gleaned from the trades of non-dealer customers, described explicitly in the following paragraphs.

Based on available information, each dealer simultaneously and independently quotes a scalar price to his customers at which he agrees to buy and sell any amount.[8] We denote dealer $i$'s round 1 price on day t as $P_{1t}^{i}$. Each dealer then receives a customer order $C_{1t}^{i}$ —unobservable to the rest of the market—that has two components. (The dealer sees only his own customer order in total, not the components separately.) Both components are executed at his quoted price $P_{1t}^{i}$. (Let $C_{1t}^{i}<0$ denote net customer selling—dealer $i$ buying.) The second component is present only if a macro announcement has occurred (indicator variable $A_t$=1):

$$C_{1t}^{i} = \tilde{C}_{1t}^{i} + \hat{C}_{1t}^{i} \qquad \text{with } \hat{C}_{1t}^{i} = 0 \text{ if } A_t = 0$$

The first component $\tilde{C}_{1t}^{i}$ is distributed Normal$(0, s_{\tilde{C}}^{2})$. It is unrelated to the occurrence of an announcement, uncorrelated across dealers, and uncorrelated with the payoff increment $\Delta R_t$. We refer to this first component as "portfolio shifts" of the non-dealer public (a la Evans and Lyons 1999). This component could

---

[7] In the background is an optimizing model along the following lines. There exists a costly information technology that, for agents choosing to pay the cost, produces a more precise payoff expectation from a common piece of data (the announcement). This is conceptually akin to models like Grossman and Stiglitz (1980) in which agents can pay a fixed cost to observe a signal distributed around the payoff ($\Delta R_{t+1}$ in our case).

[8] While it is true that a bid-ask spread of zero would not induce entry into dealing, 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. In equilibrium, expected utility of the commission will just balance the utility cost of marketmaking under risk aversion and asymmetric information. The model's simultaneous-move nature is in the spirit of simultaneous-move games more generally (versus sequential-move games).

arise due, for example, to shifting hedging demands or shifting transaction demands.

The second component of the customer order $\hat{C}^i_{1t}$ is non-zero only on days when an announcement occurs. It is distributed Normal($0, s^2_{\hat{C}}$). It is uncorrelated across dealers and uncorrelated with the first component received by any dealer. Unlike the first component, it is correlated with future payoffs. Specifically, its realization has correlation $r$ with the following day's payoff increment, $\Delta R_{t+1}$. This is important to the model: it implies that order flow induced by the announcement is informative of future payoffs (and dealers will factor this into their processing of order flow signals). Notice from this specification that the variance of customer order flow is higher on days with an announcement.[9]

In round 2, each dealer simultaneously and independently quotes a scalar price $P^i_{2t}$ to other dealers at which he agrees to buy and sell any amount. These interdealer quotes are observable and available to all dealers. Dealers then simultaneously and independently trade on other dealers' quotes. (Orders at a given price are split evenly across all dealers who have quoted that price.) Let $T^i_{2t}$ denote the interdealer trade initiated by dealer $i$ in round 2 (negative for dealer-$i$ net selling). At the close of round 2, all agents observe the interdealer order flow from that period:

$$\Delta x_t = \sum_{i=1}^{N} T^i_{2t} \,. \tag{2}$$

[9] How do dealers respond to announcements in practice? For major scheduled announcements (which represent only a subset of our data), bank economists produce a forecast of the announcement's value (e.g., the employment report) that is distributed to that bank's dealers that morning. The sheet with the forecast will often include a prediction for the exchange rate impact, but dealers are left to their own judgment regarding how to respond. Because spreads are so low in the major markets (often 1/100 of 1 percent), dealers cannot move their post-announcement prices without regard to the prices of others for fear of being arbitraged. Dealers will often shade post-announcement prices slightly to induce trades that help build a desired position. But these price effects are quite small. Our best qualitative reading from observing the process is that some price adjustment does indeed occur rapidly across the market, with little apparent role for flow. But announcements often induce follow-on trading later in the day from customers whose trading response is not instantaneous. It is these customer trading responses that our model is designed to capture.

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 simultaneously and independently quotes a scalar price $P_{3t}^{i}$ at which he agrees to buy and sell any amount. These quotes are observable and available to the public. We assume that total public demand for the risky asset in round-3, denoted $C_{3t}$, is less than infinitely elastic. With our earlier assumptions (CARA-Normal framework with non-strategic customers), this allows us to write total public demand in round 3 as a linear function of the expected return:

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

where the positive coefficient $g$ captures the public's aggregate risk-bearing capacity, and $W_3$ is the available public information (includes all payoff increments $\Delta R_t$ , announcements, and interdealer flows $\Delta x_t$ through day t).

The equilibrium relation between interdealer order flow and price adjustment follows directly from results established for other simultaneous-trade models (Lyons 2001). First, within a given round, all dealers quote a common price (this is necessary for ruling out arbitrage opportunities). It follows that this price is conditioned on common information only. Though each day's payoff increment $\Delta R_t$ and announcement information $B_t$ (if an announcement) are known publicly at the beginning of round 1, interdealer order flow $\Delta x_t$ is not observed until the end of round 2. The price for round-3 trading, $P_{3t}$, reflects the information in all three of these variables, $\Delta R_t$ , $B_t$, and $\Delta x_t$.

The information impounded in the current payoff increment $\Delta R_t$ is straightforward: the current increment is relevant for all future payoffs because increments persist through subsequent days. The information in $B_t$–the common-knowledge part of an announcement—is also straightforward: it is correlated with the next day's payoff increment $\Delta R_{t+1}$ (and by our specification, $B_t$ denotes the actual size of the exchange rate impact). The information conveyed by interdealer order flow $\Delta x_t$ is of two kinds. The first operates only on announcement days and relates to the next day's payoff increment $\Delta R_{t+1}$: there is independent information in customer orders beyond that in $B_t$ that is impounded in interdealer order flow $\Delta x_t$. On announcement days, dealers recognize that underlying customer orders convey information about $\Delta R_{t+1}$, and they adjust the signal extraction coefficients they apply to interdealer flow $\Delta x_t$ accordingly.[10]

The second type of information conveyed by order flow is unrelated to future payoffs, relating instead to portfolio-balance effects. (Evans and Lyons 1999 focus on this type of information.) To understand this second type, consider the special case in which there are never days with an announcement (i.e., $A_t$=0 always). Evans and Lyons (1999) show that in equilibrium, each dealer's interdealer trade, $T^i_{2t}$, is proportional to the first-round customer order $C^i_{1t}$ he receives. This implies that when dealers observe $\Delta x_t$ at the end of round 2 (equation 2), they can infer the aggregate portfolio shift on the part of the public in round 1, $\sum_{i=1}^{N} C^i_{1t}$ (henceforth denoted $C_{1t}$). Dealers also know that for a risk-averse public to re-absorb this portfolio shift willingly in round 3—i.e., to achieve stock equilibrium—price must adjust. In particular, price adjusts in round 3 so that $C_{1t}+C_{3t}=0$, where $C_{3t}$ is given by equation (3).

Given our model is analytically similar to that in Evans and Lyons (1999), we offer only the solution, with reference to their appendix for details. (The only substantive change is the adjustment to the signal extraction coefficients to reflect the added information in order flow on days where $A_t$=1.) The resulting price changes and interdealer flows (end day t-1 to end day t) can be written as:

$$\Delta P_t = (b_1 + b_2 A_t)\Delta x_t \; + \; A_t B_t \; + \; b_3 \Delta R_t \tag{4}$$

$$\begin{aligned} \Delta x_t &= b_4(\tilde{C}^1_t + \hat{C}^1_t) \qquad \text{if } A_t = 1 \\ &= b_5 \tilde{C}^1_t \qquad\qquad\quad \text{if } A_t = 0 \end{aligned}$$

---

[10] With respect to market efficiency, note that agents in the model are not making systematic mistakes, so in this sense there is no inefficiency. At the same time, the exchange rate does not instantly convey all dispersed information, so the market does not qualify as strong-form efficient.

where $b_1$ through $b_5$ are positive constants. The order flow coefficients in the price equation, $b_1$ and $b_2$, depend on $g$ (the public's aggregate risk-bearing capacity from equation 3), the variances $s^2_R$, $s^2_{\tilde{C}}$, $s^2_{\hat{C}}$, and the correlation $r$ between $\hat{C}^i_{1t}$ and $\Delta R_{t+1}$.[11]

## 2. Empirical Model for Disentangling Sources of Price Adjustment

The theoretical model above shows that macroeconomic information can affect prices through two channels: directly via $DR_t$ and $A_tB_t$ (as in traditional exchange rate models) and indirectly via order flow $\Delta x_t$. To examine the relative importance of these channels we consider the following empirical model:

$$\Delta p_t = f(\Delta x_t) + x_t + v_t, \tag{5}$$

$$\Delta x_t = e_t + h_t, \tag{6}$$

where $\Delta p_t$ is the change in the log spot rate (DM/\$) from the end of day t-1 to the end of day t, and $\Delta x_t$ is order flow during the corresponding period.[12]

The function $f(.)$ identifies the price impact of order flow. Prices and order flow are subject to four shocks representing different sources of information hitting the market. These shocks are mean zero, mutually uncorrelated, and serially uncorrelated. The $x_t$ and $v_t$ shocks represent information directly impounded into price. The first of these, $x_t$, represents news associated with macroeconomic announcements (the $A_tB_t$ term in equation 4). The shock $v_t$ represents other news directly impounded into prices that is unexplained by

---

[11] For intuition on equation (4), consider what it implies for the level of price at the end of day t, $P_t$. $P_t$ reflects the sum of (1) past interdealer order flows $\Delta x_t$, past announcement information $A_tB_t$, and (3) past payoff increments $\Delta R_t$. The sum of past order flows remains relevant to the level of the price because they represent changes in the effective stock of risky assets (the sum of public portfolio shifts $\Sigma^T_{t=1}C_{1t}$) that the public must be induced to re-absorb. Because the beginning-of-day portfolio shifts do not mean revert, these cumulative shifts have permanent effects on price.

[12] Notice that we replace the dependent variable with the change in the log spot rate, $\Delta p_t$. This substitution makes our empirical specification comparable to standard macro models. Estimates using $\Delta P_t$ produce nearly identical results to those we report ($R^2$s, coefficient significance, autocorrelation levels, etc.).

order flow or announcement flow (the $b_3DR_t$ term in equation 4). Order flow is driven by the $e_t$ and $h_t$ shocks. The $e_t$ shocks identify the order flow effects from macroeconomic announcements (the additional order flow term $\hat{C}^i_{1t}$ that appears in equation 4 when there is an announcement, $A_t$=1). Shocks to order flow unrelated to macro announcements (i.e., portfolio shifts that are unaccounted for) are represented by the $h_t$ shocks.

Our estimation strategy relies heavily on identification through conditional heteroskedasticity. The key conditioning variable is the number macroeconomic announcements made between the end of day t-1 and the end of day t, $N_t$. In particular, the variance of the $x_t$ and $e_t$ shocks increases with the number of announcements (consistent with our model):

$$Var(e_t) = s\,(N_t) \quad \text{and} \quad Var(x_t) = w(N_t) \tag{7}$$

where $s\,(0) = w(0) = 0$, with $s\,'(.) \geq$ and $w'(.) \geq 0$. This specification implies that neither $x_t$ or $e_t$ shocks affect prices on days when there are no macroeconomic announcements. The shocks $v_t$ and $h_t$ are independent of macroeconomic announcements, so their variances are unrelated to $N_t$. For simplicity, we assume these variances are constant:

$$Var(v_t) = s_v \quad \text{and} \quad Var(h_t) = s_h \;. \tag{8}$$

To estimate the model described in (5) – (8) we specify linear forms for the price-impact function, $f(.)$, and the variance functions, $s\,(.)$ and $w(.)$. (Following the presentation of results, we subject these specification assumptions to diagnostic testing.) The parameters of these functions along with $s_v$ and $s_h$ can then be estimated by the Generalized Method of Moments, with conditions on second moments playing a central role in identification (moment condition details appear in results tables below).

## 3. Data

The dataset contains time-stamped, tick-by-tick observations on actual transactions for the largest spot market – DM/$ – over a four-month period, May 1 to August 31, 1996 (full 24-hour trading day). These data are the same as those used by Evans (2001), and we refer readers to that paper for additional detail. The 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 would not 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. One drawback is that it is not possible to identify the size of individual transactions. For model estimation, order flow $\Delta x_t$ is therefore measured as the difference between the number of buyer-initiated trades and the number of seller-initiated trades.[14]

---

[13] At the time of our sample, interdealer transactions accounted 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 accounted for about 60 percent of interdealer trade and brokered trading accounted for about 40 percent. For more detail on the Reuters Dealing 2000-1 System see Lyons (2001) and Evans (2001).

[14] In direct interdealer trading, orders sizes are standardized, so variation in size is much smaller than variation in the size of individual customer-dealer trades. See Jones, Kaul, and Lipson (1994) for evidence that trade size in equity markets contains no information beyond that in the number of transactions. Our data set does include total dollar volume over our sample, which allows us to calculate an average trade size; we use this below to interpret the estimated coefficients.

The variables in our model are measured daily (at 4 pm GMT). The change in the spot rate (DM/$), $\Delta p_t$, is the log change in the purchase transaction price. When a purchase transaction does not occur precisely at 4 pm, we use the immediately preceding purchase price (with roughly 1 million transactions per day, the preceding transaction is generally within a few seconds of 4 pm). Order flow, $\Delta x_t$ is the difference between the *number* of buyer-initiated trades and the number of seller-initiated trades (in hundred thousands, negative sign denotes net dollar sales).

Our announcement data come from the Reuter's News service (source: Olsen Associates). We compute $N_t$ as the number of announcements relating to US or German news items between 4:01 pm GMT on day t-1 and 4 pm GMT on day t. Table 1 presents descriptive statistics for the three central variables in our analysis: daily announcement flow $N_t$, daily (log) price changes $Dp_t$, and daily interdealer order flows $Dx_t$. Note that the median number of announcements each day is 11, 8 of which are German announcements and 3 are U.S. announcements. These announcements include employment, GDP, trade balance, durable goods, PPI, retail sales, housing starts, leading indicators, initial jobless claims, factory orders, German M3 figures, and releases following the biweekly Bundesbank meeting (the first four of which were identified as more important for exchange rate volatility by Andersen and Bollerslev 1998).

Our four-month sample sharply constrains our ability to work with announcements on a disaggregated basis. One partition for which we do have some statistical power is U.S. versus German announcements. Results for this partitioning are presented in the next section. We did not have sufficient power to address finer partitions (e.g., separating announcements found to be important in previous work, separating announcements by day of week, and separating announcements by predictability of timing). In addition, we recognize that it is the surprise component of announcements that should affect the exchange rate. But the large set of announcements we consider here (many of which are unscheduled) make it unlikely that reliable expected values can be computed. This is a long-standing challenge for the literature. We have chosen to focus on announcement flow as the number of announcements, allowing the data to speak to whether there is news in the announcements (on average), and whether that news is common-knowledge news or dispersed information that affects price through order flow.

Because our GMM estimation is based on moment conditions that include variances, it is important that we measure variances precisely. In particular, measures of variance (e.g., the variance of $\nu_t$ in equation 5) for a given day based only on beginning and end-of-day values are quite noisy. Using measures of integrated variance (see, e.g., Andersen et al. 2001) increases the precision of daily estimates considerably. As is common, we use a five-minute sampling frequency to estimate daily (integrated) variances. This applies both to return variances and to order flow variances. (See also the table notes for further detail.)

## 4. Empirical Results

We turn to the empirical model in equations (5) and (6) for sorting out the direct versus indirect (via order flow) effects of announcements. We choose linear specifications for the variance functions, $s(N_t)$ and $w(N_t)$ in equation (7) (a specification choice that is tested below). The full empirical model takes the form:

$$\Delta p_t = a\,\Delta x_t + x_t + v_t \quad \text{with} \quad x_t \sim (0, wN_t), \quad v_t \sim (0, s_v), \tag{9}$$

$$\Delta x_t = e_t + h_t \quad \text{with} \quad e_t \sim (0, s\,N_t), \quad h_t \sim (0, s_h). \tag{10}$$

This model embeds five parameters, $a, w, s, s_v$ and $s_h$. The parameter $a$ captures the price impact of order flow (as in Evans and Lyons 1999). The parameter $w$ governs the direct effect of announcements on price: the greater the number of announcements $N_t$ the higher the variance of this component of daily returns. The parameter $s$ governs the direct effect of announcements on order flow: the greater the number of announcements $N_t$ the higher the variance of this component of daily order flow.[15] The variance $s_v$ reflects the component of daily returns unexplained by order flow or announcements. The variance $s_h$ reflects the component of daily order flow unexplained by announcements.

Table 2 reports GMM parameter estimates together with standard errors calculated from the asymptotic covariance matrix (allowing for heteroskedasticity).[16] The table reports estimates for specifications using the flow of all announcements lumped together (labeled All News Together) and also with the U.S. and German announcements introduced separately. In both specifications the estimate of the price-impact parameter $a$ is highly statistically significant. (Its size corresponds to a price impact of roughly 50 basis points per $1 billion in order flow, which is comparable to the estimate in Evans and Lyons 1999.) With all the announcements lumped together, both of the key parameters $w$ and $s$ are significant and correctly signed (positive), implying that both direct and indirect effects of announcements on price are present. When U.S. and German announcements are introduced separately, both of the $w$ estimates are significant and correctly signed (positive; $w_1$ is U.S. and $w_2$ is Germany). Also in this specification, the effect of German announcements on order flow volatility ($s_2$) is positive and significant. The effect of U.S. announcements is of similar size ($s_1$ slightly larger than $s_2$) but is not significant at the five percent level. (Recall that the median number of daily German announcements is nearly three times that for U.S. announcements.) In both specifications, the two (unconstrained) variance parameters $s_v$ and $s_h$ are quite significant. The next rows of the table report Wald statistics for various parameter restrictions. The restriction that the $w$ and $s$

[15] Though our model derivation for equation (9) does not call for order flow lags, whether they are relevant empirically is a legitimate robustness concern. The daily analysis in Evans and Lyons (1999) shows, however, that order flow lags are insignificant when included in this specification. Indeed, our order flow measure shows no persistence in daily data, nor do daily returns, so we are not omitting any variables that a non-structural VAR approach would identify as significant.

[16] Though not reported, we estimated the model allowing the price-impact parameter $a$ to vary linearly with the number of announcements: the coefficient on the incremental announcement effect was positive but not significant. This corresponds to a finding that $\beta_2$ in equation (4) is zero. We also tested whether announcement effects on the variance of order flow might be non-linear (by modeling the variance of $e_t$ as $s N_t^g$ ), but found no evidence of this either (not reported).

coefficient pairs are equal (in the specification that estimates them separately) cannot be rejected.

Important summary statistics are provided in the final rows of Table 3. The rows labeled Variance Ratios provide our point estimates of the share of announcement-driven variance in total variance (for both the order flow and price equations, respectively). Roughly one-third of order flow variance is attributable to announcement flow; roughly one-quarter of the *conditional* variance of price is attributable to announcement flow via the direct channel. The final three rows present the fraction of *total* variance in price that can be attributed on average to the direct and indirect effects of announcements. Here we see that the indirect channel (through order flow) is roughly twice the size of the direct effect.

To summarize, our empirical results provide clear answers to the questions posed by this paper. Order flow induced by macroeconomic announcements is an important source of price variation: about 20 percent of the daily movements in spot rates can be linked to macroeconomic announcements through this indirect order-flow channel. Through the direct (traditional) channel, macroeconomic announcements account for only about 10 percent of daily price movements. Finally, the third source of price variation in our model—order flow unrelated to macro announcements—accounts for another (roughly) 30 percent of price variation.

Specification Checks

We include two additional lines of analysis as specification checks on our baseline model (i.e., the Table 2 model). Both are quite straightforward. First, we provide scatter-plots of the raw data. The first plot in Figure 2 shows the number of announcements $N_t$ on the horizontal axis and the (integrated) variance of daily returns on the vertical axis. The second plot in Figure 2 shows the same number of announcements $N_t$ on the horizontal axis and the (integrated) variance of daily order flow on the vertical axis. For our purposes, the latter plot is the most central to our thesis: an increased flow of announcements clearly increases the variance of order flow. That the induced order flow has subsequent impact on price (per our model estimates in Table 2) is consistent with the first of these two plots.

Our second line of specification check uses regression analysis to examine the sources and form of heteroskedasticity. We consider both (1) the extent to which announcement flow can account for the conditional variance of returns and order flow (equation 7) and (2) whether the linear specifications of the variance functions $s(N_t)$ and $w(N_t)$ in our baseline model are warranted. To account for the conditional variance of returns, we perform a first-step regression to measure the conditional mean: we regress (OLS) exchange-rate returns $Dp_t$ on order flow $Dx_t$ (using a five-minute sampling frequency). We then test for announcement-related heteroskedasticity by regressing the daily variance of the first-step residuals (measured as the integrated variance over the five-minute observations) on a constant and the daily flow of announcements. To account for the conditional variance of order flow, we regress the integrated variance of order flow (also over five-minute observations) on a constant and the daily flow of announcements; it is not necessary to perform a first step regression in this case because estimates of the conditional mean of daily order flow are quite small and statistically insignificant.

Table 3 reports results from these two regressions. (The regressor $N_t$ denotes the total number of announcements over the day; the superscripts "us" and "g" denote US and Germany, respectively.) As the upper panel of the table shows, there is evidence of a direct announcement effect on the conditional variance of prices (i.e., after the conditional mean effects from order flow have been netted). The lower panel shows that announcement flow also affects the conditional variance of order flow. These results are consistent with the visual evidence in the scatter plots in Figure 2. Note too the column labeled "Non-linear." This column reports the p-values for a chi-squared LM test of whether squared announcement terms in these regressions can be excluded. None of the tests are significant, providing support for the linear variance functions of our baseline model. The tests for serial correlation in the residuals of these regressions (both first and fifth order) show some evidence of autocorrelation. There is also evidence of heteroskedasticity in the order flow equation. The reported t-statistics are corrected for both (White correction). In sum, these results support the linear specifications we adopted in our benchmark GMM model and corroborate that both direct and indirect price effects from announcements are present.

Causality

Though the direction of causality in our model runs strictly from order flow to price (as is true of microstructure theory generally), there is a popular alternative hypothesis that involves reverse causality, namely feedback trading. This alternative is addressed in some detail in Evans and Lyons (1999). Let us consider it here as well, albeit briefly.

We begin with some perspective. Most models of feedback trading are based on non-rational behavior of some kind, making them less appealing to many economists on a priori grounds. Models of feedback trading that do not rely on non-rational behavior generally require that returns be forecastable using the first lag of returns. But this is not a property of major floating exchange rates (whether daily or intradaily, based on transactions data). Accordingly, the class of feedback trading models that might be relevant here is the non-rational class.[17]

Existing empirical evidence on feedback trading in foreign exchange is scant. Valid instruments for identifying returns-chasing order flow have not been employed and it is not clear which variables would qualify. One piece of relevant evidence is provided by Killeen et al. (2001). Using daily data on foreign exchange order flow, they find that order flow Granger causes returns but returns do not Granger cause order flow. This evidence is purely statistical, however, and applies at the daily frequency, so its message (though suggestive) is not definitive.

Our analytical approach to causality in Evans and Lyons (1999) begins with the following question: Suppose intraday (i.e., contemporaneous in daily data) positive-feedback trading is present, under what conditions could it account for the key moments of our daily data? To address this question, we decompose measured order flow $\Delta x_t$ into two parts, an exogenous part from portfolio shifts (corresponding to the customer flow component $\tilde{C}_{1t}$ in the model of this paper)

[17] Whether the non-rational class is intellectually appealing is not an issue we could hope to resolve here. We simply offer the fact that immense amounts of money are at stake when dealing in foreign exchange at major banks (the source of our data). These banks take the evaluation of traders' performance and decision making very seriously.

and a contemporaneous part due to feedback trading (i.e., $\Delta x_t^{fb} = f\Delta p_t$). The sign of the parameter $f$ that most people have in mind for explaining our results is positive, i.e., positive feedback trading (based on the positive price-impact coefficient $a$ we find on contemporaneous order flow). The punch line of that bias analysis is that for reasonable parameter values, the feedback trading needed to produce the key moments of the data is actually negative. For positive feedback trading to account for the key moments, the direct effect of macroeconomic information would have to be one to two orders of magnitude more important than the indirect effect (through order flow). In our judgment this would be extreme indeed.

In Evans and Lyons (2001b) we take up causality again, in this case using hourly data, which allows the feedback trading parameter $f$ to be estimated from the data. In broad terms, the feedback trading alternative predicts that (1) the estimated coefficient $f$ will be positive and significant and (2) the coefficient on the first part of order flow (the truly exogenous part) will be smaller than the $a$ we estimate, if not zero. In the end, estimates of the feedback parameter, $f$, are negative and statistically insignificant. Thus, insofar as there is any empirical evidence of feedback trading in these data, it points to the presence of negative rather than positive feedback trading. Moreover, estimates of the coefficient on the exogenous part of order flow in the return equation are larger (slightly) than the $a$ estimates from the null without feedback trading, and they remain highly statistically significant, in contrast to what the feedback trading alternative predicts.

## 5. Conclusions

This paper addresses whether the underlying cause of order flow is traditional macroeconomic information, a hypothesis that contrasts with that in the model of Evans and Lyons (1999). We do so by examining the links between order flow and macro announcements. Our model distinguishes three sources of exchange-rate variation. The first source mirrors traditional models—public announcement information that is impounded in price immediately and directly (i.e., with no role for order flow). The second source is an indirect effect of public announcement information that operates via induced order flow. The third source of exchange rate variation mirrors Evans and Lyons (1999), namely order flow that is not related to public announcement information. Our paper is the first to attempt to disentangle these sources.

We find that all three sources of price variation are significant. The indirect effect of announcements (via order flow) accounts for twice as much price variation as the direct effect of announcements (not involving order flow). Thus, even when one would expect order flow's role to be muted, flow is still the key driver of price variation. Together, these two announcement effects account for about 30 percent of price variance. The third source—order flow unrelated to announcements—accounts for another (roughly) 30 percent of price variance.

Though the literature linking exchange rates and announcements is well developed, to our knowledge this paper is the first to make use of order flow to sort out the relationship. Our use of order flow also allows us to draw inferences in areas where past empirical work has lacked power, for example, in providing evidence that market participants are using different exchange-rate models. Order flow appears to convey information about differing individual assessments of announcements' relevance.

Why, one might ask, does this paper find that roughly 30 percent of volatility comes from direct and indirect effects of announcements when past work found that announcement effects are impounded rapidly in price and account for less than 10 percent of total volatility? A possible answer consistent with our results is that the order flow (from non-dealer participants) induced by announcements takes time (hours). In this setting, short event windows (as are used in many papers) cannot capture the total impact. Non-event-study approaches may attribute these knock-on flow effects to other factors (e.g., time-of-day effects). At this stage, though, these are but conjectures.

Future work might examine the use of order flow as an instrument for determining the directional information in announcements (i.e., good or bad news, as opposed to simply news, which has been our focus here). Past work on how announcements affect volatility suggest that announcements do indeed convey important information. However, work on how announcements affect the *direction* of exchange-rate changes has been less successful. There may be little precision in identifying the direction of news in an announcement based on, say, consensus forecasts. In contrast, if investors are selling after an announcement, this provides a backed-by-money indication that their assessment of the news is negative.

Future work might also analyze price responses in other asset markets (e.g., equity and bond markets) using our order flow approach. Though much work has been done in these markets that links announcements to volumes and volatility, none (to our knowledge) uses order flow data to sort out differential interpretation of public news. This is true despite the fact that theory on differential interpretation of public news is rather well developed (see, e.g., Harris and Raviv 1993 and Kandel and Pearson 1995, among many others). Theoretical work in this area does not in general specify enough institutional richness that implications for order flow can be derived, hence the empirical disconnect.[18]

Finally, if our "anticipation conjecture" in the introduction is correct (i.e., that order flow is how changes in expected future fundamentals $f^e_{t+1}$ are conveyed to the market), then order flow should provide better forecasts of subsequent macro variables than past macro variables. Moreover, if this hypothesis is to explain the negative results of Meese and Rogoff (1983a), then the macro variables being forecast need to be largely beyond the one-year horizon. This is testable.[19]

[18] That this disconnect arises is another reminder that order flow and demand are not one to one. To clarify, consider a simple counter-example. In the common-knowledge tradition of exchange rate economics, when positive public news arrives, demand increases, causing price to increase—without any order flow occurring or needing to occur. This is incompatible with demand and order flow being one to one: the demand shift occurs without the occurrence of order flow.

[19] Consider the following (flawed) logic: "Suppose a regression of exchange rate changes on order flow produced an $R^2$ statistic of 1. We know that a regression of exchange rate changes on macro fundamentals is basically 0. In that case, we would know that order flow does not reflect macro." Missing from this view is the possibility that "macro" variables, as specified by traditional empirical models, provide poor measures of changing macro expectations (the sine qua non of pricing in forward looking asset markets).

**Table 1: Data Distributions**

| | | | | | Announcements. | | |
|---|---|---|---|---|---|---|---|
| | $\Delta p_t$ | $\Delta x_t$ | $\lvert\Delta p_t\rvert$ | $\lvert\Delta x_t\rvert$ | US: $N_t^{us}$ | German: $N_t^{g}$ | Total: $N_t$ |
| % | | | | | | | |
| 0.05 | -7.78 | -308 | 0.07 | 5 | 0 | 3 | 5 |
| 0.25 | -2.48 | -61 | 0.92 | 28 | 1 | 6 | 9 |
| 0.50 | 0.20 | 8 | 2.43 | 83 | 2 | 8 | 11 |
| 0.75 | 2.22 | 91 | 3.91 | 140 | 5 | 12 | 16 |
| 0.95 | 4.55 | 186 | 7.78 | 319 | 7 | 19 | 22 |

Notes: $\Delta p_t$ is 1000 times the change in log price between 4:00 pm on day *t* and day t-1. $\Delta x_t$ is total interdealer order flow over the same time interval. Sample: daily observations from May 1 to August 31, 1996.

## Table 2: GMM Estimates

$$\Delta p_t = a\Delta x_t + x_t + v_t, \qquad Var(x_t) = wN_t, \qquad Var(v_t) = s_v$$

$$\Delta x_t = e_t + h_t, \qquad Var(e_t) = s\ N_t, \qquad Var(h_t) = s_h$$

| | **All News Together** | | **US and German News Separately** | |
|---|---|---|---|---|
| | Est. | t-stat. | Est. | t-stat. |
| Parameters | | | | |
| $a$ | 0.02 | (12.22) | 0.02 | (12.22) |
| $s_v$ | 42.00 | (7.64) | 41.55 | (8.00) |
| $s_h$ | 3.63 | (5.07) | 3.62 | (5.18) |
| $w_1$ | 1.22 | (3.84) | 2.07 | (2.12) |
| $w_2$ | | | 0.98 | (2.03) |
| $s_1$ | 0.15 | (3.08) | 0.17 | (1.41) |
| $s_2$ | | | 0.15 | (2.32) |
| Wald Tests: | | | | |
| $w_1 = w_2 = 0$ | | | 29.48 | (<0.01) |
| $s_1 = s_2 = 0$ | | | 9.87 | (<0.01) |
| $w_1 = w_2, s_1 = s_2$ | | | 1.04 | (0.60) |
| Variance Ratios | | | | |
| $Var(e_t)/Var(e_t + h_t)$ | 0.328 | | 0.329 | |
| $Var(x_t)/Var(x_t + v_t)$ | 0.256 | | 0.263 | |
| $R^2_{\Delta p}(direct)$ | 0.097 | | 0.100 | |
| $R^2_{\Delta p}(indirect)$ | 0.203 | | 0.204 | |
| $R^2_{\Delta p}(total)$ | 0.300 | | 0.304 | |

Parameters are estimated with GMM (asymptotic t-statistics in parentheses). Sample: daily observations from May 1 to August 31, 1996. Wald statistics are for the null hypothesis listed (p-values in parentheses). $Var(e_t)/Var(e_t + h_t)$ is the mean fraction of the (integrated) variance of order flow due to announcements, and $Var(x_t)/Var(x_t + v_t)$ is the mean fraction of the (integrated) variance of return residuals due to announcements. (We compute the integrated variance of return residuals, i.e., $Var(x_t + v_t)$, using the estimate of α from daily data.) For all variances, integrated daily variance is calculated using a 5-minute sampling frequency. $R^2_{\Delta p}(direct)$, and $R^2_{\Delta p}(indirect)$ are the mean fractions of the daily variance in prices attributed to announcement via the direct and indirect channels. The moments used in estimation are:

$E[(x_t + v_t)\Delta x_t] = 0,$

$E[Var(v_t) + Var(x_t) - s_v - wN_t] = 0,$ $\qquad E[(Var(v_t) + Var(x_t) - s_v - wN_t) \otimes N_t] = 0,$

$E[Var(h_t) + Var(e_t) - s_h - s\ N_t] = 0,$ $\qquad E[(Var(h_t) + Var(e_t) - s_h - s\ N_t) \otimes N_t] = 0.$

## Table 3: Variance Regressions

| Equation | Const. | $N_t$ | $N_t^{us}$ | $N_t^g$ | $R^2$ | SEE | Non-linear | Serial | Hetero |
|---|---|---|---|---|---|---|---|---|---|
| | **Regressors** | | | | **Diagnostics** | | | | |
| Price change | | | | | | | | | |
| | 34.76 | 1.69 | | | 0.14 | 24.81 | 0.47 | 0.04 | 0.53 |
| | (5.57) | (4.56) | | | | | | 0.46 | 0.98 |
| | 34.35 | | 2.65 | 1.41 | 0.15 | 24.83 | 0.72 | 0.05 | 0.57 |
| | (5.77) | | (2.26) | (2.61) | | | | 0.50 | 0.99 |
| Order flow | | | | | | | | | |
| | 3.12 | 0.18 | | | 0.11 | 3.06 | 0.41 | <0.01 | <0.01 |
| | (4.13) | (3.64) | | | | | | 0.07 | 0.04 |
| | 3.11 | | 0.21 | 0.18 | 0.11 | 3.08 | 0.77 | <0.01 | <0.01 |
| | (4.25) | | (1.52) | (2.53) | | | | 0.06 | 0.04 |

The table reports OLS regression coefficients (t-statistics in parenthesis). Sample: daily observations from May 1 to August 31, 1996. Regressions take the form $Var(u_t) = bz_t + w_t$ where $Var(u_t)$ is the (integrated) variance of the estimated residuals from the price change equation $\Delta p_t = a\Delta x_t + u_t$, or the (integrated) variance from the residuals in the order flow equation $\Delta x_t = u_t$. (In both cases, integrated variance is calculated using a 5-minute sampling frequency.) $z_t$ is a vector of regressors including; a constant, the number of US announcements $N_t^{us}$, German announcements, $N_t^g$, and total announcements $N_t$, all on day $t$. The Non-linear column presents the p-value of a chi-squared LM test for exclusion of the squared announcement term (or terms, in the case where country announcements are separated). The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and fifth-order (bottom row) serial correlation in the residuals. The Hetero column presents the p-value of a chi-squared LM test for first-order (top row) and fifth–order (bottom row) ARCH in the residuals.

**Figure 2**

**Number of Announcements (horizontal axis)**
**and Daily Variance (vertical axis)**

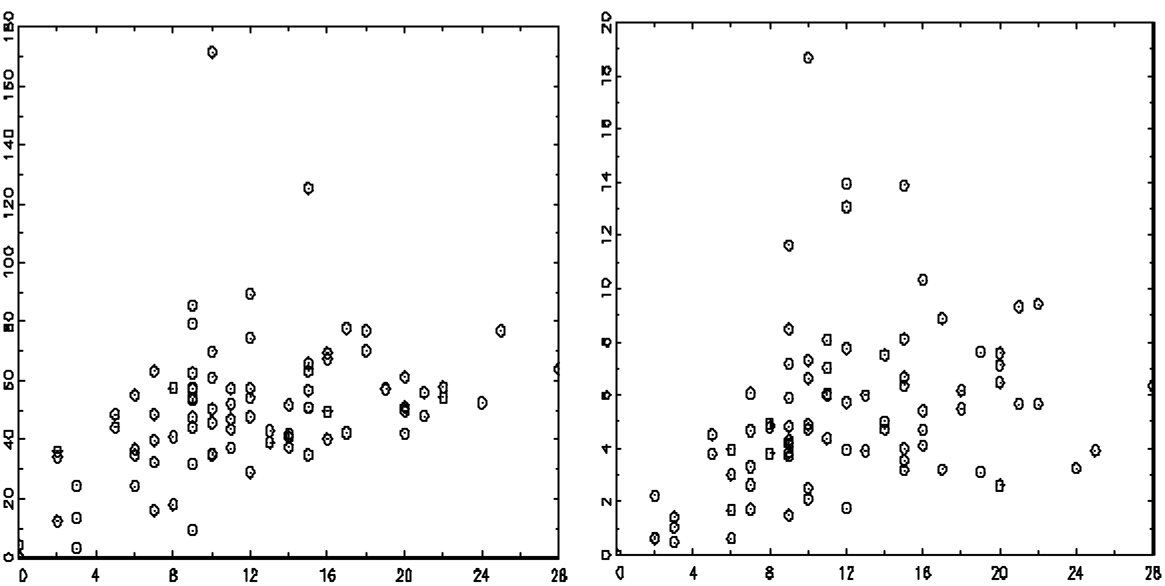


$Var(\Delta p_t)$ on vertical axis $\qquad$ $Var(\Delta x_t)$ on vertical axis

The daily variances of the log price change and order flow (vertical axes) are integrated variances, calculated using a five-minute sampling frequency. Sample: daily observations from May 1 to August 31, 1996.

## Appendix: The Dealer's Problem

As noted in the text, solving the dealer's problem in the model of this paper corresponds closely to solutions of related models (e.g., in Evans and Lyons 1999 and Evans and Lyons 2001a), so we provide only an overview of that problem here (details available on request). Within a given day t, let $W_j^i$ denote the end-of-round j wealth of dealer i, using the convention that $W_0^i$ denotes wealth at the end of day t-1. (We suppress notation to reflect the day t where clarity permits.) With this notation, and normalizing the gross return on the riskless asset to one, we can define the dealers' problem over the four choice variables described in section 1, namely, the three scalar quotes $P_j^i$, one for each round j, and the outgoing interdealer trade $T_2^i$:[20]

$$\max_{\{P_1^i,\, P_2^i,\, P_3^i,\, T_2^i\}} \quad E\left[-\exp(-qW_3^i \mid \Omega^i)\right] \tag{A1}$$

s.t.

$$W_3^i = W_0^i + C_1^i(P_1^i - \overline{P}_2^i) + \overline{T}_2^i\left(P_2^i - P_3^i\right) + \left(T_2^i - C_1^i\right)\left(P_3^i - \overline{P}_2^i\right)$$

Dealer i's wealth over the three-round trading day is affected by positions taken two ways: incoming random orders and outgoing (deliberate) orders. The incoming random orders include the public order $C_1^i$ and the incoming interdealer order $\overline{T}_2^i$ (tilde distinguishes incoming *interdealer* orders and prices from outgoing). The outgoing order is the interdealer trade $T_2^i$. $\overline{P}_j^i$ denotes an incoming interdealer quote received by dealer i in round j. As an example, the second term in the budget constraint reflects the position from the public order $C_1^i$ received in round one at dealer i's own quote $P_1^i$ and subsequently unwound at

[20] This daily problem is a valid sub-problem of the dealers' infinite horizon problem because dealers carry no overnight positions in the model and there are no shocks to the investment opportunity set over time (i.e., no intertemporal hedging demands). The former feature of the model comes from our assumption that the mass of customers is large in a convergence sense relative to the number of dealers, so that dealers have a comparative disadvantage in holding overnight positions.

the incoming interdealer quote $\bar{P}_2^i$ in round-two. (Recall that the sign of dealer i's position is opposite that of $C_1^i$, so a falling price is good for dealer i if the public order $C_1^i$ is a buy, i.e., positive. The dealer's speculative positioning based on information in $C_1^i$ is reflected in the final term of the budget constraint.) Term three reflecting the incoming (random) dealer orders is analogous to term two.

Term four of the budget constraint reflects the dealer's speculative and hedging demands. The outgoing interdealer trade in round 2 has three components:

$$T_2^i = C_1^i + D_2^i + E[\bar{T}_2^i \mid \Omega_{2T}^i] \tag{A2}$$

where $D_2^i$ is dealer i's speculative demand in round 2, and $E[\bar{T}_2^i \mid \Omega_{2T}^i]$ is the dealer's hedge against incoming orders from other dealers (this term is zero in equilibrium given the distribution of the $C_1^i$'s). The dealer's total demand (speculative plus hedging) can be written as follows:

$$D_2^i + E[\bar{T}_2^i \mid \Omega_{2T}^i] = T_2^i - C_1^i$$

which corresponds to the position in term four of the budget constraint.

As shown in appendixes of the papers noted above, constant absolute risk aversion and conditionally Normally distributed returns produces a second-round interdealer trade $T_2^i$ that is proportional to the round-one customer order received by dealer i, $C_1^i$. This in turn implies (from equation 2) that public observation of interdealer order flow *Dx* is a sufficient statistic for the aggregate customer order flow in round one, $C_1$. Given that $C_1$ is a sum of Normally distributed random variables, inference based on this variable takes the standard linear form, producing the parameters $b_1$ and $b_2$ that appear in equation (4). That the order flow process in equation (4) is conditional on the occurrence of an announcement ($A_t$) then follows directly from our specification of round-one customer order flow as depending on the occurrence of an announcement.

## References


Andersen, T., and T. Bollerslev (1998), Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance*, 53: 219-265.

Andersen, T., T. Bollerslev, F. Diebold, C. Vega (2001), Micro effects of macro announcements: Real-time price discovery in foreign exchange, typescript, Northwestern University, September.

Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2001), The distribution of realized exchange rate volatility, *Journal of the American Statistical Association*, 96: .

Balduzzi, P., E. Elton, and C. Green (2001), Economic news and bond prices, *Journal of Financial and Quantitative Analysis*, forthcoming.

Chinn, M. (1991), Some linear and non-linear thoughts on exchange rates, *Journal of International Money and Finance*, 10: 214-230.

Cornell, B. (1982), Money supply announcements, interest rates, and foreign exchange, *Journal of International Money and Finance*, 1: 201-208.

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.

Ederington, L., and J. Lee (1995), The short-run dynamics of price adjustment to new information, *Journal of Financial and Quantitative Analysis*, 30: 117-134.

Engel, C., and J. Frankel (1984), Why interest rates react to money announcements: An answer from the foreign exchange market, *Journal of Monetary Economics*, 13: 31-39.

Evans, M. (2001), FX trading and exchange rate dynamics, NBER Working Paper 8116, February, *Journal of Finance*, forthcoming.

Evans, M., and R. Lyons (1999), Order flow and exchange-rate dynamics, *Journal of Political Economy*, forthcoming (available at www.haas.berkeley.edu/~lyons).

Evans, M., and R. Lyons (2001a), Portfolio balance, price impact, and secret intervention, NBER Working Paper 8356, July.

Evans, M., and R. Lyons (2001b), Time-varying liquidity in foreign exchange, typescript, U.C. Berkeley, October.

Fleming, M., and E. Remolona (1999), Price formation and liquidity in the U.S. treasury market, *Journal of Finance*, 54: 1901-1915.

Flood, R., and A. Rose (1995), Fixing exchange rates: A virtual quest for fundamentals, *Journal of Monetary Economics* 36, 3-37.

Frankel, J., and K. Froot (1990), Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review*, 80: 181-185.

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.

Goodhart, C., S. Hall, S. Henry, and B. Pesaran (1993), News effects in a high-frequency model of the sterling-dollar exchange rate, *Journal of Applied Econometrics*, 8: 1-13.

Grossman, S., and J. Stiglitz (1980), On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.

Hakkio C., and D. Pearce (1985), The reaction of exchange rates to economic news, *Economic Inquiry*, 23: 621-635.

Hardouvelis, G. (1988), Economic news, exchange rates, and interest rates, *Journal of International Money and Finance*, 7: 23-25.

Harris, M., and A. Raviv (1993), Differences of opinion make a horse race, *Review of Financial Studies*, 6: 473-506.

Hasbrouck, J., and D. Seppi (2001), Common factors in prices, order flows, and liquidity, *Journal of Financial Economics*, 59: 383-411.

Isard, P. (1995), *Exchange Rate Economics*, Cambridge University Press: Cambridge, UK.

Ito, T., and V. Roley (1987), News from the U.S. and Japan: Which moves the yen/dollar exchange rate? *Journal of Monetary Economics*, 19: 255-277.

Jones, C., G. Kaul, and M. Lipson (1994), Transactions, volume, and volatility, *Review of Financial Studies*, 7: 631-651.

Kandel, E., and N. Pearson (1995), Differential interpretation of public signals and trade in speculative markets, Journal of Political Economy, 103: 831-872.

Killeen, W., R. Lyons, and M. Moore (2001), Fixed versus flexible: Lessons from EMS order flow, NBER Working Paper 8491, September.

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

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

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.