

Understanding Order Flow

October 2005

Martin D. D. Evans¹
Georgetown University and NBER
Department of Economics
Washington DC 20057
Tel: (202) 338-2991
evansm1@georgetown.edu

Richard K. Lyons
U.C. Berkeley and NBER
Haas School of Business
Berkeley, CA 94720-1900
Tel: (510) 643-2027
lyons@haas.berkeley.edu

Abstract

This paper develops a model for understanding end-user order flow in the FX market. The model addresses several puzzling findings. First, the estimated price-impact of flow from different end-user segments is, dollar-for-dollar, quite different. Second, order flow from segments traditionally thought to be liquidity-motivated actually has power to forecast exchange rates. Third, about one third of order flow's power to forecast exchange rates one month ahead comes from flow's ability to forecast future flow, whereas the remaining two-thirds applies to price components unrelated to future flow. We show that all of these features arise naturally from end-user heterogeneity, in a setting where order flow provides timely information to market-makers about the state of the macroeconomy.

Keywords: Exchange rates, forecasting, microstructure, order flow.

JEL Codes: F3, F4, G1.

¹Corresponding author: Martin Evans. Both authors thank the National Science Foundation for financial support, including a clearing-house for micro-based research on exchange rates (at georgetown.edu/faculty/evansm1 and at faculty.haas.berkeley.edu/lyons).

Introduction

This paper addresses order flow heterogeneity and its empirical implications. By order flow heterogeneity, we mean transactions initiated by agents of different types (e.g., non-financial corporations versus hedge funds versus mutual funds). Recent theoretical work on exchange rates stresses the analytical importance of heterogeneity across agents, stemming from both dispersed information and non-informational shocks to asset demands (e.g., Bacchetta and van Wincoop 2003, Hau and Rey 2002, Dunne, Hau, and Moore 2004, Evans and Lyons 2004a,b). Empirical predictions from these models are borne out: trades have causal and persistent effects on price, a finding that runs counter to textbook exchange rate models (e.g., Evans and Lyons 2002a,b, Payne 2003, among many others).² Theory generally assumes, however, that agents are symmetrically heterogeneous; that is to say, they differ, but in the same way. In contrast, trades in the foreign exchange (FX) market come from categories of agents that are quite different: they have different motivations, different attitudes toward risk, and different horizons. Extant theory provides little guidance for empiricists seeking to exploit transaction data that is disaggregated into segments.³

Several puzzling empirical findings suggest that new modeling is needed. First, the price impact of order flow is, dollar-for-dollar, quite different across end-user segments. This is not what one would expect to see if agents are symmetrically heterogeneous. Second, order flow from segments traditionally thought to be liquidity-motivated actually has power to forecast exchange rates. Third, about one third of order flow's power to forecast exchange rates comes from flow's ability to forecast future flow, whereas the remaining two-thirds applies to price components unrelated to future flow. In this paper, we present a model where all these features arise naturally from heterogeneity across end-user segments, in a setting where order flow provides timely information to market-makers about the state of the macroeconomy.

A one-period version of the Kyle (1985) model illustrates why intuition based on standard microstructure models is an unreliable guide to empirical work using order flow data from different end-user (henceforth "customers") segments. In this model, the change in price quoted by a market-maker (i.e., the change in the log exchange rate s_t) depends on the total order flow arriving in the market, x_t :

$$\Delta s_t = \beta x_t,$$

where β , the price-impact coefficient, governs the sensitivity of the market-maker's price quote to order flow. By assumption, the market-maker cannot distinguish different order flow components, but must instead respond to the aggregated total flow x_t . Expected order flow in this model is zero. Even in dynamic versions of the Kyle model, the market-maker expects order flow to be zero in each period (a property that also applies

²In textbook models, all of which assume that variables relevant to exchange rates are common knowledge (CK), trades per se have no causal effect on price. Demand shifts have causal effects on price. But because demand shifts come from public information only, prices adjust before transactions occur, so no causality from trades to price is present. Put differently, in CK environments demand and order flow are quite different, whereas in purely non-CK environments they are the same.

³Recent research using transaction data from different segments includes Froot and Ramadorai (2002), Carpenter and Wang (2003), Mende and Menkhoff (2003), Bjønnes, Rime, and Solheim (2004), Marsh and O'Rourke (2005), and Sager and Taylor (2005).

to standard versions of the other canonical information model, the Glosten and Milgrom 1985 model). This conditional i.i.d. structure is an analytically attractive feature of these models.

Now suppose that a researcher has a dataset that breaks total FX customer flow at a given bank into three segments, say short-term investors (e.g., hedge funds), long-term investors (e.g., mutual funds), and international-trade-based (e.g., non-financial corporations); x_t^i for $i = 1, 2, 3$. The researcher then runs the regression:

$$\Delta s_t = \beta_1 x_t^1 + \beta_2 x_t^2 + \beta_3 x_t^3 + \epsilon_t. \quad (1)$$

With the perspective of the Kyle model, it would be natural to interpret the β_i coefficients as price-impact parameters. However, this would be problematic for three reasons: First, the three regressors are not likely to be independent intratemporally. Indeed, in our data the order flows from different segments are significantly correlated. If the x_t^i s are known to covary, and each of the β 's are non-zero, then no one coefficient summarizes the total price-impact of changes in a single flow segment. Second, the three regressors are not likely to be independent intertemporally. In fact, the flow segments in our dataset are significantly auto-correlated. In this context, the specification in (1) is a reduced-form for potentially complex microeconomic dynamics. As non-structural estimates, the β 's are not reliable measures of the price-impact of incremental trades. Third, the specification misses the fact that regressors in equation (1) come from a single bank, whereas the flows that move the exchange rate are the market-wide flows from all segments. This is problematic in terms of price-impact parameters: positive correlation between flow segments across banks means that when the regression omits other-bank flows, the flows from the source bank are getting too much price-impact credit (a form of omitted variable bias). More fundamentally, the FX market is not transparent, at least not with respect to customer flows, so that the flows that drive Δs_t proximately are in fact the interdealer flows. The exchange rate reflects the information in any individual bank's customer flows only when other dealers learn that information.

We present both simulation results and empirical estimates. The simulation results address the relation in our model between exchange rates and customer order flows. The empirical estimates are based on roughly six years of customer transaction data from Citibank. Our simulations show that:

- customer flows provide more precise information about fundamentals when the mix of customers is tilted toward longer-horizon participants;
- flows from customer segments can produce negative coefficients in contemporaneous return regressions, even when positively correlated with fundamentals; and
- customer flows forecast returns because they are correlated with the future market-wide information flow that dealers use to revise their FX prices.

Based on our empirical analysis we find that:

- both the aggregate and disaggregated customer flows received by Citibank are positively auto-correlated;

- contemporaneous correlations across flow segments are low at the daily frequency, but high at the monthly frequency;
- some customer segments do produce negative coefficients in contemporaneous return regressions;
- the proportion of excess return variation that segment flows explain rises with the horizon; and
- about one-third of order flow’s power to forecast exchange rates one month ahead comes from flow’s ability to forecast future flow, with the remaining two-thirds applying to price components unrelated to future flow.

The remainder of the paper is in three sections. Section 1 presents the model. Section 2 describes our data, and presents our empirical analysis and the model simulations. Section 3 concludes. We provide technical details on how we solve the model in an appendix.

1 Model

The model we develop is based on Evans and Lyons (2004a,b). These papers embed the salient features of the spot FX market in a general equilibrium setting to study how information concerning the macroeconomy is transmitted to exchange rates via trading. The model we present here considers this transmission process in greater detail, emphasizing the role of end-user heterogeneity.

Our model focuses on the behavior of two groups of participants in the spot FX market: dealers and end-user customers. Dealers act as financial intermediaries. They quote prices at which they are willing to trade, and they initiate trade with each other. All non-dealer market participants are termed customers. This group comprises individuals, firms, and financial institutions such as hedge and mutual funds. Customers have the opportunity to initiate trade with dealers at the prices they quote. The resulting pattern of trade defines customer order flow. In particular, positive customer order flow occurs at a given bank when the value of customer orders to purchase foreign currency at the quoted spot rate exceeds the value of orders to sell. Customer order flow is only observed by the recipient dealer. Any information contained in customer order flow only becomes known to dealers across the market as the result of interdealer trading. This information aggregation process was the focus of earlier models of FX trading (e.g., Evans and Lyons 2002a, and 2004a). In this paper we focus on how the information contained in customer order flow is related to customer-type. For this purpose we distinguish between liquidity-motivated traders, short-term investors, and long-term investors. We then examine how differences across customer types affects the information contained in customer order flow, and how this, in turn, affects the joint dynamics of order flows and exchange rates.

1.1 Dealers

There are a large number of dealers who act as intermediaries in the spot FX market. As such, each dealer quotes prices at which he stands ready to buy or sell foreign currency to customers and other dealers. Dealers also have the opportunity to initiate transactions with other dealers at the prices they quote. Thus, unlike

standard international macro models, the behavior of the spot exchange rate is determined entirely by the FX prices dealers quote.

In Evans and Lyons (2004a) we derive an expression for the foreign currency prices dealers quote from their preferences and the trading rules in the market. These rules determine the sequence of quoting and trading during each period. We are interested in the transmission of information over periods of a month or more, so we refer to each period as a month. At the start of each month dealers quote FX prices to customers. Under our trading rules these prices are good for any amount and are publicly observed. Each dealer then receives orders from a distinct subset of customers. Customer orders are only observed by the recipient dealer and so may represent a source of private information. Dealers next quote prices in the interdealer market. These prices, too, are good for any quantity and are publicly observed. Dealers then have the opportunity to trade among themselves. Interdealer trading is simultaneous and trading with multiple partners is feasible.

This sequencing of events obviously oversimplifies actual trading in the market. In reality dealers quote prices to customers and other dealers, receive orders from customers and other dealers, and initiate their own trades on a continual basis. Our simplified sequencing allows us to focus on how information conveyed by customer orders is transmitted among dealers and embedded in spot rates. In this sense, trading between dealers and customers, and among dealers in our model, are metaphors for the many rounds of quoting and trading that take place during an actual month. As in Evans and Lyons (2004a), we allow for greater transparency in interdealer trading than is present in the actual FX market to compensate for the simplified sequencing.

In our trading environment, optimal quote decisions take a simple form; all dealers quote the same price for foreign currency to both customers and other dealers. We represent the month- t quote as

$$s_t = (1 - b)\mathbb{E}_t^D \sum_{i=0}^{\infty} b^i f_{t+i}, \quad (2)$$

where $1 > b > 0$. s_t is the log price of foreign currency quoted by all dealers, and f_t denotes exchange rate fundamentals. While equation (2) takes the present value form familiar from standard international macro models, here it represents how dealers quote the price for foreign currency in equilibrium. In the market setting we consider, all dealers choose to quote the same price because doing otherwise opens them up to arbitrage, a costly proposition. Consequently, the month- t quote must be a function of information known to all dealers. We incorporate this requirement in (2) with the expectations operator \mathbb{E}_t^D , that denotes expectations conditioned on information common to all dealers at the start of month t , Ω_t^D . This is not to say that all dealers have the same information. Far from it. The customer order flows received by individual dealers represent an important source of private information, so there is likely to be a good deal of information heterogeneity across dealers at any one time. Equation (2) implies that dealer quotes only reflect that part of fundamental information that is common to all dealers.

The definition of fundamentals in equation (2) depends on the macroeconomic structure of the model. For example, f_t may include home and foreign money supplies and household consumption, as in Evans and Lyons (2004b), or “output gaps” and inflation rates as in Engel and West (2005). For our purpose here we

need not take a stand on the composition of fundamentals. It suffices that all dealers agree on the elements of f_t , and its equilibrium process. For simplicity, we assume that this process is a random walk:

$$f_t = f_{t-1} + u_t + v_t, \tag{3}$$

where $u_t \sim i.i.d.N(0, \sigma_u^2)$ and $v_t \sim i.i.d.N(0, \sigma_v^2)$. Innovations to fundamentals come from two sources. u_t shocks are immediately observed by all dealers and customers (i.e., u_t is publicly observed at the start of month t). These shocks represent public news that has an immediate and known consequence for fundamentals. The v_t shocks to fundamentals are not immediately observed. Individual dealers and customers may have some information on v_t as it occurs, but the information is imprecise and private. This information may come from private research, in the case of hedge or mutual funds, or from orders for goods and services in case of individual firms (i.e., in the form of micro information). Over time, information on v_t is collected and aggregated by government agencies.

The distinction between the u_t and v_t shocks plays an important role in our model. The month- t spot exchange rate is determined by the foreign currency price that all dealers quote at the start of month t . This price can incorporate information contained in the u_t shock because it is observed by all dealers and has a known consequence for fundamentals. By contrast, v_t shocks will not affect the month- t spot rate because they are not known to all dealers. Our focus is on how information concerning the v_t shocks becomes known to dealers, and how this information is reflected in the dynamics of spot rates.

Dealers learn about the v_t shocks from two sources. The first is via the scheduled macroeconomic announcements made by government agencies. These announcements provide information about past values of v_t because it takes time for the agencies to compile and aggregate the information. We shall assume that macro announcements at the start of month t reveal the true value of v_{t-2} . This means that the true state of fundamentals in month $t - 2$, f_{t-2} , is common knowledge after the month t announcement. The second source of information to dealers comes from order flow. Recall that customer order flow is a source of private information to dealers. In general, customer order flows will not be perfectly correlated across dealers because customers have heterogeneous reasons for trading. Consequently, information contained in individual customer order flows only becomes known to all dealers if and when it is reflected in interdealer trading. The process by which private information from customer order flows becomes common knowledge to all dealers is analyzed in detail in Evans and Lyons (2004a). Here we make a reduced-form assumption about this information aggregation process. In particular, we assume it takes a month of interdealer trading before all dealers know the aggregate of customer orders hitting the market. Let x_t denote aggregate customer order flow from month- t trading (i.e., x_t is the sum of the signed customer orders placed with individual dealers at the start of month t). We assume that interdealer trading in month- t is sufficiently informative to reveal the value of x_t to all dealers by the start of month $t + 1$. Interdealer trading does not reveal the constituent components of x_t .

In summary, the information available to all dealers at the start of each month is given by:

$$\Omega_{t+1}^D = \{u_{t+1}, v_{t-1}, x_t\} \cup \Omega_t^D. \quad (4)$$

Since the spot exchange rate in month t is simply the FX price quoted by all dealers given common information Ω_t^D , the evolution of Ω_t^D shown here plays a key role in driving exchange rates.

1.2 Customers

The customer orders received by actual foreign exchange dealers come from a variety of sources: individuals, non-financial firms, and financial firms such as pension funds, mutual funds and hedge funds. We will not attempt to model each separately. Instead, we note that the customer pool comprises heterogeneous agents with different preferences. We represent this heterogeneity by splitting the customer pool into three groups: short-term investors, long-term investors, and international-trade-based agents.

Both short- and long-term investors choose to allocate their wealth between domestic and foreign bonds based on their private forecasts of returns. In particular, we assume that there is a continuum of short-term investors on the unit interval indexed by i . The desired foreign bond position for investor i is given by

$$a_{i,t}^S = \gamma^S \mathbb{E}_{i,t} [s_{t+1} - s_t], \quad (5)$$

where $\gamma^S > 0$ and $\mathbb{E}_{i,t}$ denotes expectations conditioned on short-term investor i 's information at the start of month t , $\Omega_{i,t}^S$. For simplicity we assume that domestic and foreign interest rates are equal so $s_{t+1} - s_t$ represents the excess return on foreign bond holdings. Of course, if investors held equities or other risky assets in addition, returns on these assets would also affect their desired portfolio position. Again we ignore this possibility for the sake of clarity. The γ^S parameter measures the sensitivity of the investor's portfolio choice to expected returns. Greater risk aversion and/or uncertainty concerning future returns is represented by a smaller value for γ^S .

Long-term investors have a two- rather than one-month horizon. In equilibrium, it takes two months before shocks to fundamentals are fully and permanently reflected in the spot rate. A two-month horizon is therefore sufficient to represent the preferences of far-sighted investors. Again, there is a continuum of these investors on the unit interval. The desired foreign bond position of long-term investor j is given by

$$a_{j,t}^L = \gamma^L \mathbb{E}_{j,t} [s_{t+2} - s_t], \quad (6)$$

where $\gamma^L > 0$ and $\mathbb{E}_{j,t}$ denotes expectations conditioned on $\Omega_{j,t}^L$. As above, we ignore the possibility that long-term investors could hold other assets. The portfolio position of long-term investors may be more or less sensitive to variations in expected returns depending of the value of γ^L .

Changes in the desired portfolio positions drive the customer order flows from both short- and long-term

investors. In particular, the customer order flows from each group in month t are given by

$$x_t^S = \int_0^1 (a_{i,t}^S - a_{i,t-1}^S) di, \quad (7a)$$

$$x_t^L = \int_0^1 (a_{j,t}^L - a_{j,t-1}^L) dj. \quad (7b)$$

Notice that both order flows aggregate the expectations of individual investors concerning the future path of the exchange rate. As such, they aggregate the private information that individuals use in forming these expectations. We assume that this private information comes in the form of signals concerning v_t . In particular, at the start of month t , each short-term investor i and long-term investor j receives a noisy signal concerning the value of v_t , $\tilde{v}_{i,t}$ and $\tilde{v}_{j,t}$, where

$$\tilde{v}_{i,t} = v_t + \xi_{i,t}, \quad (8a)$$

$$\tilde{v}_{j,t} = v_t + \xi_{j,t}. \quad (8b)$$

The idiosyncratic noise terms, $\xi_{i,t}$ and $\xi_{j,t}$, are i.i.d. normally distributed shocks with zero means and variances equal to σ_ξ^2 , and $\int_0^1 \xi_{i,t} di = \int_0^1 \xi_{j,t} dj = 0$. Unlike dealers, investors do not observe order flows. Rather they observe spot rates, public news, macro announcements and private signals. The information used in their portfolio decisions is therefore characterized by:

$$\Omega_{i,t+1}^S = \{u_{t+1}, v_{t-1}, s_{t+1}, \tilde{v}_{i,t+1}\} \cup \Omega_{i,t}^S, \quad (9a)$$

$$\Omega_{j,t+1}^L = \{u_{t+1}, v_{t-1}, s_{t+1}, \tilde{v}_{j,t+1}\} \cup \Omega_{j,t}^L. \quad (9b)$$

The third component of order flow comes from the agents engaged in international trade in goods and services. We do not model this source of order flow in any detail but rather assume a specification with two key features.⁴ First, order flows from this source are persistent, and second they are correlated with exchange rate fundamentals. Specifically, we assume that the order flow associated with real trade, hereafter trade-based order flow, x_t^T , follows an AR(1) process:

$$x_t^T = \rho x_{t-1}^T + \delta v_t + \varepsilon_t, \quad (10)$$

with $|\rho| < 1$, where $\varepsilon_t \sim i.i.d.N(0, \sigma_\varepsilon^2)$. When δ is non-zero, changes in real economic activity that affect exchange rate fundamentals also affect the demand for foreign currency by trade-based agents, which in turn shows up order flow. Our specification also includes a liquidity shock ε_t , a component of trade-based order flow unrelated to changing fundamentals.

Equations (7) and (10) describe the three components that comprise aggregate customer order flow during

⁴See Osler (1998) for a model in which the currency demands of exporters and importers are explicitly considered.

month t :

$$x_t = x_t^S + x_t^L + x_t^T. \quad (11)$$

A key assumption in our model is that interdealer trading allows all dealers to learn the value of x_t by the end of month t trading, but not the constituent components, x_t^S , x_t^L and x_t^T .

1.3 Solving the Model

An equilibrium in our model comprises: (i) a process for spot rates consistent with the quote equation (2) given the equilibrium evolution of common dealer information Ω_t^D in (4), and (ii) a process for aggregate customer order flow in (11) consistent with investor expectations. Finding these processes is complicated by the fact that the information conveyed to dealers by order flow is itself a function of how investors use their private information in forecasting the future exchange rate.

We solve the model by the ‘‘conjecture and verification’’ method. Specifically, we first conjecture the form for equilibrium quotes consistent with (2) and an assumption about the form of Ω_t^D . Next, we compute the components of customer order flow consistent with the exchange rate process implied by these conjectures. This provides us with an expression for aggregate customer order flow, x_t . Finally, we verify that our initial conjecture concerning dealer information, (i.e., the evolution of Ω_t^D) can be supported by an inference problem based on public news, macro announcements and the process for x_t .

We start with the process for the exchange rate. Our specification for the fundamentals process in (3) implies that the quote equation in (2) simplifies to $s_t = \mathbb{E}_t^D f_t$. Thus, the log price of FX quoted by all dealers is simply their estimate of current fundamentals based on common information. This implies that the equilibrium depreciation rate can be written as $\Delta s_{t+1} = \mathbb{E}_{t+1}^D f_{t+1} - \mathbb{E}_t^D f_t$. Now equation (3) also implies that $\mathbb{E}_t^D f_{t+1} = \mathbb{E}_t^D f_t$, so combining these expressions gives:

$$\Delta s_{t+1} = (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) f_{t+1}. \quad (12)$$

Thus, the equilibrium depreciation rate between t and $t+1$ is solely determined by the flow of new information to all dealers concerning the state of fundamentals in $t+1$. We term this, the market-wide information flow. One important implication of (12) is that the depreciation rate must be uncorrelated with any elements in Ω_t^D . This rules out serial correlation in Δs_t , among other things.

The next step is to break down the market-wide information flow into its component parts. For this purpose, we apply the expectations operators \mathbb{E}_{t+1}^D and \mathbb{E}_t^D to the identity $f_{t+1} = f_{t-2} + u_{t+1} + u_t + u_{t-1} + \delta(v_{t+1} + v_t + v_{t-1})$. This gives

$$\begin{aligned} \Delta s_{t+1} &= (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) f_{t-2} + (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) (u_{t+1} + u_t + u_{t-1}) + (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) (v_{t+1} + v_t + v_{t-1}), \\ &= u_{t+1} + (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) v_t + (v_{t-1} - \mathbb{E}_t^D v_{t-1}). \end{aligned} \quad (13)$$

The simplifications in the last line follow from three important aspects of the model. First, public news

and macro announcements make the true state of fundamentals common knowledge with a two month delay. This means that $\mathbb{E}_{t+1}^D f_{t-2} = \mathbb{E}_t^D f_{t-2} = f_{t-2}$. Second, the public news shocks, u_t , are by definition common knowledge and unanticipated so $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D)(u_t + u_{t-1}) = 0$, $\mathbb{E}_t^D u_{t+1} = 0$, and $\mathbb{E}_{t+1}^D u_{t+1} = u_{t+1}$. Third, dealer estimates of the v_t shocks combine their priors with their observations of order flow and macro announcements. At the start of month t , dealers know the order flows from earlier trading months (i.e., x_{t-1}, x_{t-2}, \dots), and the history of macro announcements (i.e., v_{t-2}, v_{t-3}, \dots). These sources are uninformative about the value of v_t or v_{t+1} because customers were unaware of these shocks when making their trading decisions in month $t - 1$ and earlier. As a consequence, $\mathbb{E}_{t+1}^D v_{t+1} = \mathbb{E}_t^D v_{t+1} = 0$, and $\mathbb{E}_{t+1}^D v_{t-1} = v_{t-1}$.

Equation (13) is helpful for understanding the joint dynamics of exchange rates and order flow. The equation provides a decomposition of the new information dealers use in revising the FX prices they quote between t and $t + 1$. As one would expect, the public news component, u_{t+1} , has an immediate impact on the depreciation rate because it leads all dealers to revise their estimate of fundamentals, f_{t+1} . The other two components of the information flow come from revisions in dealer estimates of v_t and v_{t-1} . These revisions come from two sources: the macro announcement on the value of v_{t-1} made at the start of $t + 1$, and the customer order flow from month t trading, x_t . The $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D)v_t$ component identifies what dealers learn about v_t before it becomes known publicly via the macro announcement in $t + 2$. The last component, $v_{t-1} - \mathbb{E}_t^D v_{t-1}$, identifies the new information in the macro announcement made at the start of $t + 1$. The relative importance of these last two components depends on the information content of customer order flow.

The information in order flow depends on the actions of customers. In particular, investor order flows, x_t^S and x_t^L , will only contain information on v_t if investors find it useful to use their private information on v_t to forecast returns. To analyze the information content in customer order flow we must therefore examine how investors allocate their portfolios. For this purpose, we need a conjecture for the equilibrium depreciation process. This conjectured process needs to be consistent with (13) and so relies on an assumption about dealer expectations which must also be verified. We conjecture (and verify below) that the equilibrium process for the depreciation rate can be written as

$$\Delta s_{t+1} = u_{t+1} + \lambda v_t + (1 - \lambda)v_{t-1} + \pi \varepsilon_t - \pi \varepsilon_{t-1}, \quad (14)$$

where λ and π are parameters that need to be determined.

Recall that both short- and long-term investors receive a private signal concerning the value of v_t at the start of month t . These signals are described in equation (8). This specification implies that for short-term investor i , $\mathbb{E}_{i,t} v_t = \theta \tilde{v}_{i,t}$, and long-term investor j , $\mathbb{E}_{j,t} v_t = \theta \tilde{v}_{j,t}$ where $\theta \equiv \sigma_v^2 / (\sigma_v^2 + \sigma_\xi^2)$. Investor expectations concerning the value of v_{t-1} are similarly given by $\mathbb{E}_{i,t} v_{t-1} = \theta \tilde{v}_{i,t-1}$ and $\mathbb{E}_{j,t} v_{t-1} = \theta \tilde{v}_{j,t-1}$. It might seem strange that these estimates incorporate no information from month t variables. However, recall from equation (12) that the depreciation rate depends on the new information received during month t by all dealers. Thus, variables known to dealers at the start of t will be of no use to investors in forecasting Δs_{t+1} . This leaves investors' month t signals, $\tilde{v}_{i,t}$ and $\tilde{v}_{j,t}$, but these variables are uncorrelated with v_{t-1} , so the month $t - 1$ estimate remains the best estimate of v_{t-1} for each investor in month t . Investor expectations

concerning the ε_t shocks are more straightforward. Only dealers observe customer order flow, so individual investors have no information on the values of ε_t and ε_{t-1} at the start of month t .⁵ Investor expectations concerning the rate of depreciation can therefore be computed from (14) as:

$$\begin{aligned}\mathbb{E}_{i,t}[s_{t+1} - s_t] &= \lambda\theta\tilde{v}_{i,t} + (1-\lambda)\theta\tilde{v}_{i,t-1}, \\ \mathbb{E}_{j,t}[s_{t+2} - s_t] &= \theta\tilde{v}_{j,t} + (1-\lambda)\theta\tilde{v}_{j,t-1}.\end{aligned}$$

Using these expressions to identify investor expectations in (5) and (6), and combining the results with (7) and (8), gives us the following expressions for the investor components of customer order flow:

$$x_t^S = \gamma^S\theta(\lambda v_t + (1-2\lambda)v_{t-1} - (1-\lambda)v_{t-2}), \quad (15)$$

$$x_t^L = \gamma^L\theta(v_t - \lambda v_{t-1} - (1-\lambda)v_{t-2}). \quad (16)$$

Notice that in aggregate the investor order flows provide information on v_t and v_{t-1} even though the value of these shocks is not known to individual investors. Aggregation of the order flows eliminates the idiosyncratic component of investors' private information so that dealers can make more precise inferences about v_t and v_{t-1} using aggregate order flow than would be possible from the currency orders of individual customers. Equations (15) and (16) also show the investor order flows are related to fundamentals even though investors are not concerned with the behavior of fundamentals per se. Investors care about returns which are in turn determined by how dealers change their FX quotes. They only use their information on v_t insofar as it helps them predict dealer quotes. If investor behavior makes customer order flow more informative to dealers about the behavior of fundamentals, it will only be because investors found it in their best interest to use their private information on v_t in managing their portfolios.

All that now remains is to verify that equation (14) is consistent with equation (13) given dealer inferences regarding v_t and v_{t-1} based on macro announcements and their observations of aggregate customer order flow, $x_t = x_t^S + x_t^L + x_t^T$, where x_t^S , x_t^L and x_t^T follow (10), (15) and (16). This is a relatively complex problem, so we relegate the technical details to the Appendix.

1.4 Customer Characteristics and Order Flow

We now turn to the properties of the equilibrium. Equation (14) shows that the equilibrium depreciation rate Δs_{t+1} is driven by three factors: public news u_{t+1} , fundamental news v_t , and liquidity shocks, ε_t :

$$\Delta s_{t+1} = u_{t+1} + \lambda v_t + (1-\lambda)v_{t-1} + \pi\varepsilon_t - \pi\varepsilon_{t-1}. \quad (14)$$

The π parameter measures the equilibrium influence of liquidity shocks on the depreciation rate. Recall that liquidity shocks are uncorrelated with actual fundamentals, f_t , so if dealers observed f_t contemporaneously

⁵Investors eventually learn the values of ε_t using the information in news, macro announcements and the history of exchange rates.

(i.e., $\mathbb{E}_t^D f_t = f_t$), their FX quotes would not depend on liquidity shocks. Dealers can use order flow, x_t , to revise their estimates of f_{t+1} between the start of months t and $t+1$, and these revisions drive spot rates (see equation 12). So any shock that contributes to unexpected aggregate order flow, $x_t - \mathbb{E}_t^D x_t$, will affect Δs_{t+1} provided order flow contains some information on the behavior of fundamentals that cannot be learnt from other sources. In sum, therefore, π will be non-zero in any equilibrium where dealers find customer order flow informative about fundamentals. Note, however, that even in these cases, the effect of a liquidity shock on the price of FX is short-lived. Equation (14) implies that the effect of a ε_t shock increases dealer quotes for the log price of FX by $\pi\varepsilon_t$ in month $t+1$, and lowers their quotes by $\pi\varepsilon_t$ in month $t+2$. Intuitively, by the start of month $t+2$, all dealers can infer the exact value of ε_t from macro announcements and their observations on order flow. They can therefore refine their earlier inferences regarding fundamentals based on x_t .

The λ parameter indicates the speed with which dealers learn about fundamentals from order flow. In the extreme case where order flow is completely uninformative, macro announcements would be the only source of information on v_t . In this case λ would equal zero, and v_t shocks would impact dealer quotes with a two month lag, which is the reporting lag for macro announcements. At the other extreme, if observations on x_t allowed dealers to estimate the value of v_t with complete precision by the start of month $t+1$, λ would equal one. Under these circumstances, the month $t+1$ announcement concerning the value of v_{t-1} would be informationally redundant as far as dealers were concerned, so the value of v_{t-1} would not affect the change in their FX quotes, Δs_{t+1} .

We can now turn to the focus of our analysis; namely, the question of how the information content of order flow is related to the characteristics of customers. In our model, the characteristics of customer order flow depend on the investors' sensitivities, γ^S and γ^L , the precision of private information, θ , and the correlation between fundamentals and trade-based order flow, governed by δ . Ideally, we would like to study the analytical dependence of λ and π on these parameters, but the complexity of the model makes this impossible. We therefore base our analysis on calibrated solutions to the model.

Figure 1 plots the equilibrium values of λ and π for different values of γ^S and γ^L computed from solutions to the model where $\sigma_v^2 = 2$, $\sigma_u^2 = 0.01$, $\sigma_\varepsilon^2 = 10$, $\lambda = 0.3$ and $\delta = 1$.⁶ We have examined the equilibrium values of λ and π based on many other calibrations and found that the basic characteristics displayed in Figure 1 are robust. We will therefore use these values as a baseline parameterization in much of the numerical analysis of the model below. As we shall see, their use enables the model to replicate many important features of the data.

Two features stand out from the plots in Figure 1. First, the equilibrium values of π are relatively insensitive to the different values for γ^S and γ^L ; the values for π range between 0.13 and 0.23. Second, the greater the sensitivity of investors to expected returns, the faster is the speed with which dealers learn about

⁶The parameters γ^S , γ^L and θ are not separately identified because they only appear in the order flow equations as $\gamma^S\theta$ and $\gamma^L\theta$. For simplicity, we normalize θ to one so that we can discuss the results in terms of γ^S and γ^L . Of course, increasing (decreasing) the precision of the private signals increases (decreases) θ , and so is analytically equivalent to raising (lowering) γ^S and γ^L .

Figure 1A: π





Figure 1B: λ



Notes: Equilibrium values for λ and π graphed as functions of the investor sensitivity parameters $\gamma = (\gamma^s, \gamma^l)$. The solid lines represent the case where $\gamma = \gamma^s = \gamma^l$, dashed lines the case where $\gamma = \gamma^s$ and $\gamma^l = 0$, and dot-dashed lines the case where $\gamma = \gamma^l$ and $\gamma^s = 0$. The other parameters in the model are $\sigma_v^2 = 2$, $\sigma_u^2 = 0.01$, $\sigma_\varepsilon^2 = 10$, $\lambda = 0.3$, and $\delta = \theta = 1$.

fundamentals from order flow (i.e., λ is increasing in γ^s and γ^l).

To understand the economics behind these results, we first note that $\pi = \lambda \cong 0.17$ when customer order flow is driven solely by trade-based agents. In this equilibrium, order flow, x_t , is not particularly informative about fundamentals because it does not incorporate any of the private information on v_t available to investors. Consequently, most of the impact of v_t shocks on spot rates is delayed for two months until it is fully revealed by macro announcements. The equilibrium value of λ rises with increases in both γ^s and γ^l because aggregate order flow now contains private information from investors on the values of v_t . Importantly, order flow is more informative to dealers in these equilibria because investors have a greater incentive to act on their private information. Investors recognize that dealers place greater weight on order flow when revising their FX quotes, and use their private information on v_t to forecast these revisions. At the same time, their own portfolio choices are more sensitive to private information, with the result that in aggregate, order flow provides a more precise signal to dealers concerning v_t .

Figure 1 also shows how π and λ are related to the degree of forward-looking behavior driving order flow. The dashed and dashed-dot plots show the equilibrium values of π and λ for cases where investors are either all short-term (i.e., $\gamma^l = 0$) or long-term (i.e., $\gamma^s = 0$). Here we see that order flow is more informative when only long-term investors contribute to order flow. The reason is that private information has a larger impact on investor expectations concerning two-month returns than one-month returns (i.e., $\partial \mathbb{E}_{j,t} [s_{t+2} - s_t] / \partial \tilde{v}_{j,t} = \theta$ and $\partial \mathbb{E}_{j,t} [s_{t+1} - s_t] / \partial \tilde{v}_{j,t} = \lambda \theta < \theta$). Consequently, the order flows generated by long-term investors will be more sensitive to v_t shocks than the flows from short-term investors (see equations 15 and 16).

To summarize, our model shows how the mix of customer-types affects the information content of order

flow. In particular, we have demonstrated that customer order flows provide more precise and timely information concerning fundamentals to dealers when the pool of customers is more forward-looking and willing to react to private information.

2 Empirical Analysis

2.1 Data

Our empirical analysis utilizes a dataset that comprises customer order flows, spot rates and interest rates over approximately six years. The order flows come from the customer orders received by Citibank in the USD/EUR market from April 1993 to June 1999.⁷ Citibank’s share of customer orders in the USD/EUR market at that time was in the 10-15 percent range, no other bank had a larger market share in these currencies. The customer orders are aggregated at the daily frequency and measure in \$m the imbalance between customer orders to purchase and sell euros. Days “begin” at 00:00 GMT and any trades executed with customers over a weekend – a relatively rare event – are included in Monday’s order flows. Thus a trading week comprises five days. Daily order flows are split into three categories: non-financial corporations –henceforth “Corporations”, unleveraged financial institutions (e.g., mutual funds)—henceforth “Investors”, and leveraged financial institutions (e.g., hedge funds)—henceforth “Hedge”. We also distinguish between trades executed with Citibank in the US, and those elsewhere within Citibank’s global trading operation (referred to as “Non-US”). Thus, our customer order flows are partitioned into six non-overlapping segments corresponding to three participant types and two trade locations. We take the day d spot rate as the offer rate (USD/EUR) quoted by Citibank at the end of trading (approximately 17:00 GMT) on day $d - 1$. Excess returns are computed from these quotes using Euro deposit rates of the appropriate maturity from Datastream.

Table 1 presents descriptive statistics on depreciation rates and the customer order flows. Since overnight interest rates are approximately equal in our sample, the daily depreciation rate behaves very similarly to excess returns, and neither displays any significant serial correlation. By contrast, both the aggregate and disaggregated customer order flows received by Citibank are serially correlated. The estimated autocorrelation coefficients are quite small, but many are positive and highly statistically significant. These statistical patterns are repeated at the weekly and monthly frequency: Depreciation rates and excess returns are serially uncorrelated while some of the order flow segments display a small but significant degree of autocorrelation. (We do not report these lower frequency results to conserve space.) Our theoretical model is consistent with these features of the data. The depreciation rate is serially uncorrelated by construction, as we noted above. Aggregate customer order flow is slightly positively autocorrelated under our baseline parameterization with $\gamma^s = \gamma^l = 1$: The first order autocorrelation in x_t is equal to 0.05.

The lower portion of Table 1 reports the contemporaneous order flow correlations at the daily and monthly

⁷Before January 1999, data for the euro are synthesized from data in the underlying markets against the dollar, using weights of the underlying currencies in the euro.

Table 1: Sample Statistics

	mean	std	skew	kurt.	ρ_1	ρ_2	ρ_3
Depreciation Rate							
Δs	-0.009	0.544	0.149	4.691	-0.022 (0.448)	0.003 (0.918)	-0.011 (0.658)
Order Flows							
Corporate							
US: x^C	-3.424	45.555	-1.650	19.692	0.075 (0.018)	0.053 (0.060)	-0.035 (0.255)
Non-US: \hat{x}^C	-11.879	81.666	0.573	11.440	0.033 (0.192)	0.045 (0.091)	0.028 (0.286)
Hedge							
US: x^H	-0.783	138.745	0.502	15.426	0.114 (0.003)	0.046 (0.076)	0.001 (0.967)
Non-US: \hat{x}^H	2.257	82.462	0.400	7.598	-0.025 (0.506)	-0.023 (0.397)	-0.029 (0.212)
Investors							
US: x^I	3.821	59.977	-1.946	33.632	0.063 (0.048)	0.023 (0.437)	0.031 (0.188)
Non-US: \hat{x}^I	3.170	112.391	2.472	40.441	0.068 (0.003)	0.038 (0.070)	0.027 (0.175)
Aggregate							
x	-4.940	226.073	0.677	9.418	0.098 (0.001)	0.059 (0.016)	0.026 (0.310)
Correlations							
Daily\Monthly	x^C	\hat{x}^C	x^H	\hat{x}^H	x^I	\hat{x}^I	x
		0.960	0.370	-0.948	-0.923	-0.620	0.781
	\hat{x}^C	0.020	0.261	-0.943	-0.951	-0.645	0.748
	x^H	0.033	-0.048	-0.323	-0.056	-0.483	0.549
	\hat{x}^H	-0.014	-0.038	-0.012	0.889	0.578	-0.713
	x^I	-0.031	0.014	-0.063	0.021	0.467	-0.682
	\hat{x}^I	-0.013	-0.022	0.038	-0.011	0.067	-0.257
	x	0.209	0.319	0.601	0.339	0.265	0.521

Notes: The depreciation rate, Δs , are calculated as the daily change in the natural log of the spot price (USD/EUR) x 100. Order flows are daily aggregates for the euro in \$m. The statistics reported below ρ_i are the sample autocorrelations at lag i . P-values for the null hypothesis of no autocorrelation are reported in parenthesis. The lower panel reports correlations between the order flows at the daily frequency below the leading diagonal and at the monthly frequency above the leading diagonal. The sample spans the period 4/01/93 – 6/30/99, and includes observations on 1682 trading days.

frequency. At the daily frequency, the correlations between the flow segments are quite small, but at the monthly frequency they range from approximately -0.95 to 0.95. This difference in the correlation structure is important for understanding the results presented below.

Table 2: Contemporaneous Return Regressions

Horizon	Corporate		Hedge		Investors		R^2	χ^2
	US	Non US	US	Non US	US	Non US		
1 day	-0.155	-0.240					0.015	15.133
	(0.113)	(0.067)						(0.001)
			0.174	0.204			0.024	21.791
			(0.055)	(0.060)				(<0.001)
					-0.047	0.369	0.044	38.261
				(0.120)	(0.060)		(<0.001)	
	-0.147	-0.214	0.153	0.194	-0.029	0.353	0.078	75.465
	(0.107)	(0.064)	(0.054)	(0.056)	(0.121)	(0.059)		(<0.001)
1 week	-0.118	-0.469					0.061	32.07
	(0.138)	(0.083)						(<0.001)
			0.349	0.114			0.077	27.965
			(0.069)	(0.096)				(<0.001)
					-0.005	0.523	0.105	37.728
				(0.154)	(0.086)		(<0.001)	
	-0.167	-0.358	0.275	0.069	-0.051	0.447	0.195	111.527
	(0.133)	(0.077)	(0.064)	(0.090)	(0.143)	(0.080)		(<0.001)
1 month	0.065	-0.594					0.129	22.434
	(0.266)	(0.126)						(<0.001)
			0.389	0.166			0.103	8.75
			(0.135)	(0.225)				(0.013)
					-0.091	0.719	0.205	34.636
				(0.215)	(0.119)		(<0.001)	
	0.120	-0.376	0.214	-0.074	0.000	0.583	0.299	58.424
	(0.185)	(0.102)	(0.137)	(0.196)	(0.208)	(0.130)		(<0.001)

Notes: The table reports OLS estimates of the coefficients in the regression of excess returns, er_{d+h} , on the customer order flow segments that aggregate net orders for the euro in \$m on days d to $d + h - 1$. Estimates are computed at the daily frequency, with $h = 5$ and 20 for the 1-week and 1-month horizon regressions. The table reports asymptotic standard errors corrected for heteroskedasticity in parentheses. For the 1-week and 1-month results, standard errors are also corrected for the induced MA($h-1$) process in from overlapping observations. The right hand column reports Wald tests and p-values for the null that all the coefficients on the order flows are zero.

2.2 Customer Orders and Returns

We begin our empirical analysis of the relationship between exchange rates and customer order flows in Table 2. Here we report the results of regressing excess returns on Citibank's customer flows at the one day, one week and one month horizon. Excess returns are computed as $er_{d+h} \equiv s_{d+h} - s_d + \nabla_d^i h$ where s_d is the log of the quote at the start of day d , and $\nabla_d^i h$ is the interest differential on day d for h day deposits. Order flows for each segment are aggregated from day d to day $d + h - 1$. Thus the order flows cover the same period as the revision in dealer quotes determining excess returns.

The results in Table 2 contain several noteworthy features. First, the coefficients on the order flow seg-

ments are quite different from each other. Some are positive, some are negative, some are highly statistically significant, others are not. Second, while the coefficients on order flow are jointly significant in every regression we consider, the proportion of the variation in excess returns that they account for rises with the horizon. For example, using all six order flow segments, the R^2 statistics rise from 8 percent at one day, to 30 percent at the one month horizon. Third, the explanatory power of the order flows shown here is much less than that reported for interdealer order flows. Evans and Lyons (2002a), for example, report that interdealer order flow accounts for approximately 60 percent of the variations in the \$/DM at the daily frequency.

These findings appear puzzling when judged against the perspective of a standard market-making model, like Kyle (1985). In particular, the negative coefficients in Table 2 seem to indicate, rather counter-intuitively, that market-makers view orders to purchase the euro from some customers as an indicator that the euro is overvalued. Similarly, it is hard to understand why orders from some customers appear to carry information at one frequency but not at others. Remember that there is no serial correlation in returns, so quote revisions at the monthly frequency are just the sum of daily revisions.

Our model provides the key to understanding the results in Table 2. Recall from the model that the depreciation rate reflects the revision in dealer quotes driven by new information concerning fundamentals. This information arrives in the form of public news, macro announcements and interdealer order flow, but not the customer order flows of individual dealers such as Citibank: Any information concerning fundamentals contained in the customer flows received by individual banks affects the FX price quoted by dealers only once it is inferred from the interdealer order flows observed by all dealers. This distinction between individual customer order flows and aggregate interdealer order flow is crucial to understanding the results in Table 2.

Our model simulations are useful in quantify this distinction. Specifically, we first solve the model for our baseline parameterization with $\gamma^s = \gamma^l = 1$. We then use the equilibrium values of λ and π to simulate data on the depreciation rate, Δs_{t+1} , aggregate order flow, x_t , and its components, x_t^T , x_t^S and x_t^L for a sample spanning 10,000 months. Table 3 reports the results from regressing the depreciation rate on the customer order flows in this simulated data.

Rows (i) – (iii) of Table 3 show that the revision in dealer quotes between the start of months t and $t + 1$ is positively correlated with the individual components of customer order flow. Notice that none of these individual flows accounts for more than one third of the variation in quotes. This is not due to the presence of public news shocks, they account for less than one percent of the variance in Δs_{t+1} in our simulations.⁸ Rather, the low R^2 statistics indicate that none of the order flow components is strongly correlated with the flow of information dealers use to revised their quotes. Of course dealers in our model are assumed to observe aggregate customer order flow, x_t , rather than its components, but in row (iv) we see that x_t accounts for only one third of the variation in quotes. Individual order flows are not strongly correlated with the flow of information dealers use to revise quotes.

⁸This implication of our simulations is consistent with empirical evidence on the exchange rate effects of macro news. For example, while Andersen et al. (2003) find that news items can account for a significant proportion of the variation in exchange rates in the five-minutes following their release, they also note that news items account for only 0.2 percent of exchange rate variations over their full sample. For further discussion, see Evans and Lyons (2003).

Table 3: Simulated Return Regressions

	x_t^T	x_t^S	x_t^L	x_t	R^2
(i)	0.225 (0.075)				0.329
(ii)		0.506 (0.013)			0.138
(iii)			0.220 (0.008)		0.075
(iv)				0.160 (0.002)	0.333
(v)	0.203 (0.003)	1.073 (0.028)	-0.479 (0.017)		0.426

Notes: The table reports OLS estimates of the coefficients in regressions of the depreciation rate, Δs_{t+1} , on aggregate order flow, x_t , trade-based order flow, x_t^T , and the order flows from short- and long-term investors, x_t^S and x_t^L . The regressions are estimated with 10,000 months of simulated data computed from the solution of our model using the baseline parameterization with $\gamma^S = \gamma^L = 1$. OLS standard errors are reported in parenthesis.

In row (v) we report the results from regressing the depreciation rate on all three components of customer order flow. Here we see two important implications of the model. First, the coefficients on the order flow components are quite different from each other. The coefficients on the short-term component, x_t^S , is twice the size of the estimate from the univariate regression in row (ii), and the coefficient on the long-term component, x_t^L , falls from the univariate estimate of 0.22 to -0.48. Thus, the model can produce heterogeneity of the kind observed in the coefficient estimates on Citibank’s customer order flows. The second implication concerns the R^2 statistic of 0.43. The individual components of customer order flow in our model account for almost 30 percent more of the variation in dealer quotes than aggregate customer order flow alone.

The results in Table 3 are robust to variations in the parameter values used to solve the model. They arise from two key features: First, realizations of aggregate order flow, x_t , are much less informative about shocks to fundamentals than the flow of information reaching dealers between the start of month t and $t + 1$. In our model dealers make inferences regarding fundamentals by combining their observation on x_t with the history of macro announcements (i.e., v_{t-1}, v_{t-2}, \dots) and earlier flows (i.e., x_{t-1}, x_{t-2}, \dots). The resulting dealer estimates of v_t are much more precise than any estimate of v_t based solely on x_t . Consequently, variations in x_t account for a relatively small fraction of the variation in depreciation rates because they poorly approximate the actual flow of information on fundamentals dealers use in revising the FX quotes. The second key feature concerns the correlations between the order flow components. In our simulations the correlation between the investor components is approximately 0.9, and the correlation between the trade-based and either investor component is approximately 0.3. These positive correlations are a robust feature of the model because, as equations (10), (15) and (16) show, the components of order flow are driven by

common shocks to fundamentals. At the same time, the components are less than perfectly correlated thanks to differences in investor preferences and the presence of liquidity shocks. These characteristics imply that variations in the three order flow segments are jointly a better approximation to the flow of information affecting dealer quotes than a single order flow, even the aggregate order flow observed by dealers. This is why the R^2 statistic in row (v) of Table 3 is almost 30 percent higher than the statistics from regressions using a single order flow. The correlation across order flow segments also means that the coefficients on the order flow components have no structural interpretation. Realizations of v_t , v_{t-1} , ε_t and ε_{t-1} characterize the flow of information from aggregate order flow and macro announcements dealers use to revise quotes (see equation 14). The different coefficients on the order flows simply reflect how these shocks are reflected in the depreciation rate and the order flow components.⁹

Although there is no exact mapping from Citibank’s customer flows to the components of customer order flow in our model, the discussion above provides theoretical perspective on the empirical findings in Table 2. Citibank’s customer flows are representative of the customer orders received by other (large) dealers, but they are not the vehicle through which information is impounded into FX quotes. Rather they represent one factor driving interdealer order flow that is itself an important source of information to dealers setting FX quotes across the market. Essentially the same relationship exists in our model between the components of customer order flow and the aggregate order flow used by dealers to revise quotes. The one difference is that Citibank’s customer orders undoubtedly contain idiosyncratic elements that are absent from x_t^T , x_t^S and x_t^L . According to this view then, the results in Table 2 reflect the fact that Citibank’s customer flows *approximate* the flow of information dealers are using across the market to revise their FX quotes. The individual coefficients have no structural interpretation in terms of measuring the price-impact of different customer orders, they simply map variations in customer flows into an estimate of the information flow being used by dealers across the market. This interpretation can also account for the increase in explanatory power of customer flows as we move from horizons of one day to one month. As the horizon lengthens, the idiosyncratic elements in Citibank’s customer flows become relatively less important, with the result that the flows are more precise proxies for the market-wide flow of information driving quotes. This interpretation is also consistent with the observation from Table 1 that the correlations between customer flow segments are much higher at the monthly frequency than at the daily frequency.

2.3 Forecasting Returns

One novel implication of our model is that it takes time before shocks to fundamentals are fully reflected in exchange rates. This fact is clear from the equation for the equilibrium depreciation process:

$$\Delta s_{t+1} = u_{t+1} + \lambda v_t + (1 - \lambda)v_{t-1} + \pi\varepsilon_t - \pi\varepsilon_{t-1}. \quad (14)$$

⁹For example, liquidity shocks, ε_t , contribute one-for-one to trade-based order flow, x_t^T , and increases the log FX quote between t and $t + 1$ (see equation 14). Since liquidity shocks have no effect on the other order flow components, regression analysis tells us that the coefficient on x_t^T in row (v) of Table 3 reflects these co-movements.

Recall that in the extreme case where order flow is completely uninformative about fundamentals, λ equals 0, so v_t shocks affect dealer quotes with a two month lag. At the other extreme, if observations on order flow allow dealers to make completely precise estimates of v_t by the start of month $t + 1$, λ equals one. In equilibrium, the value of λ lies between these extremes because order flow is somewhat informative. In fact, our numerical results in Figure 1 showed that the equilibrium value of λ is higher when the pool of customers placing orders with dealers is more forward-looking.

In this section, we examine the implications of equation (14) for the intertemporal relationship between customer order flows and returns. We begin by examining the relationship between future excess returns and current customer order flows. In particular, Table 4 reports the results of regressing excess returns between month t and month $t + 1$, er_{t+1} , on Citibank's customer order flows during month $t - 1$. Notice that unlike the regressions in Table 2, the customer order flows come from trading prior to the period covering the revision in quotes.

Table 4: Forecasting Regressions

Corporate		Hedge		Investors		R^2	χ^2
US	Non US	US	Non US	US	Non US		
1.179	-0.051					0.119	18.041
(0.306)	(0.133)						(<0.001)
		0.090	0.135			0.010	1.116
		(0.160)	(0.173)				(0.572)
				-0.965	0.131	0.110	15.434
				(0.264)	(0.109)		(<0.001)
0.985	-0.008	0.001	-0.038	-0.762	0.146	0.185	33.629
(0.259)	(0.137)	(0.136)	(0.182)	(0.242)	(0.128)		(<0.001)

Notes: The table reports OLS estimates of the coefficients in the regression of monthly excess returns, er_{d+h} , on the customer order flow segments from trading on days $d - h$ to $d - 1$. Estimates are computed at the daily frequency with $h = 20$ trading days. The table reports asymptotic standard errors in parentheses corrected for heteroskedasticity and an MA($h-1$) process. The right hand column reports Wald tests and p-values for the null that all the coefficients on the order flows are zero.

The results in Table 4 show that customer order flows have remarkable forecasting power for future quote revisions. The six flow segments jointly account for approximately 19 percent of the variation in future excess returns. This level of forecasting power is an order of magnitude above that usually found in exchange rate equations. For example, the forecast power of interest differentials for monthly excess returns is only in the 2 - 4 percent range. We also note, once again, that there are large differences in the estimated coefficients across flow segments and across specifications.

To gain theoretical perspective on these results, we return to simulations of our model. Table 5 reports the results of regressing the depreciation rate, Δs_{t+1} , on the aggregate and disaggregate customer order

flows from trading in month $t - 1$, $\{x_{t-1}^T, x_{t-1}^S, x_{t-1}^L\}$. As above, the regressions are estimated from 10,000 monthly observations of simulated data computed from a solution to the model using our benchmark parameterization.

Row (i) of Table 5 shows that aggregate customer order flow has no forecasting power for the future depreciation rate; both the coefficient estimate and R^2 statistic are very close to zero. By contrast, as the results in row (ii) show, the components of customer order flow have substantial forecasting power. In this case the R^2 statistic is approximately 0.25 – a figure somewhat higher than we found using all six segments of Citibank’s customer flows. The estimated coefficients on the three order flow segments are also quite different from each other, but each is highly statistically significant.

Table 5: Simulated Forecasting Regressions

	x_t^T	x_t^S	x_t^L	x_t	R^2
(i)				-0.001 (0.003)	<0.001
(ii)	-0.151 (0.004)	0.191 (0.033)	0.250 (0.019)		0.254

Notes: The table reports OLS estimates of the coefficients in the regression of the depreciation rate, Δs_{t+1} , on the aggregate and disaggregate customer order flows from trading in month $t - 1$, $\{x_{t-1}^T, x_{t-1}^S, x_{t-1}^L\}$. The regressions are estimated with 10,000 months of simulated data computed from the solution of our model using the baseline parameterization with $\gamma^S = \gamma^L = 1$. OLS standard errors are reported in parenthesis.

The results in Table 5 are easily understood using the equation for the equilibrium depreciation rate, equation (14). The equilibrium values for λ and π are approximately 0.57 and 0.22 under our baseline parameterization. These values imply that order flow is informative about fundamentals, but dealer inferences are not so precise that all the effects of v_t shocks are immediately reflected in FX quotes. In short, it takes time in this equilibrium for dealers to fully learn about shocks to fundamentals. As consequence, the market-wide flow of information driving the revision in dealer quotes between t and $t + 1$ depends on the history of v_t and ε_t shocks before they become common knowledge to dealers. This history-dependence shows up in the presence of the v_{t-1} and ε_{t-1} terms in equation (14). Now, insofar as order flow segments from month $t - 1$ trading are correlated with v_{t-1} and ε_{t-1} , these flows will be correlated with the market-wide information flow driving quote revisions, and so they will have forecasting power for the future depreciation rate as we see in Table 5.

It should be stressed that this result does not rely on inefficient inference by dealers. In our model, dealers make optimal use of their observations on aggregate order flow and macro announcements to revised their FX quotes. Rather, the forecasting power from the flow segments reflects the fact that they jointly contain more information about fundamentals than the observed history of aggregate order flow and macro

announcements up to the start of month t . Aggregating the segments destroys this information advantage, with the result that aggregate order flow should have no forecasting power for the future depreciation rate. Indeed, this is exactly what we see in row (i) of Table 5.

We conjecture that this same mechanism accounts for the forecasting power of Citibank’s customer flows. If these order flows contain more information about fundamentals than was contemporaneously known to dealers across the market, and some of this information became widely known via the ensuing interdealer order flow and incorporated into quotes, then the customer flows should have forecasting power for future excess returns. Once again, the estimated coefficients on the flow segments do not have a structural interpretation, they simply map variations in the flow segments into the flow of information concerning fundamentals that has yet to be fully assimilated by dealers across the market.

2.4 Returns, Order Flows and the Pace of Information Aggregation

We have argued that Citibank’s customer orders have forecasting power for future returns because they are correlated with the market-wide information flow that dealers use to revise their FX quotes. Moreover, this correlation arises because it takes time for information about fundamentals to be fully assimilated across the market. If, on the contrary, information about the latest change in fundamentals becomes available to dealers quickly because interdealer order flow is very informative, customer flows would have no forecasting power for future returns.

If the quality of market-wide information means that dealers are slow to fully learn about changes in fundamentals, customer order flows will contain private information on not just the latest changes in fundamentals, but also past changes that are not yet common knowledge. In our model, the revision in quotes between t and $t + 1$ depends on both v_t and v_{t-1} when $1 < \lambda < 0$ (see equation 14). This means that investors use their private information on both v_t and v_{t-1} when forecasting returns and placing their month- t FX orders. Aggregate customer order flow in month t will therefore aggregate private information on not just the most recent shock to fundamentals, v_t , but also the earlier shock, v_{t-1} , which is not yet known to all dealers. Notice that this backward-looking feature of order flow does not arise because investors are backward-looking. It occurs, instead, because forward-looking investors recognize that information aggregation across dealers is not instantaneous.

These observations provide a new perspective on the contemporaneous relationship between returns and customer order flows shown in Table 2. In particular, they raise the possibility that Citibank’s customer flows proxy for the market-wide flow of information received by dealers concerning a *history* of changes in fundamentals. If this is indeed the case, the projection of returns on contemporaneous customer flows should be forecastable with lagged flows. To investigate this possibility, we first compute the fitted values and residuals, $\hat{e}r_{t+1}$ and $\hat{\xi}_{t+1}$, from the contemporaneous regression for excess returns between month t and $t + 1$ on the six customer flow segments reported in the lower panel of Table 2. We then regress $\hat{e}r_{t+1}$ and $\hat{\xi}_{t+1}$ on the six flow segments from trading during month $t - 1$. The regression estimates are reported in Table 6.

The results in row (i) show that Citibank’s customer flows have significant forecasting power for the

Table 6: Order Flow Decompositions

	Corporate		Hedge		Investors		R^2	χ^2
	US	Non US	US	Non US	US	Non US		
(i)	0.202 (0.169)	-0.073 (0.069)	0.052 (0.047)	0.091 (0.112)	-0.239 (0.098)	0.010 (0.067)	0.070	13.383 (0.037)
(ii)	0.706 (0.244)	0.091 (0.112)	0.035 (0.100)	-0.252 (0.174)	-0.492 (0.186)	0.188 (0.123)	0.130	20.897 (0.002)

Notes: The table reports OLS estimates of the coefficients of regressions of $\hat{e}r_{d+20}$ (row i) and $\hat{\xi}_{d+20}$ (row ii) on the six order flow segments from trading on days $d - 20$ to $d - 1$, where $\hat{e}r_{d+20}$ and $\hat{\xi}_{d+20}$ are the fitted values and residuals from the regression of er_{d+20} on the six flows from days d to $d + 19$ (see the last row in Table 2). Estimates are computed at the daily frequency. The table reports asymptotic standard errors in parentheses corrected for heteroskedasticity and an MA(19) process. The right hand column reports Wald tests and p-values for the null that all the coefficients on the order flows are zero.

projection of returns on contemporaneous flows estimated by $\hat{e}r_{t+1}$. This finding is consistent with the idea that the market-wide information flow during month t contains information about earlier changes in fundamentals that are not yet common knowledge to dealers. Row (ii) of Table 6 reports the results of regressing the residuals from the contemporaneous regression, $\hat{\xi}_{t+1}$, on the lagged customer flows. Here too the flows have significant forecasting power. Over our sample they account for 13 percent of the variations in the residuals.

These findings are broadly consistent with our model. In Table 7 we report the results from regressing $\widehat{\Delta}s_{t+1}$ and $\hat{\zeta}_{t+1}$ on the lagged customer flow components, $\{x_{t-1}^T, x_{t-1}^S, x_{t-1}^L\}$ where $\widehat{\Delta}s_{t+1}$ and $\hat{\zeta}_{t+1}$ are the fitted values and residuals from the contemporaneous regression of Δs_{t+1} on x_t^T , x_t^S , and x_t^L shown in row (v) of Table 3. As in the Citibank data, we see that lagged customer flows have forecasting power for both the projection of the depreciation rate on the contemporaneous flow components, and the associated projection error. The only significant difference between these findings and the results in Table 6 is that the R^2 statistics are approximately 3.5 times larger – a point we shall return to.

The results in row (i) of Table 7 indicate that about one quarter of the market-wide information flow captured by variations in the contemporaneous projection relates to historical changes in fundamentals. Recall that these results are derived from simulations in which the equilibrium value of λ is approximately 0.57, so dealer learning about earlier changes in fundamentals makes a significant contribution to quote revisions. Citibank’s customer flows are a less precise proxy for the market-wide information flow than the customer flow segments in our model, so the R^2 statistic of 0.07 in row (i) of Table 6 almost surely understates the importance of dealer learning in the actual market.

Dealer learning also explains the results in row (ii) of Table 7. Here we see that lagged customer flows have forecasting power for the component of the depreciation rate that is uncorrelated with contemporaneous

Table 7: Simulated Forecast Decompositions

	x_t^T	x_t^S	x_t^L	R^2
(i)	0.049 (0.002)	-0.637 (0.021)	0.531 (0.013)	0.247
(ii)	-0.200 (0.002)	0.828 (0.020)	-0.281 (0.012)	0.490

Notes: The table reports OLS estimates of the coefficients from regressions of $\widehat{\Delta s}_{t+1}$ (row i) and $\hat{\zeta}_{t+1}$ (rows ii) on the lagged customer flow components, $\{x_{t-1}^T, x_{t-1}^S, x_{t-1}^L\}$, where $\widehat{\Delta s}_{t+1}$ and $\hat{\zeta}_{t+1}$ are the fitted values and residuals from the regression of Δs_{t+1} on x_t^T , x_t^S , and x_t^L shown in row (v) of Table 3. OLS standard errors are reported in parenthesis.

customer flows. Public news shocks, u_{t+1} , contribute to this component in our model (see equation 14), but they are uncorrelated with order flows by construction and so cannot account for the results in the table. Instead the forecasting power comes from elements in the market-wide information flow driving quote revision that are not proxied by current customer flows. Recall that dealers use observations on aggregate customer flow and macro announcements to revise their quotes. When λ is less than unity, these variables capture information from $v_t, v_{t-1}, \varepsilon_t$ and ε_{t-1} , an information flow that cannot be precisely represented by the projection of Δs_{t+1} on x_t^T, x_t^S and x_t^L . The associated projection error will therefore contain both v_{t-1} and ε_{t-1} , elements that are correlated with lagged customer flows. Thus, the three customer flow segments contain more information about fundamentals than aggregate flow, but they cannot be used to completely characterize the market-wide information flow when it takes time for dealers to learn about changes in fundamentals. Similarly, Citibank’s customer flows are not known across the market but appear to provide a partial contemporaneous characterization of the market-wide information flow. Other elements in this market-wide flow relate to earlier changes in fundamentals that are not yet common knowledge to dealers but are correlated with lagged Citibank’s customer flows. The R^2 statistic in row (ii) of Table 6 is lower than its counterpart in Table 7 because Citibank’s customer orders are less informative about the market-wide information flow than the components of customer order flow in our model.

3 Conclusion

This paper addressed order flow heterogeneity and its empirical implications. It is common for empiricists working in this area to apply results from stylized microstructure models when interpreting their results. Our analysis highlights the many pitfalls. For example, when employing datasets with flows from multiple end-user segments, some researchers have neglected to account for the contemporaneous correlation among the flow regressors. In this setting, estimated coefficients are not unbiased reflections of the total price-impact of

order flow from a given segment. Another pitfall stems from the fact that flow regressors are correlated across time. In this case, specifications that include contemporaneous flows only are reduced-forms for complex microeconomic dynamics, and cannot produce structural estimates of the price-impact of incremental trades. A third pitfall is that the data available to the researcher may not be available across the market in real time. In the FX market, for example, data on customer flows from a given bank do not capture the proximate flow driver of prices, interdealer order flow.

To clarify these matters, we compared simulated results from a model with customer-type heterogeneity with empirical estimates based on Citibank’s customer flows. Our simulations showed that: (1) customer flows provide more precise information about fundamentals when the mix of customers is tilted toward longer-horizon participants; (2) flows from customer segments can produce negative coefficients in contemporaneous return regressions, even when positively correlated with fundamentals; and (3) customer flows forecast returns because they are correlated with the future market-wide information flow that dealers use to revise their FX prices. Of these four simulation results, (2) is perhaps the most counter-intuitive. It arises because disentangling liquidity-motivated order flow from informative order flow in a dynamic setting is quite complex. On the empirical side, we showed that: (1) both the aggregate and disaggregated customer flows received by Citibank are positively auto-correlated; (2) contemporaneous correlations across flow segments are low at the daily frequency, but high at the monthly frequency; (3) some customer segments do produce negative coefficients in contemporaneous return regressions; (4) the proportion of excess return variation that segment flows can account for rises with the horizon; and (5) about one-third of order flow’s power to forecast exchange rates one month ahead comes from flow’s ability to forecast future flow.

An important direction for future work is to move toward true structural estimation. Part of the difficulty here is that there is no exact mapping between Citibank’s customer flows and the components of customer order flow in our model. The components of customer flow in our model are market-wide flows, whereas Citibank’s customer flows are but one sample from that market-wide flow. Moreover, Citibank’s segment flows are surely correlated with those from other banks, in part because many participant types split their orders across multiple banks (to reduce trading costs). This correlation introduces the potential for omitted-variable bias into any structural estimates. More fundamentally, a structural model must recognize that the customer orders received by Citibank (and other large dealers) are not the vehicle through which information is impounded into FX quotes. Rather, they correlate with price contemporaneously, and with future price, only to the extent that they correlate with the flow information that does drive quoted price proximately, namely market-wide interdealer flow. Put differently, flow information from a given bank is both more informative than the interdealer flow available to the whole market (in the sense that other banks do not observe Citibank’s flow mix), and less informative because the market has access to both interdealer flow and other sources of information that are not available to researchers. Any structural model needs to account for this complex information structure.

References

- Bacchetta, P., and E. van Wincoop (2003), Can information dispersion explain the exchange rate disconnect puzzle? NBER Working Paper 9498, February, forthcoming, *American Economic Review*.
- Bjønnes, G., D. Rime, and H. Solheim (2004), Liquidity provision in the overnight foreign exchange market, *Journal of International Money and Finance*, forthcoming.
- Carpenter, A., and J. Wang (2003), Sources of private information in FX trading, typescript, University of New South Wales, January.
- Dunne, P., H. Hau, and M. Moore (2004), Macroeconomic order flows: Explaining equity and exchange rate returns, typescript, November.
- Engel, C., and K. West (2005), Exchange rates and fundamentals, *Journal of Political Economy*, 113: 485-517.
- Evans, M., and R. Lyons (2002a), Order flow and exchange rate dynamics, *Journal of Political Economy*, 110: 170-180.
- Evans, M., and R. Lyons (2002b), Informational integration and FX trading, *Journal of International Money and Finance*, 21: 807-831.
- Evans, M., and R. Lyons (2003), How is Macro News Transmitted to Exchange Rates? NBER Working Paper 9433, January.
- Evans, M., and R. Lyons (2004a), A New Micro Model of Exchange Rates, NBER Working Paper 10379, March.
- Evans, M., and R. Lyons (2004b), Exchange Rate Fundamentals and Order Flow, typescript, Georgetown University, available at <http://www.georgetown.edu/faculty/evansm1/>.
- Evans, M., and R. Lyons (2005), Meese-Rogoff Redux: Micro-Based Exchange Rate Forecasting, *American Economic Review*, forthcoming, available at <http://www.georgetown.edu/faculty/evansm1/>.
- Froot, K., and T. Ramadorai (2002), Currency returns, institutional investor flows, and exchange rate fundamentals, NBER Working Paper 9101, August, forthcoming, *Journal of Finance*.
- Glosten, L., and P. Milgrom (1985), Bid, ask, and transaction prices in a specialist market with heterogeneously informed agents, *Journal of Financial Economics*, 14: 71-100.
- Hau, H., and H. Rey (2002), Exchange rates, equity prices, and capital flows, NBER Working Paper 9398, December, *Review of Financial Studies*, forthcoming.
- Kyle, A. (1985), Continuous auctions and insider trading, *Econometrica*, 53, 1315-1335.

- Marsh, I., C. O'Rourke (2005), Customer Order Flow and Exchange Rate Movements: Is There Really Information Content? Typescript, April.
- Mende, A., and L. Menkhoff (2003), Different counterparties, different foreign exchange trading? The perspective of a median bank, March.
- Osler, C. (1998), Short-term speculators and the puzzling behavior of exchange rates, *Journal of International Economics*, 45: 37-57.
- Payne, R. (2003), Informed trade in spot foreign exchange markets: An empirical investigation, *Journal of International Economics*, 61: 307-329.
- Rime, D. (2001), Private or public information in foreign exchange markets? An empirical analysis, typescript, Central Bank of Norway.
- Sager, M., and M. Taylor (2005), Order Flow and Exchange Rate Movements, Typescript, University of Warwick.

Appendix

This appendix provides technical description of how we solve the model. We proceed in three steps:

1. Guess the form of the equilibrium depreciation process. This is equation (14) in the text.
2. Derive an equation for customer order flows consistent with (14) and the information available to investors. This gives us the equations for the components of customer order flows in (10), (15) and (16).
3. Use the results from step 2 to show that equation (14) is consistent with the dynamics of the dealer quotes shown in equation (12) given the evolution of dealers' common information in (4).

Steps 1 and 2 are described in the text. To complete step 3, we first write the model in state space form:

$$\begin{bmatrix} f_{t+1} \\ x_{t+1}^T \\ v_{t+1} \\ u_{t+1} \\ x_t^T \\ v_t \\ v_{t-1} \\ v_{t-2} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} f_t \\ x_t^T \\ v_t \\ u_t \\ x_{t-1}^T \\ v_{t-1} \\ v_{t-2} \\ v_{t-3} \end{bmatrix} + \begin{bmatrix} 1 & 1 & 0 \\ \delta & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} v_{t+1} \\ u_{t+1} \\ \varepsilon_{t+1} \end{bmatrix},$$

$$\begin{bmatrix} u_t \\ x_{t-1} \\ v_{t-2} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \theta(\gamma^S \lambda + \gamma^L) & \theta[(1-2\lambda)\gamma^S - \lambda\gamma^L] & -\theta(1-\lambda)(\gamma^S + \gamma^L) \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} f_t \\ x_t^T \\ v_t \\ u_t \\ x_{t-1}^T \\ v_{t-1} \\ v_{t-2} \\ v_{t-3} \end{bmatrix},$$

or, more compactly

$$\begin{aligned} Y_{t+1} &= AY_t + BV_{t+1}, \\ Z_t &= CY_t. \end{aligned}$$

Note that Z_t is the vector containing the most recent information to all dealers at the start of month t .

Applying the steady state kalman filtering algorithm to this system gives:

$$\begin{aligned} Y_{t+1|t+1} &= AY_{t|t} + K(Z_{t+1} - Z_{t+1|t}), \\ \Sigma &= A((I - KC)\Sigma)A' + B\Omega_v B', \\ K &= \Sigma C' (C\Sigma C')^{-1}, \end{aligned}$$

where $Y_{t|t}$ denote estimates of Y_t conditioned on $\{Z_t, Z_{t-1}, \dots\} = \Omega_t^D$ and $\Omega_v = E[V_t V_t']$.

Now let ι be a 1×8 vector that picks out the first element in Y_t . We can write the log price of FX quote at the start of month $t + 1$ as

$$\begin{aligned} s_{t+1} &= \mathbb{E}_{t+1}^D f_{t+1} \\ &= \iota Y_{t+1|t+1} \\ &= \iota AY_{t|t} + \iota K(Z_{t+1} - Z_{t+1|t}) \end{aligned}$$

Next, note that $\iota A = \iota$, so the first term on the right is $\iota Y_{t|t} = s_t$. Making this substitution and using the equation for Z_t gives us

$$s_{t+1} = s_t + \iota KCBV_{t+1} + \iota KCA\hat{Y}_t, \quad (18)$$

where $\hat{Y}_t \equiv Y_t - Y_{t|t}$. Using the filter equations and the state space form we can show that this vector of estimation errors follows

$$\begin{aligned} \tilde{Y}_t &= AY_{t-1} + BV_t - AY_{t-1|t-1} - KC(Y_t - Y_{t|t-1}), \\ &= \Gamma A\tilde{Y}_{t-1} + \Gamma BV_t, \end{aligned}$$

where $\Gamma \equiv I - KC$. Combining this expression with equation (18) above gives

$$\Delta s_{t+1} = \iota KCBV_{t+1} + \iota KCA \sum_{i=0}^{\infty} \Pi_i V_{t-i},$$

where $\Pi_i \equiv (\Gamma A)^i \Gamma B$. It is straightforward to check that the Π_i are null vectors for $i > 1$, so this equation simplifies to

$$\Delta s_{t+1} = \iota KCBV_{t+1} + \iota KCA(\Pi_0 V_t + \Pi_1 V_{t-1}).$$

Direct calculations also verify that $\iota KCBV_{t+1} = u_{t+1}$, $\iota KCA\Pi_0 V_t = \lambda v_t + \pi \varepsilon_t$, and $\iota KCA\Pi_1 V_{t-1} = (1 - \lambda)v_{t-1} - \pi \varepsilon_{t-1}$. Hence, (18) has the same form as equation (14) in the text.