



American Finance Association

FX Trading and Exchange Rate Dynamics

Author(s): Martin D. D. Evans

Source: *The Journal of Finance*, Vol. 57, No. 6 (Dec., 2002), pp. 2405-2447

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/3094532>

Accessed: 24/11/2009 11:44

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

FX Trading and Exchange Rate Dynamics

MARTIN D. D. EVANS*

ABSTRACT

I examine the sources of exchange rate dynamics by focusing on the information structure of FX trading. This structure permits the existence of an equilibrium distribution of transaction prices at a point in time. I develop and estimate a model of the price distribution using data from the Deutsche mark/dollar market that produces two striking results: (1) Much of the short-term volatility in exchange rates comes from sampling the heterogeneous trading decisions of dealers in a distribution that, under normal market conditions, changes comparatively slowly; (2) public news is rarely the predominant source of exchange rate movements over *any* horizon.

THE ORIGINS OF NOMINAL EXCHANGE RATE DYNAMICS remain elusive. In particular, there is no widely accepted explanation for the sizable short- and medium-term movements in the dollar during the floating-rate period. More generally, theoretical models relating exchange rates to macroeconomic fundamentals are still outperformed by simple time-series models in forecasting spot rates over short- and medium-term horizons (Frankel and Rose (1995)).

This paper provides new perspective on the poor performance of exchange rate models by focusing on the information structure of trading between FX dealers in the spot market. I argue that equilibrium in this market at a point in time may be characterized by a distribution of transaction prices arising from dealers' trading decisions based on incomplete and heterogeneous information. I then develop an empirical model to study the equilibrium distribution of FX transaction prices. The paper undertakes this investigation with the aid of a new data set that details trading activity in the FX market. This analysis reveals a striking new perspective on the source of exchange rate movements over all horizons. In particular, I find that much of the short-term volatility in exchange rates comes from sampling the equilibrium distribution of transaction prices that, under normal market condi-

* Department of Economics at Georgetown University and the N.B.E.R. I am very grateful to Richard Lyons for the many suggestions he made on an earlier draft of this paper. I also thank the seminar participants at Vanderbilt University, the editor, Richard Green, and two anonymous referees for their valuable comments. Part of the research reported on here was conducted while I was visiting the Financial Markets Group at the L.S.E. and The Bank of England. I am grateful for the hospitality I received at both institutions and acknowledge the financial support of the National Science Foundation under Grant 26-3282, and Houblon-Norman Fund at The Bank of England. I also thank the staff at The Bank and Reuters for allowing me to collect the data used in the study. Any errors or omissions are my responsibility.

tions, changes comparatively slowly. I also find that public news is rarely the predominant source of long-term exchange rate movements, a result that contrasts with the assumptions of traditional macro models.

I begin by describing the institutional structure of the spot FX market. This market differs in several respects from other securities markets. In market microstructure terms, it is a decentralized multiple-dealer market in which three trading mechanisms operate simultaneously: one for direct interdealer trade, one for brokered interdealer trade, and one for non-bank customer-dealer trade. An important feature of this market structure is that direct interdealer trades and customer-dealer trades lack transparency. That is to say, the details of the trades (e.g., the bid and ask quotes, the amount, and direction of the trade) are observed only by the counterparties. This means that any information they convey diffuses slowly across the market. I argue that this lack of transparency (and differences between the mechanisms for direct and brokered interdealer trading) allows for the existence of an equilibrium price distribution without introducing arbitrage opportunities.

The information structure of the market also has implications for the dynamics of prices and the pattern of trade. I argue that the price distribution changes in response to the arrival of new information in two forms: common-knowledge (CK) news and non-common-knowledge (NCK) news. CK news is characterized by the simultaneous arrival of new information to all market participants and their homogeneous interpretation of its implications for equilibrium prices. Clearly, CK news shocks are akin to the public news shocks found in macro exchange rate models. I use the different terminology to stress the fact that the news must be heard simultaneously and must be homogeneously interpreted. NCK news can come from a private or public source. Customer orders represent an example of private NCK news to dealers from outside the market. Direct interdealer transactions are a private source of NCK news from inside the market. A macroeconomic announcement may be a source of public NCK news when there is no consensus among dealers about its implications for the exchange rate. The key difference between CK and NCK news is that CK news immediately shifts the whole distribution of transaction prices but has no effect on trading patterns. By contrast, NCK news affects both the transaction price distribution and the pattern of trading, as measured by interdealer order flow. This variable measures the direction of trade between dealers and is a proximate determinant of equilibrium prices in many trading models (see, e.g., Evans and Lyons (2002)).

Based on these observations, I develop an empirical model that can be used to examine the origins of exchange rate dynamics. I focus on direct interdealer trading that accounts for approximately 63 percent of total trading in major currency markets (see Bank for International Settlements (1998)). The model decomposes observed changes in transaction prices into three components: a CK news component, a NCK news component, and a sampling component. The former incorporates the traditional macro view of exchange rate dynamics in which all innovations in spot rates are attributable to the

arrival of public news in a CK framework. The second component identifies the price effects of NCK news from public sources, such as macroeconomic announcements, and private sources such as customer orders. This component identifies the degree to which asymmetric information affects equilibrium transaction prices. The sampling component arises from the fact that there is an equilibrium transaction price distribution at each point in time. The dispersion in this distribution reflects the heterogeneous trading decisions dealers make in a market that lacks transparency.

The model is estimated using a data set that details trading activity in the spot FX market over a four-month period, May 1 to August 31, 1996. These data are unique in that they provide information on trading between FX dealers around the world.¹ In particular, they allow us to track the pattern of trade and FX prices in the direct interdealer market on a transaction-by-transaction basis. As such, the data series constitute sequences of irregularly spaced observations on a continuous trading process. I describe how GMM estimation methods can be adapted to deal with the irregular spacing problem. Estimates are obtained from the data series sampled over a fixed five-minute observation window, allowing for the fact that the window may correspond to varying spans of "market time," the time scale at which market processes evolve at a constant rate.

The empirical results are based on trading data for the Deutsche mark/dollar (DM/\$), the most heavily traded currency pair. These data reveal that there is considerable variation in the intensity at which trading takes place between dealers. Some of the variation has a well-defined intraday pattern that appears consistent with dealers based in different geographical locations entering and leaving the market. However, actual trading intensity can differ greatly from the intraday pattern on any particular day. Transaction prices and order flow also display some interesting statistical characteristics. Price changes observed over a five-minute interval appear to be nonnormally distributed and display a significant degree of negative serial correlation. Order flow, by contrast, is positively autocorrelated and highly persistent.

The paper's central empirical results come from the model estimates and may be summarized as follows:

- The origins of exchange rate movements vary considerably according to the state of the market, measured by transaction intensity.

¹ They differ from the FX quotes shown on the screens of specialist information providers, such as Reuters, Telerate, Minex, and Quotron. These quotes represent indicative prices rather than the firm prices at which a dealer will enter into a transaction. Their relationship to the transaction price data used here is discussed in Evans (1997). There exists a large literature studying the quote data because this was the only source of market-wide information on FX trading. A partial list of papers includes: Engle, Ito, and Lin, (1990), Baillie and Bollerslev (1991), Bollerslev and Domowitz (1993), Dacorogna et al. (1993), Goodhart and Giugale (1993), Bollerslev and Melvin (1994), and Guillaume et al. (1994).

- Under normal market conditions, CK news accounts for approximately 15 percent of the variance in short-term price changes. The contribution of CK news rises to a peak of approximately 40 percent as trading intensity increases.
- The sampling component accounts for approximately 80 percent of the variance in short-term price changes under normal market conditions. When trading intensity is very high, the sampling component's contribution falls to 17 percent.
- Long-term price movements originate from both CK and NCK news. (By definition, long-term price movements are not affected by the sampling component.) When the transaction intensity is very low, more than 90 percent of the variance of permanent shocks to the price level are attributable to CK news shocks. When the intensity is high, the CK contribution falls below 20 percent. In these market states, approximately 80 percent of the variance in permanent price shocks comes from NCK news.
- NCK news affects transaction prices at least 20 minutes before inter-dealer order flow. Its peak effect on changing the price distribution occurs approximately 15 minutes after the shock and lasts for approximately 30 minutes. The strength of these effects increases with transaction intensity. When intensity is high, NCK news accounts for more than 50 percent of the variance in price changes over a 30-minute to two-hour horizon.
- NCK news has both temporary and permanent effects on the level of prices in all market states.

These results provide two new perspectives on exchange rate dynamics. First, they provide strong empirical support for the idea that equilibrium in the FX market is described by a distribution of prices rather than a particular price level. The existence of this distribution is key to understanding the short-term dynamics of exchange rates because so much of the variance in observed price changes is attributable to the sampling component. Second, they challenge the traditional macro view that emphasizes the role of public news as the primary source of exchange rate movements in a CK environment. CK news shocks are rarely the predominant source of exchange rate changes over both long or short horizons. Moreover, the contribution of NCK news to permanent price movements points to a source of exchange rate dynamics, over macro relevant horizons, that has been overlooked by traditional models. The concluding section of the paper discusses how these observations may lead to more empirically successful macro models of exchange rates.

The remainder of the paper is organized as follows. Section I describes the institutional structure of the spot FX market and the data set. Section II presents the FX trading model and describes the econometric techniques used to study the data. Section III presents the model estimates, and Section IV examines their implications. Section V concludes.

I. Market Structure and Data

The institutional structure of the spot FX market differs in several respects from other securities markets. In market microstructure terms, it is a decentralized multiple-dealer market in which multiple trading mechanisms operate simultaneously. This complex institutional structure has implications for the behavior of prices and the pattern of trade that I incorporate in the model.

A. Market Structure

Trading in the spot FX market takes place among dealers, and among dealers and non-bank customers. According to the Bank for International Settlements (1998), interdealer trading accounts for about 63 percent of total trading in major spot markets (the remaining 37 percent is among dealers and non-bank customers) and breaks into two transaction types, direct and brokered. Direct trading between dealers accounts for about 50 percent of interdealer trade and brokered trading accounts for about 50 percent.

Direct interdealer trades result from bilateral conversations between dealers typically over a sophisticated e-mail system run by Reuters (see Section I.B). A conversation is initiated when a dealer calls another dealer on the system to request a quote. Quote requests may also be accompanied by information on the size (but not the sign) of the desired trade when it differs from the standard amount of \$10 million. Users of the system 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). Thus, trade follows the so-called dealer protocol where quotes precede orders. The system allows a large number of dealer pairs to hold conversations at the same time, so many transactions can take place simultaneously. Importantly, details of each conversation, such as the quotes and the decision of the initiating dealer, are known to only the counterparties. They are never transmitted via the system to other dealers in the market. This lack of transparency means that any information conveyed by a transaction diffuses slowly across the market.

Interdealer trade also takes place through brokers. In recent years, electronic systems run by EBS and Reuters largely have supplanted traditional voice brokers.² These systems operate a continuous auction market to which dealers can submit limit or market orders that are matched electronically. They differ in several respects from direct interdealer trading. First, all dealers using the system can observe the most competitive buy and sell prices (known as “the touch”) that have been submitted, along with the quantity available at these prices. Second, all data displayed by the system is in anon-

² Between 1995 and 1998, the fraction of brokered trades passing through electronic systems increased from 14 percent to 40 percent (see Bank for International Settlements (1998)). Since then, much higher levels of trading activity have been observed on these systems, leading market participants to believe that they now dominate brokered interdealer trading. For further information on the Reuters electronic brokerage system, see Goodhart, Ito, and Payne (1996) and Payne (1999).

ymous form. So, it is not possible to identify the dealers with orders at "the touch." When orders are matched, the brokerage informs each dealer involved of the counterparty to their trade. Third, orders can be matched only between two dealers if the transaction falls within the credit line that they have already negotiated. This means that the best limit buy and sell prices displayed by the system are not the prices at which *any* dealer can trade. These differences make the direct and brokered dealing systems imperfect substitutes.³

Transactions between dealers and non-bank customers also take the form of bilateral conversations between the customer and the dealer's bank. The bank provides indicative quotes before the customer places an order to buy or sell a specified amount of foreign currency. The bank passes the customer order to the dealer who fills it at the current market price (which may differ from the indicative quote). Customer orders are a source of private information to banks and the dealers that work for them. Although they have no information regarding other banks' customer orders, FX dealers cite customer-dealer transactions as an important source of information (Lyons (2001)).

Several characteristics of this market structure are important for the analysis that follows. First, there is a lack of transparency in direct interdealer trading because details of trades are only observed by the counterparties. Second, direct interdealer trading can occur simultaneously between many dealer pairs. Third, brokered trading is not a perfect substitute for direct interdealer trading.

B. The Data

The analysis below uses new data on trading activity on the Reuters D2000-1 system in the DM/\$ spot FX market over a four-month period, May 1 to August 31, 1996. This is the most widely used direct electronic dealing system. According to Reuters, more than 90 percent of the world's direct interdealer transactions took place through the system during the sample period. Excluding weekends and a feed interruption caused by a power failure, there are 79 full trading days in the sample and 255,497 trades. Trades on the Reuters D2000-1 system take the form of the bilateral conversations described above. Every time an electronic conversation on D2000-1 resulted in a trade, a customized feed provided the database with a time-stamped record of the transaction price, a bought or sold indicator.

³ The posting of a quote to a brokerage system differs from making a firm direct quote in two respects. First, it provides a wider advertisement of a willingness to trade than bilateral direct quoting. Second, the brokerage system provides pretrade anonymity to the dealer. There are also differences from the perspective of a dealer viewing quotes. A bilateral direct quote may be for trade sizes beyond the quantities posted on the brokerage system. The bilateral quote also is firm in the sense that the dealer stands ready to provide the credit line necessary to trade at that price. By contrast, there is no guarantee that a dealer can trade at the best limit prices displayed by some brokerage systems; he may not have a credit line with the dealer posting the limit order.

Two features of the data are particularly noteworthy. First, they provide transaction information for the whole interbank market over the full 24-hour trading day. This contrasts with earlier transaction data sets covering single dealers over some fraction of the trading day (e.g., Lyons (1995), Yao (1997a, 1997b), and Bjornes and Rime (1998)). The data set makes it possible, for the first time, to analyze trading patterns and prices at the level of “the market.” The only other multiple-dealer data set in the literature covers brokered interdealer transactions (the electronic system examined by Goodhart et al. (1996) and Payne (1999)). The system they examine, however, accounts for only a small fraction of daily trading volume.⁴

Second, the data cover a relatively long time span (four months) compared with other micro data sets. This span provides a vast number of minute-by-minute observations on trading activity across a wide variety of “market states.” In the analysis that follows, I use the data collected between 00:00:01 hours British Summer Time (BST) on Monday to 24:60:60 hours BST on Friday. (All time is measured relative to BST, which corresponds to Greenwich Mean Time plus one hour.) This time interval appears to span the week of trading in the DM/\$ fairly well. Although the D2000-1 system runs 24 hours a day, seven days a week, it rarely records trades outside this interval.

Although the data has greater coverage and span than other dealer data sets, it also has some limitations. In particular, Reuters could not provide information on the identity of the counterparties involved in each trade for confidentiality reasons. Consequently, the analysis below focuses on aggregate trading behavior rather than the actions of individual dealers. The database also lacks information on the size of each trade. I therefore identify aggregate interdealer trading patterns below using the number rather than the value of transactions. However, this may not represent a significant loss of information on trading patterns across the market. Jones, Kaul, and Lipson (1994) find evidence that the size of trades has no information content for measuring trading patterns beyond that contained in the number of transactions.

II. The FX Trading Model

A. The Model

The model uses data on: (1) transaction prices, (2) interdealer order flow, and (3) trade intensity from direct interdealer trading in the DM/\$. Transactions take two forms in the data. If a dealer initiating a conversation buys dollars, the transaction price equals the ask quote in DMs per dollar offered by the other dealer. I refer to this as the DM purchase price for dollars, p_t^a . If the dealer initiating a conversation sells dollars, the transaction price will equal the bid quote given by the other dealer. I refer to this as the DM sale

⁴ There also is evidence that dealers attach more informational importance to direct interdealer order flow than to brokered interdealer order flow; see Bjornes and Rime (1998).

price for dollars, p_t^b . Thus, the designation of a transaction price as a purchase or sale price depends on who initiates the transaction. Buyer-initiated trades take place at the purchase price while seller-initiated trades take place at the sale price. Order flow measures the direction of trade as the net of buyer-initiated orders and seller-initiated orders. Interdealer order flow is positive (negative) when a dealer initiating a bilateral conversation purchases (sells) foreign exchange at the ask (bid) quote. In the model, I focus on aggregate direct interdealer order flow, x_t , defined as the difference between the number of buyer-initiated orders and seller-initiated orders per period. Trade intensity, n_t , is defined as the number of trades per period.

Dealers receive information in two forms: common-knowledge (CK) news and non-common-knowledge (NCK) news. CK news is characterized by the simultaneous arrival of new information to all market participants and their homogeneous interpretation of its implications for equilibrium prices. In the FX market, CK news could come from brokers or in the form of a macroeconomic announcement. In the former case, all dealers would have to subscribe to the brokerage system and hold the same view of the relationship between the information displayed and market-wide transactions prices from direct interdealer trade. In the case of macroeconomic announcements, the news must be heard simultaneously and must be homogeneously interpreted. Dealers have access to NCK news from both public and private sources. When the news comes from a public source, it is simultaneously heard by all dealers, but they interpret its implications for equilibrium prices differently.⁵ Many macroeconomic announcements may fall into this category because there is little consensus among practitioners or researchers on the link between exchange rates and macroeconomic fundamentals.⁶ Direct interdealer transactions and customer orders are examples of NCK news that come from private sources.

One important implication of the market's structure is that dealers do not have CK information about direct interdealer trading. This lack of transparency in direct interdealer trading has an important implication for the behavior of transaction prices: At any moment, equilibrium in the market may be characterized by a distribution of prices for direct interdealer transactions. Standard models (see, e.g., Perraudin and Vitale (1995), Lyons (1997), and Evans and Lyons (2002)) rule out this possibility by assuming that each dealer quotes a publicly observed price at which she will trade any amount with any number of other dealers. In equilibrium, this leads dealers to quote a common price to avoid being arbitrated. In actual direct dealing, however,

⁵ Brennan and Cao (1997) provide one theoretical rationale for this phenomena. They study a model where investors hold different priors about the value of assets. When public news arrives, better-informed investors change their valuation of the asset by more than the less well-informed investors.

⁶ Consistent with this view, Cheung and Wong (2000) report survey data showing that market participants have diverse opinions on the importance of economic factors in the determination of exchange rates at different horizons.

a dealer can quote a price to only one other dealer at a point in time. This restriction, and the lack of transparency, makes it possible for different dealers to quote different prices without opening themselves to arbitrage.⁷

To formalize this idea, the model assumes that equilibrium in the market at a point in time is described by a distribution of purchase prices and a distribution of sales prices. Let p_t^a and p_t^b denote observed prices drawn randomly from the respective distributions of purchase and sales prices at time t . These observed prices are related to the average transaction price, p_t , by

$$p_t^o = p_t + \omega_t^o, \quad (1)$$

for $o = \{a, b\}$. The idiosyncratic shocks ω_t^a and ω_t^b identify the degree to which observed prices differ from the market-wide average. Their size depends on the identity of the dealers whose prices we observe. The model assumes that observed prices are drawn randomly and independently from the cross-sectional distributions of purchase and sale prices every period so that ω_t^a and ω_t^b are serially uncorrelated and independently distributed. This is a reasonable assumption in a market where there are a large number of dealers who can execute transactions at any time.

Shifts in the price distributions come from CK and NCK news. As in many trading models, including the multiple-dealer models of Lyons (1997), Evans and Lyons (2002), and Evans (2001), CK news shocks have an immediate and one-for-one effect on all transaction prices, but no impact on order flow. Specifically, I assume that CK news, ε_t , and NCK news, v_t , change average prices and aggregate interdealer order flow according to

$$\Delta p_t = B(L)v_t + \varepsilon_t, \quad (2)$$

and

$$x_t = C(L)v_t, \quad (3)$$

where $B(L)$ and $C(L)$ are polynomials in the lag operator, L .⁸ The CK and NCK news shocks, ε_t , and v_t , are assumed to be mutually independent, serially uncorrelated, and independent from the idiosyncratic shocks, $\{\omega_t^a, \omega_t^b\}$.

⁷ Of course, the degree to which quotes differ at any point in time will depend on several factors, including the heterogeneity of customer order flows, dealer risk aversion, and the degree to which dealers view brokered trading as a substitute for direct trading. In Evans (2001), I present a model where heterogeneity in the customer orders received by different dealers, combined with the lack of transparency, leads to an equilibrium distribution of FX prices at which direct interdealer trading takes place.

⁸ To derive an estimable model, I assume that the order flow process (3) is invertible so that NCK news shocks represent fundamental innovations to order flow (see Hamilton (1994)).

The intuition behind the effects of CK news, ε_t , on prices and order flow is straightforward. In the face of a CK news shock, any dealer failing to move quotes one-for-one will find himself subject to large inventory imbalances as other dealers take advantage of the misalignment between the quoted prices and their new valuation of the asset (at the margin). Thus, the risk associated with large inventory imbalances (due to unexpected future price changes) creates an incentive for all dealers to move their quotes one-for-one. This leaves the difference between the quote and the value of the asset to the initiating dealer unchanged. As a result, the distribution of transaction prices and quotes shift together because all transactions take place at the quoted prices, and there is no change in order flow. Equations (1) and (2) imply that CK news shocks have an immediate and one-for-one impact on all transactions prices.⁹ Equation (3) shows interdealer order flow driven solely by NCK shocks.¹⁰

The effects of NCK news, v_t , on prices and order flow are more complex. NCK news will generally affect both prices and order flow with the exact dynamics, represented by $B(L)$ and $C(L)$ in (2) and (3), depending on the complex interaction of dealers. This is most easily understood by considering some illustrative examples. Suppose dealer A receives a customer purchase order for \$10 million of foreign currency that the dealer fills from his existing inventory. If the dealer believes that the order contains no price-relevant information (see below), he may respond in one of three ways. He can (1) simply wait for another customer with an offsetting sell order, (2) actively replenish his inventory by either initiating a direct bilateral trade or by submitting an order to a brokerage system, or (3) passively replenish his inventory by raising the prices he quotes when another dealer initiates a bilateral trade. If he follows option (1), the customer order has no impact on interdealer order flow or transaction prices.¹¹ Under options (2) and (3), the dealer's actions generally will lead to order flow and price movements. In

⁹ The specification of (2) also assumes that all CK shocks permanently affect the level of prices. This is not a completely innocuous assumption, because, as macro models show, it is possible for public news to have a temporary effect on spot rates. However, given that the estimated speed of mean reversion in spot rates following such shocks is typically measured in months or longer, and (2) is designed to describe price changes over intervals of a few minutes, the assumption is reasonable from an empirical standpoint.

¹⁰ The assumed absence of CK shocks from the order flow equation mirrors the identification scheme used by Hasbrouck (1991) in his study of specialist quotes and order flow on the NYSE. While Hasbrouck notes that specialist quote smoothing, the reporting of stale quotes, and the presence of stale limit orders could all invalidate this assumption in practice, none of these problems apply to data I use here. Another possibility is that lagged CK shocks affect order flow through the presence of feedback traders because their contribution to order flow depends systematically on past prices. There is, however, no direct empirical evidence in the literature of feedback trading in the FX market. Thus, the absence of CK shocks from (3) represents a weaker assumption in the context of the interdealer FX market than elsewhere.

¹¹ This option is unlikely to be part of an optimal trading strategy unless customer order flow is negatively autocorrelated. Empirical evidence in Lyons (1995) and Bjonnes and Rime (1998) indicates that dealers unwind inventory positions rather quickly.

particular, if dealer A initiates a bilateral trade with dealer B and buys \$10 million of foreign currency at B's ask quote, A's inventory imbalance will have been passed to B. At the same time, the transaction will have contributed \$10 million to interdealer order flow. If dealer B decides to replenish his inventory passively according to option (3), and dealer C is willing to sell him \$10 million of foreign currency at his higher bid quote, this transaction will lead to higher prices and negative interdealer order flow (because dealer C initiated the transaction). In sum, therefore, dealer A's customer order will have generated positive and then negative interdealer order flow and a higher transaction price. To see how an alternative response pattern could arise, suppose dealer A decided to replenish his inventory passively according to option (3). If dealer B initiates bilateral trade with A and sells him \$10 million of foreign currency at his higher bid quote, the higher transaction price will be accompanied by negative order flow. Subsequent interdealer order flow will depend on dealer B's actions. For example, if B follows option (2) and passes his inventory imbalance to dealer C, who just happens to have a customer wanting to sell \$10 million of foreign currency, the interdealer order flow following the price change will be positive.

These examples illustrate how one source of NCK news, a customer order, can generate different response patterns for prices and interdealer order flow according to the actions taken by a sequence of dealers attempting to manage their inventories. Examples of response patterns can be similarly constructed for cases where customer orders are the source of NCK information about the future path of prices. This occurs when a dealer believes that his customer order signals a shift in aggregate customer demand for foreign currency. For example, Lyons (1997) notes that orders from customers engaged in international trade can be private signals on the shift in demand for foreign currency due to trade flows because the orders are received in advance of trade statistics.¹² Under these circumstances, a dealer receiving a customer order will revise his view about the prices at which he can trade with dealers and customers in the future. The model in Evans (2001) describes how customer orders carrying price-relevant NCK news affect prices and interdealer order flow (and produces specific forms for $B(L)$ and $C(L)$). Macroeconomic announcements also can be a source of NCK news to dealers because they hold differing views about their implications for equilibrium prices. In this case, the announcement will lead dealers to change their quotes and desired inventory positions by differing amounts. This will trigger order flow and price changes until inventories have been redistributed among dealers in a manner consistent with their new views about equilibrium prices.

¹² Foreign exchange intervention by central banks represents another example of price-relevant NCK news. Because few central banks deal directly in the interbank market, most interventions are made by placing orders with FX dealers at one or more commercial banks (Peires (1997)). Such an order would represent private information to the dealer about future interest rates, say, and hence, the future path of prices.

Equations (1) to (3) can be used to examine the source of price movements. Combining (1) and (2) gives the period-by-period change in the observed price as

$$\Delta p_t^o = \varepsilon_t + B(L)v_t + \omega_t^o - \omega_{t-1}^o, \quad (4)$$

for $o = \{a, b\}$. This equation decomposes the period-by-period change in purchase or sales prices into three components: (1) the CK news component ε_t , (2) the NCK news component $B(L)v_t$, and (3) the sampling component $\omega_t^o - \omega_{t-1}^o$. The first component incorporates the traditional macro view of exchange rate determination in which all innovations in spot rates are attributable to the arrival of public news in a CK framework (i.e., all information relevant for exchange rate determination is CK, and the mapping from information to equilibrium prices is also CK). The second component identifies the price effects of NCK news from public sources, such as macro announcements, and private sources, such as customer orders. The sampling component arises from the presence of an equilibrium distribution of transaction prices at a point in time: the purchase price distribution and the sales price distribution.

Although the three components of price changes are mutually independent, they cannot be separately identified from equation (4) alone. To derive an estimable model for price changes, I, therefore, combine (3) and (4) to give

$$\begin{bmatrix} \Delta p_t^a \\ \Delta p_t^b \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} D(L)x_t + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t + \begin{bmatrix} \omega_t^a - \omega_{t-1}^a \\ \omega_t^b - \omega_{t-1}^b \end{bmatrix}, \quad (5)$$

where $D(L) = B(L)C(L)^{-1}$. The first two terms on the right-hand side represent the dynamics of the average price change, Δp_t , which are common to both the purchase and sales price changes. The third term is the sampling component of observed price changes that arises from the presence of the transaction price distribution. Notice, too that the polynomial $D(L)$ may take many forms depending on the dynamic responses of prices and order flow to NCK shocks (via $B(L)$ and $C(L)$). In particular, if prices respond more quickly to NCK shocks than order flow, $D(L)$ may contain both leads and lags of L . In the empirical analysis that follows, I will estimate versions of (5) with quite general specifications for $D(L)$. The parameter estimates thus obtained will then be used to decompose price changes into their various components.

Equations (3) and (5) provide an empirical framework for studying the mechanism through which information becomes embedded in prices. In particular, it allows us to differentiate between the effects of CK and NCK news. This framework does not inform us about the subject of the news. The subject of CK news may be a traditional "macro fundamental" such as output,

money supplies, or interest rates, or it may be something quite different, such as a surprise political announcement. The only requirement of the model is that all market participants receive the news simultaneously and interpret its implications for prices homogeneously. Similarly, the subject of NCK news may be a traditional macro fundamental or a political development when the source is public, or customer orders when the source is private. Equations (3) and (5) allow us to examine the information transmission mechanism without specifying the subject matter.

B. Estimation

The spot FX market is open continuously in the sense that dealers can trade with one another via the D2000-1 system 24 hours a day, seven days a week. The system also allows conversations between many dealer pairs to take place at the same time. Thus, it is possible for multiple trades to be concluded, and recorded, at the same instant. The resulting data set constitutes a sequence of irregularly spaced observations on a continuous trading process. At some points in the sample, the gaps between successive trades span many minutes, while at others, several trades appear with the same second-by-second time stamp. I will not attempt to directly model these irregular timing patterns in the analysis below. Instead, I will use prices, order flow, and trade intensity measured relative to a fixed five-minute observation interval. Hence, p_t^a and p_t^b are, respectively, the last dollar purchase and sale price recorded during interval t ; x_t is the difference between the number of buyer-initiated and seller-initiated trades during interval t ; and n_t is the number of transactions during interval t .

One drawback to adopting a fixed observation interval is that there are periods of the day when no transactions take place during the interval. I designate the price and order flow observations from these periods as "missing." All the statistics and estimates reported below are calculated without the use of these observations. For example, in computing the first order autocorrelation coefficient in the Δp_t^a series, I use only consecutive observations on Δp_t^a and Δp_{t-1}^a for which none of the values for p_t^a was "missing." More generally, I employ the GMM estimation procedure described below.

All the statistics and empirical models considered below can be written in the state-space form:

$$\begin{aligned}\xi_t &= A\xi_{t-1} + \zeta_t, \\ y_t &= C\xi_t,\end{aligned}\tag{6}$$

where ξ_t is a q -dimensioned state vector, and y_t is an r -dimensioned vector of observed variables. ζ_t is a q -dimensioned vector of shocks with zero means that are uncorrelated with ξ_{t-1} , serially uncorrelated, and have covariance matrix Ω . Although the form of A , C , and Ω vary according to the particular application, in all cases the eigenvalues of A lie inside the unit circle so that

ξ_t and y_t follow stationary processes. Hence, (6) implies that the unconditional means of ξ_t and y_t are respectively equal to a q and r -dimensioned vector of zeros. Equation (6) also implies that the covariance of the states, $\Gamma(k) \equiv \text{Cov}(\xi_t, \xi'_{t-k})$, is computed as $\Gamma(k) = A\Gamma(k-1)$ with $\Gamma(0) = \text{vec}^{-1}[(I - A \otimes A)^{-1} \text{vec}(\Omega)]$. The covariance of the observed variables is therefore given by

$$\text{Cov}(y_t y'_{t-k}) \equiv \gamma(k) = C\Gamma(k)C'. \quad (7)$$

In some applications, I also make use of a j -dimensioned vector of instruments, z_t , with the property $\text{cov}(y_t, z'_{t-i}) = 0$ for $i \geq 0$.

Let θ represent the vector of parameters to be estimated. As in the standard GMM case, I consider orthogonality conditions of the form

$$E[m_t(k; \theta)] = 0, \quad (8)$$

where

$$m_t(k; \theta) = \mathcal{D}(k) \begin{bmatrix} \text{vec}(y_t z'_{t-k}) \\ \text{vec}(y_t y'_{t-k} - \gamma(k; \theta)) \end{bmatrix}$$

for $k = 0, 1, \dots, K$. The vector $\mathcal{D}(k)$ is a vector of ones and zeros that selects the moments to be included in $m_t(k; \theta)$. Equation (8) gives a maximum of $rj + r^2$ independent conditions when $k > 0$, and $rj + r(r+1)/2$ conditions when $k = 0$.

To compute the GMM estimates, let $m_t(\theta) = [m_t(0; \theta), m_t(1; \theta) \dots m_t(K; \theta)]'$ be a vector of selected moment conditions. While all the elements of $m_t(\theta)$ can be computed for any period t , if a particular element involves a value for y_t or y_{t-k} designated as a "missing" observation, the result also is designated "missing." This holds true irrespective of the value of θ so the set of "missing" elements in $m_t(\theta)$ will not vary with θ for a particular t . Let $\Lambda = \{t_1^*, t_2^*, \dots, t_T^*\}$ denote the set of observations for which none of the elements in $m_t(\cdot)$ is missing. The estimates of θ are found by minimizing

$$Q(\theta) = m_{T^*}(\theta)' W^{-1} m_{T^*}(\theta), \quad (9)$$

where

$$m_{T^*}(\theta) = \frac{1}{T^*} \sum_{\Lambda} m_t(\theta),$$

with T^* equal to the number of observations in Λ . I follow the standard practice of first setting the weighting matrix W equal to the identity to obtain consistent estimates of θ . These estimates, $\tilde{\theta}$, then are used to compute

a consistent estimate of the optimal weighting matrix, $\tilde{W}$. The form of this weighting matrix varies across applications according to whether elements of $m_t(\theta)$ are serially correlated under the null hypothesis of a correctly specified model. The most general weighting matrix I consider follows the form proposed by Newey and West (1987). The GMM estimates, $\hat{\theta}$, are found by minimizing (9) with $W = \tilde{W}$. The asymptotic covariance matrix of the resulting estimates is $\hat{V} = [\hat{G}\tilde{W}^{-1}\hat{G}']^{-1}$, where $\hat{G} = \partial m_{T^*}(\hat{\theta})/\partial \theta'$.

Several facets of this estimation technique may be illustrated by considering its application to (5) for the case of purchase prices, p_t^a . In particular, suppose that $D(L) = d_1L^{-1} + d_2 + d_3L$, so that

$$\Delta p_t^a = d_1x_{t+1} + d_2x_t + d_3x_{t-1} + \varepsilon_t + \omega_t^a - \omega_{t-1}^a. \quad (10)$$

In this case, $y_t = \Delta p_t^a - d_1x_{t+1} - d_2x_t - d_3x_{t-1}$, $\xi_t' = [\varepsilon_t, \omega_t^a, \omega_{t-1}^a]$, and $\zeta_t' = [\varepsilon_t, \omega_t^a, 0]$, with

$$C = [1 \quad 1 \quad -1], \quad A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad \text{and} \quad \Omega = \begin{bmatrix} \Sigma_\varepsilon & 0 & 0 \\ 0 & \Sigma_\omega & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

The GMM estimates of $\theta = [d_1, d_2, d_3, \Sigma_\varepsilon, \Sigma_\omega]$ can be found with instruments $z_t = [1, x_{t+1}, x_t, x_{t-1}]$ and the moments

$$m_t(\theta) = [y_t \quad y_t x_{t+1} \quad y_t x_t \quad y_t x_{t-1} \quad y_t^2 - \gamma(0; \theta) \quad y_t y_{t-1} - \gamma(1; \theta) \quad \dots \quad y_t y_{t-k} - \gamma(k; \theta)].$$

It is clear in this application that the GMM technique does not necessarily provide the most efficient parameter estimates. For example, if ε_t and ω_t^a were normally distributed and the time series for y_t contained no missing observations, θ would be most efficiently estimated by maximum likelihood with the aid of the Kalman Filter. However, balanced against this, the GMM technique offers two important advantages. First, it does not require any distributional assumptions regarding the error processes: There is no reason to think that ε_t and ω_t^o for $o = \{a, b\}$ are normally distributed, for example. The second advantage stems from the presence of missing observations. Although the Kalman Filtering algorithm can be extended to deal with one or two missing observations (Harvey (1989)), dealing with many adds considerably to the complexity of the filter and makes estimation computationally burdensome. By contrast, the GMM technique can deal with many missing observations easily. Moreover, because the data set spans four months, the adoption of the five-minute observation interval provides us with a very




Figure 1. Average trade intensity. The figure shows the average number of direct inter-dealer transactions per minute over the 79 trading days in the sample plotted over 24 hours.

large number of (nonmissing) observations. As a consequence, the GMM estimates of θ are extremely precise judged by the metric of their asymptotic distribution.¹³

III. Results

A. Sample Statistics

Although the D2000-1 system runs 24 hours a day, the vast majority of interdealer transactions in the DM/\$ are concentrated during the European trading hours. This institutional feature gives rise to recurrent intraday patterns in the data. Exemplifying this phenomenon, Engle et al. (1990), Baillie and Bollerslev (1991), Bollerslev and Domowitz (1993), Goodhart and Giugale (1993), Payne (1997), and Andersen and Bollerslev (1998) all have studied the intraday patterns in the volatility of indicative quotes.

There is also a pronounced intradaily pattern in transactions. Figure 1 plots the average trade intensity during each five-minute observation interval calculated over the 79 trading days in the sample. As the figure shows, average trading activity follows a three-humped pattern. The first hump

¹³ Of course, the large number of observations in the data set does not guarantee that the asymptotic distribution of the GMM estimates is close to its true distribution in the sample. Although the sample period contains observations from a wide range of market states (see below, and Evans (1997)), it is possible that some states occur less (or more) frequently in the sample than is implied by the true distribution of the data. If the behavior of prices and order flow in these states differ significantly from their behavior in other states, the GMM estimates will suffer from a "Peso-problem;" see Evans (1996) for a survey.

occurs between 01:00 hours and 05:00 hours, with peak intensity of approximately five trades per minute.¹⁴ The second hump begins around 07:00 hours, peaks with approximately 25 trades per minute at 09:30 hours, and drops to 15 trades per minute by about 12:30 hours. Trading during this period is dominated by dealers situated in Europe. The third hump begins at 12:30 hours, and rises quickly to a peak of approximately 30 trades per minute at 14:30 hours. This increase coincides with U.S.-based dealers entering the market. Activity remains around 25 trades per minute until approximately 16:00 hours, when European based dealers typically leave the market. Thereafter, there is a gradual decline in activity with fewer than five trades per minute taking place after 19:00 hours.

It is important to stress that Figure 1 plots average trade intensity over the sample of 79 trading days. Actual intensity can vary considerably from day to day. For example, the most active period of trading in the sample occurs between 09:00 hours and 10:00 hours when more than 200 transactions per minute take place, eight times the average for that time of day. More generally, the difference between the actual and average intensity during each five-minute interval has a sample standard deviation of 58, so that trading intensity on any particular day could differ significantly from the pattern displayed in Figure 1.

Sample statistics for the change in purchase price, Δp_t^a , and order flow, x_t , are reported in Table I.¹⁵ From the statistics in the right-hand columns of the upper panel, the unconditional distributions for both variables appear nonnormal. The distribution for transaction price changes is skewed to the left and is highly leptokurtic. The distribution for order flow is also fat-tailed, but skewed to the right. The lower panel reports estimated autocorrelation coefficients together with the p -values for the null hypothesis of a zero coefficient. In the case of purchase prices, these estimates indicate the presence of a MA(1) process for Δp_t^a : There is a significant negative coefficient at lag one, while all coefficients at higher lags are insignificantly different from zero. By contrast, order flow displays positive autocorrelation that is statistically significant at the one percent level on lags one through six.

Further evidence on the dynamics of price changes and order flow is provided by Table II. Here, I report GMM estimates for various ARMA specifications using the variance and the first 12 autocorrelations as moments (see the Appendix). The number of overidentifying restrictions for each set of estimates is reported in the right-hand column together with the Hansen (1982) J -statistic and its associated p -value.

¹⁴ Trading in the DM/\$ is comparatively light during this period because only Asian-based dealers are typically in the market. It is also possible that the Reuters dealing system is used much less by Asian-based dealers than their counterparts in Europe and the United States.

¹⁵ The statistics for the difference in sales prices, Δp_t^b , are almost identical to those for Δp_t^a , and so are not reported to conserve space. All statistics on Δp_t^a are calculated without missing observations. The Appendix provides details on all the empirical results presented below.

Table I
Sample Statistics

The price is 100 times the last DM purchase price for dollars on the Reuter's D2000-1 system during observation interval t . Order flow, x_t , is the difference between the number of buyer-initiated and seller-initiated trades during observation interval t . The autocorrelations are computed by GMM. The p -values are calculated from Wald tests of the null hypothesis of a zero correlation allowing for conditional heteroskedasticity (see the Appendix for details).

	Mean	Maximum	Minimum	Std. Dev.	Skewness	Kurtosis			
Δp_t^a	0.000	0.500	−0.790	0.076 0.008	−0.194 0.020	7.291 0.039			
x_t	0.005	69.000	−72.000	5.211 0.007	0.102 0.017	14.260 0.034			
Autocorrelations (p -values)									
	lag								
	1	2	3	4	5	6	12	18	24
Δp_t^a	−0.319 (0.000)	−0.014 (0.170)	−0.005 (0.650)	0.002 (0.858)	−0.006 (0.539)	0.004 (0.732)	0.005 (0.607)	0.020 (0.037)	0.001 (0.899)
x_t	0.232 (0.000)	0.105 (0.000)	0.092 (0.000)	0.077 (0.000)	0.060 (0.000)	0.058 (0.000)	0.025 (0.007)	0.027 (0.005)	0.005 (0.578)

The upper panel of the table reports estimates of ARMA models for the change in purchase price. Consistent with the statistics in Table I, there is strong evidence of a moving average component in this series: The moving average coefficient(s) are highly significant in all but one model. Although the MA(1) model appears well specified when judged by the J -statistic, the autoregressive coefficients also appear to be significant in the ARMA(1,1) and ARMA(2,2) models. This is an interesting finding, because it provides some preliminary evidence that order flow contributes to price changes. Recall from (4) that observed price changes can be written as $\Delta p_t^o = \varepsilon_t + B(L)v_t + \omega_t^o - \omega_{t-1}^o$, where ε_t , v_t , and ω_t are serially uncorrelated. According to this equation, price changes should follow an MA(1) process if $B(L)v_t = 0$. However, the estimates in Table II suggest that price changes follow a higher-order process. According to (4), this must originate from the NCK news component, $B(L)v_t$.

ARMA model estimates for order flow are reported in the lower panel of Table II. These estimates indicate that order flow follows either an ARMA(1,2), ARMA(2,1), or ARMA(2,2) process with a high degree of persistence. Judged by the J -statistics, lower order models do not appear to capture the high degree of persistence in the process implied by the autocorrelations reported in Table I. For example, the estimates of the ARMA(1,2) model imply a root for the order flow process of 0.8417.

To establish an empirical benchmark for the results that follow, Table III examines the relationship between price changes and interest rates in daily

Table II
ARMA Models

The table reports GMM estimates and standard errors for ARMA models of the form

$$z_t = a_1 z_{t-1} + a_2 z_{t-2} + w_t + b_1 w_{t-1} + b_2 w_{t-2},$$

where $Ew_t = 0$ and $Ew_t^2 = \sigma^2$. The models are estimated from the mean, variance, and first 12 autocorrelations of the data (see the Appendix). The GMM estimates of the models for price changes and order flow respectively use data from 9,726 and 13,647 intervals to compute the sample moments. The right-hand columns report the results of Hansen (1982) J -tests for each specification. The column headed DF reports the degrees of freedom associated with each test. The variables are: $\Delta p_t^a = p_t^a - p_{t-1}^a$, where p_t^a is 100 times the last DM purchase price for dollars on the Reuter's D2000-1 system during observation interval t , and x_t is the difference between the number of buyer-initiated and seller-initiated trades during observation interval t .

	Coeffs.					J-stat.	p-value	DF
	a_1	a_2	b_1	b_2	σ^2			
Δp_t^a	-0.2347 (0.0084)				0.0728 (0.0006)	53.8420	0.0000	11
			0.3079 (0.0122)		0.0719 (0.0007)	4.9999	0.9312	11
	0.0768 (0.0269)		0.3903 (0.0308)		0.0717 (0.0007)	1.8596	0.9973	10
	0.0809 (0.0874)	0.0012 (0.0246)	0.3945 (0.0902)		0.0717 (0.0007)	1.8587	0.9935	9
	0.0972 (0.4022)		0.4108 (0.4041)	-0.0063 (0.1244)	0.0717 (0.0007)	1.8587	0.9935	9
	-1.0129 (0.2873)	0.0898 (0.0321)	-0.7000 (0.2882)	0.4305 (0.1066)	0.0717 (0.0007)	1.1029	0.9975	8
x_t	0.2602 (0.0100)				5.7451 (0.0615)	107.4869	0.0000	11
			-0.2241 (0.0128)		5.7527 (0.0612)	155.4307	0.0000	11
	0.7572 (0.0166)		0.5953 (0.0229)		5.8144 (0.0601)	32.9568	0.0003	10
	0.9851 (0.0348)	-0.1133 (0.0176)	0.7755 (0.0300)		5.8769 (0.0604)	8.9239	0.4443	9
	0.8417 (0.0176)		0.6282 (0.0207)	0.0959 (0.0141)	5.8771 (0.0604)	7.5449	0.5806	9
	0.7663 (0.1356)	0.0577 (0.1023)	0.5531 (0.1348)	0.1411 (0.0802)	5.8762 (0.0605)	7.3681	0.4975	8

data—"the standard macro approach." The dependent variable in each of the four regression models is the change in the log of the last DM purchase price for dollars before 16:00 hours on day τ , and 16:00 hours on day $\tau - 1$. The independent variables are the overnight interest rates for the dollar, i_τ , and Deutsche mark, i_τ^* , reported by Datastream at approximately 16:00 hours on day τ .

Table III
Daily Models

The table reports OLS estimates for regressions of the daily changes in the log spot rate. The dependent variable, $\Delta \ln p_t^a$, is the change in the log of the last DM purchase price for dollars before 16:00 hours on day τ , and 16:00 hours on day $\tau - 1$. The variables i_τ and i_τ^* are overnight interest rates for the dollar and Deutsche mark as reported by Datastream at approximately 16:00 hours on day τ . Standard errors are reported in parenthesis corrected for conditional heteroskedasticity.

Model	Constant	$\Delta(i_\tau - i_\tau^*)$	Δi_τ	Δi_τ^*	$i_{\tau-1} - i_{\tau-1}^*$	$i_{\tau-1}$	$i_{\tau-1}^*$	R^2	SEE
I		0.621 (0.767)						<0.001	0.004
II			0.592 (0.760)	-1.744 (1.395)				0.015	0.004
III	-0.027 (0.013)				-1.331 (0.649)			0.052	0.004
IV	-0.048 (0.055)					-1.257 (0.646)	1.696 (1.209)	0.055	0.004

According to many standard macro models, the level of the real exchange rate should be proportional to the real interest differential (Frankel and Rose (1995)). These models imply that innovations in the nominal interest differential should account for much of the daily change in nominal exchange rates (because daily rates of actual and expected inflation are small and stable). Row I of Table III empirically examines this implication using the daily change in the nominal interest differential, $\Delta(i_\tau - i_\tau^*)$, to proxy for the innovation in the interest differential. Although the positive coefficient on $\Delta(i_\tau - i_\tau^*)$ is correctly signed, it is not statistically significant.¹⁶ Moreover, as the R^2 statistic indicates, changes in the interest differential account for very little of the variance in daily exchange rate changes. Row II reports similar results when changes in interest rates, Δi_τ and Δi_τ^* , replace $\Delta(i_\tau - i_\tau^*)$. The lower portion of the table reports results based on uncovered interest parity (UIP), relating the daily change in the log spot rate to the previous day's interest differential. In row III, the coefficients on the interest differential are negative and statistically significant at the five percent level and the R^2 statistics are approximately five percent. According to these estimates, the nominal interest differential has a limited ability to forecast changes in the log spot rate over the next 24 hours, but the direction of the forecast is inconsistent with UIP. Row IV reports similar results for the case where lagged interest rates are included separately.

The results in Table III are broadly consistent with findings in the empirical exchange rate literature based on standard macro models using monthly and quarterly data (for surveys, see Frankel and Rose (1995) and Engel

¹⁶ The sign should be positive because, in the sticky-price monetary model, for example, an increase in the dollar interest rate requires an immediate dollar appreciation (i.e., increase in DM/\$) to make room for the expected dollar depreciation required by uncovered interest parity.

(1996)). They show that the “standard macro approach” meets with roughly the same degree of empirical success in explaining the daily behavior of exchange rates in this data sample as it has in other contexts.¹⁷

B. Model Estimates

I now turn to examine the origins of exchange rate movements using the FX trading model presented in Section III. Table IV reports GMM estimates of the univariate version of equation (5) for 10 different specifications for $D(L)$, the polynomial on order flow. The dependent variable in all cases is the change in purchase price, Δp_t^a . (Results using the change in sale price, Δp_t^b , are nearly identical.) The table reports the coefficient estimates on each of the order flow terms, together with their standard errors, which are corrected for conditional heteroskedasticity and an MA(1) error process. Inspection of these estimates reveals that the coefficients on x_{t-2} , x_{t-3} , x_{t+5} , and x_{t+6} are (individually) insignificantly different from zero at the five percent level whenever they are included in a specification. This evidence suggests that $D(L)$ is well characterized by a sixth order polynomial (i.e., one that includes terms in L^{-4} , L^{-3} , L^{-2} , L^{-1} , L^0 , and L). As row VII of the table shows, all the coefficient estimates in this specification are statistically significant. The importance of the leading order flow terms is further emphasized by the results in the last row. Here, the table reports estimates for the case where $D(L)$ includes only L^0 , L , L^2 , and L^3 terms. The R^2 statistic from this specification is 0.005, much lower than the R^2 statistics for the other specifications. Thus, order flow contributes most to the predictable variation in price changes through the leading terms.

The right-hand columns of the table report the estimated sum of the order flow coefficients together with some regression diagnostics. In every case, the sum of the coefficients is positive and statistically significant. Since the dependent variable is the change in price, these estimates imply that order flow variations have a permanent effect on the price level. The right-hand column of the table reports l -test statistics (Cumby and Huizinga (1992)) for the null hypothesis that the regression residuals follow an MA(1) process. For the preferred specification in row VII, the l -statistic is significant at the nine percent level.

While the results in Table IV provide strong evidence on the price impact of order flow, more can be learned about the origins of price movements from the bivariate version of (5). Intuitively, Δp_t^a and Δp_t^b represent two

¹⁷ It is worth reiterating that Table III provides an empirical benchmark for the results that follow rather than definitive evidence on the empirical performance of macro exchange rate models in the sample. Although interest rate innovations are the main engine of exchange rate variations in some macro models (cf. Dornbusch (1976)), they are clearly an incomplete measure of innovations in all the macro fundamentals that could have occurred. I use interest rates because they are available daily (unlike other standard macro fundamentals, such as real output, nominal money supplies, etc.) and because their relevance is cited most frequently in surveys of foreign exchange dealers; see Cheung and Chinn (1999).

Table IV
Decomposition Regressions

The table reports GMM estimates of the coefficients in the polynomial, $D(L)$, for the regression:

$$\Delta p_t^a = D(L)x_t + \varepsilon_t + \omega_t^a + \omega_{t-1}^a,$$

where p_t^a is 100 times the last DM purchase price for dollars on the Reuter's D2000-1 system during observation interval t , and x_t is the difference between the number of buyer-initiated and seller-initiated trades during observation interval t . The GMM estimates and standard errors allow for the presence of conditional heteroskedasticity and an MA(1) error structure (see the Appendix). The column headed $D(L)$ reports the sum of the estimated coefficients in $D(L)$ and its standard error. The right-hand column reports Cumby-Huizinga l -test statistics for the null hypothesis that the errors follow an MA(1) process. The associated p -values are reported in parentheses.

	$D(L)$ Coefficients ($\times 100$)											Diagnostics			
	$D(L)$														
	x_t	x_{t-1}	x_{t-2}	x_{t-3}	x_{t+1}	x_{t+2}	x_{t+3}	x_{t+4}	x_{t+5}	x_{t+6}	$D(1)$	R^2	SEE	l -test	
I	-0.134 (0.012)	-0.018 (0.011)			0.200 (0.016)						0.048 (0.018)	0.028	7.479	2.045 (0.360)	
II	-0.140 (0.012)	-0.028 (0.011)	0.001 (0.013)		0.167 (0.016)	0.157 (0.012)					0.156 (0.021)	0.044	7.364	2.219 (0.330)	
III	-0.145 (0.012)	-0.031 (0.011)	-0.002 (0.014)	-0.001 (0.011)	0.165 (0.016)	0.140 (0.013)	0.085 (0.012)				0.210 (0.021)	0.051	7.331	4.359 (0.113)	
IV	-0.146 (0.012)	-0.031 (0.011)	-0.004 (0.014)	-0.002 (0.011)	0.162 (0.016)	0.139 (0.013)	0.072 (0.013)	0.052 (0.011)			0.242 (0.023)	0.053	7.307	6.046 (0.049)	
V	-0.146 (0.012)	-0.031 (0.011)	-0.003 (0.014)	-0.002 (0.011)	0.161 (0.016)	0.138 (0.013)	0.072 (0.013)	0.050 (0.012)	0.011 (0.011)		0.249 (0.023)	0.054	7.299	5.287 (0.071)	
VI	-0.146 (0.012)	-0.031 (0.011)	-0.005 (0.014)	0.000 (0.011)	0.161 (0.016)	0.138 (0.013)	0.072 (0.013)	0.050 (0.012)	0.011 (0.011)	0.002 (0.010)	0.251 (0.024)	0.055	7.289	4.429 (0.109)	
VII	-0.145 (0.012)	-0.033 (0.011)			0.159 (0.016)	0.139 (0.013)	0.071 (0.012)	0.055 (0.011)			0.247 (0.021)	0.051	7.318	4.794 (0.091)	
VIII	-0.146 (0.012)	-0.033 (0.011)			0.159 (0.016)	0.138 (0.013)	0.071 (0.012)	0.052 (0.012)	0.012 (0.011)		0.253 (0.022)	0.052	7.315	6.315 (0.043)	
IX	-0.145 (0.012)	-0.033 (0.011)			0.159 (0.016)	0.138 (0.013)	0.071 (0.013)	0.051 (0.012)	0.012 (0.011)	0.003 (0.010)	0.256 (0.023)	0.053	7.297	6.046 (0.049)	
X	-0.091 (0.013)	-0.013 (0.012)	0.017 (0.014)	0.014 (0.011)							-0.073 (0.017)	0.005	7.559	0.075 (0.963)	

Table V
Estimates of Bivariate Model

The table reports GMM estimates of the Bivariate model:

$$\begin{bmatrix} \Delta p_t^a \\ \Delta p_t^b \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} D(L)x_t + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t + \begin{bmatrix} \omega_t^a - \omega_{t-1}^a \\ \omega_t^b - \omega_{t-1}^b \end{bmatrix},$$

where w_t^a , w_t^b , and ε_t are mutually independent and serially uncorrelated shocks with $E\omega_t^a = \omega^a$, $E\omega_t^b = \omega^b$, $E\varepsilon_t = 0$, and $\text{Var}(\varepsilon_t) = \Sigma_\varepsilon$, $\text{Var}(\omega_t^a) = \text{Var}(\omega_t^b) = \Sigma_\omega$. p_t^a and p_t^b are respectively 100 times the last DM purchase and sales price for dollars on the Reuters D2000-1 system during observation interval t . Order flow x_t is the difference between the number of buyer-initiated and seller-initiated trades during observation interval t . Asymptotic standard errors corrected for heteroskedasticity and serial correlation are reported below the parameter estimates (see the Appendix). The GMM procedure uses observations from 11,473 intervals to compute the sample moments. The table also reports Hansen's J -statistic with its associated p -value in parentheses.

	x_{t+4}	x_{t+3}	x_{t+2}	x_{t+1}	x_t	x_{t-1}
Coefficients in $D(L)$ ($\times 100$)	0.0494 (0.0103)	0.0700 (0.0114)	0.1488 (0.0112)	0.1613 (0.0156)	-0.1461 (0.0107)	-0.0374 (0.0101)
$(\Sigma_\varepsilon)^{1/2}$	0.0391	(0.0011)	$(\Sigma_\omega)^{1/2}$	0.0433	(0.0004)	
$D(1)$	0.2461	(0.0219)	J -statistic	7.6044	(0.8684)	

noisy observations on the behavior of Δp_t , with noise coming from independent sources (i.e., the sampling components from the purchase and sale price distributions). The estimates of (5) should, therefore, provide more precise information on the dynamics of Δp_t than the estimates above utilizing just Δp_t^a . To estimate the model, I drew on the results in Table IV by assuming that $D(L)$ is a sixth order polynomial containing terms in L^{-4} to L . I also assumed that ω_t^a and ω_t^b have the same variance, Σ_ω , and constant means.¹⁸ The GMM estimates were obtained using instruments $z_t = [x_{t+4}, \dots, x_{t-1}]$ and moments derived from the covariances of observed price changes.

The GMM estimates of (5) are reported in Table V. The order flow coefficient estimates are generally similar to those reported in row VII of Table IV, and are all highly statistically significant. The sum of the estimated coefficients is 0.2461, with a standard error of 0.0219. This estimate also is similar to the one in Table IV and implies a strong rejection of the null hypothesis that customer order shocks only have temporary effects on prices. The estimates of $\Sigma_\varepsilon^{1/2}$ and $\Sigma_\omega^{1/2}$, the standard deviations of CK news, and the idio-

¹⁸ The latter assumption implies that there is no period-by-period change in the difference between the average purchase and sales price, $\bar{p}_t^a - \bar{p}_t^b$. This is not a completely innocuous assumption because dealers can change the spread between their bid and ask quotes, which in turn could alter $\bar{p}_t^a - \bar{p}_t^b$. However, spreads in the interdealer market are extremely small (e.g., Lyons (2001) reports a median spread of DM 0.0003/\$ for the dealer he studied), so keeping the means of ω_t^a and ω_t^b constant is not unreasonable.

syncratic shocks are rather similar. Thus, the cross-sectional dispersion of transaction prices appears to be of the same order of magnitude as the variance of CK shocks.

These estimates provide empirical support for the existence of an equilibrium distribution of transaction prices at a point in time. While I argued above that an equilibrium distribution can arise from heterogeneity of dealers' information and a lack of transparency in the market, are other interpretations possible? In particular, couldn't the estimates simply imply the presence of some sort of measurement error?¹⁹ In one sense, the answer clearly is no. The data on p_t^a and p_t^b come from the Reuters trading system and constitute part of the audit trail. There is no doubt that transactions took place at exactly the prices recorded in the p_t^a and p_t^b series. Of course, it is possible that dealers make typing errors or irrational decisions when using the trading system. The prices p_t^a and p_t^b will then contain measurement errors in the sense that they do not correspond to equilibrium transaction prices in any theoretical model that excludes these possibilities. However, even in this case, the assumptions of the model are valid so long as the ω_t^o shocks are serially uncorrelated. Thus, while it may be possible to attribute some of the price distribution implied by the estimates to economic factors that have not been the focus of theoretical trading models, it cannot be attributed even partially to pure measurement error in the data.

C. Robustness

The estimates reported in Tables IV and V are based on the dynamics of price changes and order flow measured over a fixed five-minute observation window. As a check for robustness, I also estimated models with a $2\frac{1}{2}$ minute window.²⁰ Based on this window, I found that the lead-lag relationship between order flow and price changes covers the same time span as the results in Table IV. I also found that the estimates of $D(1)$, $\Sigma_\varepsilon^{1/2}$, and $\Sigma_\omega^{1/2}$ were consistent with those reported in Table V. Overall, the results in Tables IV and V appear reasonably robust to the choice of the fixed five-minute observation window.

Time deformation represents another factor that may influence the results in Tables IV and V. Time deformation occurs when a fixed interval measured in minutes corresponds to varying spans of "market time": the time scale at which market process evolves at a constant rate (Stock (1988)). If this time scale is grounded in the rate at which information becomes available to dealers, and the arrival of NCK information is an important source of interdealer trade, then the variable trading intensities we observed in Figure 1 may be signaling the presence of time deformation that affects

¹⁹ Zhou (1996) attributes the negative serial correlation displayed by indicative FX quote changes to measurement "noise" and bid-ask bounce. Bid-ask bounce cannot play a role here because we are considering time series composed solely of purchase prices, Δp_t^a , and sales prices, Δp_t^b .

²⁰ A Supplemental Appendix (available at <http://www.georgetown.edu/evansmd/datapage.htm>) reports these estimates together with other ancillary results referred to below.

Table VI
Tests for State Dependency

See Table V for the definitions of Δp^a , Δp^b , and x . The upper panel reports Wald tests for the null hypothesis of zero coefficients on all the terms listed at the head of each column in models of the form

$$z_t = D(L)x_t + D(L)x_t n_t + D(L)x_t n_t^2 + D(L)x_t n_t^3 + w_t.$$

In each case, the model including $D(L)x_t$ and the listed terms was estimated by GMM allowing for heteroskedasticity and an MA(1) error structure. These estimates then are used to construct the Wald test. The p -value is reported below each statistic. In cases where the change in price is the dependent variable, $D(L)$ takes the form of the sixth-order polynomial in the Bivariate Model. For the case of order flow, $D(L) = d_1 L + d_2 L^2 + \dots d_6 L^6$. The lower panel reports the results of tests for heteroskedasticity. The center three columns report Glesjer (1969) tests for heteroskedasticity in the variance of each shock using the variables listed at the head of each column. The shock v_t is the innovation to the ARMA(2,2) model for order flow estimated in Table II above, while ε_t , ω_t^a , and ω_t^b are the shocks from the Bivariate Model. The right-hand column reports LM statistics for first-order ARCH. In all cases, p -values for the null hypothesis of homoskedasticity are shown in parentheses.

Variable	Nonlinearity					
	$D(L)x_t n_t$	$D(L)x_t n_t$	$D(L)x_t n_t^2$	$D(L)x_t n_t$	$D(L)x_t n_t^2$	$D(L)x_t n_t^3$
Δp_t^a	97.844 (<0.001)		92.988 (<0.001)		166.284 (<0.001)	
Δp_t^b	167.426 (<0.001)		252.563 (<0.001)		296.390 (<0.001)	
x_t	11.254 (0.081)		16.868 (0.158)		28.132 (0.060)	
Shock	Heteroskedasticity					
	n	n	n^2	n	n^2	n^3
ε_t	545.127 (<0.001)	609.454 (<0.001)		662.451 (<0.001)		123.525 (<0.001)
ω_t^a	6.272 (0.013)	6.471 (0.039)		14.234 (0.003)		17.981 (<0.001)
ω_t^b	2.542 (0.117)	5.017 (0.081)		25.759 (<0.001)		41.626 (<0.001)
v_t	4,705.604 (<0.001)	5,152.894 (<0.001)		5,188.820 (<0.001)		2,097.581 (<0.001)

the estimates in Table V. To examine this possibility, I used trade intensity as a state variable to test for the presence of state dependency in the dynamics of price changes and order flow.

The upper panel of Table VI reports the results of the state-dependency tests for the dynamics of price changes and order flow. In the former case, I estimated models for price changes that include the distributed lag of order flow $D(L)x_t$, together with the interaction terms $D(L)x_t n_t$, $D(L)x_t n_t^2$, and

$D(L)x_t, n_t^3$. I then tested the null hypothesis that the estimated coefficients on all the included interactions terms are zero using a Wald test corrected for the presence of heteroskedasticity and an MA(1) residual error structure. As the table shows, these test statistics strongly reject the null. The Wald tests for state dependency in the order flow dynamics show much weaker evidence against the null; none is significant at the five percent significance level. These tests are based on AR(6) models for order flow (i.e., $D(L) = d_1L + \dots + d_6L^6$), but the results are robust to the use of higher-order models.

The lower panel of Table VI reports tests for heteroskedasticity in the variances of ε_t , ω_t^a , and ω_t^b . The center three columns report Glesjer (1969) tests for heteroskedasticity using combinations of n_t , n_t^2 , and n_t^3 . As the table shows, there is strong evidence against the null of homoskedasticity in all three cases. The right-hand column reports LM statistics for first order ARCH (Engle (1982)) for each of the three shocks. These statistics also imply a rejection of the homoskedastic null at very high significance levels. The table also reports tests for heteroskedasticity in the variance of the innovations in order flow. The tests use the innovations from the estimates of the ARMA(2,2) order flow model in Table III. All four test statistics are highly significant.

Table VII provides evidence of the degree of variability in trade intensity. The upper panel of this table reports estimates of a six-state first-order Markov process for trade intensity. The states are defined as $\{j: n \in [S_j, S_{j+1})\}$ for $j = 1, \dots, 6$ where S_j is the lower bound for state j reported in the table. Thus, an observation of $n_t = 10$, would represent state $j = 3$ of the Markov process. The body of the table reports the matrix of estimated transition probabilities, with entry i, j denoting the probability of transition from state i to state j . Two features of these estimates are noteworthy. First, the probability of remaining in the same state, reported on the leading diagonal, is less than 75 percent in every case. Second, the probability of leaving the current state is highest at intermediate levels of trading activity. These features are more pronounced in the lower panel of the table, which reports estimates of a six-state Markov process for the "deseasonalized" trading rates. In this case, states are defined as $\{j: n - \bar{n} \in [S_j, S_{j+1})\}$, where $\bar{n}$ denotes the average trade intensity for the interval from which n is observed (see Figure 1). These estimates indicate that unusually high or low trade intensities are more likely to persist than intensities that are closer to the norm for that particular time of day. They also show that trade intensities can vary over the complete range of states from period to period. In contrast to the upper panel, all the off-diagonal transition probabilities are nonzero.

Overall, the results in Table VI strongly indicate that the dynamics of price changes vary significantly with the state of the market as measured by trading intensity. Moreover, as Table VII shows, there are considerable variations in trade intensity over the sample period. Some of these variations can be attributed to a fairly well-defined intraday pattern shown in Figure 1 that appears consistent with dealers in different locations entering and leaving the market. However, on any particular day, actual trade intensity can vary considerable from this norm. While these findings do not constitute

Table VII
Markov Models for Trade Intensities

The table reports estimates of the transition probabilities for a first-order, six-state Markov process for trade intensities. Panel A reports estimates based on the raw intensities, while panel B shows estimates for the deseasonalized intensities, where the latter are computed as $n_t - \bar{n}_t$, with $\bar{n}_t$ denoting the sample average rate for observation period t (plotted in Figure 1). The table reports the lower bounds that define the six states, and the estimated unconditional (ergodic) probability of each state occurring. The transition probabilities are estimated as the relative frequency that a particular transition occurred over the sample.

Panel A: With Seasonals						
State	Transition Probabilities					
	States					
	1	2	3	4	5	6
1	0.743	0.188	0.028	0.001	0.000	0.001
2	0.215	0.568	0.209	0.008	0.001	0.000
3	0.040	0.237	0.574	0.301	0.087	0.022
4	0.001	0.006	0.160	0.496	0.453	0.201
5	0.000	0.001	0.019	0.126	0.274	0.245
6	0.000	0.000	0.010	0.068	0.186	0.532
Ergodic probabilities	0.210	0.244	0.283	0.160	0.053	0.050
Lower bounds	0	1	4	14	26	36

Panel B: Without Seasonals						
State	Transition Probabilities					
	States					
	1	2	3	4	5	6
1	0.725	0.110	0.068	0.108	0.105	0.036
2	0.108	0.497	0.266	0.182	0.081	0.034
3	0.066	0.265	0.466	0.278	0.129	0.047
4	0.060	0.104	0.167	0.297	0.288	0.164
5	0.021	0.016	0.020	0.086	0.232	0.221
6	0.019	0.007	0.013	0.049	0.165	0.498
Ergodic probabilities	0.248	0.251	0.250	0.150	0.050	0.050
Lower bounds	-28.672	-3.644	-0.974	1.775	8.472	16.582

direct evidence of time deformation, they certainly are consistent with the idea that speed of market processes varies through time.

D. The State-Dependent Model

The results above point to the need to incorporate state dependency into the model. To this end, I consider an extension of (5) in which $D(L)$ is replaced by $D(L, n_t)$, a state-dependent sixth order polynomial

$$D(L, n) = d_1(n)L^{-4} + d_2(n)L^{-3} + \dots + d_5(n) + d_6(n)L, \quad (11)$$

with state-dependent coefficients $d_i(\cdot)$. As above, ε_t , ω_t^a , and ω_t^b are mutually independent and serially uncorrelated shocks, but their variances are now state dependent: $\text{Var}(\varepsilon_t) = \Sigma_\varepsilon(n)$ and $\text{Var}(\omega_t^a) = \text{Var}(\omega_t^b) = \Sigma_\omega(n)$. I model state dependency in the coefficients and variances as

$$\begin{aligned} d_i(n) &= d_i(0)\exp(-n/\gamma) + d_i(\infty)(1 - \exp(-n/\gamma)), \\ \Sigma_i(n) &= \Sigma_i(0)\exp(-n/\gamma) + \Sigma_i(\infty)(1 - \exp(-n/\gamma)), \end{aligned} \tag{12}$$

where $d_i(0)$, $d_i(\infty)$, $\Sigma_i(0)$, and $\Sigma_i(\infty)$ are all parameters to be estimated. These functional forms make $d_i(n)$ and $\Sigma_i(n)$ smooth monotonic functions of the transaction rate and are similar to the transition functions used in nonlinear time series models (Potter (1999)). They bound the coefficients between $d_i(0)$ and $d_i(\infty)$, and the variances between $\Sigma_i(0)$ and $\Sigma_i(\infty)$ as the transaction rate varies between zero and infinity. The positive scaling parameter γ governs the rate at which $d_i(n)$ and $\Sigma_i(n)$ vary with the transaction rate (a larger value for γ reduces the effect of a given change in n). For the sake of parsimony, the value of γ is common to all functions.

Table VIII reports the GMM estimates of the state-dependent model with γ set equal to 100. Attempts to estimate the model with γ unrestricted gave estimates close to 100, but the standard errors on all the other parameters were much larger than the values reported in the table. (Reestimating the model with γ set to different values of 90 and 110 had negligible effects on the estimated parameters.) The upper portion of the table reports the estimated bounds in the order flow polynomial, $D(L, n)$. When compared to the estimates in Table V, we see that the estimates of $d_i(0)$ are generally smaller in absolute value than their state-independent counterparts, while the estimates of $d_i(\infty)$ are generally larger. The estimated range for the individual $d_i(\cdot)$ functions appears quite large. A Wald test for the null hypothesis of $d_i(0) = d_i(\infty)$ for $i = \{1, 2, \dots, 6\}$, reported in the lower panel, is highly significant. This test statistic supports the presence of state dependency in the price change dynamics. The $d_i(\cdot)$ estimates also imply significant state dependence in $D(1, n)$, which measures the long run impact of order flow on the price level. The estimated lower and upper limits are -0.1546 and 1.3169 , with standard errors of 0.0300 and 0.1020 as n ranges from zero to infinity.

The center of the table reports estimates for the variance parameters. In models where all four parameters were left unrestricted, the estimates of $\Sigma_\varepsilon(0)$ and $\Sigma_\omega(\infty)$ were very close to zero (i.e., <0.0001), so the table reports estimates where these parameters are restricted to zero. With these restrictions, $\text{Var}(\varepsilon_t) = \Sigma_\varepsilon(\infty)(1 - \exp(-n/100))$ and $\text{Var}(\omega_t^a) = \text{Var}(\omega_t^b) = \Sigma_\omega(0)\exp(-n/100)$. The estimated value for $\Sigma_\omega(0)$ implies that the standard deviation for the idiosyncratic shocks slowly falls from 0.047 to approximately 0.006 as n varies from 0 to 200 . The estimate for $\Sigma_\varepsilon(\infty)$ implies that the standard deviation of CK shocks is smaller than $\text{Var}(\omega_t)^{1/2}$ for n less

Table VIII
Bivariate Model with State-Dependency

The model takes the form

$$\begin{bmatrix} \Delta p_t^a \\ \Delta p_t^b \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} D(L, n_t) x_t + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t + \begin{bmatrix} \omega_t^a - \omega_{t-1}^a \\ \omega_t^b - \omega_{t-1}^b \end{bmatrix},$$

where ω_t^a , ω_t^b , and ε_t are mutually independent and serially uncorrelated shocks with $E\omega_t^a = \omega^a$, $E\omega_t^b = \omega^b$, $E\varepsilon_t = 0$. The variances are given by $\text{Var}(\varepsilon_t) = \Sigma_\varepsilon(n_t)$, $\text{Var}(\omega_t^a) = \text{Var}(\omega_t^b) = \Sigma_\omega(n_t)$, with $\Sigma_\varepsilon(n) = \Sigma_\varepsilon(0)\exp(-n/100) + \Sigma_\varepsilon(\infty)(1 - \exp(-n/100))$. The state-dependent polynomial is $D(L, n_t) = d_1(n_t)L^{-4} + d_2(n_t)L^{-3} + \dots + d_6L$. Variables are defined in Table V. Panel A shows GMM estimates and asymptotic standard errors corrected for heteroskedasticity and serial correlation (see the Appendix). The GMM procedure uses observations from 11,473 intervals to compute the sample moments. Panel B reports the Hansen J -statistic, a Wald statistic for the null that $D(L, 0) = D(L, \infty)$, and LM-type statistics for misspecification state-dependent polynomial and variances (see the Appendix). The lower portion of the table reports autocorrelations, and standard errors, in the estimated standardized shocks.

Panel A: Estimates							
	$\mathbf{x}_{t+4}$	$\mathbf{x}_{t+3}$	$\mathbf{x}_{t+2}$	$\mathbf{x}_{t+1}$	$\mathbf{x}_t$	$\mathbf{x}_{t-1}$	
Coefficients in $D(L,0)$	0.0195	0.0154	0.0075	−0.0579	−0.0963	−0.0428	
(Std. Errs.) (×100)	(0.0169)	(0.0160)	(0.0167)	(0.0182)	(0.0171)	(0.0165)	
Coefficients in $D(L,\infty)$	0.1575	0.2555	0.5384	0.6424	−0.2697	−0.0072	
(Std. Errs.) (×100)	(0.0659)	(0.0567)	(0.0563)	(0.0594)	(0.0426)	(0.0517)	
	$\Sigma_\varepsilon(0)^{1/2}$	$\Sigma_\varepsilon(\infty)^{1/2}$	$\Sigma_\omega(0)^{1/2}$	$\Sigma_\omega(\infty)^{1/2}$	$D(1,0)$	$D(1,\infty)$	
	0.000	0.0900	0.0479	0.0000	−0.1546	1.3169	
	(N/A)	(0.0005)	(0.0001)	(N/A)	(0.0300)	(0.1020)	
Panel B: Diagnostics							
	Statistic					p -value	
J -test for overidentifying restrictions	9.315					(0.968)	
Wald test for $D(L,0) = D(L,\infty)$	255.247					(<0.001)	
LM Test for misspecification in $D(L,n)$	0.726					(0.999)	
LM Test for misspecification in $\Sigma_\varepsilon(n)$	0.128					(0.721)	
LM Test for misspecification in $\Sigma_\omega(n)$	0.094					(0.760)	
Panel C: Residual Autocorrelations (std. errs.)							
Residual	lag = 1	2	3	4	5	6	12
$\varepsilon_t^2/\Sigma_\varepsilon(n_t)$	0.1051 (0.0222)	0.0087 (0.0120)	0.0243 (0.0098)	0.031 (0.0148)	0.024 (0.0158)	0.0312 (0.0191)	0.0034 (0.0081)
$(\omega_t^a)^2/\Sigma_\omega(n_t)$	0.0605 (0.0233)	0.0152 (0.0138)	0.0174 (0.0171)	−0.0085 (0.0111)	0.0008 (0.0095)	0.02 (0.0121)	−0.0023 (0.0094)
$(\omega_t^b)^2/\Sigma_\omega(n_t)$	0.0543 (0.0230)	0.0219 (0.0097)	0.0103 (0.0109)	0.0023 (0.0103)	0.0035 (0.0086)	0.0037 (0.0096)	0.0170 (0.0112)

than 35. As the transactions rate rises beyond 35, the standard deviation increases slowly from 0.042 toward 0.090.²¹

The lower portion of the table reports the results of various diagnostic tests. The Hansen (1982) J -test for the overidentifying restrictions of the model has a p -value of 0.968. The table also reports the results of LM-type tests for misspecification in the estimated $D(L, n)$, $\Sigma_\varepsilon(n)$, and $\Sigma_\omega(n)$ functions. None of these statistics is statistically significant. This suggests that the model did manage to incorporate most of the state dependency in price dynamics. The model is less successful in accounting for all the heteroskedasticity exhibited by price changes. The bottom of the table reports autocorrelations for the standardized estimates of the shocks. If the model had completely captured the heteroskedasticity in price changes, these autocorrelations should all be close to zero. As the table shows, this is not the case. In particular, there appears to be significant first-order serial correlation in the estimates of $\varepsilon_t^2/\Sigma_\varepsilon(n_t)$.

Overall, the results in Table VIII show that the state-dependent dynamics of price changes are reasonably well characterized by the state-dependent model. While the model does not identify *all* the sources of heteroskedasticity in prices, it does appear to capture the role played by changing trade intensity.

IV. Origins of Price Changes

The results above allow us to examine the origins of price changes in several different ways. In particular, we can (1) study the dynamic response of prices to CK and NCK shocks, (2) decompose the variance of observed price changes into different theoretical components, and (3) examine the sources of seasonality in price heteroskedasticity.

Recall from Section II.A that equilibrium order flow is represented as $x_t = C(L)v_t$, where v_t is a CKN news shock. Although this process also could be state dependent with the coefficients in $C(L)$ functions of the transaction rate, the test statistics in the upper panel of Table VI suggest that order flow follows a reasonably stable process. Thus, as an empirical matter, we can substitute $C(L)v_t$ for x_t in the state-dependent model (5) to give

$$\Delta p_t^o = B(L, n_t)v_t + \varepsilon_t + \omega_t^o - \omega_{t-1}^o, \quad (13)$$

where $B(L, n) = D(L, n)C(L)$ for $o = \{a, b\}$.

²¹ I also compared the estimates of the sampling variance implied by the GMM estimates, $\hat{\Sigma}_\omega(n_t)$, with measures of the cross-sectional dispersion in prices calculated directly from the data. The latter measures are computed as the sample variance of all purchase prices (denoted $\hat{\Omega}_t^a$) or sales prices (denoted $\hat{\Omega}_t^b$) during each five-minute observation interval t . The regression evidence reported in the Supplemental Appendix shows that, while there is a statistically significant relationship between the dispersion measures and sampling variance, they are not linked tightly.




Figure 2. Impulse response functions. Panels A, B, and C plot $B(L, n) = D(L, n)C(L)$ where $D(L, n)$ is the estimated state-dependent polynomial on order flow from Table VIII and $C(L)$ is the polynomial implied by the estimated ARMA model for order flow in Table II. Panel D plots $B(L, n) = D(L)C(L)$ where $D(L)$ is estimated the polynomial on order flow from Table V (without state dependence). The solid, dashed, and dotted lines, respectively, show the impulse response based on ARMA(2,2), ARMA(2,1), and ARMA(1,2) order flow models.

Equation (13) provides us with the means to empirically examine the origins of price changes. In particular, the state-dependent polynomial $B(L, n)$ identifies the impulse response of prices to a one unit NCK news shock while trade intensity remains at n . To calculate these impulse responses, I first used the estimated ARMA models for order flow in Table II to calculate the coefficients in $C(L)$. I then combined these values with the estimates of $D(L, n)$ from Table VIII to compute $B(L, n)$. Panels A, B, and C of Figure 2 report the results for $n = 5, 20$, and 40 . As there is little to choose among the ARMA(2,1), ARMA(1,2), and ARMA(2,2) estimates for order flow, the figure plots the impulse responses implied by each specification. For completeness, panel D plots the impulse responses implied by the estimates of $D(L)$ in Table V (i.e., without state dependency).

Figure 2 displays three noteworthy features. First, the impulse responses appear fairly robust to the choice of ARMA specification used in calculating $C(L)$. Second, the dynamic response of prices seems to vary considerably with trade intensity. When the state of the market is characterized by low

trade intensity, NCK news shocks have comparatively small effects on prices. Prices rise for the first three periods following the shock and then fall back toward their original level. In fact, the total effect on the price level (i.e., $B(1,5)$) is less than 0.002 in all three cases. Thus, NCK news shocks have small and transitory effects on prices when transaction rates are low. Panels B and C show very different responses. Here prices rise for approximately seven periods (35 minutes) following the shock. The cumulative effect on the price level is approximately 0.40 when $n = 20$ and 0.63 when $n = 40$. Thus, NCK news shocks have much larger and long-lasting price effects when trade intensity is high.

The third noteworthy feature concerns the time of the peak response. In all cases, NCK news shocks have their largest (positive) effect on price changes during the third period, 15 minutes after the shock. This suggests that it may take some time before the transmission of information through transaction prices affects the price-setting decisions of a significant number of dealers. Consequently, the delayed peaked response may well reflect the importance of transaction prices as a medium for information flows between dealers and the relatively slow speed at which such information diffuses across the whole market.

We can also use equation (13) to decompose the variance of observed price changes into different theoretical components. In particular, consider the k -period price change, $\Delta^k p_t^o \equiv \sum_{i=0}^{k-1} \Delta p_{t+i}^o$, for $o = \{a, b\}$. Substituting for Δp_t^o with (13), gives

$$\Delta^k p_t^o = \omega_t^o - \omega_{t-k}^o + \sum_{i=0}^{k-1} \varepsilon_{t-i} + B(L, k, n_t) v_t, \quad (14)$$

where $B(L, k, n_t) = \sum_{i=0}^{k-1} B(L, L^i n_t) L^i$. Since the v_t , ε_t , and ω_t^o shocks are mutually independent and serially uncorrelated, we can use (14) to write the variance of price changes as

$$\text{Var}(\Delta^k p_t^o) = \Sigma_\omega(n_t) + \Sigma_\omega(n_{t-k}) + \sum_{i=0}^{k-1} \Sigma_\varepsilon(n_{t-i}) + \text{Var}(B(L, k, n_t) v_t). \quad (15)$$

The first two terms identify the contribution of the idiosyncratic sampling shocks to the variance. These terms can be calculated from the estimates of the $\Sigma_\omega(\cdot)$ function in Table VIII. The contribution of CK news, given by the third term, similarly can be calculated from the estimates of $\Sigma_\varepsilon(\cdot)$. The fourth term identifies the contribution of the NCK news shocks. This term can be decomposed further as

$$\text{Var}(B^*(L, k, n_t) v_t) + \text{Var}(B(1, k, n_t) v_t) + 2B^*(L, k, n_t) B(1, k, n_t) \text{Var}(v_t), \quad (16)$$

where $B^*(L, k, n_t) \equiv B(L, k, n_t) - B(1, k, n_t)$. $B(1, k, n_t)$ denotes the permanent effect of a one-unit NCK news shock on prices over k -periods, so $B^*(L, k, n_t)$ identifies the transitory price effect of a NCK news shock at

lag L . The first term in the expression above gives us the transitory contribution of NCK news shocks to the variance of k -period price changes. The permanent contribution of NCK news shocks is shown by the second term. The last term identifies twice the covariance between $B^*(L, k, n_t)v_t$ and $B(1, k, n_t)v_t$.

To compute these components, we need estimates of $B(L, k, n_t)$ and the variance of the v_t shocks. In light of the state-dependency results reported in Table VI, it is clearly inappropriate to assume that v_t is homoskedastic. I, therefore, reestimated the ARMA(2,2) model for order flow allowing for a state-dependence variance, $\Sigma_v(n) = \Sigma_v(0)\exp(-n/\gamma) + \Sigma_v(\infty)(1 - \exp(-n/\gamma))$. With γ set equal to 100 (as above), the GMM estimates of the AR and MA coefficients are almost identical to those reported in Table II, while the estimates of $\Sigma_v(0)$ and $\Sigma_v(\infty)$ are 0.001 and 0.022, respectively.²²

Table IX reports variance decompositions based on (15) for various horizons, k , and trade intensities, n . The upper panel reports the fraction of the price change variance attributable to sampling, $R_\omega(k, n) \equiv 2\Sigma_\omega(n)/\text{Var}(\Delta^k p_t^o)$. According to these estimates, most of the short-term variability in prices is attributable to sampling, unless trade intensity is very high. For example, $R_\omega(1, 10)$ is approximately 84 percent. Although $R_\omega(k, n)$ falls at all intensities as the horizon increases, sampling continues to contribute more than 17 percent at the two-hour horizon for intensities of 10 or less. These results indicate that, under most conditions, a majority of the high frequency variations in observed transaction prices result from the presence of significant dispersion in the equilibrium price distribution.

The middle panel of Table IX reports the contribution of the NCK news shocks to the variance of price changes: $R_v(k, n) \equiv \text{Var}(B(L, k, n_t)v_t)/\text{Var}(\Delta^k p_t^o)$. These estimates show that NCK news shocks contribute more to price volatility in states of the market with higher trade intensities. NCK news shocks also make a larger contribution to the price variance as the horizon rises. Across all intensities, NCK news contributes approximately 40 percent of the variance in prices at the two-hour horizon. The contribution of NCK news shocks is even higher in states where trade intensities are at least 40.

These findings are inconsistent with the traditional macro view that stresses the importance of CK news. According to this view, the estimates of both $R_\omega(k, n)$ and $R_v(k, n)$ should be close to zero. Of course, the macro view still may be accurate at much longer horizons. To address this possibility, the right-hand column of the center panel reports the contribution of the NCK news shocks to the variance of the permanent innovations in prices. According to (13), observed prices can be written as $p_t^o = \bar{p}_t + I(0)$ terms, where $\Delta\bar{p}_t = \varepsilon_t + B(1, n)v_t$. Hence, the relative contribution of order flow can be calculated as $R_v(\infty, n) \equiv B(1, n)^2\Sigma_v/\text{Var}(\Delta\bar{p}_t)$. The right-hand column of the table shows that the estimates of $R_v(\infty, n)$ follow a U-shaped pattern, with a

²² As a robustness check, I also computed variance decompositions assuming that v_t was homoskedastic with Σ_v set equal to the variance of the innovations from the ARMA(2,2) model. The results are similar to Table IX and can be found in the Supplemental Appendix.

Table IX
Variance Ratios

Variance decompositions derived from estimates of the state-dependent model in Table VIII and the ARMA(2,2) model for order flow with state-dependent heteroskedasticity. The column headings show the horizon k measured in minutes. The ratios $R_\omega(k, n)$ and $R_v(k, n)$ respectively measure the contribution of sampling and NCK shocks to the variance of observed price changes. The ratio $R_{v\cdot}(k, n)$ measures the contribution of NCK shocks that only affect the price level temporarily.

$n \setminus k$	5	30	60	120	∞
Panel A: $R_\omega(k, n) = \text{Var}(\omega_t^o - \omega_{t-k}^o) / \text{Var}(\Delta^k p_t^o)$					
2	96.39%	81.78%	69.45%	53.36%	0.00%
5	91.37%	64.22%	47.64%	31.42%	0.00%
10	83.70%	45.89%	29.94%	17.66%	0.00%
20	69.61%	24.65%	13.79%	7.33%	0.00%
30	56.60%	13.48%	6.88%	3.48%	0.00%
40	45.04%	7.73%	3.75%	1.85%	0.00%
60	27.65%	3.07%	1.42%	0.68%	0.00%
80	17.13%	1.49%	0.68%	0.32%	0.00%
All	67.11%	20.13%	10.72%	5.54%	0.00%
Panel B: $R_v(k, n) = \text{Var}(B(L, k, n)v_t) / \text{Var}(\Delta^k p_t^o)$					
2	0.18%	0.77%	0.91%	1.08%	1.67%
5	0.38%	1.00%	0.74%	0.49%	0.00%
10	0.80%	3.11%	3.54%	3.87%	4.34%
20	3.25%	17.68%	21.71%	24.10%	26.82%
30	8.53%	36.71%	42.22%	45.07%	47.98%
40	15.95%	52.09%	57.28%	59.75%	62.16%
60	32.32%	70.31%	73.97%	75.60%	77.11%
80	45.91%	79.19%	81.80%	82.93%	83.97%
All	5.53%	30.62%	36.81%	40.22%	43.87%
Panel C: $R_{v\cdot}(k, n) = \text{Var}(B^*(L, k, n)v_t) / \text{Var}(\Delta^k p_t^o)$					
2	0.17%	0.53%	0.38%	0.29%	0.00%
5	0.38%	1.00%	0.74%	0.49%	0.00%
10	0.74%	2.54%	2.06%	1.22%	0.00%
20	2.22%	8.72%	6.49%	3.59%	0.00%
30	4.52%	14.45%	10.15%	5.52%	0.00%
40	6.68%	18.04%	12.37%	6.70%	0.00%
60	9.43%	21.46%	14.44%	7.83%	0.00%
80	10.73%	22.86%	15.28%	8.29%	0.00%
All	3.13%	12.16%	8.64%	4.75%	0.00%

minimum value of zero when $n = 5$. At high trade intensities, the estimates of $R_{v\cdot}(\infty, n)$ are more than 80 percent. These estimates are consistent with the regression finding in Evans and Lyons (2002) that order flow accounts for more than 60 percent of daily price changes. They also stand in sharp contrast to the traditional macro view concerning the origins of exchange rate movements. (Recall from Table III that interest rate innovations account for very little of the variance in daily exchange rate changes in this sample.)




Figure 3. Variance decompositions over a typical trading day. Solid lines plot the contribution of the sampling component to the variance of observed prices, $R_\omega(k, n_t)$, against the left-hand axis. Dashed lines plot the contribution of NCK shocks to the variance of observed prices, $R_v(k, n_t)$, against the right-hand axis. The horizontal axis covers 24 hours. The plotted contributions are calculated using n_t equal to the average trade intensity rate for the sample during each five-minute interval.

Further information on the importance of NCK news shocks comes from the lower panel of Table IX. Here, I report the contribution of NCK news shocks that temporarily affect the price level, $R_v^*(k, n) \equiv \text{Var}(B^*(L, k, n_t)v_t)/\text{Var}(\Delta^k p_t^o)$. Once again, the estimates of $R_v^*(k, n)$ increase as n rises. The largest estimates of more than 22 percent are found at high trade intensities with the 30-minute horizon. Across all intensities, the temporary NCK news component contributes approximately 12 percent to the variance of the price changes over 30 minutes. This pattern is consistent with the delayed peak in the impulse responses plotted in Figure 2.

We also can use the variance components to examine sources of seasonal heteroskedasticity. Recall from Figure 1 that there are pronounced intraday patterns in trade intensities. Figure 3 combines these patterns with the estimates of $R_\omega(k, n)$ and $R_v(k, n)$ to show how different shocks contribute to the variance of price changes over a typical 24 hour period. Three features of these plots stand out. First, sampling variability is a very significant source of price variance outside of European trading hours. For example, panel A shows that the estimates of $R_\omega(1, n)$ are more than 80 percent before 7:00

hours and after 17:00 hours. Although sampling contributes less to the price variance over longer horizons, $R_\omega(k, n)$ remains above 30 percent even at the two-hour horizon. Second, although sampling contributes less to the variance of price changes during European trading, it continues to be an important source at the five-minute horizon; the estimates of $R_\omega(1, n)$ remain above 60 percent. Third, NCK news shocks contribute most to price variance during European trading. The peak values for $R_v(k, n)$ range from 10 percent at the five-minute horizon to approximately 50 percent at two hours. When combined with the values for $R_\omega(k, n)$, these estimates imply a peak contribution for CK news to the price variance, defined as $R_\varepsilon(k, n) \equiv 1 - R_\omega(k, n) - R_v(k, n)$, of approximately 30 percent at the five-minute horizon, and 50 percent at two hours.

V. Conclusion

The primary aim of this paper has been to provide a new perspective on the source of exchange rate dynamics. The perspective comes from considering how trading in the FX market actually takes place. I argued that the lack of transparency in dealer–customer and direct interdealer transactions could give rise to an equilibrium distribution of transaction prices rather than a single price level. Based on this observation, I then developed an empirical framework to study exchange rate dynamics. Applying this framework to transaction data for the DM/\$, several striking results emerge. First, there is strong evidence supporting the presence of equilibrium price distribution. Second, CK news shocks are rarely the predominant source of exchange rate changes over both long and short horizons. Third, NCK news shocks are an empirically important source of long-term exchange-rate dynamics.

The first of these findings is key to understanding the short-term dynamics of exchange rates. Unless trading between dealers is extremely active, the dispersion in the equilibrium distribution is large enough to account for most of the observed variance in high frequency price changes. This finding puts a new perspective on the high frequency volatility of exchange rates. It implies that much of the volatility we observe comes from sampling the heterogeneous trading decisions of dealers in an equilibrium distribution that, under normal market conditions, changes comparatively slowly.

The second finding speaks more directly to assumptions that lie at the heart of the traditional macro view of exchange rate dynamics. Recall that this view assumes (1) all information relevant for exchange rate determination is CK, and (2) the mapping from information to equilibrium prices is also CK. CK news shocks meet both of these requirements but account for only 56 percent of the persistent movements in exchange rates across all market states and 20 percent when trading activity is high. Thus, a key implication of my results is that models based on assumptions (1) and (2) are

simply too restrictive to account for all the persistent exchange rate movements we observe. In this sense, my findings do direct attention away from the common knowledge framework modeled in traditional macro exchange rate models.

Where should our attention be redirected? The third finding points toward NCK news shocks. Macro announcements represent one possible source of NCK news to dealers when they hold differing views about an announcement's implications for equilibrium prices. Another source of NCK news to dealers are customer orders, which are determined endogenously by agents outside the market. Although general equilibrium models have yet to be developed identifying the determinants of customer orders, there are suggestions in the literature about their origins. For example, Evans and Lyons (2002) point out that orders can embody NCK information about valuation numerators (i.e., future interest differentials) and denominators (i.e., anything that affects discount rates). In the former case, customer orders could reflect changes in the expected future path of interest rates. Moreover, there is mounting microeconomic evidence that private information plays an important role in FX trading (Lyons (1995), Yao (1997a), Bjonnes and Rime (1998), Cheung and Wong (2000), Ito et al. (1998), Covrig and Melvin (1998), and Payne (1999)).

In summary, my results direct attention away from exchange rate models with a common-knowledge environment dominated by a small number of macro variables. They point, instead, toward models with richer informational structures in which the sources of NCK news can be identified.

Appendix: Estimation Details

The appendix describes how the GMM technique is applied to produce the estimation results reported in the tables.

A. Autocorrelations

To estimate $\text{Corr}(w_t, w_{t-i})$, let $y_t = w_t - \alpha_0 - \alpha_1 w_{t-1}$ and $z'_t = [1, w_t]$. The GMM estimates of α_0 and α_1 are then computed using $m_t(\theta) = y_t z'_{t-i}$, with weighting matrix Γ_{0T} . Table II reports the estimate of α_1 , together with the p -value for the Wald test of the null hypothesis of $\alpha_1 = 0$.

B. ARMA Models

To illustrate how the ARMA model estimates in Table II are calculated, consider the ARMA(2,2) case:

$$w_t = \alpha_1 w_{t-1} + \alpha_2 w_{t-2} + \nu_t + \beta_1 \nu_{t-1} + \beta_2 \nu_{t-2}. \quad (\text{A1})$$

The first step is to write the model in the state space form of (6):

$$\begin{bmatrix} w_t \\ w_{t-1} \\ \nu_t \\ \nu_{t-1} \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 & \theta_1 & \theta_2 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} w_{t-1} \\ w_{t-2} \\ \nu_{t-1} \\ \nu_{t-2} \end{bmatrix} + \begin{bmatrix} \nu_t \\ 0 \\ \nu_t \\ 0 \end{bmatrix},$$

$$y_t = \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} w_t \\ w_{t-1} \\ \nu_t \\ \nu_{t-1} \end{bmatrix}, \quad \Omega = \begin{bmatrix} \Sigma_\nu & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \Sigma_\nu & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (\text{A2})$$

The estimates are then computed with $m_t(k; \theta) = y_t y_{t-k} - \gamma(k; \theta)$, for $k = 0, 1, \dots, 12$, where $\gamma(k; \theta) \equiv \text{Cov}(y_t, y_{t-k})$ is calculated from the state space form. The weighting matrix is Γ_{0T^*} where

$$\Gamma_{v,T^*} = \frac{1}{T^*} \sum_{\Lambda} m_t(\tilde{\theta}) m_{t-v}(\tilde{\theta})'. \quad (\text{A3})$$

C. Decomposition Regressions

Table IV reports GMM estimates of

$$\Delta p_t^a = \sum_{i=\underline{k}}^{i=\bar{k}} \alpha_i x_{t-i} + \nu_t, \quad (\text{A4})$$

where ν_t follows an MA(1) process that is independent from all leads and lags of x_t . Let $y_t = \Delta p_t^a - \sum_{i=\underline{k}}^{i=\bar{k}} \alpha_i x_{t-i}$, and $z_t' = [x_{t-\underline{k}} \dots x_{t-\bar{k}}]$. The GMM estimates of the α_i 's are computed using $m_t(\theta) = y_t z_t'$ and the Newey–West weighting matrix, $\Gamma_{0T^*} + \frac{1}{2}(\Gamma_{1T^*} + \Gamma_{1T^*}')$. Notice that this choice allows for both heteroskedasticity and an MA(1) process in ν_t .

D. Tests for State Dependency

The upper panel of Table VI reports Wald tests for nonlinearity in models of the form

$$\Delta p_t^a = D_0(L)x_{t-i} + D_1(L)x_t n_t + D_2(L)x_t n_t^2 + D_3(L)x_t n_t^3 + \nu_t, \quad (\text{A5})$$

where ν_t follows an MA(1) process that is independent from all leads and lags of x_t and $D_i(L) = d_1^i L + d_2^i L^0 + d_3^i L^{-1} + \dots d_6^i L^{-4}$. Each cell reports the

Wald statistic and p -value for the null hypothesis of $d_i^j = 0$ for $i = \{1, 2 \dots 6\}$ with $j = 1$ (left-hand column), $j = 1$ and 2 (center column), and $j = 1, 2$, and 3 (right-hand column). In each case, the d_i^j coefficients are estimated by GMM for the specification, including $D_0(L)x_t$ and the regressors listed at the head of each column along the lines described in Section C.

The lower panel of Table VI reports tests for heteroskedasticity in the shocks to the Bivariate Model and ARMA(2,2) order flow model. Estimates of the shocks from the Bivariate Model are obtained as

$$(\tilde{\varepsilon}_t)^2 = \hat{y}_{1t}\hat{y}_{2t} = (\hat{\varepsilon}_t + \hat{\omega}_t^a - \hat{\omega}_{t-1}^a)(\hat{\varepsilon}_t + \hat{\omega}_t^b - \hat{\omega}_{t-1}^b) = (\hat{\varepsilon}_t)^2 + \zeta_t^e,$$

$$(\tilde{\omega}_t^a)^2 = -\hat{y}_{1t+1}\hat{y}_{1t} = -(\hat{\varepsilon}_{t+1} + \hat{\omega}_{t+1}^a - \hat{\omega}_t^a)(\hat{\varepsilon}_t + \hat{\omega}_t^a - \hat{\omega}_{t-1}^a) = (\hat{\omega}_t^a)^2 + \zeta_t^a,$$

$$(\tilde{\omega}_t^b)^2 = -\hat{y}_{2t+1}\hat{y}_{2t} = -(\hat{\varepsilon}_{t+1} + \hat{\omega}_{t+1}^b - \hat{\omega}_t^b)(\hat{\varepsilon}_t + \hat{\omega}_t^b - \hat{\omega}_{t-1}^b) = (\hat{\omega}_t^b)^2 + \zeta_t^b, \quad (\text{A6})$$

where “hats” denote the GMM estimates. The estimated innovations to the ARMA(2,2) order flow model are found from (A1) as

$$(\tilde{v}_t)^2 = (w_{t+2} - \hat{\alpha}_1 w_{t+1} - \hat{\alpha}_2 w_t)(w_t - \hat{\alpha}_1 w_{t-1} + \hat{\alpha}_2 w_{t-2})/\hat{\beta}_2 = (\hat{v}_t)^2 + \zeta_t^v. \quad (\text{A7})$$

Under the null of a correctly specified model, all the error terms, ζ_t^i , have mean zero and are serially uncorrelated. To implement the Glesjer (1969) tests, I estimate $\varpi_t^2 = \alpha_0 + \underline{n}_t' \alpha + \xi_t$ for each shock ϖ_t , where the vector $\underline{n}_t$ includes the terms listed at the head of each column. The GMM estimates of α_0 and α are then computed using $m_t(\theta) = y_t z_t'$, where $y_t = \varpi_t^2 - \alpha_0 - \underline{n}_t' \alpha$ and $z_t' = [1, \underline{n}_t']$ with weighting matrix Γ_{0T} . The table reports the Wald test for $\alpha = 0$ based on these estimates.

E. Bivariate Models

The Bivariate Model in (5) can be written in the state space form of (6) as

$$\begin{bmatrix} \omega_t^a \\ \omega_t^b \\ \omega_{t-1}^a \\ \omega_{t-1}^b \\ \varepsilon_t \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \omega_{t-1}^a \\ \omega_{t-1}^b \\ \omega_{t-2}^a \\ \omega_{t-2}^b \\ \varepsilon_{t-1} \end{bmatrix} + \begin{bmatrix} \omega_t^a \\ \omega_t^b \\ 0 \\ 0 \\ \varepsilon_t \end{bmatrix},$$

$$\Omega(n_t) = \begin{bmatrix} \Sigma_\omega(n_t) & 0 & 0 & 0 & 0 \\ 0 & \Sigma_\omega(n_t) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \Sigma_\varepsilon(n_t) \end{bmatrix},$$

$$y_t \equiv \begin{bmatrix} \Delta p_t^a - \sum_{i=1}^{i=-4} d_i(n_t) x_{t-i} \\ \Delta p_t^b - \sum_{i=1}^{i=-4} d_i(n_t) x_{t-i} \end{bmatrix} = \begin{bmatrix} 1 & 0 & -1 & 0 & 1 \\ 0 & 1 & 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} \omega_t^a \\ \omega_t^b \\ \omega_{t-1}^a \\ \omega_{t-1}^b \\ \varepsilon_t \end{bmatrix}. \quad (\text{A8})$$

Table V reports estimates of the Bivariate Model without state dependence where $\Sigma_\omega(n_t) = \Sigma_\omega$, $\Sigma_\varepsilon(n_t) = \Sigma_\varepsilon$, and $d_i(n_t) = d_i$ are the parameters in $D(L)$. To obtain the GMM estimates, I use instruments $z'_t = [x_{t-1} \dots x_{t+4}]$, the Newey–West weighting matrix $\Gamma_{0T^*} + \frac{1}{2}(\Gamma_{1T^*} + \Gamma'_{1T^*})$, and moments $m_t(k; \theta)$ for $k = 0, 1$, and 2 (see equation (8)).

Table VIII reports estimates of the Bivariate Model with state dependence where $d_i(n_t)$ and $\Sigma_i(n)$ take the functional forms shown in Section III.D. The GMM estimates are obtained using instruments $z'_t = [x_{t-1}, \dots, x_{t+4}, x_{t-1} \exp(-n_t/100), \dots, x_{t+4} \exp(-n_t/100)]$, the Newey–West weighting matrix $\Gamma_{0T^*} + \frac{1}{2}(\Gamma_{1T^*} + \Gamma'_{1T^*})$, and moments $m_t(k; \theta)$ for $k = 0, 1$, and 2 , where

$$\gamma_t(0; \theta) = \begin{bmatrix} \Sigma_\omega(n_t) + \Sigma_\omega(n_{t-1}) + \Sigma_\varepsilon(n_t) & \Sigma_\varepsilon(n_t) \\ \Sigma_\varepsilon(n_t) & \Sigma_\omega(n_t) + \Sigma_\omega(n_{t-1}) + \Sigma_\varepsilon(n_t) \end{bmatrix}$$

$$\gamma_t(1; \theta) = \begin{bmatrix} -\Sigma_\omega(n_{t-1}) & 0 \\ 0 & -\Sigma_\omega(n_{t-1}) \end{bmatrix}, \quad \gamma_t(2; \theta) = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}. \quad (\text{A9})$$

The Wald test for $D(L, 0) = D(L, \infty)$ is computed as $\nabla \hat{\alpha}'(\hat{\Omega}_{\nabla \alpha})^{-1} \nabla \hat{\alpha}$, where

$$\nabla \hat{\alpha} = [\hat{\alpha}_1(0) - \hat{\alpha}_1(\infty), \dots, \hat{\alpha}_{-4}(0) - \hat{\alpha}_{-4}(\infty)] \quad (\text{A10})$$

and $\hat{\Omega}_{\nabla \alpha}$ is the estimated asymptotic covariance matrix of $\nabla \alpha$. To test for misspecification in the $\alpha_i(n)$ and $\Sigma_i(n)$ functions, I use the GMM version of the LM test developed by Newey and West (1987). In the case of the $\alpha_i(n)$ functions, I consider alternative specifications of the form $\tilde{\alpha}_i(n) = \alpha_i(n) + \varphi_i n$. To test the null hypothesis that $\varphi_i = 0$ for all i , I used the two-step

procedure suggested by Greene (1997). First, I computed the derivative for the GMM criterion function $Q(\theta)$ with $\tilde{\alpha}_i(n)$ replacing $\alpha_i(n)$ at the GMM estimates with $\varphi_i = 0$. I then calculated the Wald statistic for the null hypothesis that this vector of derivatives equals zero. In the case of the variance functions $\Sigma_\omega(n)$ and $\Sigma_\varepsilon(n)$, the alternative specifications take the form of $\tilde{\Sigma}_\omega(n) = \Sigma_\omega(n) + \varphi n$ and $\tilde{\Sigma}_\varepsilon(n) = \Sigma_\varepsilon(n) + \varphi n$.

The lower panel of Table VIII reports autocorrelations for the estimated shocks. The shocks are calculated from the GMM estimates as in (A6), and standardized as $\tilde{\omega}_t^2 = \varpi_t^2 [\hat{\Sigma}_\omega(n_t)]^{-1}$, where ϖ_t denotes the shock in question. I then estimate $\tilde{\omega}_t^2 = \alpha_0 + \alpha_i \tilde{\omega}_{t-i}^2 + \xi_t$ by GMM using $m_t(\theta) = y_t z'_{t-i}$, where $y_t = \tilde{\omega}_t^2 - \alpha_0 - \alpha_i \tilde{\omega}_{t-i}^2$ and $z'_t = [1, \tilde{\omega}_t^2]$ with weighting matrix Γ_{0T}^* . The table reports the estimate of α_i and GMM standard error allowing for heteroskedasticity. Under the null hypothesis that $\text{Var}(v_t) = \Sigma_v(n_t)$, $\alpha_i = 0$ for all lags i .

REFERENCES

- Andersen, Torben G., and Tim Bollerslev, 1998, Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies, *Journal of Finance* 53, 219–266.
- Baillie, Richard T., and Tim Bollerslev, 1991, Intraday and intermarket volatility in foreign exchange rates, *Review of Economic Studies* 58, 565–585.
- Bank for International Settlements, 1998, *Central Bank Survey of Foreign Exchange and Derivative Market Activity*, Bank for International Settlements.
- Bjornnes, Geir H., and Dagfinn Rime, 1998, FX trading live: Impact of new trading environments, mimeo, Norwegian School of Management.
- Bollerslev, Tim, and Ian Domowitz, 1993, Trading patterns and prices in the interbank foreign exchange market, *Journal of Finance* 48, 1421–1443.
- Bollerslev, Tim, and Michael Melvin, 1994, Bid-ask spreads and volatility in the foreign exchange market, *Journal of International Economics* 36, 355–372.
- Brennan, Michael, and Henry H. Cao, 1997, International portfolio investment flows, *Journal of Finance* 52, 1851–1880.
- Cheung, Yin-Wong, and Menzie D. Chinn, 1999, Traders, market microstructure and exchange rate dynamics, mimeo, Department of Economics, UC Santa Cruz.
- Cheung, Yin-Wong, and Clement Wong, 2000, A survey of market practitioners' view on exchange rate dynamics, *Journal of International Economics* 51, 401–419.
- Covrig, Vincentiu, and Michael Melvin, 1998, Asymmetric information and price discovery in the FX market: Does Tokyo know more about the Yen? mimeo, Arizona State University.
- Cumby, Robert, and John Huizinga, 1992, Investigating the correlation of unobserved expectations: Expected returns in equity and foreign exchange markets and other examples, *Journal of Monetary Economics* 30, 217–253.
- Dacorogna, Michel M., Ulrich A. Muller, Richard Nagler, Richard B. Olsen, and Olivier V. Pictet, 1993, A geographical model for the daily and weekly seasonal volatility in the FX market, *Journal of International Money and Finance* 12, 413–438.
- Dornbusch, Rudiger, 1976, Expectations and exchange rate dynamics, *Journal of Political Economy* 84, 1161–1176.
- Engel, Charles, 1996, The forward discount anomaly and the risk premium: A survey of recent evidence, *Journal of Empirical Finance* 3, 123–192.
- Engle, Robert, 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, *Econometrica* 50, 987–1008.

- Engle, Robert, Takatoshi Ito, and Wen-Ling Lin, 1990, Meteor showers or heat waves? Heteroskedastic intradaily volatility in the foreign exchange market, *Econometrica* 58, 525–542.
- Evans, Martin D., 1996, Peso problems: Their theoretical and empirical implications, in G. S. Maddala and C. R. Rao eds.: *Handbook of Statistics: Statistical Methods in Finance* (Elsevier Science, Amsterdam, The Netherlands).
- Evans, Martin D., 1997, The microstructure of foreign exchange dynamics, mimeo, Georgetown University.
- Evans, Martin D., 2001, FX trading and exchange rate dynamics, NBER Working paper No. 8116.
- Evans, Martin D., and Richard K. Lyons, 2002, Order flow and exchange rate dynamics, *Journal of Political Economy* 110, 170–180.
- Frankel Jeffrey A., and Andrew K. Rose, 1995, Empirical research on nominal exchange rates, in Gene Grossman and Kenneth Rogoff eds.: *Handbook of International Economics*, Vol. 3 (Elsevier Science, Amsterdam, The Netherlands).
- Glesjer, H., 1969, A new test for heteroskedasticity, *Journal of the American Statistical Association* 64, 316–323.
- Goodhart, Charles, and Mark Giugale, 1993, From hour to hour in the foreign exchange market, *The Manchester School* 61, 1–34.
- Goodhart, Charles, Takatoshi Ito, and Richard Payne, 1996, One day in June 1993: A study of the working of Reuters 2000-2 electronic foreign exchange trading system, in J. Frankel, J. Galie, and A. Giovannini eds.: *The Microstructure of Foreign Exchange Markets* (University of Chicago Press, Chicago, IL).
- Greene, William H., 1997, *Econometric Analysis* (Prentice Hall, London).
- Guillaume, Dominique, Michel M. Dacorogna, Richard Dave, Ulrich A. Muller, Richard B. Olsen, and Olivier V. Pictet, 1994, From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets, Discussion paper, the O&A Research Group, Zurich.
- Hamilton, James D., 1994, *Time Series Analysis* (Princeton University Press, Princeton, NJ).
- Hansen, Lars, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.
- Harvey, Andrew, 1989, *Forecasting Structural Time Series Models and the Kalman Filter* (Cambridge University Press, Cambridge).
- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 66, 179–207.
- Ito, Takatoshi, Richard K. Lyons, and Michael Melvin, 1998, Is there private information in the FX market? The Tokyo experiment, *Journal of Finance* 53, 1111–1130.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson, 1994, Transactions, volume, and volatility, *Review of Financial Studies* 7, 631–651.
- Lyons, Richard K., 1995, Tests of microstructure hypotheses in the foreign exchange market, *Journal of Financial Economics* 39, 321–351.
- Lyons, Richard K., 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275–298.
- Lyons, Richard K., 2001, *The Microstructure Approach to Exchange Rates*, MIT Press, forthcoming.
- Newey, Whitney, and Kenneth West, 1987, A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Payne, Richard, 1997, Announcement effects and seasonality in the intraday foreign exchange market, LSE Financial Markets Group Discussion Paper No. 238.
- Payne, Richard, 1999, Informed trade in spot foreign exchange markets: An empirical investigation, mimeo, LSE.
- Peires, Bettina, 1997, Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market, *Journal of Finance* 52, 1589–1614.
- Perraudin, William, and Paolo Vitale, 1995, Inter-dealer trade and information flows in a decentralized foreign exchange market, mimeo, Cambridge University.
- Potter, Simon M., 1999, Nonlinear time series modelling: An introduction, mimeo, Federal Reserve Bank of New York.

- Stock, James, 1988, Estimating continuous time processes subject to time deformation, *Journal of the American Statistical Association* 83, 77–85.
- Yao, Jian M., 1997a, Market making the interbank foreign exchange market, Working Paper S-98-3, Salomon Center, New York University.
- Yao, Jian M., 1997b, Spread components and dealer profits in the interbank foreign exchange market, Working Paper S-98-4, Salomon Center, New York University.
- Zhou, Bin, 1996, High frequency data and volatility in foreign exchange rates, *Journal of Business and Economic Statistics* 14, 45–52.