

Time-Varying Liquidity in Foreign Exchange

Martin D. D. Evans*

Richard K. Lyons

26 October 2001

Abstract

This paper addresses whether currency trades have greater price impact during periods of rapid public information flow. Central bankers often suggest that expectations are at times “ripe” for coordinated adjustment, and that periods of rapid information flow are such a time. We develop an optimizing model to account for the joint behavior of order flow and returns around announcements. Using transaction data made available by electronic trading, we estimate the price impact of trades in the DM/\$ market precisely. We then test whether trades during periods with macroeconomic announcements have higher price impact. They do. We also test for dependence of liquidity on trading volume and return volatility (two other prominent state variables in the literature on liquidity variation). We do not find any evidence that liquidity depends on these variables. The findings provide policy-makers with guidance for the timing and magnitude intervention.

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 National Science Foundation for financial assistance and seminar participants at UC Berkeley for comments.

Time-Varying Liquidity in Foreign Exchange

This paper addresses time-varying liquidity in the foreign exchange market and its relevance for policy. Liquidity is defined by the price impact of trades: trades have more price impact when markets are less liquid, other things equal. (This usage is standard in theoretical work on trading; see, e.g., Kyle 1985.) Our particular interest is whether currency trades have more price impact during periods of rapid public information flow. This is important for policy: if the answer is yes, it provides a means of increasing the efficacy of intervention (e.g., by selectively timing intervention trades). Central bankers have long suggested that market expectations are at times “ripe” for coordinated adjustment, particularly when new public information arrives. Implicit in this view is the idea that agents draw different conclusions from common macroeconomic data, making transactions at these times especially valuable for measuring expectations.

We estimate the price impact of trades using transaction data recently made available by electronic trading. These data allow precise tracking of how the market absorbs actual trades and any information conveyed by them. Though the data represent private trades rather than central bank trades, they are relevant for measuring the price impact of central bank trades as long as interventions are sterilized and convey no monetary policy signal. These conditions insure that central bank trades are uncorrelated with current and future monetary policy, which is also a property of private trades (so long as exchange rates are floating). Our analysis is relevant to the price impact of these trades per se, not to the price impact of accompanying intervention announcements. Incremental effects from intervention announcements are best analyzed using more traditional methods and data (see Dominguez and Frankel 1993).¹

We begin by developing a trading model for structural guidance on the

¹ Though intervention in practice is often coupled with an intervention announcement, this is not always the case; i.e., stealth intervention is definitely a policy option.

joint behavior of order flow and returns.² The model is sufficiently rich to allow estimation using the particular type of order flow data available (data on dealer-to-dealer trades). Two features of the model stand out in terms of testable implications. First, it includes macroeconomic announcements whose implications for exchange rate fundamentals are partially gleaned from order flow. This allows order flow to mediate part of these announcements' price impact and provides a structural understanding for why liquidity is reduced. Second, the model produces so-called "hot potato" trading, i.e., the passing of positions from dealer to dealer for risk management purposes (Flood 1994, Lyons 1997). Hot potato trading imposes testable restrictions on the joint behavior of order flow and returns that we examine in the data.

Our estimation strategy has two stages. First, we estimate a linear, constant-coefficient model with two equations, one for returns and one for order flow. That the system includes an equation for order flow (signed trades) is important: macro announcements may affect the trades process itself (versus affecting the size of a given trade's price impact). Though the constant-coefficient model does not allow for state-dependent liquidity, it does allow testing of several of our trading model's predictions. For example, it allows testing of predictions about which variables should drive both returns and order flow (including the relevance of various lags).

The second stage of our estimation strategy examines nonparametric (kernel) regressions of our two trading model equations. Kernel estimation allows us to address directly the state dependence of liquidity. (Kernel estimation is feasible here due to the large quantity of available transaction data; one could not effect similar analysis with actual intervention trades due to much smaller samples.) The state variable that is our primary focus is the flow of public macroeconomic information, measured here from the flow of macro announce-

² Order flow is a term from microstructure finance that refers to the net of buyer- and seller-initiated trades. (It is therefore not synonymous with trading volume.) In markets organized like foreign exchange (i.e., dealer markets), the initiating side of the trade is the side that decides whether to trade at the dealer's posted quote. (This measure has no counterpart in rational expectations models of trading because in that setting all trades are symmetric—there is no "initiator.") As an empirical matter, order flow plays a central role in aggregating dispersed information in many different securities markets (e.g., equities, bonds, and foreign exchange; see Lyons 2001).

ments. The kernel results (using hourly data from the largest spot market, DM/\$) show a clear positive impact of announcement flow on market liquidity (i.e., the elasticity of price with respect to order flow). Thus, we find that order flow does indeed mediate the price impact of announcements, in keeping with our trading model. We also test for dependence of liquidity on trading volume and return volatility (two other prominent state variables in the literature on liquidity variation). We do not find any evidence that liquidity depends on these variables. Finally, we find evidence of hot potato trading. Specifically, there is momentum in flows (positive flows follow positive flows) and price impact is limited to initial flow innovations.

Our approach to analyzing the price impact of unannounced intervention illuminates a promising direction for future research: central banks with precise knowledge of their own trades—e.g., time of day, method of execution (e.g., brokers versus dealers), stealth level, etc.—can estimate the influence of these various parameter settings. With the types of data now available, central banks can learn exactly how trading is affected, including the induced flow of orders on each side of the market and the process by which price adjusts. It is something like a doctor who determines how a digestive system is functioning by having a patient ingest blue dye—the whole process becomes transparent. Such is the future of empirical work on this topic.

The remainder of the paper is in five sections. The next section develops a trading model for understanding the joint behavior of dealer-to-dealer order flow and exchange rate returns. Section 2 describes the dealer-to-dealer trading data. Section 3 presents our empirical results for both the constant-coefficient model and the kernel regression model. Section 4 revisits the direction of causality in the model. Section 5 concludes.

1. Model

Our model has two main features that distinguish it from earlier empirical models of currency order flow (e.g. Evans and Lyons 1999). First, the model includes macroeconomic announcements from which implications for the exchange rate are not inferable from the announcement alone. This allows order flow to convey incremental information about agents' changing expectations.

Second, the model allows for hot potato trading. If hot potato trading is present in the data, failing to account for it at the modeling stage will result in mischaracterizing the joint behavior of order flow and returns.

Before introducing specifics, let us provide a brief sketch of the model. At the beginning of each day, the public (i.e., non-dealers) place orders in the foreign exchange market. These orders are stochastic and are not publicly observed. Initially, dealers take the other side of these trades—shifting their portfolios accordingly. To compensate the (risk-averse) dealers for the risk they bear, an intraday risk premium arises, producing some intraday mean reversion in price. At the end of each trading day dealers unload intraday positions back onto the public, whose greater risk bearing capacity (in aggregate) gives them comparative advantage in holding overnight positions. Because the public's (non-stochastic) demand at the end of the day is not perfectly elastic, the public's orders at the beginning of the day have portfolio-balance effects that persist beyond the day. (Think of the public's orders at the beginning of the day as representing, say, shifts in hedging demand, whereas public demand at the end of the day is purely speculative, being driven by sufficient expected return per unit risk borne.) The portfolio balance effects arise because even when these risky positions are shared marketwide, their price impact is not diversified away.³

Specifics

Consider an infinitely lived, pure-exchange economy with two assets, one riskless and one risky, the latter representing foreign exchange.⁴ Each day, foreign exchange earns a payoff R , publicly observed, which is composed of a series of random increments:

$$(1) \quad R_t = \sum_{i=1}^t \Delta R_i$$

The increments ΔR are i.i.d. $\text{Normal}(0, \mathbf{s}_R^2)$. We interpret the increments as

³ Note that the size of the order flows the DM/\$ spot market needs to absorb are on average more than 10,000 times those absorbed in a representative U.S. stock (e.g., the average daily volume on individual NYSE stocks in 1998 was about \$9 million, whereas the average daily volume in DM/\$ spot was about \$300 billion).

⁴ Some of the model's basic structure is shared with the model in Evans and Lyons (1999), so we omit some details here where overlap is strongest.

interest rate changes.

The foreign exchange market has two participant types, customers and dealers. There is a continuum of customers, indexed by $z \in [0,1]$, each customer having constant absolute risk aversion (CARA) and maximizing utility of the following form:

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

where E_t is the expectations operator conditional on the customers' information at time t , and c_{t+s} is consumption in period $t+s$. We assume that all customers have the same time discount factor δ and risk aversion parameter θ . There are N dealers, indexed by i . Dealers also have CARA utility. The problem dealers solve is described below, following specification of the trading environment.

Within each day t there are four rounds of trading:


Round 1: Dealers trade with customers (the public).

Round 2: Dealers trade among themselves (to share inventory risk).

Round 3: R_t is realized and dealers trade among themselves a second time.

Round 4: Dealers trade again with the public (to share risk more broadly).

The timing of events within each day is summarized below, along with some notation.



Trading Round 1

At the beginning of each day, nature chooses whether to produce a macro announcement. We denote this announcement with the indicator variable A_t , which equals 1 if there is an announcement, and 0 if not. To some customers, the announcement conveys information about the payoff increment to be realized later that day, ΔR_t .⁵ This is important to the model: it implies that the order flow induced by observation of this announcement will be informative of future payoffs (and dealers will factor this into their processing of the order flow signals).

The next event each day is dealer quoting. Each dealer simultaneously and independently quotes a scalar price to the public.⁶ We denote this round-1 price of dealer i as P_1^i . (We suppress notation for day t ; as we shall see, it is the within-day rounds—the subscripts—that capture the model's economics.) This price is conditioned on all information available to dealer i .

The final event each day in round 1 is customer (i.e., public) trading. Each of the N dealers receives a customer order C_1^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 of these components are executed at his quoted price P_1^i . (Let $C_1^i < 0$ denote net customer selling—dealer i buying.) The second component is non-zero only if a macro announcement has occurred (indicator variable $A_t=1$):

$$C_1^i = \tilde{C}_1^i + \hat{C}_1^i \quad \text{with } \hat{C}_1^i = 0 \text{ if } A_t = 0$$

The first component \tilde{C}_1^i is distributed $\text{Normal}(0, \mathbf{s}_{\tilde{C}}^2)$. The first component is also

⁵ It is important that dealers learn about this information from order flow, rather than from the announcement itself. The 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 will involve costs, 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).

⁶ 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. The simultaneous-move nature of the model is in the spirit of simultaneous-move games more generally (versus sequential-move games).

uncorrelated across dealers and uncorrelated with the payoff increment ΔR at all leads and lags. (This first component reflects “portfolio shifts” of the non-dealer public; though we leave their precise source unspecified, they could come from changing hedging demands, changing transactional demands, or changing risk preferences—a la Evans and Lyons 1999.) On announcement days, the second component of the customer order \hat{C}_1^i is distributed $\text{Normal}(0, \mathbf{s}_c^2)$. It is positively correlated with the payoff increment to be realized later that day, DR_t , but is uncorrelated with the first component of customer order flow \tilde{C}_1^i (across all dealers). For the analysis below, it is useful to define the aggregate customer demand in round 1 as:

$$C_1 = \sum_{i=1}^N C_1^i$$

Trading Round 2

Round 2 is the first of two interdealer trading rounds. Each dealer simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell (any amount), denoted P_2^i . These interdealer quotes are observable and available to all dealers. Dealers then simultaneously and independently trade on other dealers’ quotes. If more than one dealer posts a quote at which a dealer wants trade, the dealer’s desired trade is allocated to the dealer whose index i is the nearest preceding index (letting index 1 wrap to index N).⁷

Let T_2^i denote the net interdealer trade initiated by dealer i in round two. At the close of round 2, all agents observe a noisy signal of interdealer order flow from that period:

$$(2) \quad \Delta x_2 = \sum_{i=1}^N T_2^i + \mathbf{n}$$

where \mathbf{n} is distributed $\text{Normal}(0, \mathbf{s}_n^2)$, independently across days. The model’s difference in transparency across trade types corresponds well to institutional

⁷ This trade allocation rule is important for generating hot potato trading. The hot potato result does not require that a dealer’s trade be routed to only one other dealer at the given price; it requires only that the dealer’s trade is not evenly split across all other dealers at the given price.

reality: customer-dealer trades in major foreign-exchange markets (round 1) are not generally observable, whereas interdealer trades do generate signals of order flow than can be observed publicly.⁸

Trading Round 3

Round 3 is the second of the two interdealer trading rounds. At the outset of round 3 the payoff increment ΔR_t is realized and the daily payoff R_t is paid (both observable publicly). Like in round 2, each dealer then simultaneously and independently quotes a scalar price to other dealers at which he agrees to buy and sell (any amount), denoted P_3^i . These interdealer quotes are observable and available to all dealers in the market. Each dealer then simultaneously and independently trades on other dealers' quotes. If more than one dealer posts a quote at which a dealer wants trade, the allocation rule is the same as that for the first interdealer trading round.

Let T_3^i denote the net interdealer trade initiated by dealer i in round 3. At the close of round 3, all agents observe interdealer order flow from that period:

$$(3) \quad \Delta x_3 = \sum_{i=1}^N T_3^i$$

This specification with noiseless observation of round-three order flow captures dealer learning and the increasing precision of their order-flow beliefs.

Trading Round 4

In round 4, dealers share overnight risk with the non-dealer public. Unlike round 1, the public's trading in round 4 is non-stochastic. Initially, each dealer simultaneously and independently quotes a scalar price P_4^i at which he agrees to buy and sell any amount. These quotes are observable and available to the public.

The mass of customers on the interval $[0,1]$ is large (in a convergence sense) relative to the N dealers. This implies that the dealers' capacity for

⁸ The screens of interdealer brokers (such as EBS) are an important source of these interdealer order-

bearing overnight risk is small relative to the public's capacity. With this assumption, dealers set prices optimally such that the public willingly absorbs dealer positions, and each dealer ends the day with no position (which is common practice among actual spot foreign-exchange dealers). These round-4 prices are conditioned on the interdealer order flow Δx_3 , described in equation (3). We shall see that this interdealer order flow informs dealers of the size of the total position that the public needs to absorb (to bring the dealers back to a position of zero).

To determine the round-4 price—the price at which the public willingly absorbs the dealers' aggregate position—dealers need to know the risk-bearing capacity of the public. We assume it is finite. Specifically, given that customers' intertemporal utility is CARA (coupled with daily returns being i.i.d.), the public's total demand for foreign exchange in round-4 of day t , denoted C_4 , is proportional to the expected return on foreign exchange conditional on public information:

$$(4) \quad C_4 = g(E[P_{4,t+1} + R_{t+1} \mid \Omega_{4,t}] - P_{4,t})$$

where the positive coefficient γ captures the aggregate risk-bearing capacity of the public ($\gamma = \infty$ is infinitely elastic demand), and $\Omega_{4,t}$ includes all public information available for trading in round 4 of day t .

The Dealer's Problem

The dealer's problem is defined over six choice variables, the four scalar quotes P_1^i , P_2^i , P_3^i , and P_4^i , and the two dealer's interdealer trades T_2^i and T_3^i . The appendix provides the full specification of the dealer's problem and model's solution. Here we provide some intuition.

Consider the four quotes P_1^i , P_2^i , P_3^i , and P_4^i . No arbitrage ensures that at any given time all dealers quote a common price: quotes are executable by

flow signals.

multiple counterparties, so any difference across dealers would provide an arbitrage opportunity. Hereafter, we write P_1 , P_2 , P_3 , and P_4 in lieu of P_1^i , P_2^i , P_3^i , and P_4^i . It must also be the case that if all dealers quote a common price, then that price must be conditioned on common information only. Common information arises at the end of round 2 (interdealer order flow Δx_2), at the beginning of round 3 (payoff increment DR), and at the end of round 3 (interdealer order flow Δx_3). The price for round-4 trading, P_4 , reflects the information in all three of these sources. (Recall that the customer orders engendered by macro announcements are not common knowledge.)

Given our model is analytically similar to that in Evans and Lyons (1999), we relegate solution details to the appendix. The resulting price changes and interdealer flows (end hour $h-1$ to end hour h) for a representative hour can be written as:

$$(5) \quad \Delta P_h = (\mathbf{b}_1 + \mathbf{b}_2 A_h) \Delta x_h - \mathbf{b}_3 \Delta P_{h-1} + \mathbf{h}_h^p$$

$$(6) \quad \Delta x_h = \mathbf{b}_4 \Delta x_{h-1} + \mathbf{b}_5 \Delta P_{h-1} + \mathbf{h}_h^x$$

where $\mathbf{h}_h^p = \mathbf{b}_6 \Delta R_h$, $\mathbf{h}_h^x = (\mathbf{b}_7 + \mathbf{b}_8 A_h) C_{1h}$ and the constants β_1 through β_8 are positive. (The variable C_{1h} denotes the aggregate round-one order flow in hour h .) The order flow coefficients in the price equation, β_1 and β_2 , depend on γ (the public's aggregate risk-bearing capacity from equation 4), the variances \mathbf{s}_R^2 , \mathbf{s}_C^2 , \mathbf{s}_C^2 , and the correlation between announcement engendered order flow \hat{C}_{1h} and ΔR_h .

The intuition for each of the terms in equations (5) and (6) is as follows. The first term in the return equation is the price impact of order flow. The contemporaneous price impact of order flow increases when a macro announcement immediately precedes it (picked up by the indicator variable A_h). The second term in the return equation reflects the fact that price effects from order flow have a mean reverting component, due to the transitory intraday risk premia

that arise in the model. The residual in the return equation reflects the flow of unmeasured macroeconomic information (orthogonal to order flow Dx_h). The first two terms in the order flow equation capture hot potato trading: interdealer order flow in the model is positively autocorrelated as positions are passed from dealer to dealer for risk management purposes. Hot potato trading covaries positively with lagged returns as well (due to contemporaneous price impact) and because order flow is measured with noise. The residual in the order flow equation reflects the flow of unmeasured order flow from customers (contemporaneous).

2. 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. 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. According to Reuters, over 90 percent of the world's direct interdealer transactions took place through the system.⁹ 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

⁹ 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).

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 Δx_t is therefore measured as the difference between the *number* of buyer-initiated and seller-initiated trades.

The variables in our empirical model are measured hourly. We take the spot rate, as the last purchase-transaction price (DM/\$) in hour h , P_h . (With roughly 1 million transactions per day, the last purchase transaction is generally within a few seconds of the end of the hour. Using purchase transactions eliminates bid-ask bounce.) Order flow, Δx_h , is the difference between the number of buyer- and seller-initiated trades (in thousands, negative sign denotes net dollar sales) during hour h . We also make use of three further variables to measure the state of the market: the number of macroeconomic announcements a_h ; trading intensity n_h , measured by the gross number of trades during hour h ; and price volatility s_h , measured by the standard deviation of all transactions prices during hour h . The macroeconomic announcements comprise all those reported over the Reuter's News service that relate to macroeconomic data for the U.S. or Germany (Money Market Headline News). The source is Olsen Associates (Zurich); for details, see, e.g., Andersen and Bollerslev (1998).

Although trading can take place on the D2000-1 system 24 hours a day, 7 days a week, the vast majority of transactions in the DM/\$ take place between 6 am and 6 pm, London time, Monday to Friday. The results we report below are based on this sub-sample. (They are similar to results based on the 24-hour trading day.) This sub-sample includes a vast number of trades, providing us with considerable power to testing the state dependence of liquidity.

3. Results

Stage 1: Constant Coefficient Model

Table 1 presents results for our first-stage estimation, the constant-coefficient model with two equations, one for returns and one for order flow

(hourly data). In the returns equation, both of the variables our model predicts should be relevant are significant and correctly signed. The magnitude of the coefficient on order flow Δx_h implies that the contemporaneous impact of order flow on price is about 60 basis points per \$1 billion.¹⁰ The coefficient on lagged returns implies that about three quarters of order flow's impact effect persists indefinitely.

Two other important facets of the returns-equation results warrant attention. First, note from row (iii) that there is no mean reversion in price unconditionally. This is consistent with the exchange rate following a martingale in hourly data. It is only when order flow is included that price exhibits some conditional mean reversion. (This conditional mean reversion result is not a violation of market efficiency, however, because our data were not available to market participants in real time.) Second, note from row (v) that lags beyond those predicted by our model are not significant in the returns equation. (This holds for further lags of both variables as well; not reported.)

Estimates for the order flow equation accord with the model as well. Both variables predicted to be relevant are significant and properly signed, whereas lags beyond those predicted are not significant. In this equation, as in the returns equation, there is evidence of heteroskedasticity, so our standard errors are adjusted for this. These results indicate that there is momentum in order flows: positive flows follow positive flows, which is consistent with the presence of hot potato trading.

Stage 2: Kernel Regression Estimates of State Dependent Liquidity

To test for state dependence in liquidity, we consider nonparametric regressions of the form:

$$\Delta p_h = \mathbf{p}(\Delta x_h, \Delta p_{h-1}, s_{h-1}) + \mathbf{h}_h^p,$$

$$\Delta x_h = \mathbf{m}(\Delta x_{h-1}, \Delta p_{h-1}, s_{h-1}) + \mathbf{h}_h^x,$$

¹⁰ This is based on an average trade size in our sample of \$3.9 million. (This average trade size is available despite individual trade sizes not being available.) We use log price change as our dependent

where $p(.)$ and $m(.)$ are arbitrary fixed, unknown, and nonlinear functions of the variables shown, and h_t is a mean zero i.i.d. error. (The vector s_{h-1} is a vector of state variables, specified below.) Our strategy here is to estimate these functions by kernel regression and then test whether our estimates of order flow's price impact are influenced by the flow of macroeconomic announcements. As noted in the introduction, the period following announcements is commonly thought to be a time when market expectations are ripe for coordinated adjustment (and our trading model provides a structural understanding of how this might work). To insure that our estimates for the announcement state variable are robust, we consider two additional potential state variables suggested by theory (see, e.g., Easley and O'Hara 1992): price volatility and trading volume.¹¹

Tables 2 and 3 present our kernel regression results for the returns and order flow equations, respectively (hourly data). For the returns equation (Table 2), \hat{p}_j denotes the derivative of the estimated function $\hat{p}(\cdot)$ with respect to j 'th variable. Thus, the first three rows indicate how the price impact of order flow (the first argument in the \hat{p} function) varies with the variables shown. Note the significant effect from the announcements state variable a_{h-1} . This variable is defined as the number of macro announcements in hour $h-1$. Accordingly, it implies that the price impact of order flow is about 10 percent higher with each macro announcement in the previous hour ($0.024/0.26$ being about 10 percent, where 0.26 comes from the linear price impact estimate in Table 1). Thus, liquidity in this market depends on the pace of public information flow. The other two potential state variables do not appear to matter for the price impact of order flow: σ_{h-1} and n_{h-1} are insignificant. (σ_{h-1} is the standard deviation of all transaction prices in the previous hour and n_{h-1} is the number of transactions in the previous hour.) Note too that the time-of-day dummies are insignificant (see table

variable because this is a more common measure of returns in the empirical literature. Use of the raw (unlogged) price change has no qualitative effect on our results.

¹¹ The announcement and volume state variables are also considered in the empirical analysis of central bank intervention by Dominguez (2001). Using intervention trade data, she finds that interventions that occur near macro announcements and during heavy trading volume are the most likely to have large effects. In her analysis, "heavy volume" means occurring when both London and New York are trading, and "near macro announcements" means that the public announcement of interventions occurred near the time of the macro announcement.

notes for definition); thus, within the 6 am to 6 pm London time period, time does not appear to define distinct states once we control for the other variables in our specification. Finally, note that there is no evidence of non-linearity in order flow's price effect, once these state variables are included.

From the lower panel of Table 2, one sees that all three of the state variables influence the persistence of price effects. Higher trading volume increases the persistence of hourly price movements, whereas price volatility and announcement flow decrease the persistence of hourly price movements. From this lower panel, note too the significant non-linearity in the Δp_{h-1} variable. This implies that larger moves are less persistent (controlling for the other variables).

Table 3 presents kernel regression results for the order flow equation. The three state variables and the time-of-day dummies are insignificant across the board in this case. There is some slight evidence of non-linearity of in the Δx_{h-1} variable, but the magnitude of the effect is not large economically. Bottom line: the order flow process does not appear to be state or time dependent as specified here.

Impulse Responses

Impulse responses provide an efficient way to summarize the effects of announcement flow on the dynamic, two-equation system examined in Tables 2 and 3. These impulse responses are presented in Figures 1-2. Figure 1 illustrates the effect of order flow shocks on both returns and subsequent order flow, including how announcement flow influences the dynamics. Figure 2 illustrates the effect of return shocks on both returns and subsequent order flow, also including how announcement flow influences the dynamics. These figures complement the evidence in Tables 2 and 3 in that they provide an illustration of the effects of announcements over time. For completeness, Figures 3-6 address the effects of the two additional state variables, trading volume and price volatility. (See the appendix for computation details.)

The strongest state-variable effects are manifest in figures 1 and 2, which illustrate the effects of announcement flow. (That these effects should be strongest is consistent with the kernel regression results reported above, given their

complementarity.) The experiment in figure 1 is the following. Suppose the announcement flow in the previous hour is high (one deviation above the sample average) and is expected to remain high forever: how does this change the effect of an innovation in order flow on price and subsequent order flow? The increased price impact from announcement flow is clear from the upper panel of Figure 1. Note too that the order flow and return responses persist. (This is conditional on the announcement flow persisting, so it should not be viewed as an indication of a non-stationary system.) Note too the price shock effect on the next hour's order flow from Figure 2. This is consistent with hot potato trading, i.e., one average much of price variation is due to order flow innovations, which on average are followed by subsequent flows in the same direction (though the price impact of those subsequent flows is indistinguishable from zero).

The impulse responses in figures 3-6 are broadly consistent with the results on liquidity state dependence from the kernel regressions. The two additional state variables of price volatility and trading volume have little effect on the price impact of order flow innovations. The order flow process is more significantly affected. Figure 3, for example, exhibits differential effects on order flow due to high trading volume (number of trades). High trading volume also significantly alters the order flow response to price shocks, as shown in figures 4 and 6.

4. Bias Analysis

Though the direction of causality in our model runs 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 section examines whether feedback trading can account for our results.

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, which is not a property of major floating exchange rates (and is not a property of our hourly data either—see Table 1, row iii). Accordingly, the

class of feedback trading models that might be relevant here is the non-rational class.¹²

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 for the issue in this paper.

Our approach here is to pose and address the following question: Suppose intra-hour (i.e., contemporaneous in hourly data) positive-feedback trading is present, under what conditions could it account for the key moments of our data? To address this question, we decompose measured order flow Δx_h into two components:

$$(7) \quad \Delta x_h = \Delta x_h^* + \Delta x_h^{fb}$$

where Δx_h^* denotes exogenous order flow from portfolio shifts as identified in our model, and Δx_h^{fb} denotes contemporaneous order flow due to feedback trading, where:

$$(8) \quad \Delta x_h^{fb} = f \Delta p_h .$$

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 coefficients on contemporaneous order flow in the price-change equations in Table 1).

Next, suppose the true structural model can be written as:

¹² 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.

$$(9) \quad \begin{aligned} \Delta p_h &= \mathbf{b}_{11} \Delta x_h^* + \mathbf{b}_{12} \Delta p_{h-1} + \mathbf{h}_h^p \\ \Delta x_h^* &= \mathbf{b}_{21} \Delta x_{h-1}^* + \mathbf{b}_{22} \Delta p_{h-1} + \mathbf{h}_h^x \end{aligned}$$

Notice that the equations in (9) are valid reduced-forms from our model that could be estimated by OLS if one had data on Δx_h^* . However, if feedback trading is present (i.e., $f \neq 0$), estimates of (9) using measured order flow, Δx_h , will suffer from simultaneity bias.

We can evaluate the size of this bias by estimating (7) – (9) as a whole system of equations. Specifically, we can combine (7) – (9) into a bivariate system for price changes and measured order flow. The dynamics of this system depend on seven parameters: $f, \mathbf{b}_{11}, \mathbf{b}_{12}, \mathbf{b}_{21}, \mathbf{b}_{22}$, and the variances of \mathbf{h}_h^p and \mathbf{h}_h^x . These parameters can be estimated by GMM using sample estimates of the covariance matrix for the vector $[\Delta p_h, \Delta x_h, \Delta p_{h-1}, \Delta x_{h-1}]$, as described more fully in the appendix. In broad terms, the feedback trading alternative predicts that (1) the coefficient \mathbf{b}_{11} on exogenous order flow will be smaller than those in Table 1, if not zero, and (2) the coefficient on feedback trading f will be positive and significant.

The GMM estimates are reported in Table 4. The last row of the table reports the estimate of the feedback parameter, f : the estimate is negative and statistically insignificant. Thus, insofar as there is any empirical evidence of feedback trading in our data, it points to the presence of negative rather than positive feedback trading. Moreover, estimates from the price equation show that our causal interpretation of the order flow's impact on price (based on our model) remains intact: the estimate of \mathbf{b}_{11} is slightly larger than those in Table 1 and remains highly statistically significant, in contrast to what the feedback trading alternative predicts.

5. Conclusions

We are, for the first time, at the point where we can measure liquidity in the FX market, and why it varies over time. Our results show that liquidity in

currency markets depends on the pace of public information flow. Liquidity does not, however, appear to depend on trading volume or return volatility (two other prominent state variables in the literature on liquidity variation). These results provide policymakers with concrete guidance for increasing the efficacy of intervention by selectively timing their trades. The trading model we develop provides a structural interpretation for why the effect from public information flow should be present in the data. In the model, the implications of new macro-economic data for the exchange rate are not inferable from the macro data alone, which allows order flow to convey incremental information about agents' changing expectations.

Our trading model provides testable implications beyond those involving public information flow. For example, it provides a structural account for the joint behavior of order flow and returns. Empirically, the model accords well with the data. The variables the model predicts should be relevant are indeed significant. The variables (and lags of variables) the model predicts should be insignificant are insignificant. Our estimate of the contemporaneous impact of order flow on price is about 60 basis points per \$1 billion. Of this 60 basis points, roughly 80 percent persists indefinitely. Our model also produces hot potato trading, and the data accord well with predictions in this respect as well. Specifically, we find momentum in order flows (positive flows follow positive flows in hourly data), though only the initial innovations in flow appear to have impact on prices.

Though our policy focus in this paper has been central bank intervention, order flow analysis is relevant for other policy issues as well. Consider two examples. The first is the liquidity hole that occurred in the dollar-yen market in October 1998 (the immediate aftermath of the LTCM crisis). Within about a day, the yen/\$ rate fell from about 132 to about 117 and bid-ask spreads rose to nearly one yen (i.e., rose to about 30 times their typical size in the interbank market). This was not a trivial event in terms of resource allocation: the new lower level of the exchange rate was persistent. Why market liquidity dried up so fast and so drastically is still a puzzle, one that order flow analysis may help to resolve (see, e.g., the analysis of Citibank's customer trades through October 1998 in Lyons 2001, including those for hedge funds). A second policy area where order flow

analysis may prove valuable is currency market design, especially in emerging markets. With appropriate data, one could estimate how price impact in these markets changes as a function of the market state (devaluation likelihood, etc.). Also, one could determine whether customer forward trades have the same price impact as customer spot trades of similar size. If not, one could quantify the difference. (Many developing countries restrict or even forbid forward trading on the belief that such trading is more “speculative” in nature than spot trading and is therefore more destabilizing.) One might also compare price impact across countries, in an effort to determine which institutional structures are better at promoting liquidity.¹³

The issue of price impact is related to the issue of market stability. Policymakers in some developing countries appear to believe that additional liquidity is destabilizing. In theory, it is less liquidity that is destabilizing, not more liquidity: the less the liquidity, the larger the price impact, and the more prices move (other things equal). To make the case that other things are not equal, in a way that might reverse the relationship between liquidity and stability, one could use the discipline of microstructure trading models to identify the countervailing forces.

¹³ One paper that addresses speculative attacks in Mexico using a microstructure approach is Carrera (1999). For theoretical work on the design of currency markets in developing countries see Kirilenko (1997).

Appendix A: Kernel Regressions

We consider nonparametric regressions of the form:

$$y_t = m(z_t) + h_t$$

where $m(\cdot)$ is an arbitrary fixed but unknown nonlinear function of the variables in the vector z_t , and h_t is a mean zero i.i.d. error. An estimate of the $m(\cdot)$ function is estimated by kernel regression as:

$$\hat{m}(z_t) = \frac{\sum_{j=0, j \neq t}^T K_b(z_t - z_j) y_j}{\sum_{j=0, j \neq t}^T K_b(z_t - z_j)}$$

where $K_b(u) = b^{-1} K(u/b)$ with $K(u) \geq 0$ and $\int K(u) du = 1$. In this application, we use the multivariate Gaussian kernel $K(u) = (2\pi)^{-d/2} \exp(-u'u/2)$ where $d = \dim(u)$. The bandwidth parameter, b , is chosen by cross-validation. That is to say, b minimizes:

$$\frac{1}{T} \sum_t (y_t - m(z_t))^2 w_t$$

where w_t is a weighting function that cuts off 5% of the data at each end of the data interval as in Hardle (1990), p. 162. We follow the common practice of including the standardized value of each of these variables in the Gaussian kernel (i.e., each element of z_t is divided by its sample standard deviation).

Asymptotic theory for kernel regressions in the time series context appear in Bierens (1983) and Robinson (1983). Robinson shows that consistency and asymptotic normality of the estimator can be established when the data satisfy α -mixing with mixing coefficients $\alpha(k)$ that obey the condition $T \sum_t^\infty \alpha(k)^{1-2/d} = O(1)$ and $E |y_t|^d < \infty$, with $d > 2$.

Appendix B: Impulse Response Functions

We use our kernel estimates of the price change and order flow equations

$$\Delta p_h = \mathbf{p}(\Delta x_h, \Delta p_{h-1}, s_{h-1}) + \mathbf{h}_h^p,$$

$$\Delta x_h = \mathbf{m}(\Delta x_{h-1}, \Delta p_{h-1}, s_{h-1}) + \mathbf{h}_h^x,$$

where $s_{h-1} = [a_{h-1}, \mathbf{s}_{h-1}, n_{h-1}]$ is the vector of state variables, to compute two sets of impulse responses as follows.

Responses to order flow shocks

The impulse response to an order flow shock of \mathbf{d} in hour h , is defined as

$$I_{\Delta p}(\mathbf{t}) = E[\Delta p_{h+t} | \Omega_{h-1}, \Delta x_h = \mathbf{d}] - E[\Delta p_{h+t} | \Omega_{h-1}, \Delta x_h = 0],$$

$$I_{\Delta x}(\mathbf{t}) = E[\Delta x_{h+t} | \Omega_{h-1}, \Delta x_h = \mathbf{d}] - E[\Delta x_{h+t} | \Omega_{h-1}, \Delta x_h = 0],$$

where Ω_{h-1} is the information set containing the history of price changes, order flows, and the state variables until the start of hour h :

$$\Omega_{h-1} = \{\Delta x_{h-1}, \Delta x_{h-2}, \dots, \Delta p_{h-1}, \Delta p_{h-2}, \dots, s_{h-1}, s_{h-2}, \dots\}.$$

$I_{\Delta p}(\mathbf{t})$ identifies the change in expectations regarding Δp_{h+t} if order flow in hour h changes from zero to \mathbf{d} , given the particular history of past order flow, price changes, and the state variables in Ω_{h-1} . Similarly $I_{\Delta x}(\mathbf{t})$ identifies the change in expectations regarding Δx_{h+t} if order flow in hour h changes from zero to \mathbf{d} , given Ω_{h-1} . Notice that given the specification of the price and order flow equations, $I_{\Delta x}(0) = \mathbf{d}$ and $I_{\Delta p}(0) = \mathbf{p}(\mathbf{d}, \Delta p_{h-1}, s_{h-1}) - \mathbf{p}(0, \Delta p_{h-1}, s_{h-1})$ so the impact of the order flow shock can be directly calculated from the kernel regression estimates.

Calculating the change in expectations for $t > 0$ is much more complicated because current price changes and order flow depend nonlinearly on past prices and order flow. Specifically, consider the one-period-ahead price change forecast. By iterated expectations, this forecast can be expressed as

$$E[\Delta p_{h+1} | \Omega_{h-1}, \Delta x_h = \mathbf{d}] = E[E[\Delta p_{h+1} | \Omega_{h-1}, \Delta x_h = \mathbf{d}, \Delta x_{h+1}, \Delta p_h] | \Omega_{h-1}, \Delta x_h = \mathbf{d}].$$

The inner conditional expectation can be calculated from the kernel estimates of $\mathbf{p}(\Delta x_{h+1}, \Delta p_h, s_h)$ given realizations of $\Delta x_{h+1}, \Delta p_h$, and s_h . To compute $E[\Delta p_{h+1} | \Omega_{h-1}, \Delta x_h = \mathbf{d}]$ we therefore have to calculate the expected value of $\mathbf{p}(\Delta x_{h+1}, \Delta p_h, s_h)$ using the joint distribution of $\Delta x_{h+1}, \Delta p_h$, and s_h conditional on $\{\Omega_{h-1}, \Delta x_h = \mathbf{d}\}$. When $\mathbf{p}(\cdot)$ is linear, this is a straightforward calculation. When $\mathbf{p}(\cdot)$ is nonlinear, the expectation must be calculated by numerical simulation (described below).

History dependency also complicates the calculation of impulse response functions in nonlinear models. Notice that the impulse responses specify both the shock and history Ω_{h-1} . In linear models, the affect of a shock does not depend on the history of past shocks so it is unnecessary to specify Ω_{h-1} . Here we must specify Ω_{h-1} because the price impact of order flow in hour h may (in principle) vary according to the value of Δp_{h-1} , and the state variables, z_{h-1} . The results presented in the paper assume a particular history Ω_{h-1}^0 , where $\Delta p_{h-1}^0 = 0$, $\Delta x_{h-1}^0 = 0$ and s_{h-1} equals a vector of constants, s^0 (specified below). With this history, $E[\Delta p_{h+t} | \Omega_{h-1}^0, \Delta x_h = 0] = E[\Delta x_{h+t} | \Omega_{h-1}^0, \Delta x_h = 0] = 0$.

The Impulse Responses are computed as follows:

1. Compute the vector of price and order flow residuals from the kernel regressions: $\mathbf{h}_h^p = \Delta p_h - \mathbf{p}(\Delta x_h, \Delta p_{h-1}, s_{h-1})$, and $\mathbf{h}_h^x = \Delta x_h - \mathbf{m}(\Delta x_{h-1}, \Delta p_{h-1}, s_{h-1})$. These residual vectors have mean zero.
2. Set $\Delta x_h = \mathbf{d}$, and compute $E[\Delta p_h | \Omega_{h-1}^0, \Delta x_h = \mathbf{d}] = \mathbf{p}(\mathbf{d}, 0, s^0)$ from the kernel estimates and $I_{\Delta x}(0) = \mathbf{d}$. Compute $\Delta \tilde{p}_h = E[\Delta p_h | \Omega_{h-1}^0, \Delta x_h = \mathbf{d}] + \tilde{\mathbf{h}}_h^p$ where $\tilde{\mathbf{h}}_h^p$ is a random drawing from the vector of price residuals.
3. Add the realizations $\{\Delta \tilde{p}_h, \mathbf{d}\}$ to the data set and compute new kernel estimates of the $\mathbf{p}(\cdot)$ and $\mathbf{m}(\cdot)$ functions using the optimized bandwidth parameters calculated in the data sample, $\tilde{\mathbf{p}}_h(\cdot)$ and $\tilde{\mathbf{m}}_h(\cdot)$.
4. Compute $E[\Delta x_{h+1} | \Omega_h^0] = \tilde{\mathbf{m}}_h(\Delta \tilde{x}_h, \Delta \tilde{p}_h, z^0)$ and $\Delta \tilde{x}_{h+1} = E[\Delta x_{h+1} | \Omega_h^0] + \tilde{\mathbf{h}}_{h+1}^x$ where $\tilde{\mathbf{h}}_{h+1}^x$ is a random drawing from the vector of order flow residuals, and $\Omega_h^0 = \Omega_{h-1}^0 \cup \{\Delta \tilde{p}_h, \mathbf{d}\}$.
5. Compute $E[\Delta p_{h+1} | \Omega_h^0, \Delta \tilde{x}_{h+1}] = \tilde{\mathbf{p}}_h(\Delta \tilde{x}_{h+1}, \Delta \tilde{p}_h, z^0)$ and $\Delta \tilde{p}_{h+1} = E[\Delta p_{h+1} | \Omega_h^0, \Delta \tilde{x}_{h+1}] + \tilde{\mathbf{h}}_{h+1}^p$ where $\tilde{\mathbf{h}}_{h+1}^p$ is a random drawing from the vector of price residuals.
6. Add the realizations $\{\Delta \tilde{p}_{h+1}, \Delta \tilde{x}_{h+1}\}$ to the data set and compute new kernel estimates of the $\mathbf{p}(\cdot)$ and $\mathbf{m}(\cdot)$ functions using the optimized bandwidth parameters calculated in the data sample, $\tilde{\mathbf{p}}_{h+1}(\cdot)$ and $\tilde{\mathbf{m}}_{h+1}(\cdot)$.
7. Repeat steps 4, 5, and 6, t times. We now have a t -period history of price changes and order flow driven by the \mathbf{d} shock to order flow in period h and subsequent random price and order flow shocks. We also have the sets of forecasts $E[\Delta p_{h+i} | \Omega_{h+i-1}^0, \Delta \tilde{x}_{h+i}]$ and $E[\Delta x_{h+i} | \Omega_{h+i-1}^0]$.
8. Calculate 1000 t -period histories of price changes and order flow using the steps 1 - 7 above and save the forecasts as $E[\Delta p_{h+i} | \Omega_{h+i-1}^0, \Delta \tilde{x}_{h+i}]^j$ and $E[\Delta x_{h+i} | \Omega_{h+i-1}^0]^j$ together with the residuals $[\tilde{\mathbf{h}}_{h+i-1}^p]^j$ and $[\tilde{\mathbf{h}}_{h+i-1}^x]^j$ for $j = 1, 2, \dots, 1000$.
9. Compute

$$E[\Delta p_{h+i} | \Omega_{h-1}^0, \Delta x_h = \mathbf{d}] = E\left[E[\Delta p_{h+i} | \Omega_{h+i-1}^0, \Delta x_{h+i}] \middle| \Omega_{h-1}^0, \Delta x_h = \mathbf{d}\right]$$

$$E[\Delta x_{h+i} | \Omega_{h-1}^j, \Delta x_h = \mathbf{d}] = E[E[\Delta x_{h+i} | \Omega_{h+i-1}^0] | \Omega_{h-1}^0, \Delta x_h = \mathbf{d}]$$

for $i = 1, 2, \dots, t$. An estimate of the expectation in the first equation of step 9 is found as the constant in a regression of $E[\Delta p_{h+i} | \Omega_{h+i-1}^0, \Delta \tilde{x}_{h+i}]^j$ on $[\tilde{\mathbf{h}}_{h+i-1}^p]^j$ for $j=1, 2, \dots, 1000$. An estimate of the expectation in the second equation of step 9 is similarly found as the constant in a regression of $E[\Delta x_{h+i} | \Omega_{h+i-1}^0]^j$ on $[\tilde{\mathbf{h}}_{h+i-1}^x]^j$. In both regressions, realizations of residuals are used as control variates to obtain more precise estimates of the expectations than would be obtained by simply averaging over $E[\Delta p_{h+i} | \Omega_{h+i-1}^0, \Delta \tilde{x}_{h+i}]^j$ or $E[\Delta x_{h+i} | \Omega_{h+i-1}^0]^j$. (For a discussion of the use of control variates, see Davidson and MacKinnon 1993, chapter 21.) The OLS standard error for the constant estimates the standard error of the simulated expectation.

The procedure described above allows us to calculate the impulse response and a confidence band (i.e., ± 1.96 standard errors) for an order flow shock of \mathbf{d} given a particular set of values for the state variables, \mathbf{s}^0 . In our baseline case, we set \mathbf{s}^0 equal to the sample average for \mathbf{s}_h . We compare this baseline case against alternatives in which one element of \mathbf{s}^0 is changed.

Responses to price shocks

The impulse response to a price shock of \mathbf{d} in hour h , is defined as

$$I_{\Delta p}(\mathbf{t}) = E[\Delta p_{h+t} | \Omega_{h-1}, \Delta x_h = 0, \Delta p_h = \mathbf{d}] - E[\Delta p_{h+t} | \Omega_h, \Delta x_h = 0, \Delta p_h = 0],$$

$$I_{\Delta x}(\mathbf{t}) = E[\Delta x_{h+t} | \Omega_{h-1}, \Delta p_h = \mathbf{d}] - E[\Delta x_{h+t} | \Omega_{h-1}, \Delta p_h = 0],$$

Given the timing of order flow and price changes in our model, $I_{\Delta x}(0) = 0$ and $I_{\Delta p}(0) = \mathbf{d}$. The impulse responses for $\mathbf{t} > 0$ are calculated by simulation in the same manner as the order flow responses.

Appendix C: Bias Analysis Estimates

Let \mathbf{a} represent the vector of parameters to be estimated (i.e., \mathbf{f} , \mathbf{b}_{11} , \mathbf{b}_{12} , \mathbf{b}_{21} , \mathbf{b}_{22} , and the variances of \mathbf{h}_h^p and \mathbf{h}_h^x). Combining equations (7) – (9), we can write the dynamics of returns and measured order flow as

$$Y_h = AY_{h-1} + B\mathbf{h}_h$$

where $Y_h = [\Delta p_h, \Delta x_h]'$, $\mathbf{h}_h = [\mathbf{h}_h^p, \mathbf{h}_h^x]'$,

$$A = \begin{bmatrix} \mathbf{b}_{11} + \mathbf{b}_{11}\mathbf{b}_{22} - \mathbf{f}\mathbf{b}_{11}\mathbf{b}_{21} & \mathbf{b}_{11}\mathbf{b}_{21} \\ \mathbf{f}(\mathbf{b}_{11} + \mathbf{b}_{11}\mathbf{b}_{22}) + \mathbf{b}_{22} - \mathbf{f}^2\mathbf{b}_{11}\mathbf{b}_{21} - \mathbf{f}\mathbf{b}_{21} & \mathbf{f}\mathbf{b}_{11}\mathbf{b}_{21} + \mathbf{b}_{21} \end{bmatrix}, \quad B = \begin{bmatrix} 1 & \mathbf{b}_{11} \\ \mathbf{f} & 1 + \mathbf{f}\mathbf{b}_{11} \end{bmatrix},$$

and $\Theta = \text{Cov}(\mathbf{h}_h, \mathbf{h}_h')$, a diagonal matrix.

We can therefore compute the covariance of returns and measured order flow, as a function of the model parameters \mathbf{a} ; $\Gamma(k; \mathbf{a}) = \text{Cov}(Y_h, Y_{h-k})Y_h$, as

$$\Gamma(k; \mathbf{a}) = A\Gamma(k-1; \mathbf{a})$$

with

$$\Gamma(0; \mathbf{a}) = \text{vec}^{-1} \left((I - A \otimes A)^{-1} \text{vec}(\Theta) \right).$$

The GMM estimates of \mathbf{a} are based on orthogonality conditions of the form

$$E[m_h(k; \mathbf{a})] = 0$$

where

$$m_h(k; \mathbf{a}) = D(k) \text{vec} \left(Y_h Y_{h-k}' - \Gamma(k; \mathbf{a}) \right)$$

and $D(k)$ is a vector of ones and zeros that selects the unique elements in $\Gamma(k; \mathbf{a})$.

The results in Table 4 use $k = 0$ and 1, for a total of seven moments.

Appendix D: The Dealer's Problem and Model Solution (Incomplete)

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 six choice variables described in section 2, namely, the four scalar quotes P_j^i , one for each round j , and the two outgoing interdealer trades, T_2^i and T_3^i :

$$\begin{aligned}
 (A1) \quad & \text{Max} \quad E[-\exp(-qW_4^i | \Omega^i)] \\
 & \{P_1^i, P_2^i, P_3^i, P_4^i, T_2^i, T_3^i\} \\
 & \text{s.t.} \\
 & W_4^i = W_0^i + C_1^i(P_1^i - \tilde{P}_2^i) + \tilde{T}_2^i(\tilde{P}_2^i - \tilde{P}_3^i) + \tilde{T}_3^i(\tilde{P}_3^i - P_4^i) \\
 & \quad + (T_2^i - C_1^i)(\tilde{P}_3^i - \tilde{P}_2^i) + (T_3^i + T_2^i - C_1^i - \tilde{T}_2^i)(P_4^i - \tilde{P}_3^i)
 \end{aligned}$$

Dealer i 's wealth over the four-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 orders \tilde{T}_2^i and \tilde{T}_3^i (tilde distinguishes incoming interdealer orders and prices from outgoing). The outgoing orders are the two interdealer trades T_2^i and T_3^i . \tilde{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 the incoming interdealer quote \tilde{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 two terms of the budget constraint.) Terms three and four reflecting the incoming (random) dealer orders are analogous.

Terms five and six of the budget constraint reflect the dealer's speculative and hedging demands. The outgoing interdealer trade in round 2 has three components:

$$(A2) \quad T_2^i = C_1^i + D_2^i + E[\tilde{T}_2^i | \Omega_{2T}^i]$$

where D_2^i is dealer i 's speculative demand in round 2, and $E[\tilde{T}_2^i | \Omega_{2T}^i]$ 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:

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

which corresponds to the position in term five of the budget constraint. The sixth term in the budget constraint is analogous: the dealer's total demand in round three is his total trade in round three (T_3^i) plus his total demand in round two ($T_2^i - C_1^i$) less the random interdealer order he received in round two (\tilde{T}_2^i).

The conditioning information Ω^i at each decision node (4 quotes and 2 outgoing orders) is summarized below (see also the daily timing in the text).

$$\begin{aligned}\Omega_{1P}^i &= \left\{ \bigcup_{k=1}^{t-1} \{ \Delta R_k, \Delta x_{2k}, \Delta x_{3k}, P_{1k}, P_{2k}, P_{3k}, P_{4k}, A_k, A_t \} \right\} \\ \Omega_{2P}^i &= \{ \Omega_{1P}^i, P_{1t}, C_1^i \} \\ \Omega_{2T}^i &= \{ \Omega_{2P}^i, P_{2t} \} \\ \Omega_{3P}^i &= \{ \Omega_{2T}^i, \Delta x_{2t} \} \\ \Omega_{3T}^i &= \{ \Omega_{3P}^i, P_{3t}, \Delta R_t \} \\ \Omega_{4P}^i &= \{ \Omega_{3T}^i, \Delta x_{3t} \}\end{aligned}$$

At this stage it is necessary to treat each of the prices in these information sets as a vector that contains the price of each individual dealer i (though in equilibrium each of these prices is a scalar, as shown below).

Equilibrium

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

To solve for the symmetric BNE, first consider optimal quoting strategies.

PROPOSITION A1: A quoting strategy is consistent with symmetric BNE only if quotes within any single trading round are common across dealers.

PROPOSITION A2: A quoting strategy is consistent with symmetric BNE only if $P_1=P_2$ and these prices are equal to the final round price P_4 from the previous day.

PROPOSITION A3: A quoting strategy is consistent with symmetric BNE only if the common round-three quote is:

$$P_3 = P_2 + I_2 \Delta x_2$$

if there is no announcement ($A_t=0$) and

$$P_3 = P_2 + I_2^* \Delta x_2$$

if there is an announcement ($A_t=1$), where the constants $I_2 < I_2^*$ are strictly positive and Δx_2 denotes the signal of round-two interdealer order flow.

PROPOSITION A4: A quoting strategy is consistent with symmetric BNE only if the common round-four quote is:

$$P_4 = P_3 + I_3 \Delta x_3 + d\Delta R - y(P_3 - P_2)$$

if there is no announcement ($A_t=0$) and

$$P_4 = P_3 + I_3^* \Delta x_3 + d\Delta R - y(P_3 - P_2)$$

if there is an announcement ($A_t=1$), where the constants $I_3 < I_3^*$, δ , and ψ are strictly positive and Δx_3 denotes round-three interdealer order flow.

Propositions A1 through A4

The proof of proposition A1 is straightforward: That all dealers post the same quote in any given trading round is required to eliminate risk-free arbitrage. (Recall from section 2 that all quotes are scalar prices at which the dealer agrees to buy/sell any amount, and trading with multiple partners is feasible.)

The proof of proposition A2 is straightforward as well: Common prices require that quotes depend only on information that is commonly observed. In round one, this includes the previous day's round-four price. Because there is no new information that is commonly observed between round four and round two quoting the following day, the round-four price is not updated. (Recall that public trading in round four is a deterministic function of round-four prices and therefore conveys no information. Recall too that the occurrence of an announcement $A_t=1$ at the beginning of the day provides no price-relevant information that is common to dealers by itself.) Thus, dealers' round-two quotes are not conditioned on individual realizations of C_1^i .

Propositions A3 and A4 require equations that pin down the levels of the four prices. Per above, these equations are necessarily functions of public information. Naturally, they also embed the equilibrium trading rules of dealers and customers. The equations are the following:

$$(A3) \quad E[C_1 | \Omega_{1P}] + E[ND_2^i(P_1) | \Omega_{1P}] = 0$$

$$(A4) \quad E[C_1 | \Omega_{2P}] + E[ND_2^i(P_2) | \Omega_{2P}] = 0$$

$$(A5) \quad E[C_1 | \Omega_{3P}] + E[ND_3^i(P_3) | \Omega_{3P}] = 0$$

$$(A6) \quad E[C_1 | \Omega_{4P}] + E[C_4(P_4) | \Omega_{4P}] = 0$$

where C_1 denotes the sum of C_1^i over all N dealers. The first three equations

state that for each round j ($j=1,2,3$), at price P_j dealers willingly absorb the estimated demand from customers (realized at the beginning of the day, but not observed publicly). The fourth equation states that at price P_4 the public willingly absorbs the estimated beginning-of-day customer portfolio shift C_1 . These equations pin down equilibrium prices because any price other than that which satisfies each generates irreconcilable demands in interdealer trading in rounds two and three (e.g., if price is too low, all dealers know that on average dealers are trying to buy from other dealers, which is inconsistent with rational expectations; see Lyons 1997 for a detailed treatment in another model within the simultaneous trade approach).

From these equations, $P_2 - P_1 = 0$ follows directly from two facts: (1) the expected value of C_1 conditional on public information Ω_{1P} or Ω_{2P} is zero and (2) expected dealer demand D_2^i is also zero at this public-information-unbiased price. To be more precise, this statement *postulates* that the dealer's demand D_2^i has this property; derivation of the optimal trading rule shows that this is the case.

That $P_3 - P_2 = \lambda_2 \Delta x_2$ if there is no announcement (with $\lambda_2 > 0$) follows from two facts: (1) interdealer order flow Δx_2 is the only public information revealed in this interval and (2) Δx_2 is positively correlated with—and therefore provides information about—the morning portfolio shift C_1 . The positive correlation arises because each of the dealer orders T_2^i of which Δx_2 is composed is proportional to the C_1^i received by that dealer. A positive expected C_1 induces an increase in price because it implies that dealers—having taken the other side of these trades—are short and need to be induced to hold this short position with an expected downward drift (intraday) in price. In addition, when an announcement occurs, the customer order flow C_1 is especially “rich” in the sense that it contains flow positively correlated with that day's payoff increment ΔR . This increases the price impact of subsequent interdealer order flow, dollar for dollar.

The exact size of this downward drift in price depends on where price is expected to settle at the end of the day. Per proposition A4, $P_4 - P_3 = \lambda_3 \Delta x_3 + \delta \Delta R - \psi(P_3 - P_2)$ when there is no announcement. This price change depends positively on the two pieces of public information revealed in this interval, Δx_3 and ΔR .¹⁴ The logic behind the positive Δx_3 effect is the same as that behind the positive Δx_2 effect in round two: a positive average T_3^i implies that the market's estimate of C_1 from Δx_2 was too low; absorption of the additional short position requires a price increase. (That a positive average T_3^i implies this is clear from the derivation of T_3^i .) The term $\delta \Delta R$ is the perpetuity value of the change in the daily payoff R_t . Finally, the drift term $-\psi(P_3 - P_2)$ is the equilibrium compensation to dealers for having to absorb the morning portfolio shift through the interval in which ΔR (and the associated price risk) is realized. This is an intraday price effect that dissipates by the end of the day.

¹⁴ Interdealer order flow Δx_3 is observed without noise, which means it reveals the value of C_1 fully. The price in round four must therefore adjust such that equation (A6) is satisfied exactly.

Equilibrium Trading Strategies

An implication of common interdealer quotes is that in rounds two and three each dealer i receives an order from exactly one other dealer, namely dealer $i+1$ (recall the trade allocation rule in section 2). These orders correspond to the position disturbances \tilde{T}_2^i and \tilde{T}_3^i in the dealer's problem in equation (A1).

Given the quoting strategy described in propositions 1-4, the following dealer trading strategy is optimal and corresponds to symmetric linear equilibrium:

PROPOSITION A5: The trading strategy profile:

$$T_2^i = \mathbf{a} C_1^i$$

if there is no announcement ($A_t=0$) and

$$T_2^i = \mathbf{a}^* C_1^i$$

if there is an announcement ($A_t=1$), with $0 < \mathbf{a} < \mathbf{a}^*$, conforms to Bayesian-Nash equilibrium.

PROPOSITION A6: The dealer trading strategy:

$$T_3^i = \mathbf{k}_1 C_1^i + \mathbf{k}_2 \Delta x_2 + \mathbf{k}_3 \tilde{T}_2^i$$

if there is no announcement ($A_t=0$) and

$$T_3^i = \mathbf{k}_1^* C_1^i + \mathbf{k}_2^* \Delta x_2 + \mathbf{k}_3^* \tilde{T}_2^i$$

if there is an announcement ($A_t=1$), conforms to Bayesian-Nash equilibrium.

Sketch of Proofs for Propositions A5 and A6

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

Propositions A5 and A6 are special cases of the analysis in Lyons (1997), which is also set in the context of a simultaneous-trade game with two interdealer trading rounds. Accordingly, we refer readers to that analysis for details on the derivation of optimal trading rules in this setting. One difference warrants note here: the Lyons (1997) analysis also includes private and public signals (denoted s and s in that paper) beyond any signals gleaned from order flow. Signals of this kind are not present in the specification here (i.e., one sets them equal to zero when applies the results of that paper to this model).

From Model Solution to Estimable Equation

The model above is split into 4 distinct trading rounds. Actual currency markets obviously do not map directly into this four-round structure. Our empirical implementation needs to apply to the “representative” trading interval, that interval being in our case one hour. From the model, we know that the distribution of price changes in any given hour depends on which “round” of trading the market is in and the probability of transitioning to a later round. Based on these transition probabilities, and the determinants of price within any given round, one can derive a conditional distribution of hourly price changes. We have done this in an earlier paper, based on a model that does not include announcements or hot potato trading. That same analysis applies to the setting here, however, so we refer people to that earlier paper for details (Evans and Lyons 2001, pages 13-14). Application of that analysis to the present model yields the estimating equations presented in the text, equations (5) and (6):

$$(5) \quad \Delta P_h = (\mathbf{b}_1 + \mathbf{b}_2 A_h) \Delta x_h - \mathbf{b}_3 \Delta P_{h-1} + \mathbf{h}_h^p$$

$$(6) \quad \Delta x_h = \mathbf{b}_4 \Delta x_{h-1} + \mathbf{b}_5 \Delta P_{h-1} + \mathbf{h}_h^x$$

where the subscript h denotes hour, the residual $\mathbf{h}_h^p = \mathbf{b}_6 \Delta R_h$, the residual $\mathbf{h}_h^x = (\mathbf{b}_7 + \mathbf{b}_8 A_h) C_{1h}$, and the constants β_1 through β_8 are positive.

Table 1: Linear Models

$$\Delta p_h = \mathbf{b}_{11}\Delta x_h + \mathbf{b}_{12}\Delta p_{h-1} + \mathbf{h}_h^p$$

$$\Delta x_h = \mathbf{b}_{21}\Delta x_{h-1} + \mathbf{b}_{22}\Delta p_{h-1} + \mathbf{h}_h^x$$

Equation	Regressors					Diagnostics			
	Δx_h	Δp_{h-1}	Δx_{h-1}	Δp_{h-2}	Δx_{h-2}	R ²	SEE	Serial	Hetero
$\frac{\Delta p_h}{\Delta x_h}$									
(i)	0.258 (12.93)	-0.203 (-4.82)				0.212	0.001	0.437 0.238	0.071 0.000
(ii)	0.225 (11.51)					0.173	0.002	0.000 0.000	0.070 0.000
(iii)		-0.061 (-1.47)				0.003	0.002	0.150 0.218	0.271 0.000
(iv)	0.234 (11.61)		-0.041 (-2.63)			0.180	0.002	0.000 0.000	0.067 0.000
(v)	0.258 (12.74)	-0.202 (-4.56)	-0.001 (-0.08)			0.212	0.001	0.205 0.235	0.071 0.000
$\frac{\Delta x_h}{\Delta p_h}$									
(i)		0.477 (5.11)	0.098 (2.42)			0.095	0.003	0.759 0.209	0.000 0.000
(ii)			0.205 (5.08)			0.042	0.003	0.520 0.161	0.000 0.000
(iii)		0.553 (5.97)				0.087	0.003	0.009 0.019	0.000 0.006
(iv)		0.479 (5.08)	0.097 (2.34)		-0.008 (-0.23)	0.095	0.003	0.369 0.397	0.000 0.000
(v)		0.486 (4.83)	0.091 (2.04)	0.028 (0.45)		0.096	0.003	0.516 0.394	0.000 0.000

* T-statistics in parentheses are calculated with asymptotic standard errors corrected for the presence of heteroskedasticity. OLS estimates are based on hourly observations from 6:00 to 18:00 BST from May 1 to August 31, 1996, excluding weekends. Δp_h is the hourly change in the log spot exchange rate (DM/\$) times 10,000. Δx_h is the hourly interdealer order flow, measured contemporaneously with Δp_h (negative for net dollar sales, in thousands). The Serial column presents the p-value of a chi-squared LM test for first-order (top row) and sixth-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 sixth-order (bottom row) ARCH in the residuals.

Table 2: Nonparametric (Kernel) Regression of Price Change Equation

$$\Delta p_h = \mathbf{p}(\Delta x_h, \Delta p_{h-1}, \mathbf{s}_{h-1}, n_{h-1}, a_{h-1}) + \mathbf{h}_h^p$$

Dependent Variable									Diagnostics		
	Const.	Δx_h	Δp_{h-1}	a_{h-1}	n_{h-1}	\mathbf{s}_{h-1}	t_h	t_h^2	R ²	Serial	Hetero
\hat{p}_1	0.134	<0.001	<0.001	0.024	0.001	-0.014			0.071	0.866	0.556
	(15.25)	(0.34)	(1.00)	(3.79)	(0.26)	(-1.23)				0.231	0.315
	0.077			0.023	-0.001	-0.013	0.010	<0.001	0.071	0.943	0.531
	(1.99)			(3.56)	(-0.26)	(-0.98)	(1.49)	(1.37)		0.198	0.323
	0.136			0.024	0.001	-0.018			0.069	0.968	0.536
	(14.72)			(3.80)	(0.35)	(-1.43)				0.225	0.313
\hat{p}_2	-0.050	0.001	-0.002	-0.028	0.016	-0.072			0.066	0.116	<0.001
	(-3.12)	(1.55)	(2.85)	(-2.63)	(2.56)	(-3.29)				0.230	<0.001
	-0.005			-0.026	0.018	-0.055	-0.007	<0.001	0.044	0.294	<0.001
	(-0.07)			(-2.34)	(2.49)	(-1.92)	(-0.55)	(0.20)		0.147	<0.001
	-0.061			-0.027	0.014	-0.052			0.038	0.279	<0.001
	(-3.66)			(-2.53)	(2.22)	(-1.96)				0.195	<0.001

* T-statistics in parentheses are calculated with standard errors corrected for heteroskedasticity. Δp_h is the hourly change in the log spot exchange rate (DM/\$) times 10,000. Δx_h is the hourly interdealer order flow, measured contemporaneously with Δp_h (negative for net dollar sales, in thousands). a_h is the number of macroeconomic announcements, σ_h is the standard deviation of all the transactions prices, and n_h is the number of transactions, all during hour h . τ is a vector of three dummy variables, $[\tau_1 \tau_2 \tau_3]$. τ_1 equals one for hours between 6:00 am and 7:59 am, zero otherwise; τ_2 equals one for hours between 8:00am and 11:59am, zero otherwise; and τ_3 equals one for hours between 12:00 pm and 1:59 pm, zero otherwise. \mathbf{p}_j is the derivative of the estimated function $\hat{\mathbf{p}}(\cdot)$ with respect to j th variable. The nonlinear function $\mathbf{p}(\cdot)$ is estimated non-parametrically by Kernel regression (estimated by OLS). 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.

Table 3: Nonparametric (Kernel) Regression of Order Flow Equation

$$\Delta x_h = \mathbf{m}(\Delta x_{h-1}, \Delta p_{h-1}, \mathbf{s}_{h-1}, n_{h-1}, a_{h-1}) + \mathbf{h}_h^x$$

Dependent Variable									Diagnostics		
	Const.	Δx_{h-1}	Δp_{h-1}	a_{h-1}	n_{h-1}	\mathbf{s}_{h-1}	\mathbf{t}_h	\mathbf{t}_h^2	R ²	Serial	Hetero
$\hat{\mathbf{m}}_1$	0.146	0.001	-0.001	0.003	-0.008	-0.045			0.022	0.521	0.459
	(6.06)	(2.14)	(-1.30)	(0.26)	(-0.69)	(-1.65)				0.007	0.806
	-0.001			0.003	-0.013	-0.023	-0.009	0.028	0.013	0.543	0.417
	(-1.96)			(0.21)	(-0.96)	(-0.92)	(-0.11)	(1.89)		0.009	0.749
	0.139			0.005	-0.009	-0.035			0.010	0.593	0.408
	(6.56)			(0.38)	(-0.78)	(-1.50)				0.007	0.738
$\hat{\mathbf{m}}_2$	0.326	0.001	0.002	-0.029	0.019	-0.013			0.032	0.704	0.361
	(8.13)	(1.51)	(1.86)	(-1.78)	(1.18)	(-0.28)				0.178	0.943
	0.287			-0.029	0.021	-0.037	0.013	-0.001	0.011	0.437	0.443
	(2.36)			(-1.72)	(1.09)	(-0.86)	(0.61)	(-0.80)		0.154	0.951
	0.340			-0.029	0.020	-0.042			0.009	0.460	0.431
	(9.32)			(-1.74)	(1.28)	(-1.10)				0.156	0.951

* T-statistics in parentheses are calculated with standard errors corrected for heteroskedasticity. Δp_h is the hourly change in the log spot exchange rate (DM/\$) times 10,000. Δx_h is the hourly interdealer order flow, measured contemporaneously with Δp_h (negative for net dollar sales, in thousands). a_h is the number of macroeconomic announcements, σ_h is the standard deviation of all the transactions prices, and n_h is the number of transactions, all during hour h . τ is a vector of three dummy variables, $[\tau_1 \tau_2 \tau_3]$, see Table 2 for definition. \mathbf{m}_j is the derivative of the estimated function $\mathbf{m}(\cdot)$ with respect to j th variable. The nonlinear function $\mathbf{m}(\cdot)$ is estimated non-parametrically by Kernel regression (estimated by OLS). 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.

Table 4

Feedback Trading Model

$$\begin{aligned}\Delta x_h &= \Delta x_h^* + f \Delta p_h \\ \Delta p_h &= \mathbf{b}_{11} \Delta x_h^* + \mathbf{b}_{12} \Delta p_{h-1} + \mathbf{h}_h^p \\ \Delta x_h^* &= \mathbf{b}_{21} \Delta x_{h-1}^* + \mathbf{b}_{22} \Delta p_{h-1} + \mathbf{h}_h^x\end{aligned}$$


Price Change Equation Order Flow Equation

<u>Coefficient on</u>		
Δx_h^*	0.262 (14.93)	
Δp_{h-1}	-0.232 (-3.73)	0.190 (0.37)
Δx_{h-1}^*		0.147 (2.00)
$Var(\mathbf{h}_h)^{1/2}$	13.604 (14.64)	24.163 (30.52)
Feedback Parameter: f	-0.563 (-0.37)	


* T-statistics in parentheses. Estimated using GMM (see appendix for details).

Figure 1

Order Flow Shocks: High versus Average Announcement Flow



A: Price change responses




B: Order flow responses


Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation order flow shock when the number of announcements is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). Vertical axis units are variable standard deviations. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

Figure 2

Price Shocks: High versus Average Announcement Flow




A: Price change responses




B: Order flow responses

Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation price shock when the number of announcements is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). No confidence band is shown for the high transaction case in panel A. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

Figure 3
Order Flow Shocks: High versus Average Volume (# Trades)




A: Price change responses




B: Order flow responses

Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation order flow shock when the number of transactions is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). No confidence band for the high transaction case is shown in panel A for clarity. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

Figure 4
Price Shocks: High versus Average Volume (# Trades)




A: Price change responses




B: Order flow responses

Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation price shock when the number of transactions is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). No confidence band is shown for the high transaction case in panel A. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

Figure 5
Order Flow Shocks: High versus Average Price Volatility




A: Price change responses




B: Order flow responses

Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation order flow shock when price volatility is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). No confidence band for the high transaction case is shown for clarity. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

Figure 6
Price Shocks: High versus Average Price Volatility



A: Price change responses



B: Order flow responses

Notes: Average impulse response patterns and 95% confidence bands for a one standard deviation price shock when price volatility is one standard deviation above the average sample level (solid plot) and at the sample average (dashed plot). No confidence band for the high dispersion case is shown in panel A. Based on two-equation system examined in Tables 2 and 3 (hourly data). See appendix for computation details.

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.
- Balduzzi, P., E. Elton, and C. Green (2001), Economic news and bond prices, *Journal of Financial and Quantitative Analysis*, forthcoming.
- Bierens, H. (1983), Uniform consistency of kernel estimators of a regression function under general conditions, *Journal of the American Statistical Association*, 77 699-707.
- Carrera, J. (1999), Speculative attacks to currency target zones: A market microstructure approach, *Journal of Empirical Finance*, 6, 555-582.
- Cornell, B. (1982), Money supply announcements, interest rates, and foreign exchange, *Journal of International Money and Finance*, 1: 201-208.
- Davidson, R., and J. MacKinnon, (1993), *Estimation and Inference in Econometrics*, Oxford University Press.
- 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.
- Dominguez, K. (2001), The market microstructure of central bank intervention, typescript, University of Michigan, August.
- Dominguez, K., and J. Frankel (1993), *Does foreign-exchange intervention work?* Institute for International Economics, Washington, D.C.
- Easley, D., and M. O'Hara (1992), Time and the process of security price adjustment, *Journal of Finance*, 47: 577-606.
- 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 (2001), Portfolio balance, price impact, and secret intervention, NBER Working Paper 8356, July.
- Fleming, M., and E. Remolona (1999), Price formation and liquidity in the U.S. treasury market, *Journal of Finance*, 54: 1901-1915.
- Flood, M. (1994), Market structure and inefficiency in the foreign exchange market, *Journal of International Money and Finance*, 13: 131-158.
- Frankel, J., and K. Froot (1990), Chartists, fundamentalists, and trading in the foreign exchange market, *American Economic Review*, 80: 181-185.

- 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.
- Hardle W. (1990), *Applied Nonparametric Regression*, Econometric Society Monograph, Cambridge University Press.
- 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.
- Kandel, E., and N. Pearson (1995), Differential interpretation of public signals and trade in speculative markets, *Journal of Political Economy*, 103: 831-872.
- Kirilenko, A. (1997), Endogenous trading arrangements in emerging foreign exchange markets, typescript, International Monetary Fund.
- Killeen, W., R. Lyons, and M. Moore (2001), Fixed versus flexible: Lessons from EMS order flow, NBER Working Paper 8491, September.
- Koop, G., M. Pesarab, and S. Potter (1996), Impulse response analysis on nonlinear multivariates, *Journal of Econometrics* 74 pp. 119-147.
- Kyle, A. (1985), Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.
- 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, November (chapters at www.haas.berkeley.edu/~lyons).
- Robinson, P. (1983), Nonparametric estimators for time series, *Journal of Time Series Analysis*, 4, 185-207.