

Exchange Rate Fundamentals and Order Flow

October 26, 2004

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) 642-1059
lyons@haas.berkeley.edu

Abstract

This paper addresses whether transaction flows in foreign exchange markets convey information about fundamentals. We begin with a GE model in the spirit of Hayek (1945) in which fundamental information is first manifest in the economy at the micro level, i.e., in a way that is not symmetrically observed by all agents. With this information structure, induced foreign exchange transactions play a central role in the aggregation process, providing testable links between transaction flows, exchange rates, and future fundamentals. We test these links using data on all end-user currency trades received by Citibank over 6.5 years, a sample sufficiently long to analyze real-time forecasts at the quarterly horizon. The predictions are borne out in four empirical findings that define this paper's main contribution: (1) transaction flows forecast future macro variables such as output growth, money growth, and inflation, (2) transaction flows forecast these macro variables significantly better than spot rates do, (3) transaction flows (proprietary) forecast future spot rates, and (4) though proprietary flows convey new information about future fundamentals, much of this information is still not impounded in the spot rate one quarter later. A bottom line is that the significance of transaction flows for exchange rates extends well beyond high frequencies.

Keywords: Exchange Rate Dynamics, Microstructure, Order Flow.

JEL Codes: F3, F4, G1

¹We thank the following for valuable comments: Anna Pavlova, Andrew Rose, and seminar participants at the Fed (Board of Governors), UC San Diego, UC Berkeley, and the NBER (IFM). Both authors thank the National Science Foundation for financial support, which includes funding for a clearinghouse for recent micro-based research on exchange rates (at georgetown.edu/faculty/evansm1 and at faculty.haas.berkeley.edu/lyons).

Introduction

Exchange rate movements at horizons up to one year remain unexplained by observable macroeconomic variables (Meese and Rogoff 1983, Frankel and Rose 1995, Cheung et al. 2002). In their survey, Frankel and Rose (1995) describe evidence to date as indicating that "no model based on such standard fundamentals ... will ever succeed in explaining or predicting a high percentage of the variation in the exchange rate, at least at short- or medium-term frequencies." Seven years later, Cheung et al.'s (2002) comprehensive study concludes that "no model consistently outperforms a random walk."

This paper addresses this longstanding puzzle from a new direction. Rather than running a regression of exchange rates on macro variables, we address instead the microeconomic mechanism by which information is impounded in exchange rates by the market. One way to frame our approach is via the present value relation, in which the exchange rate can be expressed as the sum of two terms, one reflecting measured macro fundamentals, F_t , and the other unexplained by measured macro fundamentals, U_t :

$$s_t = F_t + U_t,$$

where s_t denotes the log nominal exchange rate. Naturally, both of these terms are discounted sums of expected current and future values, i.e.:²

$$F_t \equiv E_t \sum_{i=0}^{\infty} b^i f_{t+i}$$

and

$$U_t \equiv E_t \sum_{i=0}^{\infty} b^i u_{t+i}$$

where b^i is a discount factor, f_{t+i} is the measured macro fundamental, E_t is the expectations operator, and the u_{t+i} are unexplained by measured macro variables.

When addressing this decomposition empirically, an important element is the gap between two information sets, one that pins down the spot rate s_t and the other that pins down econometric estimates of $E_t f_{t+i}$. The spot rate is pinned down by the market, using all available information. In contrast, empirical estimates of $E_t f_{t+i}$ are a function of the econometrician's information set. As an empirical matter, then, one component of what econometricians estimate as U_t is in fact capturing departures of the econometrician's information set from that of the market.

Our approach focuses on this gap between the information sets of the econometrician and the market. Specifically, we address whether microeconomic information that is available to the market, but not available to the econometrician, is helpful in forming estimates of $E_t f_{t+i}$ (i.e., helpful for forecasting future macro variables). We recognize that a positive finding is not itself a resolution of the Meese-Rogoff determination puzzle. It is instead an investigation of what may be a missing link in that puzzle. More specifically, we

²We do not write "current and expected future values" because asset markets generally do not know the current values of macro variables (like GDP growth). See the discussion in Evans (2004).

would write the present value relation as follows:

$$s_t = F_t + X_t + Z_t,$$

where the term X_t captures information about the present value of future fundamentals that is not captured in macro-econometric measures of fundamentals F_t (and Z_t now captures shocks that are orthogonal to fundamentals). How does this information get impounded in the current price s_t ? This is where the micro view comes in: If transaction flows reaching the market (i.e., entering the price-setters' information sets) are conveying signals of future macro realizations—signals truly incremental to the public information that forms F_t —then marketmakers will impound this information in price when it is observed. This produces "unexplained" exchange rate variation relative to movements in macro-econometric measures of F_t .

We view our analysis as complementary to the recent finding by Engel and West (2004a) that the spot rate s_t and the discounted sum F_t are in fact linked in the way that theory would predict, i.e., s_t helps to forecast the future fundamentals in F_t . Our purpose is to suggest another channel through which the theory might find support, namely, that the microeconomic information available to markets for setting prices is in fact conveying information about fundamentals that has not been evident based on information available to the econometrician.

The type of fundamental information that we view F_t as missing in this new, information-theoretic approach is not concentrated "insider" information, but rather information that is dispersed around the economy and aggregated by the market (Hayek 1945). In textbook models, such information does not exist: relevant information is either symmetric economy-wide, or, sometimes, asymmetrically assigned to a single agent—the central bank. And, as a result, no textbook model predicts that marketwide order flow should drive exchange rates.³ The objective of this paper is to determine whether transaction flows do in fact convey information about macro fundamentals, and if so, to identify the specific macro variables involved.

To fulfill this objective, we need to proceed in a disciplined way. We develop a simple general equilibrium model of information aggregation that provides—in a setting of incomplete markets—a utility-based present value representation for exchange rates. (In this sense, it is an empirically friendly variation on the DGE model in Evans and Lyons 2004). The key output is a set of testable relationships between transaction flows, current and future exchange rate returns, and future fundamentals (e.g., money supplies). The framework incorporates a continuum of optimizing households that make consumption and portfolio decisions, a realistic set of asset markets, and financial intermediaries that quote security prices and fill household orders for financial assets. The model also includes sufficient goods market segmentation that firms can follow pricing policies consistent with empirical evidence.

The model's equilibrium has a number of noteworthy features. For example, the presence of dispersed information about fundamentals leads to a concurrent correlation between price change and order flow that matches the data. More novel, in the model order flow has superior forecasting power for *future* fundamentals than current spot rates. Related to this result, the model clarifies why dispersed information about

³Order flow is a measure of actual transactions, signed according to whether the initiator is on the buy side (+) or the sell side (−). In textbook models—all of which are based on information that is public—information arrival that is good for the dollar generates dollar appreciation to the market-clearing level immediately; at this new price, there should be no tilt toward transactions initiated by dollar buyers, i.e., no correlation between order flow and price changes in response to the public information.

fundamentals becomes impounded in spot rates only slowly. The equilibrium also has a feature that helps resolve the puzzling finding in past work that fundamentals and spot rates are not cointegrated. Specifically, when markets are incomplete and information is dispersed, tests of cointegration between standard measures of fundamentals and spot rates are incorrectly sized.

Our main empirical findings are consistent with our model: (1) transaction flows forecast *future* exchange rates changes (and do so more effectively than forward discounts), (2) transaction flows forecast subsequent macroeconomic variables such as money growth, output growth, and inflation, and (3) in cases where proprietary flows convey significant new information about future fundamentals, much of this information is still not impounded in the exchange rate itself three months later. All of these results are qualitative departures from findings in any micro-based empirical work thus far.

These results also shed new light from the macro perspective. For example, they direct attention away from understanding the U_t term in the present value relation as "missing fundamental variables" like the FX risk premium, and away from bubbles and behavioral explanations (without, of course, ruling these possibilities out). The results also indicate that information aggregation takes place on a macroeconomic time-scale, rather than on the ultra-high frequency time scale typically associated with the frenzied world of trading. Rather, the picture that emerges is nuanced, emphasizing flows of dispersed information, but within a framework for how exchange rates are determined that is explicitly macro. The approach is complementary to macro rather than competing.

The rest of the paper is organized as follows. Section 1 provides an overview of our modeling approach. Section 2 presents the model. Section 3 describes the data. Section 4 presents both specific empirical implications and results from tests of those implications. Section 5 concludes.

1 Modeling Approach

As noted, central to this paper is the distinction between different information sets. Engel and West (2004b) make a useful distinction along these lines between the information sets of the econometrician on the one hand, and the market on the other, where the latter is understood to contain the former. In this paper our focus is on the information set of the market, or, more precisely, the information sets of different market participants, and how those information sets interact. Decomposition of this "market" information set clarifies how information aggregation actually takes place.

There are three distinct information sets that we consider, one for each of the three agent types in our model: home households, foreign households, and dealers (i.e., marketmakers). We denote these three information sets as Ω_t^H , Ω_t^{H*} , and Ω_t^D , respectively. We distinguish between home and foreign households because we want to allow the information conveyed by actions of these two agent types to differ. (Constraining these two household information sets to be equal is simply a special case of our model.) We assume that $\Omega_t^D \subset \Omega_t^H$ and $\Omega_t^D \subset \Omega_t^{H*}$, that is, households have strictly more information than dealers. We do not take this assumption literally; it is instead something of a reduced form. In economic terms, what we have in mind is that the realized actions of households (i.e., their trades) convey incremental information about fundamentals to the dealers.⁴ It need not be the case that each of those households literally knows more than any given

⁴If the idea that households have strictly more information than dealers do is counter-intuitive, think of it this way: we

dealer. Rather, one could model this as households (or firms) that are trading purely for allocative reasons: for example, each household receives a money demand shock and is thereby privately motivated to trade foreign exchange. In this setting, none of the households would consider itself to have superior information. But the aggregate of those realized household trades would in fact convey information about the average household shock, i.e., the state of the macroeconomy. For parsimony, we do not model this heterogeneity at the home and foreign household level. Instead, we assume that households in any given country share the same information about the macroeconomy. Extending the model to capture this (admittedly more realistic) heterogeneity is a natural extension, but not one that would alter our main testable implications, so we have chosen this simpler, more transparent specification.

We consider households who have superior information of a rather general form (i.e., whose trades convey information of a rather general form). For example, we let $\Omega_t^H = \{\Omega_t^D, v_t\}$ for some vector of variables v_t . Upon observing households' trades, dealers update their estimates of fundamentals, knowing that the trades of households are based on this richer information set. This simple structure is sufficient to generate a host of readily testable implications. These include: (1) dispersed information about fundamentals should produce a strong contemporaneous correlation between order flow and spot rate changes, (2) the spot rate should be a relatively weak predictor of future fundamentals, (3) order flow should be a relatively powerful predictor of future fundamentals, and (4) tests of cointegration between spot rates and fundamentals are likely to over-reject cointegration. We should also note that our objective here is not to solve for the dynamics of every endogenous variable in our general equilibrium model. Rather, we want to put just enough structure on the problem to generate the empirical guidance that we seek.

The following two equations provide a clear picture of how these different information sets enter our model, and how they connect to the traditional present value relation presented in the introduction. These equations also provide some framing for the detail of the next section. The first of these equations is, on the face of it, a standard present-value relation:

$$s_t = \mathbb{E}_t^D \sum_{i=0}^{\infty} b^i f_{t+i}.$$

where, as in the introduction, f_{t+i} is the macro fundamental (in period $t+i$) and b^i is the discount factor. (The discount factor in our model will take the form $\left(\frac{1}{1+\eta}\right) \left(\frac{\eta}{1+\eta}\right)^i$, where η is the semi-elasticity of money demand to the nominal interest rate.) The difference here is the expectations operator, $\mathbb{E}_t^D$. The superscript D denotes conditioning on the dealers' information set, Ω_t^D . Our model thus explicitly recognizes that the exchange rate reflects information only when it becomes evident to the price-setters themselves; if the information is dispersed throughout the economy, but dealers do not yet have a sense for it, then it will not be in price.

The second equation is relevant for understanding how $\mathbb{E}_t^D (\Omega_t^D)$ evolves through time, i.e., how dealers learn. This is the equation that links the dealers' learning variable, order flow, to the differences in

are addressing here whether dealers, upon observing the transaction demands of end users, might learn something about the state of the macro-economy. In our model, order flow is observable by all agents once realized, so dealers do not observe flow more precisely than others. In actual FX markets, dealers do observe order flow more precisely than others (i.e., this type of information advantage to being a dealer exists). We chose to abstract from it to focus on what the dealers learn from the flow.

information sets described above. Specifically, order flow x_t in our model takes the form:

$$x_t \cong \phi(\mathbb{E}_t^H s_{t+1} - \mathbb{E}_t^D s_{t+1}) + \phi^*(\mathbb{E}_t^{H^*} s_{t+1} - \mathbb{E}_t^D s_{t+1})$$

Order flow, then, is a function of the gaps between household forecasts of next period's exchange rate and the dealers' forecast. As noted, in our model these gaps capture the presence of dispersed information. With these two equations, the basic story is evident: exchange rates move when dealer expectations move; and dealer expectations move when dealers see order flow, since order flow conveys information that is not yet publicly observable. Exchange rates will be adjusting to realized trades because that is the mechanism by which dispersed information about the state of the macro-economy is learned. Implications for the joint behavior of order flow, spot rates, and fundamentals will be our core empirical predictions.

2 The Model

In this section we outline the building blocks of the model. Additional analytic detail is provided in the appendix; for still further detail, see the related model in Evans and Lyons (2004). Readers more interested in our empirical results—the main contribution of the paper—might skim propositions 1 and 2 of this section and proceed to section 3 describing the data. Section 4 presents the specific empirical implications that we test.

2.1 Households and Firms

There are two countries populated by a continuum of households arranged on the unit interval $[0,1]$. We assume that half the households live in each country, and use the index $z \in [0, 1/2)$ to denote households in the home country, and $z^* \in [1/2, 1]$ to denote households in the foreign country. All households derive utility from consumption and real balances. Preferences for home consumers are given by:

$$\mathbb{U}_{z,t} = \mathbb{E}_t^H \sum_{i=0}^{\infty} \delta^i \left\{ \frac{1}{1-\gamma} C_{z,t+i}^{1-\gamma} + \frac{\chi}{1-\gamma} \left(\frac{M_{z,t+i}}{P_{t+i}} \right)^{1-\gamma} \right\}, \quad (1)$$

where $0 < \delta < 1$ is the discount factor, and $\gamma \geq 1$. $\mathbb{E}_t^H$ denotes expectations conditioned on home household information, $\Omega_{z,t}$. $M_{z,t}$ is the stock of domestic currency held by household z , and $C_{z,t}$ is a CES consumption index defined by:

$$C_{z,t} = \left((C_{z,t}^1)^{\frac{\theta-1}{\theta}} + (C_{z,t}^2)^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \quad (2)$$

where $C_{z,t}^j$, $j = \{1, 2\}$ is the consumption of good j by household z in period t . The corresponding home price index is

$$P_t = \left((P_t^1)^{1-\theta} + (P_t^2)^{1-\theta} \right)^{\frac{1}{1-\theta}}, \quad (3)$$

where, P_t^j is the home currency price of good j .

Foreign households preferences are given by:

$$\mathbb{U}_{z^*,t} = \mathbb{E}_t^{\text{H}^*} \sum_{i=0}^{\infty} \delta^i \left\{ \frac{1}{1-\gamma} C_{z^*,t+i}^{1-\gamma} + \frac{\chi}{1-\gamma} \left(\frac{M_{z^*,t+i}^*}{P_{t+i}^*} \right)^{1-\gamma} \right\}, \quad (4)$$

where $\mathbb{E}_t^{\text{H}^*}$ denotes expectations conditioned on foreign household information, $\Omega_{z^*,t}$. $M_{z^*,t}^*$ is the stock of foreign currency, and the foreign price index is given by:

$$P_t^* = \left((P_t^{*1})^{1-\theta} + (P_t^{*2})^{1-\theta} \right)^{\frac{1}{1-\theta}}, \quad (5)$$

where P_t^{*j} is the foreign currency price of good j .

Each household has access to an array of assets. In particular we assume that each household can hold domestic nominal bonds, foreign nominal bonds, and a portfolio of other assets, including both home and foreign equities. Let $B_{z,t}$ and $B_{z,t}^*$ denote the stock of home and foreign one-period nominal bonds held by household z , and let $A_{z,t}$ be the nominal value of the the portfolio of other assets, all held in period t .

There are two firms, one producing each good. Firms set the prices at which they will sell their good to households in each market. Let $Q_t^j \equiv P_t^j / S_t P_t^{*j}$ denote the relative price at which firm j sells goods in the home market relative to the foreign market. We allow for the fact the goods arbitrage is costly so that firm j may set prices such that Q_t^j differs from one. In other words, we do not assume that the law of one price (LOOP) holds good by good. Substituting for P_t^j in the home price index, we obtain:

$$\begin{aligned} P_t &= S_t \left((Q_t^1 P_t^{*1})^{1-\theta} + (Q_t^2 P_t^{*2})^{1-\theta} \right)^{\frac{1}{1-\theta}} \\ &= S_t P_t^* Q_t^1 \left(\frac{1 + (Q_t^2 P_t^{*2} / Q_t^1 P_t^{*1})^{1-\theta}}{1 + (P_t^{*2} / P_t^{*1})^{1-\theta}} \right)^{\frac{1}{1-\theta}}. \end{aligned}$$

Taking logs and a first-order approximation of the term on the right hand side around $Q_t^1 / Q_t^2 = 1$ and $P_t^{*2} / P_t^{*1} = \Upsilon$ gives:

$$p_t - s_t - p_t^* = \varphi q_t^1 + (1 - \varphi) q_t^2 = q_t$$

where $\varphi = 1 / (1 + \Upsilon^{1-\theta})$, and lowercase letters denote natural logs (for example, $p_t \equiv \ln P_t$).

This equation relates the log real exchange rate, $p_t - s_t - p_t^*$, to the pricing policies of the two firms via q_t^j . If consumption goods could be freely and instantaneously moved between countries, goods arbitrage would make it impossible for a firm to set the home and foreign currency prices of their product unless $q_t^j = 0$. By contrast, we assume that the presence of transactions and/or other costs segments goods markets allowing firms to set prices such that q_t^j differs from zero. Thus, deviations from LOOP will be the source of real exchange rate variations in the model as is consistent with the empirical evidence (see, e.g., Engel 1999).

At the beginning of each period, home households observe the return on their assets, r_t , and the goods market prices, P_t^1 and P_t^2 set by firms. They also see the home and foreign nominal interest rates, i_t and i_t^* , and the spot exchange rate that are quoted by financial intermediaries. With this information, household z , makes his consumption and portfolio allocation choices. Specifically, let $\alpha_{z,t}^{\text{B}^*} \equiv S_t P_t^{\text{B}^*} B_{z,t}^* / P_t W_{z,t}$ and $\alpha_{z,t}^{\text{A}} \equiv A_{z,t} / P_t W_{z,t}$ respectively denote the share of wealth held in foreign bonds and other assets, where

$W_{z,t}$ is the real value of wealth at the beginning of period t . Appendix A shows that the optimal portfolio choices are given by:

$$\begin{bmatrix} \alpha_{z,t}^{B*} \\ \alpha_{z,t}^A \end{bmatrix} = \frac{\rho}{\gamma} (\Xi_t^H)^{-1} \begin{bmatrix} \mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t + \frac{1}{2} \mathbb{V}_t^H(\Delta s_{t+1}) - \theta_{z,t}^s \\ \mathbb{E}_t^H r_{t+1} - i_t + \frac{1}{2} \mathbb{V}_t^H(\Delta s_{t+1}) - \theta_{z,t}^r \end{bmatrix}, \quad (6)$$

where

$$\theta_{z,t}^v = \gamma \mathbb{C} \mathbb{V}_t^H (c_{z,t+1} - w_{z,t+1}, v_{t+1}) + (1 - \gamma) \mathbb{C} \mathbb{V}_t^H (\Delta p_{t+1}, v_{t+1}),$$

for $v = s, r$ and Ξ_t^H is the conditional covariance matrix for the vector $(\Delta s_{t+1}, r_{t+1})'$. $\mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t - \theta_{z,t}^s$ and $\mathbb{E}_t^H r_{t+1} - i_t - \theta_{z,t}^r$ are the risk-adjusted expected excess returns on foreign bonds and other assets. (The variance terms arise because we are working with log excess returns.) $\theta_{z,t}^v$ identifies the consumption hedging factor associated with foreign bonds ($v = s$) and other assets ($v = r$). Optimal holdings of real balances satisfy:

$$m_{z,t} - p_t = \kappa + c_{z,t} - \eta i_t \quad (7)$$

where $\kappa \equiv \frac{1}{\gamma} \ln \chi + i \exp(i) \eta$ and $\eta = 1/\gamma(\exp(i) - 1) > 0$. Finally, the optimal choice of log consumption can be approximated by:

$$c_{z,t} - w_{z,t} = \frac{\rho k}{1 - \rho} + \left(1 - \frac{1}{\gamma}\right) \mathbb{E}_t^H \sum_{i=0}^{\infty} \rho^{i+1} (i_{t+i} - \Delta p_{t+1+i}) + \mathbb{E}_t^H \sum_{i=1}^{\infty} \rho^{i-1} (h_{t+i} - \lambda i_{t+i-1}), \quad (8)$$

where $k > 0, \lambda > 0, 1 > \rho > 0$, and h_t is the log excess return on wealth. Given expectations regarding future exchange rates, interest rates, inflation and the return on assets, equations (6), (7), (8) and the budget constraint describe the approximate behavior of home households. The behavior of foreign households is similar (see Appendix A for details).

This completes our description of the household behavior. We turn now to the central question of how securities prices are determined, which, naturally, focuses on how the information sets of price-setting agents evolve.

2.2 Financial Intermediaries

There are D dealers, indexed by d , that act as intermediaries in four financial markets: the home money and bond markets, and the foreign money and bond markets. As such, each dealer quotes prices at which they stand ready to buy or sell securities to households 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 exchange and interest rates are determined by the securities prices dealers choose to quote as the solution of a utility maximization problem. We therefore begin by considering the preferences and constraints that characterize of the optimization problem facing dealers. We do not model the entry decision that results in there being D dealers: one could add this by modeling entry among a set of "potential" dealers and introducing an ex-ante commission for intermediation.⁵

⁵To provide a bit more detail, dealers in this model would need to be compensated for the adverse selection that they will face in this model. The ex-ante commission for providing financial services would achieve this (but without adding any particular

For the sake of simplicity, we assume that all dealers are located in the home country. The preferences of dealer d are given by:

$$\mathbb{E}_t^D \sum_{i=0}^{\infty} \delta^i \frac{1}{1-\gamma} C_{d,t+i}^{1-\gamma}, \quad (9)$$

where $\mathbb{E}_t^D$ denotes expectations conditioned on the dealer's period t information, $\Omega_{d,t}$, and $C_{d,t}$ represents dealer d 's consumption of goods 1 and 2 aggregated via the CES function shown in (2). Note that dealer preferences only differ from those of households in that real balances have no utility value. As a consequence, dealers will not hold currency in equilibrium (since as an asset class currencies are dominated by bonds) and act solely as intermediaries between households and central banks in the money markets. The importance of this feature will become apparent when we derive the equations that determine exchange rates below.

Timing in the model is as follows (see Evans and Lyons 2004 for details within a similar trading structure). At the beginning of each period, each dealer d quotes a price $S_{d,t}$ at which he is willing to buy or sell foreign currency. He also quotes prices at which he is willing to buy or sell one period pure discount bonds— $P_{d,t}^B$ and $P_{d,t}^{B*}$ respectively. All prices are good for any quantity and are publicly observed. Then each dealer receives a orders for bonds and currency from a subset of households. Household orders are only observed by the recipient dealer and so represent a source of private information. Then each dealer quotes prices for foreign currency and bonds in the interdealer market. These prices, too, are good for any quantity and publicly observed, so that trading with multiple partners (e.g., arbitrage trades) is feasible. Based on these interdealer quotes, each dealer then chooses the amount of foreign currency, $T_{d,t}^{B*}$, he wishes to purchase (negative values for sales) by initiating a trade with other dealers. Interdealer trading is simultaneous and, to the extent trades are desired at a quote that is posted by multiple dealers, those trades are divided equally among dealers posting that quote. Finally, each dealer chooses the amount of home and foreign currency, $T_{d,t}^M$ and $T_{d,t}^{M*}$, they wish to purchase from the central banks.

The problem facing the dealer is to choose prices $S_{d,t}$, $P_{d,t}^B$ and $P_{d,t}^{B*}$, trades $T_{d,t}^B$, $T_{d,t}^{B*}$, $T_{d,t}^M$ and $T_{d,t}^{M*}$, and consumption $C_{d,t}$ to maximize expected utility (9). Formally, we can write the dealer's problem as:

$$\mathcal{J}_t(W_{d,t}) = \max_{\{\alpha_{d,t}^{A*}, \alpha_{d,t}^{B*}, C_{d,t}, P_{d,t}^B, P_{d,t}^{B*}, S_{d,t}\}} \left\{ \frac{1}{1-\gamma} C_{d,t}^{1-\gamma} + \beta \mathbb{E}_t^D \mathcal{J}(W_{d,t+1}) \right\}, \quad (10)$$

$$\text{s.t.} \quad W_{d,t+1} = \exp(i_t - \Delta p_{t+1}) (H_{d,t+1} W_{d,t} - C_{d,t} + \Pi_{d,t+1}), \quad (11)$$

$$\begin{aligned} \text{where } H_{d,t+1} &= 1 + (\exp(\Delta s_{t+1} + i_t^* - i_t) - 1) (\alpha_{d,t}^{B*} - \xi_t) + (\exp(r_{t+1} - i_t) - 1) \alpha_{d,t}^A, \\ \Pi_{d,t+1} &= \exp(\Delta s_{t+1} + i_t^* - i_t) (P_t^{B*} - P_{d,t}^{B*}) S_t T_t^{B*} + (S_{d,t} P_{d,t}^{B*} - S_t P_t^{B*}) T_t^{B*} + (S_{d,t} - S_t) T_t^{M*}, \\ \alpha_{d,t}^{A*} P_t W_{d,t} &= A_{d,t}, \\ \alpha_{d,t}^{B*} P_t W_{d,t} &= S_t P_t^{B*} (B_{t-1}^* + T_{d,t}^{B*} - \mathbb{E}_t^D T_t^{B*}), \text{ and} \\ \xi_t P_t W_{d,t} &= S_t P_t^{B*} (T_t^{B*} - \mathbb{E}_t^D T_t^{B*}). \end{aligned}$$

$W_{d,t}$ denotes the real wealth of dealer d at the start of period t . This comprises the dealer's holding of bonds and other assets: $W_{d,t} \equiv (B_{d,t-1} + S_t B_{d,t-1} + \exp(r_t) A_{d,t-1}) / P_t$. $\alpha_{d,t}^{B*}$ identifies the fraction of wealth dealer

richness to the aggregation problem being solved).

d wishes to hold in foreign bonds, taking into account expected orders from households and other dealers, $\mathbb{E}_t^D T_t^{B*}$. $H_{d,t+1}$ is the excess return on wealth between the start of periods t and $t+1$. This return depends upon the excess return on foreign bonds, and the actual fraction of wealth held in foreign bonds at the end of period t trading, $\alpha_{d,t}^{B*} - \xi_t$, where ξ_t represents the effects of unexpected foreign bond orders (the latter being important to the model). The dealer's trading profit is identified by $\Pi_{d,t+1}$. This is a function of the spread between the price quoted by the dealer and the price elsewhere in the market, multiplied by the flow of incoming orders for the asset. For example, $(S_{d,t} - S_t) T_t^{M*}$ identifies the profit for filling household orders for foreign currency at price $S_{d,t}$ when the price elsewhere in the market is S_t .

We solve the dealer's problem in the Appendix. In this environment, the optimal quotes for dealer d are given by:

$$\begin{aligned} S_{d,t} &= S_t = \mathcal{F}_s(\Omega_t^D) \\ P_{d,t}^B &= P_t^B = \mathcal{F}_B(\Omega_t^D), \\ P_{d,t}^{B*} &= P_t^{B*} = \mathcal{F}_{B*}(\Omega_t^D). \end{aligned}$$

where $\Omega_t^D = \cap_d \Omega_{d,t}^D$ is the information set common of all dealers at the beginning of period t . The functions $\mathcal{F}_s(\cdot)$, $\mathcal{F}_B(\cdot)$ and $\mathcal{F}_{B*}(\cdot)$ (described below) map elements of the information set into a common quote for foreign currency, home bonds and foreign bonds. In words, optimal quotes have the twin features of being common across all dealers, and a function of only common information.

To see why optimal quotes must have these features, consider how the choice of spot rate quote affects $\Pi_{d,t+1}$ via the last term $(S_{d,t} - S_t) T_t^{M*}$. Suppose dealer d quotes a price $S_{d,t} > S_t = \mathcal{F}_s(\Omega_t^D)$. Because all quotes are observable and are good for any amount, incoming orders for foreign currency will be negative ($T_t^{M*} < 0$) as dealers and households attempt to make arbitrage profits. Under these circumstances, $(S_{d,t} - S_t) T_t^{M*}$ has limiting value of $-\infty$. Similarly, if $S_{d,t} < S_t$, arbitrage trading will generate an incoming flow of foreign currency orders (i.e., $T_t^{M*} > 0$) so $(S_{d,t} - S_t) T_t^{M*}$ will again have a limiting value of $-\infty$. Thus, optimal quotes must be common across dealers to avoid the (expected utility) losses associated with arbitrage. This requires that quotes be a function of information that is known to all dealers, Ω_t^D .

Appendix B characterizes the optimal portfolio and consumption choices of the dealer by combining log linearized versions of the budget constraint with the first order conditions from (10) and (11). Dealer d 's optimal choice for the share of wealth in foreign bonds and other assets is

$$\begin{bmatrix} \alpha_{d,t}^{B*} \\ \alpha_{d,t}^A \end{bmatrix} = \left(\frac{1-\mu}{\gamma} \right) (\Xi_t^D)^{-1} \begin{bmatrix} \mathbb{E}_t^D \Delta s_{t+1} + i_t^* - i_t + \frac{1}{2} \mathbb{V}_t^D(\Delta s_{t+1}) - \theta_{d,t}^s \\ \mathbb{E}_t^D r_{d,t+1} - i_t + \frac{1}{2} \mathbb{V}_t^D(\Delta s_{t+1}) - \theta_{d,t}^r \end{bmatrix} \quad (12)$$

where

$$\theta_{d,t}^\omega = \gamma \mathbb{C} \mathbb{V}_t^D(c_{d,t+1} - w_{d,t+1}, \omega_{t+1}) + (1-\gamma) \mathbb{C} \mathbb{V}_t^D(\Delta p_{t+1}, \omega_{t+1})$$

for $\omega = s, r$ and Ξ_t^D is the conditional covariance matrix for the vector $(\Delta s_{t+1}, r_{t+1})'$. Notice that this expression takes the same general form as equation describing the portfolio choices of home households, except that the moments are conditioned on the dealer's information set, $\Omega_{d,t}$. The log consumption-wealth

ratio for dealer d is approximated by

$$c_{d,t} - w_{d,t} = \left(\frac{1}{\mu} - 1\right) k_d + \left(1 - \frac{1}{\gamma}\right) \mathbb{E}_t^D \sum_{i=0}^{\infty} (1 - \mu)^{i+1} (i_{t+i} - \Delta p_{t+1+i}) + \mathbb{E}_t^D \sum_{i=1}^{\infty} (1 - \mu)^{i-1} h_{d,t+i} \quad (13)$$

where $k_d > 0$, $1 > \mu > 0$ and $h_{d,t}$ is the log excess return on dealer's wealth. There are two differences between this expression and the equation for the log consumption wealth ratio of home households. First, dealers discount the future at rate $1 - \mu$ which is larger than the rate used by households ρ , because dealers are unconcerned with holding real balances. Second, expected excess returns on dealer wealth, $h_{d,t}$, generally differ from the excess returns expected by households because they hold different assets and have different information sets. This difference shows up in the last present value term on the right hand side.

2.3 Exchange and Interest Rate Determination

In the last two subsections, we characterized how dealers and households make consumption and portfolio decisions once prices for foreign currency and bonds have been quoted. Furthermore, we have argued that utility maximizing dealers will quote common prices for currency and bonds based information they *all* possess before trading starts. In this section, we turn to the question of how these quotes are related to that part of dealers' information that is common. This will pin down the determination of the spot exchange rate together with home and foreign interest rates.

From this point forward, our analysis will focus on specific differences between dealers and household information sets. To ease empirical implementation below, we keep things simple by assuming that all agents within a given group (dealers, home households, and foreign households) have the same information. With this simplification, we can use a representative agent within each of these three groups to describe behavior. The focus, then, will be on differences in the information sets available to home households, foreign households, and dealers.

The mapping from dealer's common information to quotes is identified by the requirements of market clearing. Because households are the sole holders of their national currencies, market clearing in the two money markets requires that the quotes for home and foreign bonds be set such that the implied interest rates satisfy:

$$-\ln P_t^B \equiv i_t = i - \frac{1}{\eta} \mathbb{E}_t^D (m_t - p_t - c_t), \quad (14)$$

$$-\ln P_t^{B*} \equiv i_t^* = i - \frac{1}{\eta} \mathbb{E}_t^D (m_t^* - p_t^* - c_t^*), \quad (15)$$

where i is a constant and quantities without a z subscript denote an aggregate (i.e., $m_t = \ln \int_z m_{z,t} dz$), or equivalently the quantity chosen by the representative home or foreign household. (Recall that η is the semi-elasticity of money demand to the nominal interest rate.) Thus m_t , m_t^* , c_t and c_t^* denote the log aggregate of home money, foreign money, home consumption and foreign consumption respectively. The right hand side of these equations identifies the interest rate that will equate the expected stock of currency with aggregate household demand. $\mathbb{E}_t^D$ now denotes the expectation of the representative dealer and is retained in both equations to allow for the possibility that dealers may not be able to precisely forecast m_t , m_t^* , c_t or c_t^* .

In addition to the requirements of money market clearing, we also impose a condition on the composition on the dealers portfolio, namely that dealers cannot hold speculative positions in foreign bonds. The actual foreign bond holdings at the end of period- t trading are given by $B_{d,t}^* = (\alpha_{d,t}^{B*} - \xi_t) P_t W_{d,t} / S_t P_t^{B*}$ and so are comprised of two components; desired holdings $\alpha_{d,t}^{B*} P_t W_{d,t} / S_t P_t^{B*}$, and unexpected holdings $\xi_t P_t W_{d,t} / S_t P_t^{B*} = T_t^{B*} - \mathbb{E}_t^D T_t^{B*}$. Our restriction on dealer portfolios sets $\alpha_{d,t}^{B*}$ to zero.⁶ Combining this restriction with (12) gives:

$$\mathbb{E}_t^D \Delta s_{t+1} + i_t^* - i_t = \beta^D \mathbb{E}_t^D er_{d,t+1},$$

where $\beta^D \equiv \mathbb{C}\mathbb{V}_t^D(r_{d,t+1}, \Delta s_{t+1}) / \mathbb{V}_t^D(r_{d,t+1})$ and $er_{d,t+1}$ is the risk adjusted excess return on dealers' other assets equal to $r_{d,t+1} - i_t - \frac{1}{2} \mathbb{V}_t^D(r_{d,t+1}) + \theta_{d,t}^r - (\theta_{d,t}^s - \frac{1}{2} \mathbb{V}_t^D(s_{t+1})) / \beta^D$. To focus on other aspects of the model we assume that $\mathbb{E}_t^D er_{d,t+1} = 0$ in the analysis below.⁷ Substituting for i_t and i_t^* with (14) and (15) and combining the result with the equation for the real exchange rate, leads to the following proposition:

Proposition 1 (spot rates as a present value) *The spot exchange rate evolves according to:*

$$s_t = \left(\frac{\eta}{1 + \eta} \right) \mathbb{E}_t^D s_{t+1} + \frac{1}{1 + \eta} \mathbb{E}_t^D f_t$$

where

$$f_t = [c_t^* - c_t + m_t - m_t^* - q_t]. \quad (16)$$

Solving this equation forward and applying the law of iterated expectations we obtain:

$$s_t = \mathbb{E}_t^D \left(\frac{1}{1 + \eta} \right) \sum_{i=0}^{\infty} \left(\frac{\eta}{1 + \eta} \right)^i f_{t+i}. \quad (17)$$

Equation (17) takes a standard form: the current spot rate is equal to the present value of fundamentals, f_t , as in traditional monetary models of the exchange rate. There are, however, two differences. First, the definition of fundamentals in (16) includes the difference between foreign and home consumption rather than income. Second, fundamentals affect the spot rate only via dealers' expectations. This is a particularly important feature of the model: Since the current spot rate is simply the common price of foreign currency quoted by dealers before trading starts, it must only be a function of information that is common to all dealers at the time, Ω_t^D . Equations (17) and (16) represent the mapping from elements of Ω_t^D to the spot exchange rate and so must include the conditional expectation operator $\mathbb{E}_t^D$.

Further insight into the behavior of the spot rate can had by decomposing fundamentals into two com-

⁶This assumption is akin to the requirement that dealers finish trading with no open forex positions (which, empirically, describes dealer behavior well; see Lyons 1995, Bjonnes and Rime 2004). Since expected excess returns are an increasing function of $\alpha_{d,t}^{B*}$, if there is no constraint on $\alpha_{d,t}^{B*}$ then dealers will quote S_t such that $\alpha_{d,t}^{B*}$ is infinitely large. Setting $\alpha_{d,t}^{B*}$ to zero rules out this form of market manipulation. Alternatively, we could place an upper bound on $\alpha_{d,t}^{B*}$ without changing our analysis because dealers would choose S_t to place $\alpha_{d,t}^{B*}$ at the bound.

⁷Allowing $\mathbb{E}_t^D er_{d,t+1}$ is a straightforward extension - we simply need to adjust the definition of fundamentals derived below. Any change in $\mathbb{E}_t^D er_{d,t+1}$ creates an incentive for dealers to take a speculative position in foreign bonds unless they set current spot rates so that the expected excess return on foreign bonds matches the change in $\beta^D \mathbb{E}_t^D er_{d,t+1}$.

ponents: $f_t = f_t^m + f_t^c$ where

$$\begin{aligned} f_t^m &= m_t - m_t^* + \left(\frac{1}{\gamma} - 1\right)q_t, \\ f_t^c &= c_t^* - c_t - \frac{1}{\gamma}q_t \end{aligned}$$

This decomposition is useful because the dynamics of f_t^m depend on monetary policy and the pricing decisions of firms while the dynamics of f_t^c are governed by the consumption decisions of households. To see this more clearly, we first use the consumption Euler equations to write

$$\begin{aligned} \Delta c_{t+1} &= \Delta c + \frac{1}{\gamma} (i_t - \Delta p_{t+1}) + c_{t+1} - \mathbb{E}_t^H c_{t+1} + \frac{1}{\gamma} (\Delta p_{t+1} - \mathbb{E}_t^H \Delta p_{t+1}), \\ \Delta c_{t+1}^* &= \Delta c + \frac{1}{\gamma} (i_t^* - \Delta p_{t+1}^*) + c_{t+1} - \mathbb{E}_t^H c_{t+1} + \frac{1}{\gamma} (\Delta p_{t+1}^* - \mathbb{E}_t^{H^*} \Delta p_{t+1}^*). \end{aligned}$$

Combining these equations with the definition of f_t^c gives

$$\begin{aligned} \Delta f_{t+1}^c &= \left(c_{t+1}^* - \mathbb{E}_t^{H^*} c_{t+1}^*\right) + \frac{1}{\gamma} (\Delta p_{t+1} - \mathbb{E}_t^H \Delta p_{t+1}) \\ &\quad + (c_{t+1} - \mathbb{E}_t^H c_{t+1}) + \frac{1}{\gamma} (\Delta p_{t+1}^* - \mathbb{E}_t^{H^*} \Delta p_{t+1}^*) \\ &\quad + \frac{1}{\gamma} (\Delta s_{t+1} - \mathbb{E}_t^D \Delta s_{t+1}) \end{aligned}$$

Under the assumption that dealers know less about future consumption and inflation than do households, the expectation of each term on the right conditioned on dealer's information, Ω_t^D , is zero, so that $\mathbb{E}_t^D \Delta f_{t+i}^c = 0$ for $i > 0$.

To make use of this observation, we first rewrite (17) as

$$s_t = \mathbb{E}_t^D f_t + \mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta}\right)^i \Delta f_{t+i}. \quad (18)$$

Combining this equation with our decomposition for fundamentals yields

$$s_t = \mathbb{E}_t^D f_t + \mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta}\right)^i \Delta f_{t+i}^m. \quad (19)$$

In terms of new economics, this is probably the most important equation of our model. Notice, once again, that expectations are defined over the marketmaker's information set, not a macro-econometric information set (the latter generally not containing information from conveyed by transaction quantities). The equation also makes clear that $s_t - \mathbb{E}_t^D f_t$ does not contain information useful for forecasting future changes in all fundamental variables, but rather a subset of those fundamentals (in this case the subset containing expectations of future Δf_{t+i}^m).

2.4 Transaction Flows

We now consider the implications of the interest and exchange rate equations for the transaction flows generated by the portfolio choices made by households. In particular, our aim is to identify the components that contribute to measures of order flow in the international bond markets.

Let x_t denote aggregate household order flow defined as the home currency value of aggregate household purchases of foreign bonds during period t trading. The contribution of home households to this order flow is $S_t (B_{z,t}^* - B_{z,t-1}^*) = \alpha_{z,t}^{B^*} P_t W_{z,t} \exp(i_t^*) - S_t B_{z,t-1}^*$ where $\alpha_{z,t}^{B^*}$ denotes the desired share of foreign bonds in the home households' wealth. Similarly, foreign households contribute $S_t (B_{z^*,t}^* - B_{z^*,t-1}^*) = \alpha_{z^*,t}^{B^*} S_t P_t^* W_{z^*,t} \exp(i_t^*) - S_t B_{z^*,t-1}^*$. Market clearing requires that aggregate holdings of foreign bonds by dealers and households sum to zero so that $B_{d,t-1}^* + B_{z^*,t-1}^* + B_{z,t-1}^* = 0$. Hence aggregate order flow can be written as

$$x_t = \left[\alpha_{z,t}^{B^*} \lambda_t + \alpha_{z^*,t}^{B^*} (1 - \lambda_t) \right] W_t \exp(i_t^*) + S_t B_{d,t-1}^*, \quad (20)$$

where $W_t \equiv W_t + S_t P_t^* W_{z^*,t}$ and $\lambda_t \equiv W_{z,t}/W_t$. Thus, order flow depends upon the portfolio allocation decisions of home and foreign households (via $\alpha_{z,t}^{B^*}$ and $\alpha_{z^*,t}^{B^*}$), the level and international distribution of household wealth (via W_t and λ_t) and the outstanding stock of foreign bonds held by dealers from last period's trading, $B_{d,t-1}^*$. These elements imply that order flow contains both pre-determined (backward-looking) and non-predetermined (forward-looking) components. The former are given by the level and distribution of wealth, the latter by the portfolio shares because they depend on households' forecasts of future returns.

To see this more clearly, we return to the determination of the home household's portfolio. Let $er_{z,t+1}$ be the risk adjusted excess return on the other assets held by home households equal to $r_{t+1} - i_t + \frac{1}{2} \mathbb{V}_t^H(\Delta s_{t+1}) - \theta_{z,t}^r$. We may now rewrite the portfolio allocation equation in (6) as

$$\alpha_{z,t}^{B^*} = \frac{\rho}{\gamma} \Theta^H \left(\mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t + \frac{1}{2} \mathbb{V}_t^H(s_{t+1}) - \theta_{z,t}^s \right) - \beta^H \Theta^H \mathbb{E}_t^H er_{z,t+1},$$

where $\beta^H \equiv \mathbb{V} \mathbb{C}_t^H(r_{t+1}, \Delta s_{t+1}) / \mathbb{V}_t^H(r_{t+1})$ and $\Theta^H \equiv \left(\mathbb{V}_t^H(\Delta s_{t+1}) - \frac{\gamma}{\rho} (\beta^H)^2 \mathbb{V}_t^H(r_{z,t+1}) \right)^{-1}$. Households know that dealers quote spot rates in accordance with (17). So the expected excess return on foreign bonds can be written as

$$\begin{aligned} \mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t &= \mathbb{E}_t^D \Delta s_{t+1} + i_t^* - i_t + \nabla \mathbb{E}_t^H s_{t+1} \\ &= \beta^D \mathbb{E}_t^D er_{d,t+1} + \nabla \mathbb{E}_t^H s_{t+1}, \end{aligned}$$

where $\nabla \mathbb{E}_t^\omega s_{t+1} \equiv \mathbb{E}_t^\omega s_{t+1} - \mathbb{E}_t^D s_{t+1}$ for $\omega = \{H, H^*\}$. Thus, $\nabla \mathbb{E}_t^H s_{t+1}$ denotes the spot rate forecast differential between home households and dealers. Combining this expression with the one above gives us

$$\alpha_{z,t}^{B^*} = \frac{\rho}{\gamma} \Theta^H \nabla \mathbb{E}_t^H s_{t+1} - \Theta^H \beta^H \mathbb{E}_t^H er_{z,t+1} - \Theta^H \left(\theta_{z,t}^s - \frac{1}{2} \mathbb{V}_t^H(s_{t+1}) \right). \quad (21)$$

Following the same steps for foreign households, we obtain

$$\alpha_{z^*,t}^{B^*} = \frac{\rho}{\gamma} \Theta^{H^*} \nabla \mathbb{E}_t^{H^*} s_{t+1} - \Theta^{H^*} \beta^{H^*} \mathbb{E}_t^{H^*} er_{z^*,t+1} - \Theta^{H^*} \left(\theta_{z^*,t}^s - \frac{1}{2} \mathbb{V}_t^{H^*}(s_{t+1}) \right) \quad (22)$$

Equations (21) and (22) show the desired portfolio shares for foreign bonds depend on: (i) the difference in expectations regarding future spot rates between the households and dealers, (ii) the risk adjusted expected excess return on other assets, and (iii) and the risk associated with holding foreign bonds.

The above equations produce another equation that is among the most important of this paper. Specifically, substituting the expressions for $\alpha_{z,t}^{B*}$ and $\alpha_{z^*,t}^{B*}$ in the order flow equation (20), and linearizing around the point where wealth is equally distributed between households and expectations are the same, yields the following proposition:

Proposition 2 (aggregate order flow and expected spot rates) *Aggregate order flow takes the form:*

$$x_t \cong \phi \nabla \mathbb{E}_t^H s_{t+1} + \phi^* \nabla \mathbb{E}_t^{H*} s_{t+1} + o_t \quad (23)$$

where o_t denotes the approximation terms involving the distribution of wealth and dealer's bond holdings.

Though the relation in proposition 1 is quite standard, this equation is quite new. It allows us to focus on how dispersed information—as manifested by the existence of the forecast differentials, $\nabla \mathbb{E}_t^H s_{t+1}$ and $\nabla \mathbb{E}_t^{H*} s_{t+1}$ —affects the joint behavior of order flow, spot rates, and fundamentals. This implied joint behavior provides our core empirical predictions, which are addressed in the next section. As an empirical matter, the terms o_t do not vary significantly from quarter to quarter under most circumstances, and so will not be the focus of the empirical analysis.

3 Data

Our empirical analysis utilizes a new data set that comprises end-user transaction flows, spot rates and macro fundamentals over six and half years. The transaction flow data is of a fundamentally different type and it covers a much longer time period than the data used in earlier work (e.g., Evans and Lyons 2002a,b). The difference in type is our shift from inter-marketmaker order flow to end-user order flow. By end users we are referring to three main segments: non-financial corporations, investors (such as mutual funds and pension funds), and leveraged traders (such as hedge funds and proprietary traders). Though inter-marketmaker transactions account for about two-thirds of total volume in major currency markets, they are largely derivative of the underlying shifts in end-user demands. Our data on the three end-user segments include all of Citibank's end-user trades in the largest spot market, the USD/EUR market, over a sample from January 1993 to June 1999.⁸ Citibank's end-user market share in these currencies is in the 10-15 percent range; no other bank has a larger market share in these currencies.

Advantages of our end-user data are many. First, as noted, the data are simply more powerful, covering much longer time spans. Second, because these trades reflect the world economy's primitive currency demands, the data provide a bridge to modern macro analysis. Third, the three segments span the full set of underlying demand types; those not covered by extant end-user data sets are empirically quite important for exchange rate determination, as we show below (c.f., Froot and Ramadorai 2002, which examines

⁸Before 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.

institutional investors only, and Osler 2003, which examines end-user stop-loss orders). Fourth, because the data are disaggregated into segments, we can address whether the behavior of these different flow measures is similar, and whether the information conveyed by each, dollar for dollar, is similar. The answer to both of these questions is no, as we shall see below; there is clear structure in the disaggregation, opening the door to structural insights not possible from aggregated transaction data.

Our forecasts of future fundamentals are based on real-time estimates of each fundamental variable. By "real time," we mean estimates that correspond to actual macroeconomic data available at any given time. It is, of course, these actual information sets, and the expectations that derive from them, that pin down asset pricing, in contrast to time-series on revised values of macro variables that become available many months later. (See Faust, Rogers, and Wright 2003 on this important distinction.) Specifically, we construct weekly real-time estimates of the three fundamentals we try to forecast that are based on the history of 35 different types of macro announcements (the three fundamentals being GDP growth, CPI inflation, and M1 money growth).⁹ The announcement data, including measures of ex-ante expectations for measuring news, are from MMS. Beyond the benefit of corresponding more closely to the information sets that pin down asset pricing, real-time estimates also have the advantage that they enable estimation of forecasting regressions at the weekly frequency, which would otherwise be impossible (given that variables like GDP growth are available only quarterly). Weekly frequency analysis greatly increases the efficiency of our forecasting analysis, which is helpful because we are interested in forecasting over horizons as long as a quarter or two, based on transactions data that span only 6.5 years.

Figure 1 provides an example of the forecasts produced by this real-time procedure. The lower panel is the one that is relevant for our analysis below of whether flows forecast fundamentals. But since that panel is derived from the upper panel, we begin with the upper panel. That upper panel plots demeaned "final" GDP growth releases (dotted line) and real-time estimates of those final releases (solid line), quarter by quarter. The final data release is available at the end of the third month following the quarter in question, making the average delay 4.5 months. So, for example, the high growth rate shown on the upper plot as released at the end of the first quarter of 1994 corresponds to the growth rate for the last quarter of 1993. Clearly, the real-time estimate does a good job anticipating those final estimates before they actually arrive (in part due to arrival of advanced and preliminary GDP growth estimates, which arrive at the ends of the first and second months following the quarter in question, respectively). (See Appendix B for further detail on doing this type of real-time estimation, which is drawn from Evans 2004 .) The lower plot is time-shifted relative to the upper plot, so that the cumulative GDP levels implied by the upper panel appear during the quarter to which they apply (e.g., the sharp upturn in growth noted in the upper panel for the first quarter of 1994 is now a rapidly rising real-time GDP level in the last quarter of 1993). It is this solid line in the lower panel that forms the basis of our forecasting analysis in the next section. Specifically, we will address whether order flows and/or spot rates can forecast changes in the level of this solid-line (up to two quarters ahead).

The real-time estimates used here are conceptually distinct from the real-time *data* series studied by Croushore and Stark (2001), Orphanides (2001) and others. A real-time data series comprises a set of

⁹Of the 35 announcements, 21 are US announcements and 14 are German announcements. In constructing the real-time forecasts of US (German) variables, we use US (German) announcements only.

historical values for a variable that are known on a particular date. This date identifies the vintage of the real-time data. For example, the March 31st vintage of real-time GDP data would include data releases on GDP growth up to the fourth quarter of the previous year. This vintage incorporates current revisions to earlier GDP releases but does not include a contemporaneous estimate of GDP growth in the first quarter. As such, it represents a subset of public information available on March 31st. By contrast, the March 31st real-time estimate of GDP growth comprises an estimate of GDP growth in the first quarter based on public information available on March 31st.

4 Empirical Implications and Results

Our model has a number of specific empirical implications, as we show below. To summarize: first, dispersed information about fundamentals across households and dealers can account for the strong concurrent correlation between order flow and spot rate changes observed in the data. Second, spot rates are unlikely to have much forecasting power for future fundamentals. Third, our model implies that tests for cointegration between spot rates and fundamentals will have the wrong size in typical data samples. Fourth, order flow should convey more precise information about future fundamentals than do current spot rates.

4.1 Order Flow's Link to Current Spot Rates

Our model provides a structural interpretation of the strong correlation observed between aggregate order flow and the change in spot rates. To see how this arises, recall that aggregate order flow depends on the forecast differentials, $\nabla \mathbb{E}_t^H s_{t+1} \equiv \mathbb{E}_t^H s_{t+1} - \mathbb{E}_t^D s_{t+1}$ and $\nabla \mathbb{E}_t^{H*} s_{t+1} \equiv \mathbb{E}_t^{H*} s_{t+1} - \mathbb{E}_t^D s_{t+1}$. We need to show why these differentials are correlated with spot rate changes.

For this purpose, let the vector of fundamentals be a linear combination of variables: $f_t = C y_t$ where y_t is a vector that follows:

$$\Delta y_{t+1} = A \Delta y_t + u_{t+1}, \quad (24)$$

with u_{t+1} a vector of mean zero i.i.d. shocks. This is a completely general linear specification for the dynamics of fundamentals that allows us to represent the behavior of the spot rate in a simple way. Specifically, combining (24) with (17) gives:

$$s_t = \pi \mathbb{E}_t^D y_t, \quad (25)$$

where $y_t' \equiv [y_t', \Delta y_t']$ and $\pi \equiv C \iota_1 + \frac{\eta}{1+\eta} C (I - \frac{\eta}{1+\eta} A)^{-1} A \iota_2$, with $y_t = \iota_1 y_t$ and $\Delta y_t = \iota_2 y_t$.¹⁰ π is a vector that relates the spot rate to dealers current estimate of the state vector y_t . We can now write the home forecast differential as:

$$\mathbb{E}_t^H s_{t+1} - \mathbb{E}_t^D s_{t+1} = \pi \left(\mathbb{E}_t^H \mathbb{E}_{t+1}^D y_{t+1} - \mathbb{E}_t^D y_{t+1} \right). \quad (26)$$

As above, we continue to assume that households know as much about the state of the economy as dealers do.

¹⁰To derive this equation, we use (24) to substitute for fundamentals in (18):

$$s_t = C \mathbb{E}_t^D y_t + C \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta} \right)^i A^i \mathbb{E}_t^D \Delta y_t.$$

This means that the term on the right-hand side is equal to $\pi \mathbb{E}_t^H (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1}$. In other words, the forecast differential for future spot rates depends on the households' expectations regarding how dealers revise their estimates of the future state, $\mathbf{y}_{t+1}$. As one might expect, this difference depends on the information sets, Ω_t^H and Ω_t^D . Clearly, if $\Omega_t^H = \Omega_t^D$, then $\mathbb{E}_t^H (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1}$ must equal a vector of zeros because $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1}$ must be a function of information that is not in Ω_t^D . Alternatively, suppose that households have superior information so that $\Omega_t^H = \{\Omega_t^D, v_t\}$ for some vector of variables v_t . If dealers update their estimates of $\mathbf{y}_{t+1}$ using elements of v_t , then some elements of $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1}$ will be forecastable based on Ω_t^H .

We can formalize these ideas using the Bayesian updating formula for an arbitrary information set Ω_t^ω :

$$\begin{aligned} \mathbb{E}[\boldsymbol{\varkappa}_{t+1} | \Omega_t^\omega, v_t] &= \mathbb{E}[\boldsymbol{\varkappa}_{t+1} | \Omega_t^\omega] + \mathbb{B}_{\boldsymbol{\varkappa}, v} (v_t - \mathbb{E}[v_t | \Omega_t^\omega]), \\ \mathbb{B}_{\boldsymbol{\varkappa}, v} &= \mathbb{V}_t^\omega (v_t)^{-1} \mathbb{C} \mathbb{V}_t^\omega (\boldsymbol{\varkappa}_{t+1}, v_t). \end{aligned}$$

Applying this equation in the case where $\boldsymbol{\varkappa}_{t+1} = \mathbb{E}[\mathbf{y}_{t+1} | \Omega_{t+1}^D]$, $\Omega_t^\omega = \Omega_t^D$, and $\Omega_t^H = \{\Omega_t^D, v_t\}$, gives:

$$\mathbb{E}_t^H \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} = \mathbb{B}_{\mathbb{E}_{t+1}^D \mathbf{y}_{t+1}, v_t} (v_t - \mathbb{E}[v_t | \Omega_t^D]).$$

In the case where $\boldsymbol{\varkappa}_{t+1} = \mathbf{y}_{t+1}$, $\Omega_t^\omega = \Omega_t^D$, and $\Omega_t^H = \{\Omega_t^D, v_t\}$ we get:

$$\mathbb{E}_t^H \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} = \mathbb{B}_{\mathbf{y}_{t+1}, v_t} (v_t - \mathbb{E}[v_t | \Omega_t^D]).$$

Combining these equations we obtain:

$$\begin{aligned} \mathbb{E}_t^H \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} &= \kappa (\mathbb{E}_t^H \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}), \\ \text{where } \kappa &\equiv \mathbb{B}_{\mathbb{E}_{t+1}^D \mathbf{y}_{t+1}, v_t} (\mathbb{B}'_{\mathbf{y}_{t+1}, v_t} \mathbb{B}_{\mathbf{y}_{t+1}, v_t})^{-1} \mathbb{B}'_{\mathbf{y}_{t+1}, v_t}. \end{aligned} \tag{27}$$

The matrix κ determines the expected speed at which households' information v_t is assimilated by dealers. We shall refer to this transmission process as information aggregation. When κ equals the identity matrix, households expect dealers to assimilate immediately all the information they used to forecast $\mathbf{y}_{t+1}$, so information aggregation is complete. When κ equals the null matrix, none of the information is assimilated and there is no information aggregation.

Information aggregation is critical to understanding the origins of the order flow spot rate correlation. To see why, we first combine (26) and (27) to give $\nabla \mathbb{E}_t^H s_{t+1} = \pi \kappa \nabla \mathbb{E}_t^H \mathbf{y}_{t+1}$. Applying the same technique to the foreign forecast differential gives $\nabla \mathbb{E}_t^{H*} s_{t+1} = \pi \kappa^* \nabla \mathbb{E}_t^{H*} \mathbf{y}_{t+1}$ where κ^* is the foreign counterpart of κ . Using these expressions, we can now express aggregate order flow as:

$$x_t = \phi \pi \kappa \nabla \mathbb{E}_t^H \mathbf{y}_{t+1} + \phi^* \pi \kappa^* \nabla \mathbb{E}_t^{H*} \mathbf{y}_{t+1} + o_t. \tag{28}$$

In the absence of information aggregation, $\kappa = \kappa^* = 0$, and order flow contains no information about the differences between households and dealers in the forecasts of the future state. In these circumstances, order flow ceases to have a forward-looking component that will give it forecasting power for both spot rates and fundamentals.

Accounting for the correlation between order flow and the spot rate is now straightforward. First, we use

(25) and the $\Delta s_{t+1} \equiv \mathbb{E}_t^D \Delta s_{t+1} + s_{t+1} - \mathbb{E}_t^D s_{t+1}$ to write the change in the spot rate as:

$$\Delta s_{t+1} = i_t - i_t^* + \pi \left(\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} \right).$$

(recall that $\mathbb{E}_t^D \Delta s_{t+1} = i_t - i_t^*$). Next, we use our updating formula to compute the change in the dealers' expectation over the period $\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}$. Setting $\mathbf{z}_{t+1} = \mathbf{y}_{t+1}$, $\Omega_t^\omega = \Omega_t^D$, and $\Omega_{t+1}^D = \{\Omega_t^D, v_t\}$, this gives:¹¹

$$\left(\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} \right) = \mathbb{B}_{\mathbf{y}_{t+1}, v_t} (v_t - \mathbb{E}[v_t | \Omega_t^D]).$$

The vector v_t denotes the new information available to dealers between the start of period t and the start of $t+1$ (by start, we mean before observing any actions within that period). Thus, period t order flow x_t is an element of v_t . We can therefore write:

$$\Delta s_{t+1} = i_t - i_t^* + b(x_t - \mathbb{E}_t^D x_t) + \zeta_{t+1}.$$

where $b = \pi \mathbb{B}_{\mathbf{y}_{t+1}, x_t}$ and ζ_{t+1} denotes the effect of other elements in v_t that are uncorrelated with order flow. The intuition for this equation is clear: ex-post departures from uncovered interest parity are made up of two kinds of news (dealers' perspective): order flow news and other news. ($\mathbb{E}_t^D x_t$ is dealers' best estimate of period t order flow at the start of t .) To see how the correlation between order flow and spot rates depends on the degree of information aggregation, we simply use (28) to substitute for x_t in the definition of $\mathbb{B}_{\mathbf{y}_{t+1}, x_t}$. This gives:¹²

$$\begin{aligned} b &= \mathbb{V}_t^D(x_t)^{-1} \left(\phi \pi \mathbb{V}_t^D \left(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1} \right) \kappa' \pi' + \phi^* \pi \mathbb{V}_t^D \left(\nabla \mathbb{E}_t^{H*} \mathbf{y}_{t+1} \right) \kappa^{*'} \pi' \right) \\ &\quad + \mathbb{V}_t^D(x_t)^{-1} \pi \mathbb{C} \mathbb{V}_t^D(\mathbf{y}_{t+1}, o_t). \end{aligned}$$

This expression shows that the observed correlation between order flow and price changes—i.e., the order flow news—can arise through two channels. First, if household transaction flows convey incremental information about the future fundamentals, and that information is transmitted to dealers via trading, then order flow will be correlated with spot rate changes as part of an information aggregation process. Second, price changes will be correlated with order flow when the latter contains information about the distribution of wealth that is useful in forecasting fundamentals.

Now we turn to the empirical evidence, Table 1 presents regression results for the relation between realized excess currency returns and concurrent end-user flow. Two points emerge from the results. First, the price impact (dollar for dollar) of trades from different sources is quite different, implying that these different trade

¹¹Setting $\Omega_{t+1}^D = \{\Omega_t^D, v_t\}$ in the example is of course strong: it says that dealers learn everything about v_t within a single period.

¹²First we write

$$\begin{aligned} \pi \mathbb{B}_{\mathbf{y}_{t+1}, x_t} \mathbb{V}_t^D(x_t) &= \phi \pi \mathbb{C} \mathbb{V}_t^D \left(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1}' \right) \kappa' \pi' \\ &\quad + \phi^* \pi \mathbb{C} \mathbb{V}_t^D \left(\nabla \mathbb{E}_t^{H*} \mathbf{y}_{t+1}' \right) \kappa^{*'} \pi' + \pi \mathbb{C} \mathbb{V}_t^D(\mathbf{y}_{t+1}, o_t) \end{aligned}$$

Then we use

$$\mathbf{y}_{t+1} = \mathbb{E}_t^D \mathbf{y}_{t+1} + \mathbb{E}_t^\omega \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} + (\mathbf{y}_{t+1} - \mathbb{E}_t^\omega \mathbf{y}_{t+1})$$

for $\omega = H, H^*$.

types have different information content. Our model provides helpful perspective on why location should matter, but it would be over-interpreting the model to suggest that it is conclusive in this regard. Second, the explanatory power of flows for concurrent returns is substantial: at the monthly frequency, the R^2 statistic when all flow segments are included is 30 percent.¹³ Froot and Ramadorai 2002 also find stronger links between end-user flows and returns as the horizon is extended to 1 month; their flow measure is institutional investors, however, not economy-wide.

The framework developed in our model establishes the theoretical link between the presence of dispersed information, the pace of information aggregation, and exchange rate dynamics, including their relation to order flow. The results in Tables 1 are consistent with that theoretical link. But we need a way to verify that the specific theoretical mechanism described by the model is empirically important. We start this task by considering the implications of our model for the relation between fundamentals and spot rates.

4.2 Linking Spot Rates to Fundamentals

We examine the link between spot rates and fundamentals in two ways. First, we examine the model's implications for forecasting fundamentals with spot rates. Second, we study whether our model can account for the apparent lack of cointegration between the spot rates and fundamentals (see Engel and West 2004a).

Our model provides two reasons for why the forecasts of fundamentals provided by spot rates are relatively poor. First, the model implies that only a subset of fundamentals should be forecastable. Recall that equation (19) implies that $s_t - \mathbb{E}_t^D f_t$ may only have forecasting power for elements of Δf_{t+i}^m (i.e., Δm_{t+i} , Δm_t^* and $(\frac{1}{\gamma} - 1)\Delta q_{t+i}$) because optimal decision-making by households implies that $\mathbb{E}_t^D \Delta f_{t+i}^c = 0$, for $i > 0$. The second reason arises from the mis-measurement of fundamentals. Mis-measurement is not in itself a new idea—many papers have recognized that an incomplete definition of fundamentals may be contributing to the poor forecasting performance of the spot rates. Our model allows us to enrich this argument with some specifics.

Suppose we correctly identified fundamentals and attempted to forecast future changes in elements of $\Delta f_{t+\tau}^m$ using $s_t - f_t$. To see the problem associated with this approach, we first consider the projection onto $s_t - \mathbb{E}_t^D f_t$:

$$\Delta f_{t+\tau}^m = \beta_s (s_t - \mathbb{E}_t^D f_t) + \varepsilon_{t+\tau}. \quad (29)$$

where $\varepsilon_{t+\tau}$ is the projection error. The value of the projection coefficient can be calculated with the aid of (19):

$$\beta_s = \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta} \right)^i \frac{\text{Cov}(\mathbb{E}_t^D \Delta f_{t+i}, \mathbb{E}_t^D \Delta f_{t+\tau}^m)}{\text{Var}(s_t - \mathbb{E}_t^D f_t)}.$$

Now consider the result from forecasting $\Delta f_{t+\tau}^m$ using $s_t - f_t$:

$$\Delta f_{t+\tau}^m = \hat{\beta}_s (s_t - f_t) + \hat{\varepsilon}_{t+\tau}. \quad (30)$$

¹³It is noteworthy that these end-user flows convey information beyond that in the inter-marketmaker flows used in Evans and Lyons (2002a,b). When both types of flow are included in a regression of daily excess returns, we are able to reject the null hypothesis that the coefficients on end-user flows are zero at the 1 percent level.

In this case, the projection coefficient is given by:

$$\hat{\beta}_s = \beta_s \left(1 + \frac{\mathbb{V}(f_t - \mathbb{E}_t^D f_t)}{\mathbb{V}(s_t - f_t)} \right)^{-1}.$$

When dealers have incomplete information about the current level of fundamentals, $\mathbb{V}(f_t - \mathbb{E}_t^D f_t) > 0$, and $\hat{\beta}_s$ will be pushed below β_s . This is a form of attenuation bias that arises because the use of f_t rather than $\mathbb{E}_t^D f_t$ in the forecasting equation introduces an errors-in-variables problem. Thus, even if we have the correct variables in our definition of fundamentals, the fact that we don't have dealer's estimates of these fundamentals implies that $s_t - f_t$ will have less forecasting ability for future fundamentals than dealers actually do.

A similar errors-in-variables problem plagues tests for the presence of cointegration between the spot rate and fundamentals. To see why, we rewrite equation (18) as:

$$s_t - f_t = \mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1 + \eta} \right)^i \Delta f_{t+i} - (f_t - \mathbb{E}_t^D f_t).$$

According to this equation, s_t and f_t should be cointegrated when: (i) f_t follows a non-stationary $I(1)$ process, and (ii) the dealers' error in estimating the current level of f_t is stationary $I(0)$. When dealers observe money stocks and the real exchange rate, $f_t - \mathbb{E}_t^D f_t = f_t^c - \mathbb{E}_t^D f_t^c$, so uncertainty about household consumption decisions may contribute to the dynamics of $s_t - f_t$. Let us therefore examine this possibility.

Recall that $\mathbb{E}_t^D \Delta f_{t+1} = 0$ if households know at least as much about consumption and inflation as dealers. Under these circumstances we can represent the dynamics of f_t^c when markets are incomplete by $\Delta f_{t+1}^c = v_{t+1}$ where v_{t+1} is an i.i.d. mean zero shock that is uncorrelated with f_t and elements of Ω_t^D . If dealers receive a noisy signal of f_t^c every period equal to $f_t^c + \zeta_t$, then their estimates of f_t^c will follow:

$$\mathbb{E}_t^D f_t^c = \mathbb{E}_{t-1}^D f_{t-1}^c + \varphi (f_t^c + \zeta_t - \mathbb{E}_{t-1}^D (f_{t-1}^c)),$$

where $\varphi^D \equiv \frac{\mathbb{V}_t^D(f_t^c)}{\mathbb{V}_t^D(f_t^c) + \mathbb{V}_t^D(\zeta_t)}$. Combining this equation with $\Delta f_{t+1}^c = v_{t+1}$, we find that the estimation error follows an AR(1) process:

$$(f_t^c - \mathbb{E}_t^D f_t^c) = (1 - \varphi) (f_{t-1}^c - \mathbb{E}_{t-1}^D f_{t-1}^c) + (1 - \varphi) v_t - \varphi^D \zeta_t.$$

Notice that the autoregressive coefficient in this process approaches unity as the variance of the noise rises. Thus, dealers' estimation errors will be more persistent when the information they receive about the current level of f_t^c is less precise (they are learning very slowly).

The implications of our model for forecasting fundamentals and the behavior of $s_t - f_t$ are now clear. Dealers have imprecise information about a component of fundamentals, f_t^c , that follows a non-stationary process when markets are incomplete. As a consequence, $f_t^c - \mathbb{E}_t^D f_t^c$ will follow an AR(1) process with a near unit root, that in turn may make $f_t - \mathbb{E}_t^D f_t$ extremely persistent and the sample variance of $f_t - \mathbb{E}_t^D f_t$ very large. As a consequence, there may be a significant degree of attenuation bias in the estimating β_s from the fundamentals forecasting projection (29). In terms of cointegration, $f_t - \mathbb{E}_t^D f_t$ is stationary when $\varphi^D > 0$, so that s_t and f_t are indeed cointegrated. However, realizations of $f_t - \mathbb{E}_t^D f_t$ in any sample may

appear non-stationary so that conventional tests reject cointegration between s_t and f_t . The failure to find cointegration stems not from using the wrong definition of fundamentals, f_t . Rather, it results from the fact that dealers use $\mathbb{E}_t^D f_t$ to set spot rates and the difference between f_t and $\mathbb{E}_t^D f_t$ can be very persistent when markets are incomplete because the unobserved elements of f_t follow a random walk.

Table 2 examines the time series properties of the expectational errors associated with different fundamental macro variables. The upper panel reports the results from the cointegrating regression of the real-time estimate of the fundamental variable on its own ex-post value, i.e., it addresses the expectational error $f_t - \mathbb{E}_t^D f_t$. (The reported standard errors are computed by Dynamic OLS in daily data with 10 leads and lags to correct for finite-sample bias. Standard errors contain an MA(10) correction for residual serial correlation.) The p-values reported in parentheses are for the hypothesis that the cointegration coefficient equals unity. Note that this is rejected in five of the six cases at the one-percent level, suggesting that error-correction is difficult to detect. The lower panel reports daily autocorrelations for the real-time errors, defined as the difference between the ex-post value and real-time estimate of the fundamental variables. These remain quite high, even at the one-quarter horizon, for many of the variables, which is consistent with the persistence argument above. We view these cointegration-test results as offering promising prospects for resolving the puzzle of no-cointegration (between exchange rates and fundamentals) highlighted in past work.

4.3 Forecasting Fundamentals with Spot Rates and Order Flow

Our model shows that order flow should have forecasting power for future spot rates when households' transaction flows convey information about future fundamentals that is news to dealers. This interpretation of the order flow/price change correlation implies that order flow should have forecasting power for future fundamentals. Specifically, order flow should have incremental forecasting power for $\Delta f_{t+\tau}^m$ beyond $s_t - \mathbb{E}_t^D f_t$.

To see why, consider the projection of $\Delta f_{t+\tau}^m$ on $s_t - \mathbb{E}_t^D f_t$ and the unexpected component of order flow $x_t - \mathbb{E}_t^D x_t$:

$$\Delta f_{t+\tau}^m = \beta_s (s_t - \mathbb{E}_t^D f_t) + \beta_x (x_t - \mathbb{E}_t^D x_t) + \epsilon_{t+\tau}, \quad (31)$$

where $\epsilon_{t+\tau}$ is the projection error. Order flow has incremental forecasting power when β_x differs from zero. To show that this is indeed the case, we first note that $\beta_x (x_t - \mathbb{E}_t^D x_t) + \epsilon_{t+\tau}$ must equal the projection error in (29), $\epsilon_{t+\tau}$, because $x_t - \mathbb{E}_t^D x_t$ is uncorrelated with $s_t - \mathbb{E}_t^D f_t$. Consequently, β_s takes the same value as it did in (29) and:

$$\beta_x = \frac{\mathbb{C}\mathbb{V}(\Delta f_{t+\tau}^m, x_t - \mathbb{E}_t^D x_t)}{\mathbb{V}(x_t - \mathbb{E}_t^D x_t)}.$$

Using the identity $\Delta f_{t+\tau}^m \equiv \nabla \mathbb{E}_t^\omega \Delta f_{t+\tau}^m + \mathbb{E}_t^D \Delta f_{t+\tau}^m + (\Delta f_{t+\tau}^m - \mathbb{E}_t^\omega \Delta f_{t+\tau}^m)$ for $\omega = \{H, H^*\}$ to substitute for $\Delta f_{t+\tau}^m$, and (28) to substitute for order flow, we find that:

$$\beta_x = \frac{\phi \pi \kappa \mathbb{C}\mathbb{V}(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1}, \nabla \mathbb{E}_t^H \Delta f_{t+\tau}^m) + \phi^* \pi \kappa^* \mathbb{C}\mathbb{V}(\nabla \mathbb{E}_t^{H^*} \mathbf{y}_{t+1}, \nabla \mathbb{E}_t^{H^*} \Delta f_{t+\tau}^m) + \mathbb{C}\mathbb{V}(o_t, \Delta f_{t+\tau}^m)}{\mathbb{V}(x_t - \mathbb{E}_t^D x_t)}$$

This equation shows that order flow will have incremental forecasting power when: (i) households and dealers have different forecasts of future fundamentals (i.e., $\nabla \mathbb{E}_t^\omega \Delta f_{t+\tau}^m \neq 0$ for $\omega = H$ or H^*) and information aggregation is partial, or (ii) the distribution of wealth has forecasting power for fundamentals, i.e.,

$$\mathbb{C}\mathbb{V}(o_t, \Delta f_{t+\tau}^m) \neq 0.$$

This establishes the general conditions under which order flow has incremental forecasting power for fundamentals (beyond the power in $s_t - \mathbb{E}_t^D f_t$). Does this mean that order flow should also have incremental forecasting power relative to $s_t - f_t$? There is no unambiguous answer—the question is an empirical one. If the introduced measurement error of $f_t - \mathbb{E}_t^D f_t$ is uncorrelated with order flow, then attenuation bias will only affect β_s . Under these circumstances, the forecasting power in $s_t - \mathbb{E}_t^D f_t$ will be understated but order flow will continue to have forecasting power via β_x . There is, however, no good reason why $f_t - \mathbb{E}_t^D f_t$ and order flow should be uncorrelated. In this case, the estimates of both β_s and β_x will be affected by measurement error and it is conceivable that the estimate of β_x will be close to zero. Thus, order flow could appear not to have incremental forecasting power relative to $s_t - f_t$, even though it does relative to $s_t - \mathbb{E}_t^D f_t$.

Table 3 presents Granger Causality tests for a set of macro variables considered fundamental across a wide range of modeling traditions: relative output growth, relative money growth, and relative inflation. These results provide us with some preliminary evidence on whether order flow has incremental forecasting power for fundamentals and follow the methodology employed by Engel and West (2004a). The results in Table 3 show that order flow forecasts all three of these variables at the one-percent level. Interestingly, the spot rate itself is able to forecast only one of these three variables—relative output growth—at the one-percent level. Moreover, we find no evidence of Granger causality going the opposite direction for either order flow or the spot rate. In sum, the results in Table 3 provide an initial indication that there is fundamental information in these end-user flows.¹⁴

Table 4 addresses order flow’s ability to forecast macro variables in a richer way. Specifically, we want our analysis to answer three questions that the simple Granger causality analysis does not: (1) what is the incremental forecasting power in order flow? (2) how does this forecasting power change with the forecasting horizon? and (3) is order flow able to forecast *real-time* measures of fundamentals, that is, estimates that determine the exchange rate in real time based on macroeconomic data available at that time (rather than ex-post measures of fundamentals that were available to market participants only with considerable lags; see the Data section for more detail). The regressions are estimated in weekly data, 284 observations, with weekly real-time estimates of the forecasted fundamental that are based on the history of 35 different types of macro announcements. Real-time estimates of these variables is what enables estimation of these regressions at the weekly frequency, which would otherwise be impossible (given that variables like GDP growth are available only quarterly). Weekly frequency data greatly increases the efficiency of our forecasting analysis (though the usual care must be taken of the effects of overlapping forecasts on the calculation of standard errors).

The results in Table 4 clearly show that order flow has considerable forecasting power for all of the six macro variables, and this forecasting power is typically a significant increment over the forecasting power of the other variables considered. Take US output growth as a specific example.¹⁵ At the two-quarter forecasting horizon, order flow produces an R^2 statistic of 24.6 percent, which is significant at the one-percent level. In contrast, forecasting US output growth two months out using both past US output growth

¹⁴Note that the market at large does not have access to these disaggregated end-user data, so the possibility that the exchange rate is a less powerful predictor is not a violation of any simple efficiency criterion.

¹⁵We focus on output growth here, rather than consumption growth as in the model, for two reasons: we want comparability with past work and we have more confidence in the integrity of the available data.

and the spot rate produces an R^2 statistic of only 9.6 percent, a level of forecasting power that is insignificant at conventional levels. In general, the forecasting power of order flow is greater as the forecasting horizon is lengthened.

4.4 Estimating the Speed of Information Aggregation

The results in Table 4 support the idea that order flow contains dispersed information about the future values of fundamental macro variables. While this finding is consistent with theoretical mechanism driving exchange rate dynamics in our model, it does not tell us anything about the pace of information aggregation. Recall that it is the combination of dispersed information on the one hand, and a slow pace of information aggregation on the other, that produces the qualitatively new possibilities for exchange rate dynamics outlined in the theoretical section above.

To quantify the pace at which the market aggregates macro information, we estimate the following regression:

$$\Delta^{13}\tilde{y}_{t+13} = \gamma_0 + \sum_{i=0}^5 \alpha_i \Delta\tilde{y}_{t-i} + \sum_{i=0}^{\#w} \delta_i \Delta s_{t+i} + u_{t+13}$$

where $\Delta^{13}\tilde{y}_{t+13}$ denotes the quarterly change in a given fundamental variable $\tilde{y}$ (i.e., 13 weeks), $\#w$ denotes the number of “learning weeks”, and Δs_{t+i} denotes the change in the spot exchange rate over week $t+i$. The idea here is that exchange rate changes should progressively impound more information about the fundamental $\tilde{y}_t$.

The first column of Table 5 presents the R^2 statistic for this regression over different numbers of learning weeks (denoted $R_{\Delta p}^2$). Given that exchange rate changes should progressively impound more fundamental information, one would expect the $R_{\Delta p}^2$ to increase as the number of learning weeks is increased. This is in fact what we find for each of the six variables (though the increase is often not significant—the column labeled “Sig. I.” reports p-values for the joint significance of the coefficients, corrected for heteroskedasticity and the forecast overlap).

The incremental information in order flow is expressed in columns three and four of each panel. Column three presents the proportional increase in R^2 when order flow is added to the regression. That is, we estimate:

$$\Delta^{13}\tilde{y}_{t+13} = \gamma_0 + \sum_{i=0}^5 \alpha_i \Delta\tilde{y}_{t-i} + \sum_{i=0}^{\#w} \delta_i \Delta s_{t+i} + \sum_{j=1}^6 \omega_j \Delta^{13}x_{j,t} + u_{t+13}$$

where $\Delta^{13}x_{j,t}$ is the quarterly order flow from segment j . Column three then presents the statistic $\nabla R_{\Delta p}^2 \equiv (R_{\Delta p,x}^2/R_{\Delta p}^2)-1$. P-values for the joint significance of the coefficients in this augmented regression are reported in the column labeled “Sig. II.” (corrected for heteroskedasticity and the forecast overlap).

The null hypothesis, then, is that this third column should get smaller as fundamental information conveyed by order flow is impounded in the exchange rate. Moreover, if one expected this information conveyed by flow to be fully reflected in the exchange rate by, say, three weeks, then one would expect the column-three statistic to shrink to zero when the number of learning weeks is set to three. That column three does not shrink to zero, or even close to zero, even after 12 weeks (in every case) is the main result of this table: yes, information in order flow is getting impounded in prices over time, but a lot of that information is still not impounded a quarter later. (For example, a coefficient in the third column, 12-week row, of 1.0 would

imply that a quarter later, the exchange rate alone is impounding only half the fundamental information that order flow and the exchange rate together convey. Again, because these data are not publicly available, there is no simple efficiency criterion that would rule this out.) Given that new innovations in these macro fundamentals are arriving over time, there can be no presumption that this column would in fact shrink to zero if we had enough data to estimate it over longer horizons. Rather, the asymptote for these statistics can remain well above zero, suggesting a balance in the rate at which the underlying macro state is changing and the rate at which markets are learning about the current macro state.

Together, the results in Tables 4 and 5 provide strong evidence in support of the economic mechanisms that drive exchange rates in our model. In fact, our results sharply contradict the traditional assumption that little or no information dispersion exists. Instead, they point to the presence of dispersed, fundamental-related information in order flow, and an information aggregation process that operates on a macroeconomic time scale, not in minutes, hours, or days.

One implication of these surprising findings is the order flow should have forecasting power for *spot rate* changes. If the information aggregation process takes time, then order flows should have forecasting power for future changes in spot rates over the corresponding learning period. To examine this possibility, Table 6 reports the results of estimating forecasting regressions for excess returns. Here we see that the coefficients on US corporate and US investor flows are highly statistically significant, with an R^2 statistic of 19 percent. For an excess return prediction equation to attain this fit, even when based on private information, is rather striking. For comparison, the basis of the forward bias puzzle is the well-documented result that monthly forecasting using the beginning-of-period forward discount produces significant coefficients, but in that case R^2 statistics are generally in the 2-4 percent range. The addition of the time- t forward discount as a regressor does not affect our results; the coefficient on the discount is insignificant while all the statistics remain unchanged. There is no evidence of forward-discount bias (à la Fama) in our data once we include order flows (i.e., we cannot reject the null that the coefficient on the forward discount is one).

5 Conclusion

The results of this paper indicate that information aggregation takes place on a macroeconomic time-scale, rather than on the ultra-high frequency time scale associated with frenzied trading. The picture emphasizes flows of dispersed information, but within a framework for how exchange rates are determined that is explicitly macro. In this sense, the approach does not compete with traditional macro analysis, but is instead complementary. That said, the paper does offer a qualitatively different view of why the empirical performance of macro exchange rate models is so poor. Specifically, transaction flows in the FX market convey information about the present value of future fundamentals that is not captured in macro-econometric measures of fundamentals. If transaction flows reaching the market, i.e., entering the price-setters' information sets, are conveying signals of future macro realizations—signals that are truly incremental to the public macro information—then price will impound this information only after marketmakers observe it. In this setting, if one were to regress exchange rate changes on public macro information, a significant share of those changes would remain "unexplained."

This paper also speaks to the important question of whether transaction flows are relevant to exchange rates over the longer run. There are three existing approaches to this question in the empirical literature.

One approach estimates multi-equation systems that include order flow and exchange rate equations and then measures the effect of order flow shocks on exchange rates over the long run (e.g., Payne 2003, Froot and Ramadorai 2002). A second approach looks instead at the cointegrating relationship between the level of the exchange rate and cumulative order flow (e.g., Bjornes and Rime 2003, Killeen et al. 2001). A third approach uses time-aggregated transaction data in a regression format to test whether daily order flows explain daily exchange rate changes. (At the daily frequency, the exchange rate is very nearly a random walk; since increments to a random walk last forever, explaining those increments establishes relevance over the long run, even if the increments are relatively high frequency.) All three of these previous approaches provide positive evidence that transaction flows are relevant to exchange rates over the longer run. What they do not provide, however, is any sense for why. This paper helps to understand why: if transaction flows are conveying macro information that is otherwise not available, then those information effects on price *should* persist, just as one would expect the effects of public arrival of fundamental information to persist.

The paper’s theoretical contribution is in the simple general equilibrium model of information aggregation that provides—in a setting of incomplete markets—a utility-based present value representation for exchange rates. Key output includes a set of testable relationships between transaction flows, current and future exchange rate returns, and future fundamentals (e.g., money supplies). Noteworthy analytical results include that in equilibrium, order flow is better at forecasting future fundamentals than current spot rates. Related to this result, the equilibrium clarifies why dispersed information about fundamentals becomes impounded in spot rates only slowly. The equilibrium also has a feature that helps resolve the puzzling lack of cointegration found in past work between fundamentals and spot rates. Specifically, when markets are incomplete and information is dispersed, tests of cointegration between standard measures of fundamentals and spot rates are incorrectly sized.

The paper’s main contribution is empirical. We present four main findings, all of which are consistent with our model: (1) transaction flows forecast future macro variables such as output growth, money growth, and inflation, (2) transaction flows generally forecast these macro variables better than spot rates do, (3) transaction flows forecast future spot rates, and (4) though flows convey new information about future fundamentals, much of this information is still not impounded in the spot rate one quarter later. Together, these results have the following broad-level implication. Traditionally, people have viewed past micro-empirical findings linking transaction flows and exchange rates as reflecting a high-frequency, non-fundamental part of exchange rate determination. Our findings suggest that the significance of transaction flows is deeper: Transaction flows appear to be central to the process by which expectations of future macro variables are impounded in price. This conclusion is likely to have relevance for other asset markets as well.

A Appendix A

This appendix presents sections of the model derivation that get too detailed for presentation in the text. In particular, it presents the mechanics of (1) household consumption and portfolio choices and (2) financial intermediary price setting.

A.1 Household Consumption and Portfolio Choices

At the beginning of each period, home households observe the return on their assets, r_t , and the goods market prices, P_t^1 and P_t^2 set by firms. They also see the home and foreign nominal interest rates, i_t and i_t^* , and the spot exchange rate that are quoted by financial intermediaries. With this information, household z makes his consumption and portfolio allocation choices. The budget constraint for household z is:

$$P_t^B B_{z,t} + S_t P_t^{B^*} B_{z,t}^* + A_{z,t} + M_{z,t} = B_{z,t-1} + S_t B_{z,t-1}^* + \exp(r_t) A_{z,t-1} + M_{z,t-1} - P_t C_{z,t},$$

where S_t is the period t price of foreign currency, i.e. the spot exchange rate (home currency/foreign currency). P_t^B and $P_t^{B^*}$ are respectively the period t prices of one-period pure discount bonds that pay one unit of currency in period $t+1$. r_t denotes the nominal return on the portfolio of other assets between periods $t-1$ and t . The budget constraint for foreign households is:

$$\frac{P_t^B B_{z^*,t}}{S_t} + P_t^{B^*} B_{z^*,t}^* + A_{z^*,t} + M_{z^*,t} = \frac{1}{S_t} B_{z^*,t-1} + B_{z^*,t-1}^* + \exp(r_t^*) A_{z^*,t-1}^* + M_{z^*,t-1} - P_t^* C_{z^*,t},$$

where r_t^* is the nominal return (in foreign currency) on the foreign asset portfolio, with nominal value $A_{z^*,t}^*$ in period t .

Households solve the following dynamic programming problem:

$$\begin{aligned} J(W_{z,t}) &= \max_{\alpha_{z,t}^B, \alpha_{z,t}^A, \alpha_{z,t}^M, C_{z,t}} \left\{ \frac{1}{1-\gamma} C_{z,t}^{1-\gamma} + \frac{\chi}{1-\gamma} (\alpha_{z,t}^M W_t)^{1-\gamma} + \delta \mathbb{E}_t^H J(W_{z,t+1}) \right\} \\ &\text{s.t.} \\ W_{z,t+1} &= \exp(i_t - \Delta p_{t+1}) (H_{z,t+1}^M W_{z,t} - C_{z,t}) \\ H_{z,t+1}^M &= 1 + (\exp(\Delta s_{t+1} + i_t^* - i_t) - 1) \alpha_{z,t}^{B^*} + (\exp(r_{t+1} - i_t) - 1) \alpha_{z,t}^A \\ &\quad - \exp(-i_t) (\exp(i_t) - 1) \alpha_{z,t}^M \end{aligned}$$

where $W_{z,t}$ is the value of wealth at the beginning of period t , measured in terms of the consumption index, $C_{z,t}$.¹⁶ $i_t \equiv -\ln P_t^B$ and $i_t^* \equiv -\ln P_t^{B^*}$ are the home and foreign one-period nominal interest rates. $H_{z,t+1}^M$ is the (gross) excess return on wealth between periods t and $t+1$. This depends on the share of wealth held in foreign bonds, $\alpha_{z,t}^{B^*} \equiv S_t P_t^{B^*} B_{z,t}^* / P_t W_{z,t}$, other assets, $\alpha_{z,t}^A \equiv A_{z,t} / P_t W_{z,t}$, and real balances $\alpha_{z,t}^M \equiv$

¹⁶Formally, $W_{z,t} = \exp(i_{t-1}) B_{z,t-1} / P_t + S_t \exp(i_{t-1}^*) B_{z,t-1}^* / P_t + \exp(r_t) A_{z,t-1} / P_t + M_{t-1} / P_t$

$M_{z,t}/P_t W_{z,t}$. The first-order conditions are given by:

$$C_{z,t} : \mathbb{E}_t^H \left[\delta \left(\frac{C_{z,t+1}}{C_{z,t}} \right)^{-\gamma} \exp(i_t - \Delta p_{t+1}) \right] = 1 \quad (\text{A1})$$

$$\alpha_{z,t}^M : \left(\frac{M_t}{P_t C_t} \right)^{-\gamma} = \frac{\exp(i_t) - 1}{\chi \exp(i_t)} \quad (\text{A2})$$

$$\alpha_{z,t}^A : E_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp(r_{t+1} - i_t) \right] = 1 \quad (\text{A3})$$

$$\alpha_{z,t}^{B*} : E_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \exp(\Delta s_{t+1} + i_t^* - \Delta p_{t+1}) \right] = 1 \quad (\text{A4})$$

We now characterize the solution of the household's consumption and portfolio problem with log linearized versions of the first order conditions and budget constraint. First we combine the identity $\alpha_{z,t}^M \equiv M_{z,t}/P_t W_{z,t}$ with the first-order condition for real balances and the definition of $H_{z,t+1}^M$. The budget constraint can then be rewritten as:

$$\frac{W_{z,t+1}}{W_{z,t}} = \exp(i_t - \Delta p_{t+1}) \left(H_{z,t+1} - (1 + \Gamma(i_t)) \frac{C_{z,t}}{W_{z,t}} \right)$$

where

$$\Gamma(i) \equiv \chi^{1/\gamma} \left(\frac{\exp(i) - 1}{\exp(i)} \right)^{1 - \frac{1}{\gamma}}$$

and

$$H_{z,t+1} \equiv 1 + (\exp(\Delta s_{t+1} + i_t^* - i_t) - 1) \alpha_{z,t}^{B*} + (\exp(r_{t+1} - i_t) - 1) \alpha_{z,t}^A.$$

Notice that the coefficient on the consumption-wealth ratio includes the $\Gamma(i_t)$ function because increased consumption raises holdings of real balances. This, in turn, reduces the growth in wealth because the return on nominal balances is zero.

Taking logs on both sides of the budget constraint, and linearizing the right hand side around the point where the consumption-wealth ratio and home nominal interest rate are constant, gives:

$$\Delta w_{t+1} \cong i_t - \Delta p_{t+1} + k + \frac{1}{\rho} (h_{t+1} - \lambda i_t) - \frac{1 - \rho}{\rho} (c_t - w_t) \quad (\text{A5})$$

where $\rho = 1 - \mu(1 + \Gamma(i))$, $\lambda = \frac{\mu(\gamma-1)}{\gamma(\exp(i)-1)\exp(i)} \Gamma(i)$ and $k = \ln \rho + \left(1 - \frac{1}{\rho}\right) \ln \mu + \lambda/\rho$. The sign of the λ coefficient depends on the degree of curvature in the sub-utility function. To understand why, we need to consider the two channels through which nominal interest rates affect the return on wealth via real balances. First, an increase in the interest rate lowers the excess return on wealth when real balances are a constant fraction of wealth. This can be seen from the definition of $H_{z,t+1}^M$ above. When $\gamma < (>) 1$, the former (latter) effect dominates so the excess return on wealth is negatively (positively) related to the nominal interest rate. In the case of log utility ($\gamma = 1$) the effect exactly cancel, and $\lambda = 0$. Hereafter, we focus on the case where $\gamma > 1$, so that $\lambda > 0$ and excess returns are negatively related to the nominal interest rate.

Using the definition of $H_{z,t+1}$ above, we follow Campbell and Viceira (2002) in approximating the log

excess return on wealth by:

$$\begin{aligned} h_{t+1} \cong & \alpha_{z,t}^A (r_{t+1} - i_t) + \alpha_{z,t}^{B*} (\Delta s_{t+1} + i_t^* - i_t) + \frac{1}{2} \alpha_{z,t}^A (1 - \alpha_{z,t}^A) \mathbb{V}_t^H (r_{t+1}) \\ & + \frac{1}{2} \alpha_{z,t}^{B*} (1 - \alpha_{z,t}^{B*}) \mathbb{V}_t^H (\Delta s_{t+1}) - \alpha_{z,t}^{B*} \alpha_{z,t}^A \mathbb{C}\mathbb{V}_t^H (r_{t+1}, \Delta s_{t+1}) \end{aligned} \quad (\text{A6})$$

where $\mathbb{V}_t^H(\cdot)$ and $\mathbb{C}\mathbb{V}_t^H(\cdot, \cdot)$ denote the variance and covariance conditioned on household z 's period t information, $\Omega_{z,t}$. This second order approximation holds exactly in the continuous time limit when the spot exchange rate and the price of other assets A_t follow Wiener processes.

We can now use (A5), (A6) and the linearized first order conditions to characterize the optimal choice of consumption, real balances and the portfolio shares $\alpha_{z,t}^A$ and $\alpha_{z,t}^{B*}$. Combing the log linearized versions of (A3) and (A4) with (A5) and (A6) we obtain:

$$\begin{bmatrix} \alpha_{z,t}^{B*} \\ \alpha_{z,t}^A \end{bmatrix} = \frac{\rho}{\gamma} (\Xi_t^H)^{-1} \begin{bmatrix} \mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t + \frac{1}{2} \mathbb{V}_t^H (\Delta s_{t+1}) - \theta_{z,t}^s \\ \mathbb{E}_t^H r_{t+1} - i_t + \frac{1}{2} \mathbb{V}_t^H (\Delta s_{t+1}) - \theta_{z,t}^r \end{bmatrix}$$

where

$$\theta_{z,t}^v = \gamma \mathbb{C}\mathbb{V}_t^H (c_{z,t+1} - w_{z,t+1}, v_{t+1}) + (1 - \gamma) \mathbb{C}\mathbb{V}_t^H (\Delta p_{t+1}, v_{t+1})$$

for $v = s, r$ and Ξ_t^H is the conditional covariance matrix for the vector $(\Delta s_{t+1}, r_{t+1})'$. $\mathbb{E}_t^H \Delta s_{t+1} + i_t^* - i_t - \theta_{z,t}^s$ and $\mathbb{E}_t^H r_{t+1} - i_t - \theta_{z,t}^r$ are the risk-adjusted expected excess returns on foreign bonds and other assets. The variance terms arise because we are working with log excess returns. $\theta_{z,t}^v$ identifies the consumption hedging factor associated with foreign bonds ($v = s$) and other assets ($v = r$).

All that now remains is to characterize the demand for real balances and the consumption wealth ratio. The former is found by log linearizing (A2):

$$m_{z,t} - p_t = \kappa + c_{z,t} - \eta i_t$$

where $\kappa \equiv \frac{1}{\gamma} \ln \chi + i \exp(i) \eta$ and $\eta = 1/\gamma(\exp(i) - 1) > 0$. An approximation to the log consumption wealth ratio is found by combining (A5) with the linearized version of (A1):

$$c_{z,t} - w_{z,t} = \frac{\rho k}{1 - \rho} + \left(1 - \frac{1}{\gamma}\right) \mathbb{E}_t^H \sum_{i=0}^{\infty} \rho^{i+1} (i_{t+i} - \Delta p_{t+1+i}) + \mathbb{E}_t^H \sum_{i=1}^{\infty} \rho^{i-1} (h_{t+i} - \lambda i_{t+i-1}).$$

We can characterize the behavior of foreign households in a similar way. Specifically, the linearized budget constraint for household z^* is:

$$\Delta w_{z^*,t+1} \cong i_t^* - \Delta p_{t+1}^* + k + \frac{1}{\rho} (h_{z^*,t+1} - \lambda i_t^*) - \frac{1 - \rho}{\rho} (c_{z^*,t} - w_{z^*,t}) \quad (\text{A7})$$

where the log excess return is approximated by:

$$\begin{aligned} h_{z^*,t+1} \cong & \alpha_{z^*,t}^{A*} (r_{t+1}^* - i_t^*) + \alpha_{z^*,t}^{B*} (i_t - \Delta s_{t+1} - i_t^*) + \frac{1}{2} \alpha_{z^*,t}^{A*} (1 - \alpha_{z^*,t}^{A*}) \mathbb{V}_t^{H*} (r_{t+1}^*) \\ & + \frac{1}{2} \alpha_{z^*,t}^{B*} (1 - \alpha_{z^*,t}^{B*}) \mathbb{V}_t^{H*} (\Delta s_{t+1}^*) + \alpha_{z^*,t}^{A*} \alpha_{z^*,t}^{B*} \mathbb{C}\mathbb{V}_t^{H*} (r_{t+1}^*, \Delta s_{t+1}^*). \end{aligned} \quad (\text{A8})$$

The optimal portfolio shares are:

$$\begin{bmatrix} \alpha_{z^*,t}^B \\ \alpha_{z^*,t}^A \end{bmatrix} = \frac{\rho}{\gamma} \left(\Xi_t^{H^*} \right)^{-1} \begin{bmatrix} i_t - \mathbb{E}_t^{H^*} \Delta s_{t+1} - i_t^* + \frac{1}{2} \mathbb{V}_t (\Delta s_{t+1}) - \theta_{z^*,t}^{-s} \\ \mathbb{E}_t^{H^*} r_{t+1}^* - i_t^* + \frac{1}{2} \mathbb{V}_t (r_{t+1}^*) - \theta_{z,t}^{r^*} \end{bmatrix} \quad (\text{A9})$$

where

$$\theta_{z^*,t}^\omega = \gamma \mathbb{C} \mathbb{V}_t^{H^*} (c_{z,t+1} - w_{z,t+1}, \omega_{t+1}) + (1 - \gamma) \mathbb{C} \mathbb{V}_t^H (\Delta p_{t+1}, \omega_{t+1})$$

for $\omega_t = -s_t, r_t^*$, and $\Xi_t^{H^*}$ is the conditional covariance matrix for the vector $(-\Delta s_{t+1}, r_{t+1}^*)'$. The demand for log real balances is given by:

$$m_{z^*,t}^* - p_t^* = \kappa + c_{z^*,t} - \eta_t^* \quad (\text{A10})$$

and the log consumption wealth ratio by:

$$c_{z^*,t} - w_{z^*,t} = \frac{\rho k}{1 - \rho} + \left(1 - \frac{1}{\gamma}\right) \mathbb{E}_t^{H^*} \sum_{i=0}^{\infty} \rho^{i+1} (i_{t+i}^* - \Delta p_{t+1+i}^*) + \mathbb{E}_t^{H^*} \sum_{i=1}^{\infty} \rho^{i-1} (h_{t+i}^* - \lambda i_{t+i-1}^*) \quad (\text{A11})$$

This completes our description of the household behavior in both countries. We now ready to consider the central question of how securities prices are determined. For this we focus on the behavior of financial intermediaries.

A.2 Financial Intermediaries

As above, we characterize the optimal consumption and portfolio decisions using log linear approximations to the dealer's budget constraint and first order conditions. The exact flow constraint for trades initiated by households and other dealers is given by:

$$P_{d,t}^B T_t^B + T_t^M + S_{d,t} \left(P_{d,t}^{B^*} T_t^{B^*} + T_t^{M^*} \right) = 0 \quad (\text{A12})$$

where T_t^v denotes and incoming order to purchase asset v . For trades initiated by dealer d , the constraint is:

$$P_t^B T_{d,t}^B + T_{d,t}^M + S_t \left(P_t^{B^*} T_{d,t}^{B^*} + T_{d,t}^{M^*} \right) = 0. \quad (\text{A13})$$

Notice that the prices for bonds and foreign currency in this equation (i.e., P_t^B , $P_t^{B^*}$ and S_t) are the prices the dealer is quoted by others in the market. Let $M_{d,t}$, $M_{d,t}^*$, $B_{d,t}$ and $A_{d,t}$ respectively denote dealer d 's holding of home and foreign currency bonds, and other assets at the end of period t trading. The dynamic budget constraint of dealer d is then given by:

$$\begin{aligned} M_t + P_t^B B_t + S_t (M_t^* + P_t^{B^*} B_t^*) + A_{d,t} + P_t C_t &= B_{t-1} + M_{t-1} + S_t (B_{t-1}^* + M_{t-1}^*) + \exp(r_t) A_{t-1} \\ &\quad + T_{d,t}^M - T_t^M + S_t (T_{d,t}^{M^*} - T_t^{M^*}) \\ &\quad + P_t^B T_{d,t}^B - P_{d,t}^B T_t^B + S_t (P_t^{B^*} T_{d,t}^{B^*} - P_{d,t}^{B^*} T_t^{B^*}). \end{aligned} \quad (\text{A14})$$

The terms on the left hand side identify the value of consumption and asset holdings at the end of period t trading. The terms in the first row on the right show the value of asset holdings before period t trading,

while the remaining terms in rows two and three identify the profits from trade in currency and bonds.

The problem facing the dealer is to choose prices $S_{d,t}$, $P_{d,t}^B$ and $P_{d,t}^{B*}$, trades $T_{d,t}^B$, $T_{d,t}^{B*}$, $T_{d,t}^M$ and $T_{d,t}^{M*}$, and consumption $C_{d,t}$ to maximize expected utility (9) subject to (A12) - (A14). Although this appears a complex problem, two features of the model make it relatively tractable. First, since dealers have the ability to trade with central banks after their transactions with households and other dealers are complete, they can always achieve their desired holding of both home and foreign currency, M_t and M_t^* (i.e., zero). Second, desired holdings of M_t and M_t^* must both be zero because dealers derive no direct utility from real balances. This means that under any optimal plan $T_{d,t}^M = T_t^M$, $T_{d,t}^{M*} = T_t^{M*}$, so that $M_t = M_t^* = 0$. We combine these restrictions with (A12) - (A14) to write the dealer's problem as (10) and (11) in the text.

The first-order conditions for consumption and the portfolio shares for dealer d are given by:

$$C_{d,t} : \mathbb{E}_t^D \left[\delta V_{t+1} C_{d,t}^\gamma \exp(i_t - \Delta p_{t+1}) \right] = 1 \quad (\text{A15a})$$

$$\alpha_{d,t}^{A*} : \mathbb{E}_t^D \left[\delta V_{t+1} C_{d,t}^\gamma \exp(r_{t+1} - i_t) \right] = 1 \quad (\text{A15b})$$

$$\alpha_{d,t}^{B*} : \mathbb{E}_t^d \left[\delta V_{t+1} C_{d,t}^\gamma \exp(\Delta s_{t+1} + i_t^* - i_t) \right] = 1 \quad (\text{A15c})$$

where $V_t \equiv d\mathcal{J}_t(W_{d,t})/dW_{d,t}$ is the marginal utility of wealth that follows the recursion:

$$V_t = \mathbb{E}_t^D [\delta V_{t+1} \exp(i_t - \Delta p_{t+1}) H_{d,t+1}]. \quad (\text{A16})$$

$\mathbb{E}_t^D$ denotes expectations conditioned on dealer d 's information at the start of period t . Notice that this is the same information set available to dealers before quotes were chosen because quotes are functions of common period information, Ω_t^D ; a subset of dealer d 's information, $\Omega_{d,t}$. In the special case where dealers can perfectly predict the flow of incoming orders for foreign bonds (i.e., $T_{d,t}^{B*} = \mathbb{E}_t^D T_t^{B*}$), (A16) simplifies to $V_t = C_{d,t}^{-\gamma}$ so the consumption and portfolio decisions facing dealers take the familiar form. Under other circumstances, uncertainty about incoming affects these decisions by driving a wedge between the marginal utility of wealth and consumption.

The approximate log budget constraint is:

$$\Delta w_{d,t+1} \cong i_t - \Delta p_{t+1} + k_d + \frac{1}{1-\mu} h_{d,t+1} - \frac{\mu}{1-\mu} (c_{d,t} - w_{d,t}) \quad (\text{A17})$$

where μ is the steady state consumption to wealth ratio, and $k_d = \ln(1-\mu) - \frac{\mu}{1-\mu} \ln \mu$. $h_{d,t+1}$ is the log excess return on wealth, which we approximate by:

$$\begin{aligned} \hat{h}_{t+1} \cong & \alpha_{d,t}^{B*} (\Delta s_{t+1} + i_t^* - r_{t+1}) + \frac{1}{2} \alpha_{d,t}^{B*} (1 - \alpha_{d,t}^{B*}) \mathbb{V}_t^D (\Delta s_{t+1}) + \alpha_{d,t}^A (r_{t+1} - i_t) \\ & + \frac{1}{2} \alpha_{d,t}^A (1 - \alpha_{d,t}^A) \mathbb{V}_t^D (r_{t+1}) - \alpha_{d,t}^A \alpha_{d,t}^{B*} \mathbb{C}\mathbb{V}_t^D (r_{t+1}, \Delta s_{t+1}) - \mathbb{C}\mathbb{V}_t^D (s_{t+1}, \xi_t). \end{aligned} \quad (\text{A18})$$

Combining this equation with log linearized versions of (A15a)- (A16) gives the following approximation for the log marginal utility of wealth:

$$\ln V_t \equiv v_t \cong -\gamma c_t - \varpi \quad (\text{A19})$$

where $\varpi \equiv \mathbb{C}\mathbb{V}_t^D (s_{t+1}, \xi_t)$. Substituting for v_t in the linearized first order conditions for $\alpha_{d,t}^{A*}$ and $\alpha_{d,t}^{B*}$ gives

the expression for the portfolio shares (12). The expression for the log consumption-wealth ratio (13) is derived by combining the linearized budget constraint with the first order condition for consumption.

B Appendix B: Real-Time Inference (From Evans 2004)

The aim is to obtain high frequency real-time estimates on how the macro economy is evolving. For this purpose, it is important to distinguish between information arrival and data collection periods. Information about GDP can arrive via data releases on any day t . GDP data is collected on a quarterly basis. We index quarters by τ and denote the last day of quarter τ by $Q(\tau)$, with the first, second and third months ending on days $M(\tau, 1)$, $M(\tau, 2)$ and $M(\tau, 3)$ respectively. We identify the days on which data is released in two ways. The release day for variable $\varkappa$ collected over quarter τ is $R_\varkappa(\tau)$. Thus, $\varkappa_{R(\tau)}$ denotes the value of variable $\varkappa$, over quarter τ , released on day $R_\varkappa(\tau)$. The release day for monthly variables is identified by $R_\varkappa(\tau, i)$ for $i = 1, 2, 3$. In this case, $\varkappa_{R(\tau, i)}$ is the value of $\varkappa$, for month i in quarter τ , announced on day $R_\varkappa(\tau, i)$. The relation between data release dates and data collection periods is illustrated in Figure 2.

The Bureau of Economic Analysis (BEA) at the U.S. Commerce Department releases data on GDP growth in quarter τ in a sequence of three announcements: The “advanced” growth data are released during the first month of quarter $\tau + 1$; the “preliminary” data are released in the second month; and the “final” data are released at the end of quarter $\tau + 1$. The “final” data release does not represent the last official word on GDP growth in the quarter. Each summer, the BEA conducts an “annual” or comprehensive revision that generally lead to revisions in the “final” data values released over the previous three years. These revisions incorporate more complete and detailed micro data than was available before the “final” data release date.¹⁷

Let $x_{Q(\tau)}$ denote the log of real GDP for quarter τ ending on day $Q(\tau)$, and $y_{R(\tau)}$ be the “final” data released on day $R_y(\tau)$.¹⁸ The relation between the “final” data and actual GDP growth is given by:

$$y_{R(\tau)} = \Delta^Q x_{Q(\tau)} + v_{R(\tau)}, \quad (\text{A20})$$

where $\Delta^Q x_{Q(\tau)} \equiv x_{Q(\tau)} - x_{Q(\tau-1)}$ and $v_{R(\tau)}$ includes future revisions (i.e., the revisions to GDP growth made after $R_y(\tau)$). Notice that equation (A20) distinguishes between the end of the reporting period $Q(\tau)$, and the release date $R_y(\tau)$. We shall refer to the difference $R_y(\tau) - Q(\tau)$ as the reporting lag for quarterly data. (For data series $\varkappa$ collected during month i of quarter τ , the reporting lag is $R_\varkappa(\tau, i) - M(\tau, i)$.) Reporting lags vary from quarter to quarter because data are collected on a calendar basis but announcements are not made on holidays and weekends. For example, “final” GDP data for the quarter ending in March has been released between June 27th and July 3rd.

Real-time estimates of GDP growth are constructed using the information in a specific information set. Let Ω_t denote an information set that only contains data that is publicly known at the end of day t . The real-time *estimate* of GDP growth in quarter τ is defined as $E[\Delta^Q x_{Q(\tau)} | \Omega_{Q(\tau)}]$, the expectation of $\Delta^Q x_{Q(\tau)}$ conditional on public information available at the end of the quarter, $\Omega_{Q(\tau)}$. To see how this estimate relates to the “final” data release, y , we combine the definition with (A20) to obtain:

¹⁷For a complete description of BEA procedures, see Carson (1987), and Seskin and Parker (1998).

¹⁸Because “final” data releases are themselves sometimes subsequently revised, one must choose which final measure to use. In our empirical analysis, we use the first “final” estimate. Of course, the method can accommodate other choices.

$$y_{R(\tau)} = E [\Delta^Q x_{Q(\tau)} | \Omega_{Q(\tau)}] + E [v_{R(\tau)} | \Omega_{Q(\tau)}] + (y_{R(\tau)} - E [y_{R(\tau)} | \Omega_{Q(\tau)}]). \quad (\text{A21})$$

The “final” data released on day $R_y(\tau)$ comprises three components; the real-time GDP growth estimate, an estimate of future data revisions, $E [v_{R(\tau)} | \Omega_{Q(\tau)}]$, and the real-time forecast error for the data release, $y_{R(\tau)} - E [y_{R(\tau)} | \Omega_{Q(\tau)}]$. Under the reasonable assumption that $y_{R(\tau)}$ represents the BEA’s unbiased estimate of GDP growth, and that $\Omega_{Q(\tau)}$ represents a subset of the information available to the BEA before the release day, $E [v_{R(\tau)} | \Omega_{Q(\tau)}]$ should equal zero. In this case, (A21) becomes:

$$y_{R(\tau)} = E [\Delta^Q x_{Q(\tau)} | \Omega_{Q(\tau)}] + (y_{R(\tau)} - E [y_{R(\tau)} | \Omega_{Q(\tau)}]). \quad (\text{A22})$$

Thus, the data release $y_{R(\tau)}$ can be viewed as a noisy signal of the real-time estimate of GDP growth, where the noise arises from the error in forecasting $y_{R(\tau)}$ over the reporting lag. By construction, the noise term is orthogonal to the real-time estimate because both terms are defined relative to the same information set, $\Omega_{Q(\tau)}$. The noise term can be further decomposed as

$$y_{R(\tau)} - E [y_{R(\tau)} | \Omega_{Q(\tau)}] = \left(E [y_{R(\tau)} | \Omega_{Q(\tau)}^{\text{BEA}}] - E [y_{R(\tau)} | \Omega_{Q(\tau)}] \right) + \left(y_{R(\tau)} - E [y_{R(\tau)} | \Omega_{Q(\tau)}^{\text{BEA}}] \right), \quad (\text{A23})$$

where Ω_t^{BEA} denotes the BEA’s information set. Since the BEA has access to both private and public information sources, the first term on the right identifies the informational advantage conferred on the BEA at the end of the quarter $Q(\tau)$. The second term identifies the impact of new information the BEA collects about $x_{Q(\tau)}$ during the reporting lag. Since both of these terms could be sizable, there is no a priori reason to believe that real-time forecast error is always small.

To compute real-time estimates of GDP, we need to characterize the evolution of Ω_t and describe how inferences about $\Delta^Q x_{Q(\tau)}$ can be calculated from $\Omega_{Q(\tau)}$. For this purpose, we incorporate the information contained in the “advanced” and “preliminary” GDP data releases. Let $\hat{y}_{R(\tau)}$ and $\tilde{y}_{R(\tau)}$ respectively denote the values for the “advanced” and “preliminary” data released on days $R_{\hat{y}}(\tau)$ and $R_{\tilde{y}}(\tau)$ where $Q(\tau) < R_{\hat{y}}(\tau) < R_{\tilde{y}}(\tau)$. We assume that $\hat{y}_{R(\tau)}$ and $\tilde{y}_{R(\tau)}$ represent noisy signals of the “final” data, $y_{R(\tau)}$:

$$\hat{y}_{R(\tau)} = y_{R(\tau)} + \tilde{e}_{R(\tau)} + \hat{e}_{R(\tau)}, \quad (\text{A24})$$

$$\tilde{y}_{R(\tau)} = y_{R(\tau)} + \tilde{e}_{R(\tau)}, \quad (\text{A25})$$

where $\tilde{e}_{R(\tau)}$ and $\hat{e}_{R(\tau)}$ are independent mean zero revision shocks. $\tilde{e}_{R(\tau)}$ represents the revision between days $R_{\hat{y}}(\tau)$ and $R_y(\tau)$ and $\hat{e}_{R(\tau)}$ represents the revision between days $R_{\tilde{y}}(\tau)$ and $R_y(\tau)$. The specification of (A24) and (A25) implies that the “advanced” and “preliminary” data releases represent unbiased estimates of actual GDP growth.

The three GDP releases $\{\hat{y}_{R(\tau)}, \tilde{y}_{R(\tau)}, y_{R(\tau)}\}$ represent a sequence of signals on actual GDP growth that augment the public information set on days $R_{\hat{y}}(\tau)$, $R_{\tilde{y}}(\tau)$ and $R_y(\tau)$. In principle we could construct real-time estimates based only on these data releases as $E[\Delta^Q x_{Q(\tau)} | \Omega_t^y]$, where Ω_t^y is the information set comprising data on the three GDP series released on or before day t :

$$\Omega_t^y \equiv \{\hat{y}_{R(\tau)}, \tilde{y}_{R(\tau)}, y_{R(\tau)} : R(\tau) < t\}.$$

Notice that these estimates are only based on data releases relating to GDP growth *before* the current quarter because the presence of the reporting lags exclude the values of $\hat{y}_{R(\tau)}$, $\tilde{y}_{R(\tau)}$, and $y_{R(\tau)}$ from $\Omega_{Q(\tau)}$. As such, these candidate real-time estimates exclude information on $\Delta^Q x_{Q(\tau)}$ that is available at the end of the quarter. Much of this information comes from the data releases on other macroeconomic variables, like employment, retail sales and industrial production. Data for most of these variables are collected on a monthly basis¹⁹, and as such can provide timely information on GDP growth. To see why this is so, consider the data releases on Nonfarm Payroll Employment, z . Data on z for the month ending on day $M_z(\tau, j)$ is released on $R_z(\tau, j)$, a day that falls between the 3rd and the 9th of month $j + 1$ (as illustrated in Figure 1). This reporting lag is much shorter than the lag for GDP releases but it does exclude the use of employment data from the 3rd month in estimating real-time GDP. However, insofar as employment during the first two months is related to GDP growth over the quarter, the values of $z_{R(\tau,1)}$ and $z_{R(\tau,2)}$ will provide information relevant to estimating GDP growth at the end of the quarter.

The real-time estimates we construct below will be based on data from the three GDP releases and the monthly releases of other macroeconomic data. To incorporate the information from these other variables, we decompose quarterly GDP growth into a sequence of daily increments:

$$\Delta^Q x_{Q(\tau)} = \sum_{i=1}^{D(\tau)} \Delta x_{Q(\tau-1)+i}, \quad (\text{A26})$$

where $D(\tau) \equiv Q(\tau) - Q(\tau-1)$ is the duration of quarter τ . The daily increment Δx_t represents the contribution on day t to the growth of GDP in quarter τ . If x_t were a stock variable, like the log price level on day t , Δx_t would identify the daily growth in the stock (e.g. the daily rate of inflation). Here $x_{Q(\tau)}$ denotes the log of the flow of output over quarter τ so it is not appropriate to think of Δx_t as the daily growth in GDP.

To incorporate the information contained in the i 'th. macro variable, z^i , we project $z_{R(\tau,j)}^i$ on a portion of GDP growth:

$$z_{R(\tau,j)}^i = \beta_i \Delta^M x_{M(\tau,j)} + u_{M(\tau,j)}^i, \quad (\text{A27})$$

where $\Delta^M x_{M(\tau,j)}$ is the contribution to GDP growth in quarter τ during month j :

$$\Delta^M x_{M(\tau,j)} \equiv \sum_{i=M(\tau,j-1)+1}^{M(\tau,j)} \Delta x_i.$$

β_i is the projection coefficient and $u_{M(\tau,j)}^i$ is the projection error that is orthogonal to $\Delta^M x_{M(\tau,j)}$. Notice that equation (A27) incorporates the reporting lag $R_z(\tau, j) - M_z(\tau, j)$ for variable z which can vary in length from month to month.

The real-time estimates derived in this paper are based on a information set specification that includes the 3 GDP releases and 18 monthly macro series; $z^i = 1, 2, \dots, 18$. Formally, we compute the end-of-quarter real-time estimates as:

$$E[\Delta^Q x_{Q(\tau)} | \Omega_{Q(\tau)}], \quad (\text{A28})$$

¹⁹Data on initial unemployment claims are collected week by week.

where $\Omega_t = \Omega_t^z \cup \Omega_t^y$ with Ω_t^z denoting the information set comprising of data on the 18 monthly macro variables that has been released on or before day t :

$$\Omega_t^z \equiv \bigcup_{i=1}^{21} \left\{ z_{R(\tau,j)}^i : R(\tau,j) < t \text{ for } j = 1, 2, 3 \right\}.$$

The model presented in Evans (2004) enables us to compute the real-time estimates in (A28) using equations (A20), (A24), (A25), (A26), and (A27) together with a time-series process for the daily increments, Δx_t . That model also enables us to compute *daily* real-time estimates of quarterly GDP, and GDP growth:

$$x_{Q(\tau)|i} \equiv E[x_{Q(\tau)}|\Omega_i] \tag{A29}$$

$$\Delta^Q x_{Q(\tau)|i} \equiv E[\Delta^Q x_{Q(\tau)}|\Omega_i]. \tag{A30}$$

for $Q(\tau-1) < i \leq Q(\tau)$. Equations (A29) and (A30) respectively identify the real-time estimate of log GDP, and GDP growth in quarter τ , based on information available on day i during the quarter. $x_{Q(\tau)|i}$ and $\Delta^Q x_{Q(\tau)|i}$ incorporate real-time forecasts of the daily contribution to GDP in quarter τ between day i and $Q(\tau)$. These high frequency estimates are particularly useful in studying how data releases affect estimates of the current state of the economy, and forecasts of how it will evolve in the future. As such, they are uniquely suited to examining how data releases affect a whole array of asset prices.

References

- Aguiar, M. (2002), Informed speculation and the choice of exchange rate regime, typescript, University of Chicago, March.
- Andersen, T., T. Bollerslev, F. Diebold, and C. Vega (2003), Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review*, 93: 38-62.
- Bacchetta, P., and E. van Wincoop (2003), Can information dispersion explain the exchange rate disconnect puzzle? NBER Working Paper 9498, February.
- Bjornnes, G., and D. Rime (2003), Dealer behavior and trading systems in foreign exchange markets, *Journal of Financial Economics*, forthcoming.
- Campbell, J. Y., and L. Viceira (2002), *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*, Clarendon Lectures in Economics, Oxford University Press.
- Cai, J, Y. Cheung, R. Lee, and M. Melvin (2001), ‘Once in a generation’ yen volatility in 1998: Fundamentals, intervention, and order flow, *Journal of International Money and Finance*, 20: 327-347.
- Cao, H., M. Evans, and R. Lyons (2003), Inventory Information, NBER Working Paper 9893, August, *Journal of Business*, forthcoming.
- Chari, V., P. Kehoe, and E. McGrattan (2002), Can sticky price models generate volatile and persistent real exchange rates? *Review of Economic Studies*, 69: 533-564.
- Cheung, Yin-Wong, M. Chinn, and A. Pascual (2002), Empirical exchange rate models of the nineties: Are any fit to survive? Typescript, U.C. Santa Cruz, Department of Economics.
- Corsetti, G., P. Pesenti, and N. Roubini (2001), Does one Soros make a difference? The role of a large trader in currency crises, NBER Working Paper 8303, *Review of Economic Studies*, forthcoming.
- Covrig, V., and M. Melvin (2002), Asymmetric information and price discovery in the FX market: Does Tokyo know more about the yen? *Journal of Empirical Finance*, 9: 271-285.
- Croushore, D. and T. Stark (2001), A real-time data set for macroeconomists, *Journal of Econometrics* 105, November, pp. 111-130.
- Duarte, M., and A. Stockman (2001), Rational Speculation and Exchange Rates, NBER Working Paper 8362, July.
- Engel, C. (1999), Accounting for US real exchange rate changes, *Journal of Political Economy*, 107: 507-538.
- Engel, C., and K. West (2004a), Exchange rates and fundamentals, NBER Working Paper 10723, August.
- Engel, C., and K. West (2004b), Accounting for exchange rate variability in present value models when the discount factor is near one, NBER Working Paper 10267, January, *American Economic Review*, forthcoming.

- Evans, M. (2002), FX trading and exchange rate dynamics, *Journal of Finance*, 57: 2405-2448.
- Evans, M. (2004), Where are we now? Real-time estimates of the macro economy, typescript, Georgetown University, September (at www.georgetown.edu/faculty/evansm1).
- 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 (2004), A New Micro Model of Exchange Rate Dynamics, NBER Working Paper 10379, March.
- Faust, J., J. Rogers, and J. Wright (2003), Exchange rate forecasting: The errors we've really made, *Journal of International Economics*, 60: 35-59.
- Flood, R., and M. Taylor (1996), Exchange rate economics: What's wrong with the conventional macro approach? In *The Microstructure of Foreign Exchange Markets*, edited by J. Frankel, G. Galli, and A. Giovannini, University of Chicago Press: Chicago, 261-294.
- Frankel, J., G. Galli, and A. Giovannini (1996), Introduction, In *The Microstructure of Foreign Exchange Markets*, University of Chicago Press: Chicago, 1-18.
- Frankel, J., and A. Rose (1995), Empirical research on nominal exchange rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International Economics*, Elsevier Science, Volume III: 1689-1729.
- Froot, K., and T. Ramadorai (2002), Currency returns, institutional investor flows, and exchange rate fundamentals, NBER Working Paper 9101, August.
- Hau, H., and H. Rey (2002), Exchange rates, equity prices, and capital flows, NBER Working Paper 9398, December.
- Hayek, F. (1945), The use of knowledge in society, *American Economic Review*, 35: 519-530.
- Jeanne, O. and A. Rose (2002), Noise trading and exchange rate regimes, *Quarterly Journal of Economics*, 117: 537-569.
- Killeen, W., R. Lyons, and M. Moore (2001), Fixed versus Flexible: Lessons from EMS Order Flow, NBER Working Paper 8491, September 2001, forthcoming in the *Journal of International Money and Finance*.
- Lyons, R. (1995), Tests of microstructural hypotheses in the foreign exchange market, *Journal of Financial Economics*, 39: 321-351.
- Lyons, R. (1997), A Simultaneous Trade Model of the Foreign Exchange Hot Potato, *Journal of International Economics*, 42: 275-298.

- Lyons, R. (2001), *The Microstructure Approach to Exchange Rates*, MIT Press: Cambridge, MA, 2001.
- Meese, R. (1990), Currency fluctuations in the post-Bretton Woods Era, *Journal of Economic Perspectives*, 4: 117-134.
- Meese, R., and K. Rogoff (1983), Empirical exchange rate models of the seventies, *Journal of International Economics*, 14: 3-24.
- Obstfeld, M., and K. Rogoff (1995), Exchange rate dynamics redux, *Journal of Political Economy*, 103: 624-660.
- Obstfeld, M., and K. Rogoff (1998), Risk and exchange rates, NBER Working Paper 6694, August.
- Orphanides, A. (2001), "Monetary Rules Based on Real-Time Data", *American Economic Review*.
- Osler, C. (1998), Short-term speculators and the puzzling behavior of exchange rates, *Journal of International Economics*, 43: 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.
- Vitale, P. (2004), A guided tour of the market microstructure approach to exchange rate determination, typescript, June (at faculty.haas.berkeley.edu/lyons/NewField.html).

Table 1: Contemporaneous Regressions for Excess Returns

Horizon	Corporate		Short Term		Long Term		R^2	χ^2 (p-value)
	US	Non-US	US	Non-US	US	Non-US		
1 day	-0.155 (0.113)	-0.240 (0.067)					0.015	15.133 (0.001)
			0.174 (0.055)	0.204 (0.060)			0.024	21.791 (0.000)
					-0.047 (0.120)	0.369 (0.060)	0.044	38.261 (0.000)
	-0.147 (0.107)	-0.214 (0.064)	0.153 (0.054)	0.194 (0.056)	-0.029 (0.121)	0.353 (0.059)	0.078	75.465 (0.000)
1 week	-0.118 (0.138)	-0.469 (0.083)					0.061	32.070 (0.000)
			0.349 (0.069)	0.114 (0.096)			0.077	27.965 (0.000)
					-0.005 (0.154)	0.523 (0.086)	0.105	37.728 (0.000)
	-0.167 (0.133)	-0.358 (0.077)	0.275 (0.064)	0.069 (0.090)	-0.051 (0.143)	0.447 (0.080)	0.195	111.527 (0.000)
1 month	0.065 (0.266)	-0.594 (0.126)					0.129	22.434 (0.000)
			0.389 (0.135)	0.166 (0.225)			0.103	8.750 (0.013)
					-0.091 (0.215)	0.719 (0.119)	0.205	34.636 (0.000)
	0.120 (0.185)	-0.376 (0.102)	0.214 (0.137)	-0.074 (0.196)	0.000 (0.208)	0.583 (0.130)	0.299	58.424 (0.000)

Notes: The table reports coefficient and standard errors from regressions of excess returns measured over one day, week and month, on order flows cumulated over the same horizon. The left hand column report χ^2 statistics for the null that all the coefficients on order flow are zero. Estimates are calculated at the daily frequency. The standard errors correct for heteroskedastic and the moving average error process induced by overlapping forecasts (1 week and 1 month results).

Table 2: Real-Time Estimates and Errors						
A: Cointegration Results	US			German		
	Output	Prices	Money	Output	Prices	Money
Coefficient.	0.880	1.056	1.020	0.873	0.667	0.980
Standard Error	0.025	0.008	0.002	0.026	0.024	0.006
p-value	(0.001)	(0.001)	(0.068)	(0.000)	(0.001)	(0.002)
B: Error Autocorrelations						
Lag = 1 Day	0.980	0.950	0.896	0.984	0.987	0.947
1 Week	0.904	0.749	0.483	0.920	0.934	0.765
1 Month	0.620	0.495	0.192	0.693	0.750	0.343
2 Months	0.369	0.493	0.109	0.386	0.662	0.133
1 Quarter	0.209	0.511	0.118	0.212	0.573	0.066
Notes: The upper panel reports the results from the cointegrating regression of the real time estimate of the fundamental variable on its ex post value. The reported standard errors are computed by Dynamic OLS in daily data (1682 observations) with 10 leads and lags to correct for finite sample bias. Standard errors contain an MA(10) correction for residual serial correlation. The p-values are for the hypothesis that the cointegration coefficient equals unity. The lower panel reports daily autocorrelations for the real-time errors, defined as the difference between the ex post and real-time estimate of the fundamental variables.						

Table 3: Granger Causality Significance Levels		
Variable to be Forecast	Forecasting Variable	
	Order Flows	Exchange Rate
Money Growth—US	0.00	0.72
Output Growth—US	0.00	0.01
Inflation—US	0.47	0.09
Money Growth—Germany	0.79	0.72
Output Growth—Germany	0.44	0.96
Inflation—Germany	0.00	0.71
Notes: Table presents marginal significance levels of tests whether end-user flows Granger cause three macro variables: output growth, money growth, and inflation. The tests are based on a monthly-frequency VAR for money and inflation, and a quarterly-frequency VAR for output growth. All the VARs include one lag of each of the following: the rate of exchange-rate depreciation, the macro variable, and the 6 end-user flow segments.		

Table 4: Forecasting Fundamentals

Forecasting Variables	US Output Growth				German Output Growth			
	1 month	2 months	1 quarter	2 quarters	1 month	2 months	1 quarter	2 quarters
Output	0.002 (0.607)	0.003 (0.555)	0.022 (0.130)	0.092 (0.087)	0.004 (0.295)	0.063 (0.006)	0.089 (0.009)	0.006 (0.614)
Spot Rate	0.001 (0.730)	0.005 (0.508)	0.005 (0.644)	0.007 (0.650)	0.058 (0.002)	0.029 (0.081)	0.003 (0.625)	0.024 (0.536)
Output and Spot Rates	0.003 (0.802)	0.007 (0.710)	0.031 (0.287)	0.096 (0.224)	0.059 (0.007)	0.083 (0.021)	0.099 (0.024)	0.033 (0.709)
Order Flows	0.032 (0.357)	0.080 (0.145)	0.189 (0.002)	0.246 (0.000)	0.012 (0.806)	0.085 (0.227)	0.075 (0.299)	0.306 (0.000)
All	0.052 (0.383)	0.086 (0.195)	0.199 (0.011)	0.420 (0.000)	0.087 (0.021)	0.165 (0.037)	0.156 (0.130)	0.324 (0.000)
Forecasting Variables	US Inflation				German Inflation			
	1 month	2 months	1 quarter	2 quarters	1 month	2 months	1 quarter	2 quarters
Inflation	0.003 (0.461)	0.024 (0.146)	0.005 (0.487)	0.053 (0.213)	0.007 (0.402)	0.037 (0.067)	0.053 (0.040)	0.024 (0.232)
Spot Rate	0.005 (0.351)	0.007 (0.419)	0.013 (0.391)	0.016 (0.457)	0.081 (0.000)	0.000 (0.962)	0.000 (0.858)	0.033 (0.305)
Inflation and Spot Rates	0.007 (0.505)	0.028 (0.352)	0.015 (0.636)	0.060 (0.441)	0.088 (0.002)	0.038 (0.214)	0.053 (0.112)	0.051 (0.364)
Order Flows	0.025 (0.773)	0.050 (0.629)	0.116 (0.052)	0.212 (0.000)	0.050 (0.429)	0.116 (0.010)	0.178 (0.025)	0.271 (0.000)
All	0.031 (0.788)	0.082 (0.151)	0.124 (0.010)	0.240 (0.000)	0.127 (0.005)	0.158 (0.021)	0.258 (0.005)	0.511 (0.000)

Table 4: Forecasting Fundamentals (cont.)

Forecasting Variables	US Money Growth				German Money Growth			
	1 month	2 months	1 quarter	2 quarters	1 month	2 months	1 quarter	2 quarters
Money Growth	0.071 (0.009)	0.219 (0.000)	0.253 (0.000)	0.329 (0.000)	0.050 (0.023)	0.111 (0.005)	0.122 (0.017)	0.041 (0.252)
Spot Rate	0.021 (0.054)	0.001 (0.778)	0.003 (0.732)	0.005 (0.619)	0.002 (0.558)	0.044 (0.031)	0.036 (0.123)	0.065 (0.343)
Money Growth and Spot Rates	0.086 (0.002)	0.220 (0.000)	0.267 (0.000)	0.333 (0.000)	0.050 (0.075)	0.130 (0.004)	0.129 (0.040)	0.080 (0.403)
Order Flows	0.034 (0.466)	0.119 (0.239)	0.280 (0.026)	0.424 (0.000)	0.026 (0.491)	0.082 (0.147)	0.152 (0.037)	0.578 (0.000)
All	0.096 (0.056)	0.282 (0.000)	0.417 (0.000)	0.540 (0.000)	0.074 (0.244)	0.175 (0.020)	0.284 (0.001)	0.624 (0.000)

Notes: The table reports the R^2 from the forecasting regression for the fundamental listed in the header of each panel using the forecasting variables reported on the left. The regressions are estimated in weekly data (284 observations). Significance levels for χ^2 statistics testing the null hypothesis of no predictability (corrected for heteroskedasticity and the forecast horizon overlap) are reported in parentheses. The weekly estimates of fundamentals are real time estimates based on the history of macro announcements.

Table 5: Tests for the Speed of Information Aggregation

	US Output Growth				German Output Growth			
	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II
Learning Weeks								
0	0.157	(0.119)	0.981	(0.020)	0.077	(0.003)	1.745	(0.016)
3	0.187	(0.009)	0.860	(0.016)	0.078	(0.917)	1.697	(0.023)
6	0.201	(0.003)	0.805	(0.008)	0.080	(0.986)	1.683	(0.018)
9	0.203	(0.010)	0.794	(0.004)	0.080	(0.999)	1.672	(0.019)
12	0.219	(0.001)	0.743	(0.000)	0.084	(0.998)	1.556	(0.018)
	US Inflation				German Inflation			
	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II
Learning Weeks								
0	0.021	(0.303)	5.728	(0.014)	0.012	(0.144)	23.035	(0.001)
3	0.022	(0.978)	5.648	(0.012)	0.012	(0.997)	22.747	(0.001)
6	0.032	(0.760)	3.751	(0.014)	0.031	(0.356)	8.743	(0.001)
9	0.066	(0.077)	1.580	(0.030)	0.034	(0.514)	8.451	(0.001)
12	0.080	(0.026)	1.155	(0.073)	0.051	(0.353)	5.381	(0.000)
	US Money Growth				German Money Growth			
	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II	$R^2_{\Delta p}$	Sig. I	$\nabla R^2_{\Delta p}$	Sig. II
Learning Weeks								
0	0.256	(0.086)	0.540	(0.030)	0.147	(0.011)	1.505	(0.002)
3	0.263	(0.513)	0.577	(0.024)	0.155	(0.343)	1.370	(0.005)
6	0.275	(0.244)	0.563	(0.017)	0.168	(0.118)	1.195	(0.009)
9	0.289	(0.069)	0.528	(0.018)	0.203	(0.000)	0.840	(0.016)
12	0.290	(0.152)	0.539	(0.013)	0.247	(0.000)	0.527	(0.048)

Notes: $R^2_{\Delta p}$ denotes the R^2 statistic from the regression

$$\Delta^{13} y_{t+13} = \gamma_0 + \sum_{i=0}^5 \alpha_i \Delta y_{t-i} + \sum_{i=0}^{\#w} \delta_i \Delta s_{t+i} + u_{t+13}$$

where $\Delta^{13} y_{t+13}$ denotes the quarterly change in the fundamental (listed in the header of each sub-panel) and $\#w$ denotes the number of “learning weeks”. P-values for the joint significance of the δ_i coefficients (corrected for heteroskedasticity and the forecast overlap) are reported in the column headed Sig. I. $\nabla R^2_{\Delta p}$ shows the proportional increase in R^2 when order flow is added to the regression. Specifically, let $R^2_{\nabla p,x}$ denotes the R^2 statistic from the regression

$$\Delta^{13} y_{t+13} = \gamma_0 + \sum_{i=0}^5 \alpha_i \Delta y_{t-i} + \sum_{i=0}^{\#w} \delta_i \Delta s_{t+i} + \sum_{j=1}^6 \omega_j \Delta^{13} x_{j,t} + u_{t+13}$$

where $\Delta^{13} x_{j,t}$ is the quarterly order flow from segment j. $\nabla R^2_{\Delta p} = (R^2_{\nabla p,x} / R^2_{\Delta p}) - 1$. P-values for the joint significance of the ω_j coefficients (corrected for heteroskedasticity and the forecast overlap) are reported in the column headed Sig. II.

Table 6: Forecasting Regressions for Excess Returns

Horizon	Corporate		Short Term		Long Term		R^2	χ^2 (p-value)
	US	Non-US	US	Non-US	US	Non-US		
1 week	1.119 (0.365)	-0.061 (0.170)					0.027	10.243 (0.006)
			0.045 (0.162)	0.205 (0.225)			0.003	0.983 (0.612)
					-0.652 (0.304)	0.222 (0.183)	0.015	6.003 (0.050)
	1.074 (0.363)	-0.008 (0.189)	-0.071 (0.161)	0.039 (0.228)	-0.421 (0.309)	0.247 (0.196)	0.037	16.207 (0.013)
2 weeks	1.243 (0.363)	-0.067 (0.155)					0.069	13.403 (0.001)
			0.098 (0.155)	0.230 (0.209)			0.009	1.787 (0.409)
					-0.785 (0.276)	0.203 (0.145)	0.042	9.791 (0.007)
	1.124 (0.356)	-0.004 (0.172)	-0.013 (0.143)	0.063 (0.209)	-0.536 (0.273)	0.207 (0.161)	0.092	24.352 (0.000)
3 weeks	1.262 (0.341)	-0.041 (0.142)					0.104	16.261 (0.000)
			0.097 (0.155)	0.190 (0.190)			0.011	1.532 (0.465)
					-0.864 (0.271)	0.170 (0.120)	0.071	12.247 (0.002)
	1.111 (0.313)	0.014 (0.150)	-0.005 (0.138)	0.024 (0.196)	-0.626 (0.258)	0.184 (0.143)	0.143	30.195 (0.000)
1 month	1.179 (0.306)	-0.051 (0.133)					0.119	18.041 (0.000)
			0.090 (0.160)	0.135 (0.173)			0.010	1.116 (0.572)
					-0.965 (0.264)	0.131 (0.109)	0.110	15.434 (0.000)
	0.985 (0.259)	-0.008 (0.137)	0.001 (0.136)	-0.038 (0.182)	-0.762 (0.242)	0.146 (0.128)	0.185	33.629 (0.000)

Notes: The table reports coefficient and standard errors from regressions of excess returns measured over 1, - 3 weeks and 1 month on lagged order flows cumulated over one month. The left hand column report χ^2 statistics for the null that all the coefficients on order flow are zero. Estimates are calculated at the daily frequency using 1141 trading days in the sample. The standard errors correct for heteroskedastic and the moving average error process induced by overlapping forecasts.




Figure 1: Top plot shows real-time estimates of GDP growth last quarter (solid plot), and the “final” data releases of GDP growth (dashed plot). Lower plot shows cumulative real-time log GDP (solid plot) and cumulative log GDP based on “final” data releases (dashed plot).




Figure 2: The figure illustrates the relation between data collection periods and release times for quarterly and monthly variables. The reporting lag for “final” GDP growth in quarter τ , $y_{Q(\tau)}$, is $R_y(\tau) - Q(\tau)$. The reporting lag for the monthly series $z_{M(\tau,3)}$ is $R_z(\tau,3) - M(\tau,3)$.