

NBER WORKING PAPER SERIES

EXCHANGE RATE FUNDAMENTALS AND ORDER FLOW

Martin D. D. Evans
Richard K. Lyons

Working Paper 13151
<http://www.nber.org/papers/w13151>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2007

We thank the following for valuable comments: Anna Pavlova, Andrew Rose and seminar participants at the NBER (October 2004 meeting of IFM), the Board of Governors at the Federal Reserve, the European Central Bank, the London Business School, the University of Warwick, the Graduate School of Business at the University of Chicago, UC Berkeley, the Bank of Canada, the International Monetary Fund, and the Federal Reserve Bank of New York. 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/evansml and at faculty.haas.berkeley.edu/lyons). The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Martin D. D. Evans and Richard K. Lyons. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Exchange Rate Fundamentals and Order Flow
Martin D. D. Evans and Richard K. Lyons
NBER Working Paper No. 13151
June 2007
JEL No. F31,G12,G14

ABSTRACT

We address whether transaction flows in foreign exchange markets convey fundamental information. Our GE model includes fundamental information that first manifests at the micro level and is not symmetrically observed by all agents. This produces foreign exchange transactions that play a central role in information aggregation, providing testable links between transaction flows, exchange rates, and future fundamentals. We test these links using data on all end-user currency trades received at 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 the exchange rate does, (3) transaction flows (proprietary) forecast future exchange rates, and (4) the forecasted part of fundamentals is better at explaining exchange rates than standard measured fundamentals.

Martin D. D. Evans
Department of Economics
Georgetown University
Washington, DC 20057
evansm1@georgetown.edu

Richard K. Lyons
Goldman Sachs
85 Broad Street, Floor 20
New York, NY 10004
and NBER
Richard.Lyons@gs.com

Introduction

Exchange rate movements at frequencies of one year or less remain unexplained by observable macroeconomic variables (Meese and Rogoff 1983, Frankel and Rose 1995, Cheung et al. 2005). 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 (2005) comprehensive study concludes that "no model consistently outperforms a random walk."

This paper addresses this long-standing puzzle from a new direction. Rather than attempting to empirically link macro variables to exchange rates directly, we address instead the intermediate market-based process that impounds macro information into exchange rates. Our approach is based two central ideas: First, only some of the macro information relevant for the current spot exchange rate is publicly known at any point in time. Other information is present in the economy, but it exists in a dispersed microeconomic form in the sense of Hayek (1945). The second idea relates to determination of the spot rate through the operation of the foreign exchange market. Specifically, since the spot rate literally is the price of foreign currency quoted by foreign exchange dealers, it can only reflect information that is known to dealers. Consequently, the spot rate will only reflect dispersed information once it has been assimilated by dealers, (collectively called "the market") – a process that takes place via trading. We shall argue that this trade-based mechanism is economically important because much information about the current state is dispersed, and because it takes a considerable time for dispersed information to be completely assimilated by "the market".

To make these ideas concrete, we present a two-country general equilibrium model in which the spot rate is determined via the optimal trading activities of dealers in the foreign exchange market. Our model contains three essential ingredients. First, it includes information that is not publicly observed, at least initially. Second, transaction flows are correlated with this information. Third, the equilibrium spot rate is not fully revealing. The model not only provides a theoretical rationale for the strong empirical link between spot rate changes and transaction flows (see, for example, Evans and Lyons 2002a,b), but it also delivers two new testable implications: First, transaction flows should have more power to forecast *future* fundamentals than current spot rates. Second,

insofar as the transaction flows received by individual dealers predict what the rest of “the market” will learn about fundamentals in the future, those flows should have forecasting power for *future* exchange rate returns.

We investigate these empirical predictions using a new data set that comprises USD/EUR spot rates, transaction flows and macro fundamentals over six and a half years. The transaction flows come from Citibank and represent propriety information of an important Bank in the USD/EUR market. A novel and important feature of our empirical analysis is that it utilizes high-frequency real-time estimates of macro variables. These data are estimates of the underlying macro variables based on contemporaneously available public information. As such, they provide a more precise measure of public expectations regarding fundamentals than realizations of the variables themselves. This greater precision is reflected in the strong statistical significance of our findings.

The implications of our model are strongly supported by our data. In particular we find that:

1. Transaction flows in the USD/EUR market have significant forecasting power for future output growth, money growth, and inflation in both the US and Germany.
2. Transaction flows have incremental forecasting power for macro variables beyond that contained in the history of exchange rates and the variable itself.
3. Propriety transaction flows forecast future exchange rate returns, and do so much more effectively than forward discounts.
4. The forecasting power of propriety transaction flows reflects their ability to predict how “the market” will react to the flow of subsequent information concerning macro fundamentals.

To the best of our knowledge, these are the first findings to link macro fundamentals, transaction flows and exchange rate dynamics. Taken together, they provide strong support for the idea that exchange rates vary as “the market” assimilates dispersed information regarding macro fundamentals from transaction flows.

Our analysis is related to several strands of the international finance literature. From a theoretical perspective, our general equilibrium model includes two novel ingredients: dispersed information and a micro-based rationale for trade in the foreign exchange market. Dispersed information does

not exist in textbook models: relevant information is either symmetric economy-wide, or, sometimes, asymmetrically assigned to a single agent – the central bank. As a result, no textbook model predicts that market-wide transaction flows should drive exchange rates. In recent research, Bacchetta and van Wincoop (2006) examine the dynamics of the exchange rate in a rational expectations model with dispersed information. Our model shares some of the same informational features, but derives the equilibrium dynamics from the equilibrium trading strategies of foreign exchange dealers. Our focus on the role of transaction flows as conveyors of information concerning macro fundamentals also differs from Bacchetta and van Wincoop (2006).

From an empirical perspective, our analysis is closely related to the work of Engel and West (2005). They find that spot rates have forecasting power for future macro fundamentals as textbook models predict. Indeed, our model makes the same empirical prediction. The novel aspect of our analysis, relative to Engel and West (2005), is that we investigate whether the exchange rate responds to transaction flows because they induce a change in “the market’s” expectations about future fundamentals. From this perspective, our findings should be viewed as complementing theirs. Our analysis is also related to earlier research by Froot and Ramadorai (2005), hereafter F&R. These authors examine VAR relationships between real exchange rates, excess currency returns, real interest differentials, and the transaction flows of institutional investors. In contrast to our results, they find little evidence that these flows can forecast fundamentals. Our analysis differs from F&R in three respects. First, and most substantively, transaction flows should be driven not by changes in fundamentals, but by changes in fundamentals expectations. The F&R analysis focuses on the former, whereas ours focuses on the latter. Second, we analyze transaction flows which fully span the demand for foreign currency, not just institutional investors. This facet of our flow data proves to be empirically important. Third, we require no assumption about exchange rate behavior in the long run, whereas the variance decompositions F&R use are based on long run purchasing power parity.

The rest of the paper is organized as follows. Section 1 provides an overview of our model and presents the key equations determining the spot exchange rate. Section 2 derives the theoretical link between transaction flows and exchange rate fundamentals. Section 3 describes the data. Section 4 presents our empirical analysis. Section 5 concludes.

1 The Model

Our model is a two-country, two-good dynamic general equilibrium model that incorporates explicit microfoundations of how trading takes place in the foreign exchange market. For this purpose, we need to model the behavior of households, firms, central banks and foreign exchange dealers who act as market-makers. In this section, we first present the preferences and constraints facing households and firms and describe the role of central banks. We then lay out the problem facing foreign exchange dealers and provide intuition for their equilibrium behavior. Finally we present the equilibrium equation for the spot exchange rate that plays a central role in our analysis. The Appendix describes the complete structure of the model and provides detailed mathematical derivations of our key results.

1.1 Households, Firms and Central Banks

There are two countries, each populated by a continuum of households arranged on the unit interval $[0,1]$. For concreteness, we shall refer to home and foreign countries as the US and Europe and use the index $h \in [0, 1/2)$ to denote US households and $\hat{h} \in [1/2, 1]$ to denote European households. All households derive utility from consumption and real balances. The preferences of US household h are given by:

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

where $0 < \delta < 1$ is the discount factor, $\chi > 0$ and $\gamma \geq 1$. \mathbb{E}_t^h denotes expectations conditioned on US household information, $\Omega_{h,t}$. $M_{h,t}$ is the stock of dollars held by household h , and $C_{h,t}$ is a CES consumption index defined over the two consumption goods:

$$C_{h,t} \equiv (C_{h,t}(\text{US})^{(\theta-1)/\theta} + C_{h,t}(\text{EU})^{(\theta-1)/\theta})^{\theta/(\theta-1)}, \quad (2)$$

where $C_{h,t}(i)$ is the consumption of the i -country good by household h . θ is the elasticity of substitution between the two goods, which we assume to be greater than one (see below). The price index corresponding to (2) is $P_t \equiv (P_t^{\text{US}(1-\theta)} + P_t^{\text{EU}(1-\theta)})^{1/(1-\theta)}$, where P_t^i are the prices of good i . The preferences of European households are defined in an analogous manner with respect to the foreign consumption index, $\hat{C}_{h,t}$, and real balances, $\hat{M}_{h,t}/\hat{P}_t$, where \hat{P}_t is the European price level.

Hereafter, we use “hats” to indicate European variables.

In addition to domestic currency, households can hold one-period nominal dollar bonds, B , nominal euro bonds \hat{B} , and the equities issued by US and European firms, A and \hat{A} . Let R_t and \hat{R}_t be the US and European one period gross nominal interest rates and let S_t denote the spot exchange rate, specifically, the dollar price of euros (\$/€). The budget constraint facing US household h is

$$B_{h,t} + Q_t A_{h,t} + S_t \hat{B}_{h,t} + S_t \hat{Q}_t \hat{A}_{h,t} + M_{h,t} + P_t C_{h,t} = (Q_t + D_t) A_{h,t-1} + S_t (\hat{Q}_t + \hat{D}_t) \hat{A}_{h,t-1} + R_{t-1} B_{h,t-1} + S_t \hat{R}_{t-1} \hat{B}_{h,t-1} + M_{h,t-1} \quad (3)$$

where Q_t and \hat{Q}_t are the local currency prices of US and European equities with dividends per share of D_t and \hat{D}_t respectively. The problem facing US household h in period t is to choose $B_{h,t}$, $\hat{B}_{h,t}$, $\hat{A}_{h,t}$, $A_{h,t}$, $M_{h,t}$, and $C_{h,t}(i)$ for $i = \{\text{US}, \text{EU}\}$ given prices $\{Q_t, \hat{Q}_t, P_t^{\text{US}}, P_t^{\text{EU}}\}$, dividends $\{D_t, \hat{D}_t\}$, interest rates $\{R_t, \hat{R}_t\}$, and the spot exchange rate S_t , that maximize (1) subject to (3).

There are two representative firms; a US firm producing good Y , and a European firm producing good \hat{Y} . Each firm has monopoly power in the US and European market for its good and issues equity claims to its dividend stream. To introduce consumer price-stickiness, we assume that firms set prices in local currencies before they have complete information about the state of demand in each national market.

Consider the pricing problem facing the US firm. The period- t output of the US good is $Y_t = \xi_t K_t^\nu$ with $\nu > 0$, where K_t and ξ_t denote the current stock of firm-specific capital and the state of productivity. This output can be costlessly transported to meet demand in the US and European market or used to augment the existing capital stock. Let P_t^{US} and \hat{P}_t^{US} denote the period- t dollar and euro retail prices for the US good. Given the form of household preferences, the US and European demands for the US good are given by $(P_t^{\text{US}}/P_t)^{-\theta} C_t$ and $(\hat{P}_t^{\text{US}}/\hat{P}_t)^{-\theta} \hat{C}_t$ where C_t and \hat{C}_t denote aggregate US and European consumption. We assume that prices are chosen to maximize the real value of the firm's dividend stream. If the total number of outstanding shares is normalized to unity, the pricing problem facing the US firm is

$$\mathbb{Q}_t^{\text{US}} = \max_{P_t^{\text{US}}, \hat{P}_t^{\text{US}}} \mathbb{E}_t^{\text{US}} \sum_{i=0}^{\infty} \Lambda_{t+i,t} (D_{t+i}/P_{t+i}) \quad (4)$$

subject to

$$D_t/P_t = (P_t^{\text{US}}/P_t)^{1-\theta}C_t + (S_t\hat{P}_t/P_t)(\hat{P}_t^{\text{US}}/\hat{P}_t)^{1-\theta}\hat{C}_t, \quad \text{and} \quad (5)$$

$$K_{t+1} = (1-\varrho)K_t + \xi_t K_t^\nu - (P_t^{\text{US}}/P_t)^{-\theta}C_t - (\hat{P}_t^{\text{US}}/\hat{P}_t)^{-\theta}\hat{C}_t. \quad (6)$$

where \mathbb{E}_t^{US} denotes the firm's expectations conditioned period- t information. $\Lambda_{t+i,t}$ is the stochastic discount factor between t and $t+i$ that the firm uses to value the stream of real dividends. Firms cannot hold financial assets or claims, so real dividends, D_t/P_t , must equal the the sum of US and European sales measured in terms of US aggregate consumption as shown in (5). Equation (6) describes capital accumulation with depreciation rate $\varrho > 0$.² Notice that the firm faces three (potential) sources of uncertainty when choosing period- t prices: uncertainty about aggregate consumption, C_t and \hat{C}_t ; the aggregate price levels, P_t and \hat{P}_t ; and the spot exchange rate, S_t . The European firm producing the EU good faces an analogous problem in choosing prices, P_t^{EU} and \hat{P}_t^{EU} .

The Federal Reserve (FED) and European Central Bank (ECB) play a simple role in our model. Both central banks set one period nominal interest rates so as to achieve a target level for their national money supplies. Specifically, we assume that R_t and \hat{R}_t are set at the beginning of period t such that

$$m_t^* = \mathbb{E}_t^{\text{FED}} m_t, \quad \text{and} \quad \hat{m}_t^* = \mathbb{E}_t^{\text{ECB}} \hat{m}_t$$

where $m_t \equiv \int_0^{1/2} \ln M_{h,t} dh$ and $\hat{m}_t \equiv \int_{1/2}^1 \ln \hat{M}_{h,t} dh$ are the aggregate log demands for dollars and euros and m_t^* and \hat{m}_t^* denote the targets for the US and European log money supplies. (Hereafter we denote aggregates by dropping the h subscript and use lowercase variables to denote natural logs, e.g. $s_t = \ln S_t$, $c_{h,t} \equiv \ln C_{h,t}$, etc.). Notice that interest rates are set on the basis of the FED's and ECB's expectations concerning the demand for currency, $\mathbb{E}_t^{\text{FED}} m_t$ and $\mathbb{E}_t^{\text{ECB}} \hat{m}_t$, rather than the actual demand. Insofar as central banks are unable to exactly predict the aggregate demand for currency, because individual household demands are a function of private information, excess demand is accommodated at the chosen interest rates.

1.2 Foreign Exchange Dealers

A key distinction between our model and traditional international finance models is that the spot exchange rate is determined as the foreign currency price quoted by dealers in the foreign exchange

²The firm's problem is not well-posed if the elasticity parameter θ is less than one because real dividends and future capital would be increasing functions of current relative prices.

market. We assume that there are D dealers (indexed by d) who act as market-makers in the spot market for foreign currency. As such, each dealer quotes prices at which they stand ready to buy or sell foreign currency to households and other dealers.³ Each dealer also has the opportunity to initiate transactions with other dealers at the prices they quote. We now described the decision problem facing a typical dealer in detail.

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

$$\mathbb{U}_t^d \equiv \mathbb{E}_t^d \sum_{i=0}^{\infty} \delta^i \frac{1}{1-\gamma} C_{d,t+i}^{1-\gamma}, \quad (7)$$

where \mathbb{E}_t^d denotes expectations conditioned on the dealer's period- t information, $\Omega_{d,t}$, and $C_{d,t}$ represents the dealers consumption of the 2 goods aggregated via the CES function shown in (2). Dealers have the same preferences as US households except that real balances have no utility value. As a consequence, they will not hold currency in equilibrium – a feature that proves useful in the deriving equations for the equilibrium exchange rate below. We assume that dealers are prohibited from holding equities for the same reason.

Trading in period t is split into two rounds. In round I, dealers quote prices at which they are willing to trade with households. In round II, dealers quote prices at which they will trade with other dealers and they initiate trades against other dealer's quotes. More specifically, at the start of round I, each dealer d quotes a dollar price for euros, $S_{d,t}^I$, at which he is willing to buy or sell euros. These price quotes are publicly observed and good for any quantity of euro (i.e. there is no bid-ask spread). Each dealer then receives orders for euros from a subset of households. We denote the net household order to purchase euros received by dealer d as $\mathcal{T}_{d,t}^I$. Household orders are only observed by the recipient dealer and so represent a source of private information. At the start of round II, each dealer quotes a price for euros of $S_{d,t}^{II}$. These prices, too, are good for any quantity and publicly observed, so that trading with multiple partners (e.g., arbitrage trades) is feasible. Each dealer d then chooses the quantity of euros he wishes to purchase, $T_{d,t}$, (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. We denote the net quantity of euros purchased

³More precisely, the price dealers quote is for the euro bond, which can be thought of as an interest-baring euro deposit account.

from dealer d as a result of the trades initiated by other dealers by $\mathcal{T}_{d,t}^{\text{II}}$. After round II trading is complete, dealers make their period- t consumption decisions.

Let $B_{d,t}^i$ and $\hat{B}_{d,t}^i$ denote dealer d 's holdings of dollar and euro bonds at the start of round i trading in period t . At the end of round I trading, the dealer's bond holdings are

$$B_{d,t}^{\text{II}} = B_{d,t}^{\text{I}} + S_{d,t}^{\text{I}} \mathcal{T}_{d,t}^{\text{I}}, \quad \text{and} \quad \hat{B}_{d,t}^{\text{II}} = B_{d,t}^{\text{I}} - \mathcal{T}_{d,t}^{\text{I}}, \quad (8)$$

where $S_{d,t}^{\text{I}}$ is the price quoted by dealer d , and $\mathcal{T}_{d,t}^{\text{I}}$ are the incoming household orders to purchase euros. In round II, dealer d quotes $S_{d,t}^{\text{II}}$, receives incoming order for euros of $\mathcal{T}_{d,t}^{\text{II}}$ and initiates euro purchases of $T_{d,t}$ at the price of S_t^{II} , the price quoted by other dealers. (In equilibrium all dealers quote the same price so we need not worry about the identity of the other dealers.) To finance his desired basket of consumption goods, dealer d then exchanges US bonds worth $P_t C_{d,t}$ for dollars at the US central bank, and makes his consumption purchases in the US markets for the two goods. The dealer's bond holdings at the start of period $t+1$ are therefore given by

$$\begin{aligned} \hat{B}_{d,t+1}^{\text{I}} &= \hat{R}_t(\hat{B}_{d,t}^{\text{II}} + T_{d,t} - \mathcal{T}_{d,t}^{\text{II}}), \quad \text{and} \\ B_{d,t+1}^{\text{I}} &= R_t(B_{d,t}^{\text{II}} + S_{d,t}^{\text{II}} \mathcal{T}_{d,t}^{\text{II}} - S_t^{\text{II}} T_{d,t} - P_t C_{d,t}). \end{aligned} \quad (9)$$

The problem facing dealer d at the start of round I is to choose the price quote, $S_{d,t}^{\text{I}}$, that maximizes \mathbb{U}_t^d based on current information, $\Omega_{d,t}^{\text{I}}$, subject to (8) and (9). By assumption, all dealers choose quotes simultaneously, so the choice of $S_{d,t}^{\text{I}}$ cannot be conditioned on the quotes of other dealers, i.e., $S_{n,t}^{\text{I}}$ for $n \neq d$. At the start of round II, dealer d faces the analogous problem of choosing $S_{d,t}^{\text{II}}$ that maximizes \mathbb{U}_t^d based on $\Omega_{d,t}^{\text{II}}$, subject to (9). After all the dealers have quoted their round II prices, dealer d must determine his interdealer euro order, $T_{d,t}$, to maximize \mathbb{U}_t^d based on $\Omega_{d,t}^{\text{II}}$ and $\{S_{d,t}^{\text{II}}\}_{d=1}^D$ subject to (9). Once again, the choice of $T_{d,t}$ cannot be conditioned on incoming euro orders from other dealers, $\mathcal{T}_{d,t}^{\text{II}}$, because interdealer trading takes place simultaneously. After round II trading is complete, dealer d then chooses his consumption of the US and EU goods, $C_{d,t}(\text{US})$ and $C_{d,t}(\text{EU})$, to maximize \mathbb{U}_t^d based on current information and the sequence of future constraints in (8) and (9).

1.3 The Equilibrium Exchange Rate

An equilibrium in this model is described by a set of: (i) market-clearing equity prices, (ii) consumption and portfolio rules that maximize the expected utility of households, (iii) local currency pricing rules for firms that maximize the value of their dividend streams, (iv) optimizing quote, trade and consumption rules for dealers, and (v) interest rates consistent with both central banks monetary targets. To characterize this equilibrium, we need to specify how market clearing is achieved in the equity markets and how the information used in decision-making differs across agents. For this purpose, we make the following assumptions:

A1 Households within each county have the same information.

A2 Households cannot hold the equity issued by foreign firms.

Assumption A1 rules out intranational differences in the information available to individual households. It does not rule out differences between the information available to dealers, and households, or between households in different countries. We use the index H and \hat{H} to identify a representative US and European household and denote their common information sets at the start of period t by Ω_t^H , and $\Omega_t^{\hat{H}}$ respectively. With this simplification, we can use the currency orders of representative US and European households to describe how information concerning the macroeconomy is transmitted to the exchange rate. Trade in the equity markets is ruled out by A1 and A2. Taken together, these assumptions imply that all the equities issued by US and European firms are held the domestic representative household.⁴ As a result, the market clearing real price of US equity, Q_t/P_t , must equal the value of $\mathbb{Q}_t^{\text{US}} - D_t/P_t$ under an optimal period- t pricing policy where $\Lambda_{t+i,t}$ is the discount factor of US households.⁵ The market clearing price of European equity, \hat{Q}_t/\hat{P}_t , is analogously identified from the solution to the European firm's pricing problem. Notice that all other goods and asset prices are set by either firms, central banks or dealers.

Let us now focus on the determination of the equilibrium exchange rate. For this purpose we must consider the optimal choice of dealers' quotes in the two rounds of trading. As in Lyons

⁴Obviously, this implication of A1 and A2 is at odds with the degree of international financial integration we observe in world equity markets. We use it here to avoid having to model market-making activity in both currency and equity markets – an extension we leave for future research.

⁵Note that Q_t/P_t is the ex-dividend real price of US equity in period t , while \mathbb{Q}_t^{US} is the period- t present value of current and future real dividends valued using the US household's stochastic discount factor. Hence $\mathbb{Q}_t^{\text{US}} = Q_t/P_t + D_t/P_t$.

(1997), our trading environment constitutes a game played over two trading rounds each period by the D dealers. As such, we identify optimal dealer quotes and trades by the Perfect Bayesian Equilibrium (PBE) strategies. The resulting quotes for dealer d are given by

$$S_{d,t}^I = S_{d,t}^{II} = S_t = \mathcal{F}(\Omega_t^D), \quad (10)$$

where $\Omega_t^D = \cap_d \Omega_{d,t}^I$ is the information set common to all dealers at the beginning of round I in period t .

Equation (10) shows that optimal quotes have three features: First, each dealer quotes the same prices in rounds I and II. Second, quotes are common across all dealers. Third, all quotes are a function, $\mathcal{F}(\cdot)$, of common information at the start of period t , Ω_t^D . The intuition behind these features is straightforward: Recall that round II quotes are available to all dealers, are good for any amounts, and that each dealer can initiate trades with multiple counterparties. Under these conditions, any dealer quoting a different price from S_t^{II} would expose himself to arbitrage. A similar argument applies to the round I quotes. Again, these quotes are publicly observed and households are free to place orders with several dealers. Consequently, all dealers must quote the same prices to avoid arbitrage trading losses. Dealers must also have an incentive to fill their share of incoming orders at the quoted common price (i.e., they must be willing to participate in round I). This rules out differences between the round I and round II common quote. Finally, recall that quotes must be chosen simultaneously at the beginning of each trading round. As such, round I quotes will only be common across all dealers if they depend on common dealer information, Ω_t^D . Dealers may possess private information at the start of period t , but they cannot use it in their choice of quote without exposing themselves to arbitrage losses.

The relationship between the common period- t quote, S_t , and dealers' common information, Ω_t^D , implied by the PBE of our model is identified in the following proposition:

Proposition 1 *The log spot rate implied by the PBE quote strategies of dealers in period t is*

$$s_t = \left(\frac{1}{1+\eta} \right) \mathbb{E}_t^D \sum_{i=0}^{\infty} \left(\frac{\eta}{1+\eta} \right)^i f_{t+i}, \quad (11)$$

where η is a positive constant and \mathbb{E}_t^D denotes expectations conditioned on dealers' common period- t information, Ω_t^D . f_t denotes exchange rate fundamentals, which are defined as

$$f_t \equiv \hat{c}_t - c_t + m_t^* - \hat{m}_t^* + \varepsilon_t - \eta\psi \quad (12)$$

where $\varepsilon_t \equiv \ln(S_t \hat{P}_t / P_t)$ is the log real exchange rate and ψ is a risk premium.

The Appendix provides a detailed derivation of these equations from the log linearized equilibrium conditions as well as the results reported in the propositions that follow. Here, we provide some intuition. In the equilibrium of our model, dealers must be willing to fill incoming orders for euros at the price they quote. This means that the period- t quote must be set such that the expected excess return on euros between t and $t+1$ compensates the dealers for the risk of filling incoming currency orders during period t . In other words, all dealers must quote a price, $S_t \equiv \exp(s_t)$, such that

$$\mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t = \psi, \quad (13)$$

where $\Delta s_{t+1} \equiv s_{t+1} - s_t$ and ψ is the risk premium that depends on the conditional second moments of dealers' marginal utility of wealth and the future spot rate.⁶ Notice that $\mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t$ will differ from the expectations of (log) excess returns held by an individual dealer d when he has private information about the future spot rate (i.e., $\mathbb{E}_t^d s_{t+1} \neq \mathbb{E}_t^D s_{t+1}$). Individual dealers use this private information when making the round II trading decisions, not when choosing S_t . Proposition 1 follows easily from (13) and the implications of money market clearing. In particular, our specification for household preferences implies that the expected demand for dollars conditioned on Ω_t^D is approximately $\mathbb{E}_t^D m_t = \varpi + p_t + \mathbb{E}_t^D c_t - \eta r_t$. The expected demand for euros is similarly approximated by $\mathbb{E}_t^D \hat{m}_t = \varpi + \hat{p}_t + \mathbb{E}_t^D \hat{c}_t - \eta \hat{r}_t$. Under the reasonable assumption that central banks expectations concerning aggregate money demand are at least as precise as expectations based on Ω_t^D , $\mathbb{E}_t^D m_t = \mathbb{E}_t^D m_t^*$ and $\mathbb{E}_t^D \hat{m}_t = \mathbb{E}_t^D \hat{m}_t^*$ by the law of iterated expectations. Combining these expressions with (13) gives us the equations in Proposition 1.

Equation (11) plays a central role in our analysis. It shows that the log price of euros quoted by all dealers is equal to the present value of fundamentals, f_t . There are two noteworthy differ-

⁶For the sake of clarity, we shall take this risk premium to be constant in the analysis that follows. Allowing for time-variation does not affect the focus of our analysis.

ences between this specification and the exchange rate equations found in traditional monetary models. First, the definition of fundamentals in (12) includes the difference between foreign and home consumption rather than income. This arises because household preferences imply that the demand for national currencies depends on consumption rather than income. Second, equation (11) shows that 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 euros 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 . This means that exchange rate dynamics in our model are driven by the evolution of dealers' common information.

To further emphasize the importance of dealers' information, it is useful to consider the implications of (11) for the rate of depreciation, Δs_{t+1} . Specifically, if we iterate (11) forward to get $s_t = \mathbb{E}_t^D f_t + \eta \mathbb{E}_t^D \Delta s_{t+1}$, and rearrange, we can write the depreciation rate implied by the PBE quotes as

$$\Delta s_{t+1} = \frac{1}{\eta}(s_t - \mathbb{E}_t^D f_t) + e_{t+1}, \quad (14)$$

where

$$e_{t+1} \equiv \frac{1}{1+\eta} \sum_{i=0}^{\infty} \left(\frac{\eta}{1+\eta} \right)^i (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) f_{t+i+1}. \quad (15)$$

Equation (14) shows that the evolution of dealers' information can affect the depreciation rate through two channels: First, it can affect the difference between the current spot rate and dealers' estimate of current fundamentals, $s_t - \mathbb{E}_t^D f_t$. Second, it can lead to revisions in dealers' common knowledge forecasts of future fundamentals, $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) f_{t+i+1}$ for $i \geq 0$, which as (15) shows, contribute to dealer errors in forecasting next period's spot rate, $e_{t+1} \equiv s_{t+1} - \mathbb{E}_t^D s_{t+1}$. Since the first term in (14) is multiplied by the reciprocal of the semi-interest elasticity of money demand, $1/\eta$, a small number, the second channel is more likely to be empirically relevant. Indeed, because depreciation rates are very hard to forecast over short time periods, any attempt to make progress on understanding the origins of high-frequency spot rate dynamics must focus on the second channel.⁷ This is exactly the strategy of this paper. Specifically, our aim is to investigate whether transaction flows in the foreign exchange market convey information about fundamentals to dealers that they

⁷This point holds outside the context of our specific model. Engel and West (2005) note that forecasting the depreciation rate implied by several standard models will be hard because the value of the η coefficient in the present value representation of the equilibrium exchange rate is very large. Thus, the lack of forecastability does not, in itself, imply that spot exchange rates are disconnected from fundamentals (see, also, Evans and Lyons 2005).

then incorporate into their price quotes. In other words, we ask: Do transaction flows act as a proximate driver of spot exchange rates because they convey information that leads to revisions in dealers' forecasts of fundamentals, $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D)f_{t+i+1}$?

Before we address this question in detail, it proves useful to have an overview of how information contained in customer orders becomes incorporated into the equilibrium spot rate. Recall that the customer orders received by each dealer d , $\mathcal{T}_{d,t}^1$, represent private information to the dealer. In our model, the PBE strategy for each dealer is to use this information when initiating trades with other dealers (i.e., when choosing $T_{d,t}$). As a result, interdealer trading in round II effectively aggregates the information contained in customer orders received by dealers across the market. Indeed, it is the information conveyed by interdealer trading that augments dealer's common information by the start of period $t + 1$, and hence affects dealers' PBE choice for s_{t+1} . This does not mean that dealers necessarily have complete information about the current fundamentals by the end of interdealer trading. As the model of Evans and Lyons (2004) shows, they will under some special circumstances, but in general the inference problem facing dealers is too complex for them to make precise inferences about current fundamentals from their observations of interdealer trading. We will have more to say about dealers' assimilation of information below.

Finally, a few comments about the structure of the model are in order. Our specification for the household and production sectors deliberately does not include many of the features to be found in recent two-country general equilibrium models. Our aim, instead, is to present a minimal specification that provides microfoundations for the key macroeconomic factors that affect the behavior of the spot exchange rate. These are: (i) household demands for foreign currency motivated by optimal portfolio choice, and (ii) pricing decisions by firms that imply variations in the real exchange rate. While richer specifications for preferences and the production sector would clearly improve the empirical relevance of the model along many dimensions, they would not qualitatively affect the links between exchange rates, fundamentals and transaction flows which are the focus of this paper.

2 Fundamentals and Order Flow

We now examine the link between transaction flows, fundamentals and the spot exchange rate. More specifically, our aim is to identify the conditions under which the customer order flows reaching dealers, $\mathcal{T}_{d,t}^1$, convey new information about fundamentals that dealers incorporate into their price quotes for euros. We proceed in two steps. First we identify the factors driving customer order flows. Second, we show why order flows may convey information about fundamentals.

2.1 Customer Order Flow

Let x_t denote aggregate customer order flow defined as the dollar value of aggregate household purchases of euros from dealers during period t trading. The contribution of US households to this order flow is $S_t(\hat{B}_{H,t} - \hat{B}_{H,t-1}) = \alpha_t W_{H,t} \hat{R}_t - S_t \hat{B}_{H,t-1}$ where α_t denotes the desired share of euro bonds in the US households' wealth. Similarly, European households contribute $S_t(\hat{B}_{\hat{H},t} - \hat{B}_{\hat{H},t-1}) = \hat{\alpha}_t S_t \hat{W}_{\hat{H},t} \hat{R}_t - S_t \hat{B}_{\hat{H},t-1}$ where $\hat{\alpha}_t$ is the desired share of euro bonds in European wealth. Market clearing requires that aggregate holdings of euro bonds by households and non-households (i.e., central banks and dealers) sum to zero, so that $\hat{B}_{t-1} + \hat{B}_{\hat{H},t-1} + \hat{B}_{\hat{H},t-1} = 0$ where \hat{B} denotes the aggregate holdings of non-households. Hence, aggregate order flow can be written as

$$x_t = [\alpha_t \ell_t + \hat{\alpha}_t (1 - \ell_t)] W_t \hat{R}_t + S_t \hat{B}_{t-1}, \quad (16)$$

where $W_t \equiv W_{H,t} + S_t \hat{W}_{\hat{H},t}$ is world household wealth in dollars, and $\ell_t \equiv W_{H,t}/W_t$. Thus, order flow depends upon the portfolio allocation decisions of US and European households (via α_t , and $\hat{\alpha}_t$), the level and international distribution of household wealth (via W_t and ℓ_t) and the outstanding stock of foreign bonds held by non-households from last period's trading, \hat{B}_{t-1} . These elements imply that order flow contains both pre-determined (backward-looking) and non-predetermined (forward-looking) components. The former include the level and distribution of wealth, the latter are given by the portfolio shares because they depend on households' forecasts of future returns. We formalize these observations in the following proposition.

Proposition 2 *The utility-maximizing choice of portfolios by US and European households implies that aggregate order flow may be approximated by*

$$x_t = \phi \nabla \mathbb{E}_t^H s_{t+1} + \hat{\phi} \nabla \mathbb{E}_t^{\hat{H}} s_{t+1} + o_t, \quad (17)$$

with $\phi, \hat{\phi} > 0$, where $\nabla E_t^\omega s_{t+1} \equiv E_t^\omega s_{t+1} - E_t^D s_{t+1}$ for $\omega = \{H, \hat{H}\}$ and o_t denotes terms involving the distribution of wealth, non-household bond holdings, and the consumption of European households.

Equation (17) describes the second important implication of our model. It relates order flow to the difference between households' forecasts for the future spot rate, $E_t^\omega s_{t+1}$ for $\omega = \{H, \hat{H}\}$, and dealers' forecasts, $E_t^D s_{t+1}$. In particular, there will be positive order flow for euros if households are more optimistic about the future value of the euro than dealers, so that $\nabla E_t^\omega s_{t+1} > 0$ for $\omega = \{H, \hat{H}\}$.

To understand why differences in expectations play this role, we need to focus on how households choose their portfolios. In the appendix we show that the optimal share of US household wealth held in the form of euro bonds is increasing in the expected log excess return, $\mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t$. Now, when dealers' foreign currency quotes satisfy (11) and (12), the log spot rate also satisfies $\mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t = \psi$. We can therefore write the excess return on European bonds expected by US households as

$$\mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t = \mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t + \nabla \mathbb{E}_t^H s_{t+1} = \nabla \mathbb{E}_t^H s_{t+1} + \psi.$$

Thus, when US households are more optimistic about the future value of the euro than dealers, they expect a higher excess return on euro bonds. These expectations, in turn, increase the desired fraction of US household wealth in euro bonds, so US households place more orders for euros with dealers in round 1 of period- t trading. Optimism concerning the value of the euro on the part of European households (i.e. $\nabla \mathbb{E}_t^{\hat{H}} s_{t+1} > 0$) contributes positively to order flow in a similar manner.

Of course household portfolio choices are also affected by risk. The o_t variable in (17) summarizes the effects of risk, the distribution of wealth and non-household bond holdings. These terms will not vary significantly from month to month or quarter to quarter under most circumstances, and so will not be the prime focus of the analysis below. We shall concentrate instead on how the existence

of dispersed information, manifest through the existence of the forecast differentials, $\nabla \mathbb{E}_t^H s_{t+1}$ and $\nabla \mathbb{E}_t^{\widehat{H}} s_{t+1}$, affects the joint behavior of order flow, spot rates and fundamentals.

2.2 How is Order Flow Related to Fundamentals?

To address this question, we first characterize the equilibrium dynamics of fundamentals. Let y_t denote the vector that describes the state of the economy at the start of period t . This vector includes the variables that comprise fundamentals (i.e. consumption, money targets and the real exchange rate) as well as those variables needed to describe firms' behavior, and the distribution of wealth across households and dealers. In Evans and Lyons (2004), we describe in detail the equilibrium dynamics of a model with a similar structure. Here our focus is on the empirical implications of the model, so we present the equilibrium dynamics in reduced form:

$$\Delta y_{t+1} = A\Delta y_t + u_{t+1}, \quad (18)$$

where $\Delta y_t \equiv y_t - y_{t-1}$ with u_{t+1} a vector of mean zero shocks. This specification for the equilibrium dynamics of the state variables is completely general, yet it allows us to examine the link between order flow and fundamentals in a straightforward way.

We start with the behavior of the spot exchange rate. Let fundamentals be a linear combination of the elements in the state vector: $f_t = C y_t$. When dealers quote spot rates according to (11) in Proposition 1, and (18) describes the dynamics of the state vector y_t , the spot exchange rate can be written as

$$s_t = \pi \mathbb{E}_t^D \mathbf{y}_t, \quad (19)$$

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

$$\begin{aligned} \nabla \mathbb{E}_t^H s_{t+1} &= \pi \left(\mathbb{E}_t^H \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} \right) \\ &= \pi \left(\mathbb{E}_t^H \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} \right). \end{aligned} \quad (20)$$

Suppose that US households collectively know as much about the state of the economy as dealers

do. Under these circumstances, the right hand side of (20) 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 the future spot rate depends on 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 collectively 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 formalize these ideas in the following proposition.

Proposition 3 *If US and European households are as well-informed about the state of the economy as dealers, so that $\Omega_t^D \subset \Omega_t^H$ and $\Omega_t^D \subset \Omega_t^{\hat{H}}$, then US and European forecast differentials for spot rates are*

$$\nabla \mathbb{E}_t^H s_{t+1} = \pi \kappa (\mathbb{E}_t^H \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}), \quad (21a)$$

$$\nabla \mathbb{E}_t^{\hat{H}} s_{t+1} = \pi \hat{\kappa} (\mathbb{E}_t^{\hat{H}} \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}), \quad (21b)$$

and order flow follows

$$x_t = \phi \pi \kappa \nabla \mathbb{E}_t^H \mathbf{y}_{t+1} + \hat{\phi} \pi \hat{\kappa} \nabla \mathbb{E}_t^{\hat{H}} \mathbf{y}_{t+1} + o_t. \quad (22)$$

for some matrices, κ and $\hat{\kappa}$.

The intuition behind Proposition 3 is straightforward. If US households are collectively as well-informed about the future state of the economy as dealers, then $\nabla \mathbb{E}_t^H s_{t+1} = \pi \mathbb{E}_t^H (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1}$, so the forecast differential depends on the speed at which US household expect dealers to assimilate new information concerning the future state of the economy. We term this the pace of information aggregation. If dealers learn nothing new about \mathbf{y}_{t+1} during period- t trading, $\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} = \mathbb{E}_t^D \mathbf{y}_{t+1}$. Hence, if US households expect that period- t trading will reveal nothing new to dealers, $\mathbb{E}_t^H (\mathbb{E}_{t+1}^D - \mathbb{E}_t^D) \mathbf{y}_{t+1} = 0$ and there is no difference between dealer and household forecasts of future spot rates. Under these circumstances, there is no information aggregation so κ and $\hat{\kappa}$ are equal to null matrices. Alternatively, if households expect dealers to assimilate information from period- t trading, the forecast differentials for spot rates will be non-zero. In

the extreme case where period- t trading is sufficiently informative to reveal to dealers all that households know about the future state of the economy, $(\mathbb{E}_{t+1}^D - \mathbb{E}_t^D)\mathbf{y}_{t+1}$ will equal $\mathbb{E}_t^\omega \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}$ for $\omega = \{\mathbb{H}, \widehat{\mathbb{H}}\}$. In this case, information aggregates quickly, so κ and $\hat{\kappa}$ equal the identity matrices. Under other circumstances where the pace of information aggregation is slower, the κ and $\hat{\kappa}$ matrices will have many non-zero elements. (Exact expressions for κ and $\hat{\kappa}$ are provided in the Appendix.)

Equation (22) combines (17) from Proposition 2 with (21). This equation expresses order flow in terms of forecast differentials for the future state of the economy and the speed of information aggregation. Since fundamentals represent a combination of the elements in \mathbf{y}_t , (22) also serves to link dispersed information regarding future fundamentals to order flow. In particular, if households have more information about the future course of fundamentals than dealers, and dealers are expected to assimilate at least some of this information from transaction flows each period, order flow will be correlated with variations in the forecast differentials for fundamentals.

We should emphasize that the household currency orders driving order flow in this model are driven solely by the desire to optimally adjust portfolios. Households have no desire to inform dealers about the future state of the economy, so the information conveyed to dealers via transaction flows occur as a by-product of their dynamic portfolio allocation decisions. The transaction flows associated with these decisions establish the link between order flow, dispersed information, and the speed of information shown in equation (22).

One aspect of our model deserves further clarification. Our model abstracts from informational heterogeneity at the household level, so Ω_t^H , and $\Omega_t^{\widehat{H}}$ represent the information sets of the representative US and European households. This means that the results in Proposition 3 are derived under the assumption that representative households have strictly more information than dealers ($\Omega_t^D \subset \Omega_t^H$ and $\Omega_t^D \subset \Omega_t^{\widehat{H}}$). Clearly this is a strong assumption. Taken literally, it implies that every household knows more about the current and future state of the economy than any given dealer. Fortunately, our central results do not rely on this literal interpretation. To see why, suppose, for example, that each household receives its own money demand shock and is thereby privately motivated to trade foreign exchange. In this setting, no household 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 the sake of parsimony, we have not modelled heterogeneity at the US and European household levels. Instead, we assume

that households in any given country share the same information about the macroeconomy. Extending the model to capture heterogeneity is a natural extension, but not one that would alter the main implications of our model that are the focus of the empirical analysis below.⁸

3 Data

Our empirical analysis utilizes a new data set that comprises end-user transaction flows, spot rates and macro fundamentals over six and a half years. The transaction flow data differs in two important respects from the data used in earlier work (e.g., Evans and Lyons 2002a,b). First, they cover a much longer time period; January 1993 to June 1999. Second, they come from transactions between end-users and a large bank, rather than from inter-bank transactions. Our data covers transactions with three end-user segments: non-financial corporations, investors (such as mutual funds and pension funds), and leveraged traders (such as hedge funds and proprietary traders). The data set also contains information on trading location. From this we construct order flows for six segments: trades executed in the US and non-US for non-financial firms, investors, and leveraged traders. Though inter-bank transactions account for about two-thirds of total volume in major currency markets at the time, they are largely derivative of the underlying shifts in end-user currency demands. Our data include all the end-user trades with Citibank in the largest spot market, the USD/EUR market, and the USD/EUR forward market.⁹ Citibank had the largest share of the end-user market in these currencies at the time, ranging between 10 and 15 percent. The flow data are aggregated at the daily frequency and measure in \$m the imbalance between end-user orders to purchase and sell euros.

There are many advantages of our transaction flow data. First, the data are simply more powerful, covering a much longer time span. Second, because the underlying 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. We shall see that those not

⁸As is standard in literature, we use “households” as a metaphor for a wide class of agents that constitute the private sector. In particular, households represent the class of non-dealer agents that observe some component of macro fundamentals. One way to introduce heterogeneity would be to differential between the information available to different members of this class, e.g., financial institutions and individuals.

⁹Before 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.

covered by extant end-user data sets are empirically important for exchange rate determination.¹⁰ Fourth, because the data are disaggregated into segments, we can address whether the behavior of the individual segments is similar, and whether they convey the same information concerning exchange rates and macro fundamentals.

Our empirical analysis also utilizes new high-frequency real-time estimates of macro variables for the US and Germany: specifically GDP, consumer prices, and M1 money. As the name implies, a real-time estimate of a variable is the estimated value based on public information available on a particular date. These estimates are conceptually distinct from the values that make up standard macro time-series. Importantly, because they are computed from information available to market participants contemporaneously, real-time estimates are relevant for understanding the link between the foreign exchange market (or any other financial market) and the macroeconomy.

A simple example clarifies the difference between a real-time estimate of a macro variable and the data series usually employed in empirical studies. Let \varkappa denote a variable representing macroeconomic activity during month τ , that ends on day $M(\tau)$, with value $\varkappa_{M(\tau)}$. Data on the value of \varkappa is released on day $R(\tau)$ after the end of month τ with a reporting lag of $R(\tau) - M(\tau)$ days. Reporting lags vary from month to month because data is collected on a calendar basis, but releases issued by statistical agencies are not made on holidays and weekends. (For quarterly series, such as GDP, reporting lags can be as long as several months.) The real-time estimate of \varkappa on day t in month τ is the expected value of $\varkappa_{M(\tau)}$ based on day- t information. Formally, the real-time estimate of a monthly series \varkappa is

$$\varkappa_{M(\tau)|t} \equiv \mathbb{E}[\varkappa_{M(\tau)}|\Omega_t] \quad \text{for } M(\tau - 1) < t \leq M(\tau), \quad (23)$$

where Ω_t denotes an information set that only contains data known at the start of day t . In the case of a quarterly series like GDP, the real-time estimate on day t is

$$\varkappa_{Q(i)|t} \equiv \mathbb{E}[\varkappa_{Q(i)}|\Omega_t] \quad \text{for } Q(i - 1) < t \leq Q(i), \quad (24)$$

where $Q(i)$ denotes the last day of quarter i .

¹⁰Froot and Ramadorai (2002), consider the transactions flows associated with portfolio changes undertaken by institutional investors. Osler (2003) examines end-user stop-loss orders.

Real-time estimates are conceptually distinct from the values for $\varkappa_{M(\tau)}$ or $\varkappa_{Q(i)}$ found in standard macro time series. To see why, let $v(\tau)$ denote the last day on which data on \varkappa for month τ was revised. A standard monthly time series for variable \varkappa spanning months $\tau = 1, \dots, T$ comprises the sequence $\{\varkappa_{M(\tau)|v(\tau)}\}_{\tau=1}^T$.¹¹ This latest vintage of the data series incorporates information about the value of \varkappa that was not known during month τ . We can see this more clearly by writing the difference between $\varkappa_{M(\tau)|v(\tau)}$ and real-time estimate as

$$\varkappa_{M(\tau)|v(\tau)} - \varkappa_{M(\tau)|t} = (\varkappa_{M(\tau)|v(\tau)} - \varkappa_{M(\tau)|R(\tau)}) + (\varkappa_{M(\tau)|R(\tau)} - \varkappa_{M(\tau)|M(\tau)}) + (\varkappa_{M(\tau)|M(\tau)} - \varkappa_{M(\tau)|t}). \quad (25)$$

The first term on the right hand side represents the effects of data revisions following the initial data release. We denote the value for $\varkappa_{M(\tau)}$ released on day $R(\tau)$ by $\varkappa_{M(\tau)|R(\tau)}$ so $\varkappa_{M(\tau)|v(\tau)} - \varkappa_{M(\tau)|R(\tau)}$ identifies the effects of all the data revisions between $R(\tau)$ and $v(\tau)$. Croushore and Stark (2001), Faust, Rogers, and Wright (2003) and others have emphasized that these revisions are significant for many series. The second term in (25) is the difference between the value for $\varkappa_{M(\tau)}$ released on day $R(\tau)$ and the real-time estimate of $\varkappa_{M(\tau)}$ at the end of the month. This term identifies the impact of information concerning $\varkappa_{M(\tau)}$ collected by the statistical agency before the release date that was not part of the $\Omega_{M(\tau)}$ information set. This term is particularly important in the case of quarterly data where the reporting lag can be several months. The third term on the right of (25) is the difference between the real time estimate of $\varkappa_{M(\tau)}$ at the end of month τ and the estimate on a day earlier in the month.

In this paper we construct real time estimates of GDP, consumer prices, and M1 for the US and Germany using an information set based on 35 macro data series. For the US estimates our specification for Ω_t includes the 3 quarterly releases on US GDP and the monthly releases on 18 other US macro variables. The German real-time estimates are computed using a specification for Ω_t that includes the 3 quarterly release on German GDP and the monthly releases on 8 German macro variables. All series come from a database maintained by Money Market News Services that contains details of each data release. We use the method developed in Evans (2005) to compute the real-time estimates. Specially, for each variable \varkappa we use the Kalman Filter to calculate the

¹¹For the sake of notational clarity, we have implicitly assumed that the statistical agency uses the Ω_t information set when computing data revisions. Relaxing this assumption to give the agency superior information does not affect the substance of our argument. For a further discussion, see Evans (2005).

conditional expectations in (23) and (24) from estimates of a state space model that specifies a daily time series process for \varkappa_t and its relation to the sequence of data releases (i.e. the elements of Ω_t). The Appendix provides an overview of the state space model and the estimation method.

Our real time estimates have several important attributes. First our specification insures that the information set used to compute each real-time estimate, Ω_t , is subset of the information available to participants in the foreign exchange market on day t . This means that the real-time estimate of monthly variable \varkappa , $\varkappa_{M(\tau)|t}$, can be legitimately used as a variable affecting market actively on day t . By contrast, the values for $\varkappa_{M(\tau)}$ found in either the first or final vintage of a time series (i.e., $\varkappa_{M(\tau)|R(\tau)}$ or $\varkappa_{M(\tau)|V(\tau)}$) contain information that was not known to participants on day t .

The second attribute of the real-time estimates concerns the frequency with which macro data is collected and released. Even though the macro variables are computed on a quarterly (GDP) or monthly (prices and money) basis, real-time estimates vary day-by-day as the flow of macro data releases augments the information set Ω_t . This attribute is illustrated in Figure 1, where we plot the real-time estimates of log GDP for the US and Germany. The real-time estimates (shown by the solid plots) clearly display a much greater degree of volatility than the cumulant of the data releases (shown by the dashed plots). This volatility reflects how inferences about current GDP change as information arrives in the form of monthly data releases during the current quarter and GDP releases referring to the previous quarter. A further noteworthy feature of Figure 1 concerns the difference between the real-time estimates and the ex post value of log GDP represented by the vertical gap between the solid and dashed plots. This gap should be small if the current level of GDP could be precisely inferred from contemporaneously available information. However, as the figure clearly shows, there are many occasions where the real-time estimates are substantially different from the ex post values.

A third attribute of the real-time estimates concerns their variation over our sample period. Although our data covers only six and a half years, Figure 1 shows that there is considerable variation in our GDP measures within this relatively short time span. The vertical axis shows that real-time estimates of US GDP have a range of approximately 2.4 percent around trend, while the range for German GDP is more than 4.5 percent.

Figure 2 displays the variation in the other real-time estimates. The left hand panel shows that

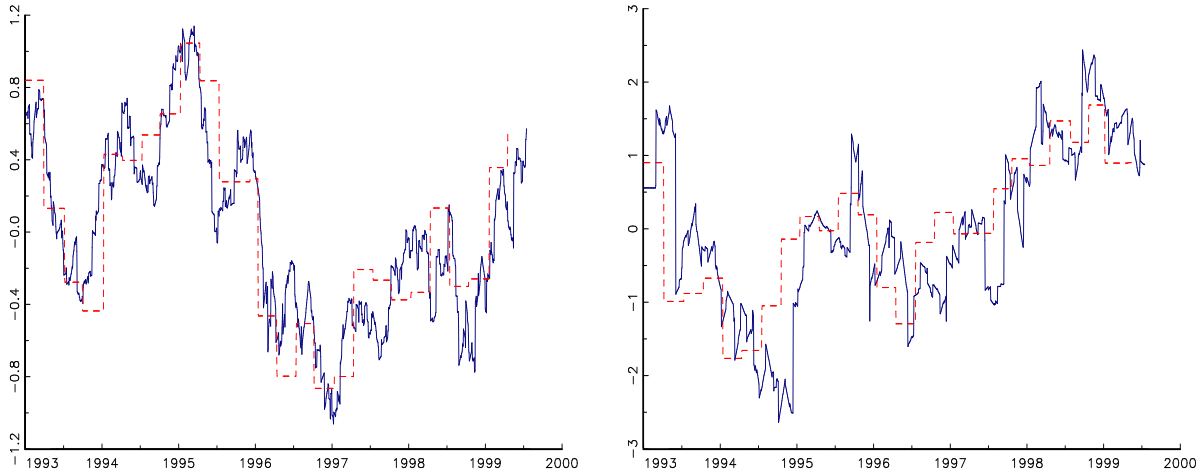


Figure 1: Real-time estimates of log GDP (solid line) and cumulant of GDP releases (dashed line). The right hand panel shows plots for US GDP, the left panel plots for German GDP. All series are detrended and multiplied by 100.

while the real-time estimates of US prices varied very little from their trend, German prices varied by almost 3 percent. In the right hand panel the real-time estimates for M1 have a range of almost 16 percent in the US and 7 percent in Germany. Because the reporting lag for both prices and money are much shorter than that for GDP, the differences between these real-time estimates and the ex-post values are much smaller than those shown in Figure 1. (We omit ex-post values from Figure 2 for clarity.) Real-time uncertainty about current consumer prices and M1 is far less than the degree of uncertainty surrounding GDP.

In sum, all but one of the real-time estimates varies significantly over our sample period. This is important if we want to study how macroeconomic conditions affect the foreign exchange market. If all of our real-time estimates were essentially constant over our sample, there would be no room for detecting how perceived developments in the macroeconomy are reflected in the foreign exchange market.

In the analysis that follows we consider the joint behavior of exchange rates, order flows and the real-time estimates of macro variables at the weekly frequency. This approach provides more precision in our statistical inferences concerning the high frequency link between flows, exchange rates and macro variables than would be otherwise possible. The weekly timing of the variables is as follows: We take the log spot rate at the start of week t , s_t , to be the log of the offer rate (USD/EUR) quoted by Citibank at the end of trading on Friday of week $t - 1$ (approximately

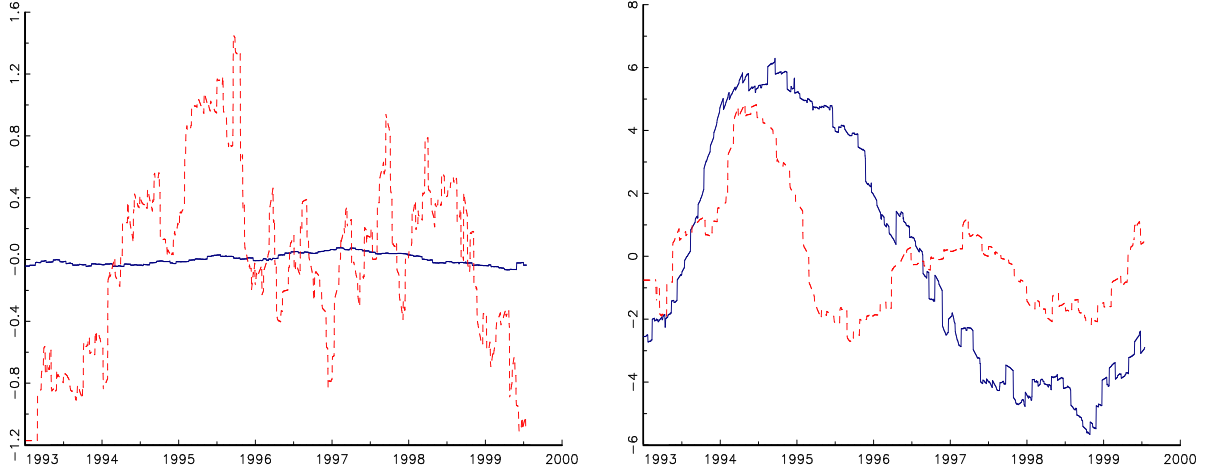


Figure 2: Left hand panel: Real-time estimates of log US consumer prices (solid line) and Germany consumer prices (dashed line). Right hand panel: Real-time estimates of US M1 (solid line) and German M1 (dashed line). All series are detrended.

17:00 GMT). This is also the point at which we sample the week- t interest rates from Datastream. The week- t flow from segment j , $x_{j,t}$, is computed as the total value in \$m of dollar purchases initiated by the segment against Citibank's quotes between the 17:00 GMT on Friday of week $t - 1$ and Friday of week t . Positive values for these order flows therefore denote net demand for euros by the end-user segment. The week- t change in the real-time estimates are computed as the difference between the Friday estimates on weeks $t - 1$ and $t - 2$. This timing insures that the week- t change in the real-time estimates are derived using a subset of the information available to foreign exchange dealers when quoting spot rates at the start of week- t trading. In other words, our timing assumptions insure that the information used to compute $\varkappa_{M(\tau)|t}$ or $\varkappa_{Q(i)|t}$ is a subset of the information available to all dealers when quoting the spot rate s_t .¹²

Summary statistics for the weekly data are reported in Table 1. The statistics in panel A show that weekly changes in the log spot rate, $\Delta s_t \equiv s_t - s_{t-1}$, have a mean very close to zero and display no significant serial correlation. These statistics are typical for spot exchange rates and suggest that the univariate process for s_t is well-characterized by a random walk. Two features stand out from the statistics on the six flow segments shown in Panel B. First, the order flows are large and volatile. Second, they display no significant serial correlation. At the weekly frequency,

¹²More precisely, our timing assumptions imply that the real-time estimates of $\varkappa_{M(\tau)|t}$ or $\varkappa_{Q(i)|t}$ incorporate macro data releases that are only a few hours old by the time dealers quote s_t .

the end-user flows appear to represent shocks to the foreign exchange market arriving at Citibank. This is not to say that flows are unrelated across segments. The (unreported) cross-correlations between the six flows range from approximately -0.16 to 0.16, but cross-autocorrelations are all close to zero.

Summary statistics for the weekly changes in the real-time estimates are reported in Panel C of Table 1. The most notable feature of these statistics concerns the estimated autocorrelations. These are generally small and insignificant at the 5% level except in the case of the M1 real-time estimates. For perspective on these findings, consider the weekly change in the monthly series \varkappa . If the weekly change falls within a single month, the change in real-time estimate is

$$\varkappa_{M(\tau)|W(j)} - \varkappa_{M(\tau)|W(j-1)} \equiv \mathbb{E}[\varkappa_{M(\tau)}|\Omega_{W(j)}] - \mathbb{E}[\varkappa_{M(\tau)}|\Omega_{W(j-1)}],$$

where $W(j)$ denotes the last day of week j . In this case the weekly change simply captures the flow of new information concerning the value of \varkappa in the current month, $\varkappa_{M(\tau)}$, and so should not be correlated with any elements of $\Omega_{W(j-1)}$, including past changes in the real-time estimates. If the weekly change occurs at the end of the month, the change in the real-time estimate can be written as

$$\varkappa_{M(\tau+1)|W(j)} - \varkappa_{M(\tau)|W(j-1)} = (\mathbb{E}[\varkappa_{M(\tau+1)}|\Omega_{W(j)}] - \mathbb{E}[\varkappa_{M(\tau+1)}|\Omega_{W(j-1)}]) + (\mathbb{E}[\varkappa_{M(\tau+1)} - \varkappa_{M(\tau)}|\Omega_{W(j-1)}]).$$

Here the first term on the right hand side represents the the flow of new information concerning $\varkappa_{M(\tau+1)}$. Once again this should not be correlated with any elements in $\Omega_{W(j-1)}$. The second term identifies initial expectations about the growth in \varkappa from month τ to $\tau + 1$. This term is a function of elements in $\Omega_{W(j-1)}$ and so may be correlated with past changes in the real-time estimates.

The autocorrelations in Table 1 are computed from all weekly changes in our sample, and so capture the characteristics of both the within and cross-month changes. The small amounts of positive serial correlation we see reflect the fact that forecasts for monthly M1 growth are positively correlated with past growth, a feature that is evident from the plots in Figure 2. That said, the over-arching implication of the estimated autocorrelations is that the weekly changes in each real-time estimates primarily reflects the arrival of new information concerning the current state of the corresponding macro variable. Our real-time estimates will therefore enable us to capture changing

Table 1: Summary Statistics							
	mean Std.	max min	skewness kurtosis	ρ_1	Autocorrelations		
					ρ_2	ρ_4	ρ_8
A: Exchange Rate							
(i) Δs_t (x100)	-0.043 1.234	3.722 -3.715	0.105 3.204	-0.061 (0.287)	0.027 (0.603)	0.025 (0.643)	-0.015 (0.789)
B: Order Flows							
(ii) Corporate US	-16.774 108.685	549.302 -529.055	-0.696 9.246	-0.037 (0.434)	-0.040 (0.608)	0.028 (0.569)	-0.028 (0.562)
(iii) Corporate Non-US	-59.784 196.089	634.918 -692.419	-0.005 3.908	0.072 (0.223)	0.089 (0.124)	-0.038 (0.513)	0.103 (0.091)
(iv) Traders US	-4.119 346.296	1710.163 -2024.275	0.026 8.337	-0.021 (0.735)	0.024 (0.602)	0.126 (0.101)	-0.009 (0.897)
(v) Traders Non-US	11.187 183.36	972.106 -629.139	0.392 5.86	-0.098 (0.072)	0.024 (0.660)	0.015 (0.747)	0.083 (0.140)
(vi) Investors US	19.442 146.627	535.32 -874.15	-1.079 11.226	0.096 (0.085)	-0.024 (0.568)	-0.03 (0.536)	-0.016 (0.690)
(vii) Investors Non-US	15.85 273.406	1881.284 -718.895	0.931 9.253	0.061 (0.182)	0.107 (0.041)	-0.030 (0.550)	-0.014 (0.825)
C: Real-Time Data							
(viii) US Output	-0.001 0.201	0.711 -0.610	0.060 0.134	0.072 (0.084)	0.107 (0.056)	-0.015 (0.788)	0.058 (0.329)
(ix) US Prices	0.000 0.030	0.250 -0.104	1.527 18.673	0.006 (0.695)	-0.034 (0.135)	0.091 (0.142)	0.004 (0.963)
(x) US Money	-0.007 1.368	5.679 -6.981	-0.230 9.160	0.076 (0.003)	0.065 (0.012)	0.132 (0.131)	0.032 (0.595)
(xi) German Output	0.002 0.514	2.840 -4.087	-0.298 20.437	0.072 (0.138)	-0.039 (0.193)	-0.009 (0.873)	0.019 (0.671)
(xii) German Prices	0.002 0.817	4.090 -3.988	0.105 8.632	0.069 (0.111)	0.005 (0.918)	0.009 (0.864)	-0.044 (0.444)
(xiii) German Money	0.022 1.421	7.447 -6.263	1.073 13.120	0.116 (0.000)	0.083 (0.000)	0.100 (0.339)	0.042 (0.473)
Notes: The table reports summary statistics for the following variables sampled at the weekly frequency between January 1993 and June 1999: (i) the weekly change in the log spot rate x100, (ii)-(vii) order flows from end-user segments cumulated over a week, and (viii)- (xiii) weekly changes in real-time estimates measured in annual percent. The last four columns on the right report autocorrelations ρ_i at lag i and p-values for the null that $\rho_i = 0$ in parentheses.							

perceptions concerning the current state of the macroeconomy rather than its actual evolution. It is the link between the changing perceptions of market participants and the behavior of exchange

rate that is the focus of our empirical analysis.

4 Empirical Analysis

In this section we examine the empirical implications of Propositions 1 - 3. First, we consider the implications of our model for the correlation between order flows and changes in spot exchange rates. Next, we examine the links between spot rates and fundamentals. Our model identifies conditions under which order flow should have incremental forecasting power beyond spot rates. We find strong empirical support for this prediction, implying that order flows convey information about macro fundamentals to the market. Finally, we investigate whether this informational role can account for the forecasting power of order flows for future changes in exchange rates.

4.1 The Order Flow/Spot Rate Correlation

Evans and Lyons (2002a,b) show that order flows account for between 40 and 80 percent of the daily variation in the spot exchange rates of major currency pairs. Propositions 1 - 3 provide a structural interpretation of this finding. Recall that when dealers' foreign currency quotes satisfy (11) and (12) in Proposition 1, the log spot rate satisfies $\mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t = \psi$. Combining this restriction with the identity $\Delta s_{t+1} \equiv \mathbb{E}_t^D \Delta s_{t+1} + s_{t+1} - \mathbb{E}_t^D s_{t+1}$ gives

$$\begin{aligned} \Delta s_{t+1} &= r_t - \hat{r}_t + \psi + s_{t+1} - \mathbb{E}_t^D s_{t+1}, \\ &= r_t - \hat{r}_t + \psi + \pi \left(\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} \right), \end{aligned} \tag{26}$$

where the second line follows from the relation between the spot rate and state vector described by equation (19). Thus, Proposition 1 implies that the rate of depreciation is equal to the interest differential, a risk premium, and the revision in dealer forecasts concerning the future state of the economy between periods t and $t + 1$. This forecast revision is attributable to two possible information sources. The first is public information that arrives right at the start of period $t + 1$, before dealers quote s_{t+1} . The second is information conveyed by the transaction flows during period t . It is this second information source that accounts for the correlation between order flow and spot rate changes in the data.

Proposition 4 *When dealer quotes for the price of foreign currency satisfy (11), and order flow follows (22), the rate of depreciation can be written as*

$$\Delta s_{t+1} = r_t - \hat{r}_t + \psi + b(x_t - \mathbb{E}_t^D x_t) + \zeta_{t+1}. \quad (27)$$

ζ_{t+1} represents the portion of $\pi(\mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1})$ that is uncorrelated with order flow, and b is a projection coefficient equal to

$$\frac{\pi \mathbb{C}\mathbb{V}(\mathbf{y}_{t+1}, o_t)}{\mathbb{V}(x_t)} + \frac{\phi \pi \mathbb{V}(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1}) \kappa' \pi'}{\mathbb{V}(x_t)} + \frac{\hat{\phi} \pi \mathbb{V}(\nabla \mathbb{E}_t^{\hat{\Pi}} \mathbf{y}_{t+1}) \hat{\kappa}' \pi'}{\mathbb{V}(x_t)}, \quad (28)$$

where $\mathbb{V}(\cdot)$ and $\mathbb{C}\mathbb{V}(\cdot, \cdot)$ denote the population variance and covariance.

Inspection of expression (28) reveals that the observed correlation between order flow and the rate of depreciation can arise through two channels. First, if the distribution of wealth and dealer bond holdings affect order flow (via o_t in equation 17) and has forecasting power for fundamentals, order flow will be correlated with the depreciation rate through the first term in (28). Since there is little variation in o_t from month to month or even quarter to quarter, it is unlikely that this channel accounts for much of the order flow/spot rate correlation we observe at a daily or weekly frequency. The second channel operates through the transmission of dispersed information. If household expectations for the future state vector differ from dealers' expectations, and information aggregation accompanies trading in period t , both the second and third terms in (28) will be positive. Notice that the depreciation rate is correlated with order flow in this case not just because households and dealers hold different expectations, but also because households expect some of their information to be assimilated by dealers from the transaction flows they observe in period t . In this sense, the correlation between order flow and the depreciation rate informs us about both the existence of dispersed information and the pace at which information aggregation takes place.

Now we turn to the empirical evidence. Table 2 presents the results of regressing currency returns between the start of weeks t and $t + \tau$ for $\tau = \{1, 4\}$ on a constant, the interest differential at the start of week t , $r_t - \hat{r}_t$, and the order flows from the six segments between the start of weeks t and $t + \tau$. These regressions are the empirical counterparts to (27) with the six flows proxying for

Table 2: Contemporaneous Return Regressions								
Horizon	Interest Differential	Corporate US Non-US		Traders US Non-US		Investors US Non-US		R ² χ ² (p-value)
1 week	-0.2 (0.391)	-0.326 (0.584)	-1.096 (0.309)					0.03 12.627 (0.002)
	-0.193 (0.364)			1.018 (0.170)	0.63 (0.350)			0.094 38.139 (0.000)
	-0.134 (0.341)					1.194 (0.576)	1.441 (0.327)	0.131 30.818 (0.000)
	-0.297 (0.325)	-0.321 (0.535)	-0.817 (0.291)	0.791 (0.170)	0.632 (0.337)	1.108 (0.572)	1.254 (0.312)	0.213 88.758 (0.000)
4 weeks	-0.182 (0.252)	-0.006 (0.165)	-0.340 (0.085)					0.058 16.101 (0.000)
	-0.168 (0.247)			0.279 (0.061)	0.11 (0.118)			0.113 23.354 (0.000)
	0.001 (0.204)					0.144 (0.121)	0.49 (0.063)	0.251 68.471 (0.000)
	-0.19 (0.209)	0.027 (0.138)	-0.202 (0.071)	0.177 (0.060)	0.046 (0.101)	0.218 (0.119)	0.41 (0.066)	0.323 109.571 (0.000)

Notes: The table reports coefficients and standard errors from regressions of returns measured over one week and one month, on a constant (estimates not reported), the lagged interest differential and order flows cumulated over the same horizon. The interest differential is computed from the one month rates on Euro Dollar and DM deposits. Estimated coefficients on the order flows are multiplied by 1000. The right hand column reports χ^2 statistics for the null that all the coefficients on order flows are zero. Estimates are calculated at the weekly frequency. The standard errors correct for heteroskedasticity and the moving average error process induced by overlapping forecasts (4 week results).

$x_t - \mathbb{E}_t^D x_t$. Several points emerge from the table. First, the coefficients on the order flow segments are quite different from each other. Some are positive, some are negative, some are highly statistically significant, others are not. Second, while the coefficients on order flow are jointly significant in every regression we consider, the proportion of the variation in returns that they account for rises with the horizon: the R^2 statistic in regressions with all six flows rises from 21 to 32 percent as we move from the 1 to 4 week horizon.¹³ Third, the explanatory power of the order flows shown

¹³Froot 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.

here is much less than that reported for interdealer order flows. Evans and Lyons (2002a), for example, report that interdealer order flow accounts for approximately 60 percent of the variations in the \$/DM at the daily frequency. Finally, we note that none of the coefficients on the interest differential are statistically significant, and many have an incorrect (i.e. negative) sign.¹⁴ This is not surprising in view of the empirical literature examining uncovered interest parity. However, the estimated coefficients on the order flows are essentially unchanged if we re-estimate the regressions with a unity restriction on the interest differential, as implied by equation (27).

The key to understanding these results lies in the distinction between unexpected order flow in the model, $x_t - \mathbb{E}_t^D x_t$, and our six end-user flows. According to the model, realized foreign exchange returns reflect the revision in dealer's quotes driven by new information concerning fundamentals. This information arrives in the form of public news, macro announcements and inter-dealer order flow, but not the end-user order flows of individual dealers such as Citibank: Any information concerning fundamentals contained in the end-user flows received by individual banks affects the FX price quoted by dealers only once it is inferred from the inter-dealer order flows observed by all dealers. In Evans and Lyons (2006) we study the relationship between end-user flows and market-wide inter-dealer order flow (i.e., the counterpart to $x_t - \mathbb{E}_t^D x_t$). This analysis shows that individual coefficients have no structural interpretation in terms of measuring the price-impact of different end-user orders, they simply map variations in end-user flows into an estimate of the information flow being used by dealers across the market. This interpretation also accounts for the pattern of explanatory power: As the horizon lengthens, the idiosyncratic elements in Citibank's end-user flows become relatively less important, with the result that the flows are more precise proxies for the market-wide flow of information driving quote revisions.

To summarize, the results in Table 2 show that end-user flows are contemporaneously linked with changes in spot rates, but the strength of the link is less than that reported elsewhere for inter-dealer order flows. Once one recognizes that Citibank's end-user flows are an imperfect proxy for inter-dealer order flows, our findings are consistent with the theoretical link between exchange rates and order flow implied by the model.

¹⁴We report results using 4 week rates on Euro-dollar and Euro-mark deposits in both panels of the table because 1 week euro-current rates were unavailable over the entire sample period. Re-estimating the regressions in the upper panel with 1 week rates when they are available over the second half of the sample gives very similar results.

4.2 Forecasting Fundamentals

According to Proposition 3, changes in the exchange rate are correlated with order flow because the latter contains information concerning fundamentals. If this is the mechanism responsible for the results reported in Table 2, order flows ought to have forecasting power for future fundamentals. We now examine whether this implication of our model applies to the end-user flows. First we derive the model's implications for forecasting fundamentals with spot rates and order flows. We then examine the forecasting power of spot rates and the end-user flows for future changes in our real-time estimates.

The model's implications for forecasting fundamentals with spot rates follow straightforwardly from Proposition 1. In particular equation (11) can be rewritten as

$$s_t = \mathbb{E}_t^D f_t + \mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta} \right)^i \Delta f_{t+i}. \quad (29)$$

Thus, the log spot rate quoted by dealers differs from dealers' current estimate of fundamentals by the present value of future changes in fundamentals. One implication of (29) is that the gap between the current spot rate and estimated fundamentals, $s_t - \mathbb{E}_t^D f_t$, should have forecasting power for future changes in fundamentals. This can be formally shown by considering the projection:

$$\Delta f_{t+\tau} = \beta_s (s_t - \mathbb{E}_t^D f_t) + \varepsilon_{t+\tau}, \quad (30)$$

where

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

and $\varepsilon_{t+\tau}$ is the projection error that is uncorrelated with $s_t - \mathbb{E}_t^D f_t$. The projection coefficient β_s provides a measure of the forecasting power of $s_t - \mathbb{E}_t^D f_t$ for the change in fundamentals τ periods ahead.

Now we turn to the forecasting power of order flow. According to Proposition 3, order flow is driven in part by differences between dealers' forecasts and household forecasts concerning future fundamentals. Consequently, if households have more precise information concerning future fundamentals than dealers, order flows should have *incremental* forecasting power beyond that contained $s_t - \mathbb{E}_t^D f_t$. We formalize this idea in the following proposition.

Proposition 5 *When dealer quotes for the price of foreign currency satisfy (11), and order flow follows (22), changes in future fundamentals are related to spot rates and order flows by*

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

where $\epsilon_{t+\tau}$ is the projection error. β_s is the projection coefficient identified in (30) and β_x is equal to

$$\frac{\mathbb{C}\mathbb{V}(o_t, \Delta f_{t+\tau})}{\mathbb{V}(x_t - \mathbb{E}_t^D x_t)} + \frac{\phi \pi \kappa \mathbb{V}(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1}) (A^{\tau-1})' C' \iota_2'}{\mathbb{V}(x_t - \mathbb{E}_t^D x_t)} + \frac{\hat{\phi} \pi \hat{\kappa} \mathbb{V}(\nabla \mathbb{E}_t^H \mathbf{y}_{t+1}) (A^{\tau-1})' C' \iota_2'}{\mathbb{V}(x_t - \mathbb{E}_t^D x_t)}.$$

The intuition behind Proposition 5 is straightforward. Recall from (29) that $s_t - \mathbb{E}_t^D f_t$ is equal to the present value of future changes in fundamentals. The first term in (31) is therefore a function of dealers' information at the start of period t , Ω_t^D . Period- t order flow will have incremental forecasting power of future changes in fundamentals, beyond $s_t - \mathbb{E}_t^D f_t$, when it conveys information about $\Delta f_{t+\tau}$ that is not already known to dealers (i.e. in Ω_t^D). The expression for β_x shows that this will happen when: (i) the distribution of wealth and dealer bond holdings affect order flow and have forecasting power for fundamentals, and (ii) when there is dispersed information concerning future fundamentals and information aggregation occurs via period- t trading. Proposition 4 showed that order flow would be correlated with the depreciation rate under these same conditions. Thus, if our theoretical rationale for the results in Table 2 holds true, we should also find that order flow has incremental forecasting power for future changes in fundamentals.

To assess the empirical evidence on this prediction, we consider forecasting regressions of the form:

$$\Delta^\tau \varkappa_{t+\tau} = a_1 \Delta^k \varkappa_t + a_2 \Delta^k s_t + \sum_{n=1}^6 \beta_j x_{j,t}^k + \eta_{t+\tau}, \quad (32)$$

where $\Delta^\tau \varkappa_{t+\tau}$ denotes the τ -week change in the real-time estimate of variable \varkappa ending at week $t + \tau$, $\Delta^k s_t$ is the rate of depreciation between weeks $t - k$ and t , and $x_{j,t}^k$ is the order flow from segment j in weeks $t - k$ to t . The first two terms on the right hand side are known to dealers at the start of week t and are used to proxy for $s_t - \mathbb{E}_t^D f_t$ in equation (31). Estimates of the β_j coefficients will reveal whether our end-user flows have incremental forecasting power for future fundamentals.

Table 3 presents the results from estimating (32) in weekly data with horizons τ ranging from one month to two quarters. We report results where k is set equal to τ , but our findings are not

sensitive to the number of cumulation weeks k . There are a total of 284 weekly observations in our sample period, so there are 11 non-overlapping observations on the dependent variable at our longest forecasting horizon (e.g. 2 quarters). In each cell of the table we report the R^2 statistic as a measure of forecasting power and the significance level of a Wald test for the joint significance of the forecasting variables. These test statistics are corrected for conditional heteroskedasticity and the moving-average error structure induced by the forecast overlap using the Newey-West estimator.

The results in Table 3 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. Consider, for example, the case of US GDP. 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 GDP two months out using both past real-time estimates of GDP 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.¹⁵

Our findings in Table 3 are robust to the inclusion of other variables as proxies for $s_t - \mathbb{E}_t^D f_t$. In particular, we have estimated versions of (32) that include multiple lags of $\Delta^k \varkappa_t$ and $\Delta^k s_t$ as well as the term spread, default spread and the commercial paper spread.¹⁶ We found that the term spread predicts US GDP and M1, and German prices and M1, while the default and commercial paper spreads predict US GDP. However, the marginal forecasting contribution of these variables is small. Moreover, in all cases, the forecasting contribution of the six flow segments remains highly significant at one and two-quarter horizons. These findings indicate that the results in Table 3 are indeed robust to the inclusion of different variables proxying for $s_t - \mathbb{E}_t^D f_t$.

Although the longest horizon we consider in Table 3 is short compared to the span of our data, our asymptotic inferences concerning forecasting power over 1 and 2 quarters may not be entirely

¹⁵Our theoretical model indicates that US and German consumption are components of fundamentals. Unfortunately, we were not able to compute real-time estimates for both consumption series because the sequence of data releases for German consumption are unavailable. We did compute real-time estimates of US consumption and found that the forecasting power of order flows is similar to that we report for US GDP.

¹⁶The term spread is the difference between the 3-month and 5-year yields on US bonds. We compute the default spread as the difference between Moody's AAA corporate bond yield and Moody's BAA corporate bond yield. The commercial paper spread is the difference between the 3-month commercial paper rate and the 3-month T-Bill rate. Before September 1997 we use the 3-month commercial paper rate, thereafter the 3-month rate for non-financial corporations. We obtained the term structure data from CRSP, and the other interest rates from the FRED database at the St Louis Fed.

Table 3: Forecasting Fundamentals

Forecasting Variables	US GDP				German GDP			
	1M	2M	1 Q	2Q	1M	2M	1Q	2Q
GDP	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)
GDP and Spot	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)
	US Prices				German Prices			
Prices	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)
Prices and Spot	0.007 (0.505)	0.028 (0.352)	0.015 (0.636)	0.06 (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)
	US Money				German Money			
Money	0.071 (0.009)	0.219 (0.000)	0.253 (0.000)	0.329 (0.000)	0.05 (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 and Spot	0.086 (0.002)	0.22 (0.000)	0.267 (0.000)	0.333 (0.000)	0.05 (0.075)	0.13 (0.004)	0.129 (0.040)	0.08 (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)	0.54 (0.000)	0.074 (0.244)	0.175 (0.020)	0.284 (0.001)	0.624 (0.000)

The table reports the R^2 statistic from the forecasting regression (32) for the real-time estimate of 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 statistics testing the null hypothesis of no predictability (corrected for heteroskedasticity and the forecast horizon overlap) are reported in parentheses. The forecasting horizons listed at the head of each column are 1 month ($\tau = 4$), 2 months ($\tau = 8$), 1 quarter ($\tau = 13$) and 2 quarters ($\tau = 26$).

reliable.¹⁷ To insure that our forecasting findings are robust, we supplemented our analysis at these two horizons with the following Monte Carlo experiment: First we estimated an AR(4) process for weekly change in the real-time estimate, $\Delta \varkappa_t$ and a fourth-order VAR for the weekly change in log spot rate, Δs_t , and the six flow segments, $x_{j,t}$. Next, we generated a pseudo data series spanning 284 weeks for $\Delta \varkappa_t$ by combining a bootstrap sample from the $\Delta \varkappa_t$ residuals with estimates of AR(4) process. Pseudo data series for Δs_t and $x_{j,t}$ are similarly generated by bootstrap sampling from the VAR residuals and estimates. Notice that under this data generation process, realizations of $\Delta \varkappa_t$ are independent from the other variables. We then used the pseudo data to estimate equation (32) at the 1 quarter ($\tau = 13$) and 2 quarter ($\tau = 26$) horizons. This process was repeated 5000 times to construct a bootstrap distribution for the regression estimates under the null hypothesis that both spot rates and order flows have no forecasting power for the real-time estimates of fundamentals.

The results of our Monte Carlo experiment are shown in Table 4. To conserve space we only report the results from regressions that include the real-time estimates, spot rates and the end-user flows. The upper entry in each cell is the OLS coefficient on the variable listed at the head of each column estimated from our data. The lower entry is the percentage contribution of the variable to the variance of the future change in the fundamental, again estimated from our data.¹⁸ The variance contribution of all six end-user flows is shown in the right hand column. We compare these estimated variance contributions to the Monte Carlo distribution of the contributions generate under the null of no forecastability, and denote by “*”, “**” and “***” cases where the variance

¹⁷Estimates of long-horizon forecasting regressions like (32) are susceptible to two well-known econometric problems. First, the coefficient estimates may suffer from finite sample bias when the independent variables are predetermined but not exogenous. Second, the asymptotic distribution of the estimates provides a poor approximation to the true distribution when the forecasting horizon is long relative to the span of the sample. Finite-sample bias in the estimates of β_j is not a prime concern here because our six flow segments display little or no autocorrelation and are uncorrelated with past changes in the real-time estimates. There should also be less of a size distortion in the asymptotic distribution than found elsewhere. For example, Mark (1997) considers a case where the data span is less than five times the length of his longest forecasting horizon. Here, we have 11 non-overlapping observations at the 2-quarter horizon.

¹⁸To compute the contribution, we take the fitted values of (32),

$$\Delta^\tau \varkappa_{t+\tau} = \hat{a}_1 \Delta^k \varkappa_t + \hat{a}_2 \Delta^k s_t + \sum_{n=1}^6 \hat{\beta}_j x_{j,t}^k + \hat{\eta}_{t+\tau}$$

multiply both sides by $\Delta^\tau \varkappa_{t+\tau}$, and take second moments:

$$\begin{aligned} \mathbb{V}(\Delta^\tau \varkappa_{t+\tau}) &= \hat{a}_1 \mathbb{C}\mathbb{V}(\Delta^k \varkappa_t, \Delta^\tau \varkappa_{t+\tau}) + \hat{a}_2 \mathbb{C}\mathbb{V}(\Delta^k s_t, \Delta^\tau \varkappa_{t+\tau}) + \sum_{n=1}^6 \hat{\beta}_j \mathbb{C}\mathbb{V}(x_{j,t}^k, \Delta^\tau \varkappa_{t+\tau}) x_{j,t}^k \\ &\quad + \mathbb{C}\mathbb{V}(\Delta^\tau \varkappa_{t+\tau}, \hat{\eta}_{t+\tau}). \end{aligned}$$

Notice that by least squares, $\mathbb{C}\mathbb{V}(\Delta^\tau \varkappa_{t+\tau}, \hat{\eta}_{t+\tau}) = \mathbb{V}(\hat{\eta}_{t+\tau})$, so we end up with a decomposition for $\mathbb{V}(\Delta^\tau \varkappa_{t+\tau})$. The variance contribution of the spot rate is therefore $\hat{a}_2 \mathbb{C}\mathbb{V}(\Delta^k s_t, \Delta^\tau \varkappa_{t+\tau}) / \mathbb{V}(\Delta^\tau \varkappa_{t+\tau})$ and so on.

Table 4: Contributions to Fundamentals' Forecasts

Forecast	Spot	Fund	Corporate		Traders		Investors		All Flows
			US	non-US	US	non-US	US	non-US	
US GDP									
1Q	-0.013	-0.111	1.476	0.329	-0.235	0.283	-0.099	-0.116	
	0.230	1.689	8.797**	3.579	2.946	0.483	0.684	1.575	18.064**
2Q	-0.109	-0.449	0.709	0.283	0.607	-0.218	-0.662	0.281	
	-2.332***	13.544	6.300	2.763	15.921**	0.458	3.600	1.839	30.882**
Prices									
1Q	0.005	0.048	-0.144	0.040	0.037	-0.068	-0.072	0.020	
	0.987**	0.330	1.874	1.829	2.406	1.955	1.301	1.695	11.060**
2Q	-0.006	0.184	0.141	-0.010	-0.045	0.083	-0.020	0.100	
	-1.247**	4.150	0.664	-0.602	1.027	1.525	0.993	17.236*	20.842
M1									
1Q	0.360	0.425	0.919	-0.957	2.129	-5.184	-9.573	-2.477	
	0.577***	21.427**	0.633	-0.826	-0.045	3.020	14.798**	2.160	19.740**
2Q	-1.039	0.439	1.391	-3.944	5.292	-1.806	-10.537	-0.056	
	1.962***	30.842**	1.259	-2.929	1.657**	0.855	19.764*	0.009	20.615*
German GDP									
1Q	0.095	-0.281	-0.983	-0.712	0.302	-0.997	-0.726	-0.426	
	0.617***	8.814*	0.257	2.108	0.409	2.646	0.368	0.394	6.181*
2Q	-0.042	-0.106	-1.677	0.260	0.024	-0.845	1.402	1.170	
	-0.859	0.837	2.630	-0.730	0.095	3.026	6.511	19.995*	31.527**
Prices									
1Q	-0.192	-0.286	2.315	0.167	-0.068	-3.479	-2.701	1.027	
	-0.307***	6.483*	2.485	0.065	0.052	9.421**	5.384	2.171	19.578**
2Q	-0.531	-0.491	1.764	0.714	0.104	-3.242	-4.703	1.394	
	7.156***	7.04*	2.263	3.142	-0.327	3.287	25.076*	3.355	36.797*
Money									
1Q	0.724	0.396	-3.224	2.408	-0.233	3.210	5.180	-5.215	
	-3.982***	13.88**	0.443	1.928	0.154	1.046	1.749	13.15**	18.469***
2Q	1.143	0.243	5.670	-3.091	0.203	8.813	-4.129	-11.436	
	-6.683***	4.935	0.938	0.349	-0.067	10.483*	1.530	50.605***	63.838***

Notes: The upper entry in each cell is the OLS coefficient in the forecasting equation computed at the 1 quarter ($\tau = 13$) and 2 quarters ($\tau = 26$) horizon. The lower entry is the percentage contribution to the variance of the forecast variable. Estimated contributions falling in the 10, 5, and 1 percent tails of the bootstrap distribution are denoted by *, **, and ***.

contribution lies in 10, 5 and 1 percent tails of the bootstrap distribution.

The results in Table 4 complement our earlier findings in two important ways. First, the Monte Carlo results confirm that our end-user flows have significant incremental forecasting power for fundamentals. Although many of the variance contributions from the individual flow segments do not appear statistically significant when compared against the Monte Carlo distribution, the joint contribution of all six flows are significant at the 5 percent level at either the one or two quarter horizon. Moreover, judged by the estimated size of the variance contributions, the order flows contain information that accounts for an economically meaningful fraction of variance in the variable being forecast. In fact, with the exception of US M1, the order flows account for more of the variance than do spot rates or fundamentals. The second noteworthy feature of Table 4 concerns the forecasting contribution of spot rates. Although the estimates are small, they are highly statistically significant at the two quarter horizon for all six variables. Consistent with the present value equation in (11), changes in the spot rates do appear to contain information about the future course of macro fundamentals. We also note from Table 4 that there is a good deal of heterogeneity in the estimated coefficients and contributions of the individual order flows.¹⁹ Imperfect classification of end-user orders into our six segments probably accounts for some of this heterogeneity. Recall that our flow segments are classified according to trade location rather than the nationality of the end-user. Nevertheless, we do note that the largest and most statistically significant contributions come from US-located trades for US variables, and non-US-located trades for German variables.

The results in Tables 3 and 4 contrast quite sharply from the findings of Froot and Ramadorai (2005). They found no evidence of a long run correlation between real interest rate differentials (their measure of fundamentals) and the transaction flows of institutional investors. One likely reason for this difference is the wider span of end-users generating the order flows in our data. The estimates in Table 4 suggest that transactions from different end-users convey different information.

Our use of the real-time estimates is also important. Recall that the change in real-time estimate comprises an ex ante forecast and an information flow. For example, the change in the real-time

¹⁹It is tempting to interpret the coefficients on the individual flows as measuring the information content of an unexpected order from a particular segment. However, as we note earlier, our six flows are correlated with one-another, so the information content of an unexpected order cannot be measured by a single coefficient.

estimate of GDP over the first quarter of the year can be written as

$$\ln GDP_{Q(2)|w(13)} - \ln GDP_{Q(1)|w(1)} = \mathbb{E} \left[\Delta^Q \ln GDP_{Q(2)} \middle| \Omega_{w(1)} \right] + \left(\ln GDP_{Q(2)|w(13)} - \ln GDP_{Q(2)|w(1)} \right),$$

where $\Delta^Q \ln GDP_{Q(2)} \equiv \ln GDP_{Q(2)} - \ln GDP_{Q(1)}$ and $w(j)$ denotes the first day on week j . Thus the change in the real-time estimate comprises the ex ante forecast of GDP growth in the first quarter, and the flow of information concerning second-quarter GDP over the first quarter. Now according to (29), variations in $s_t - \mathbb{E}_t^D f_t$ reflect changes in $\mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta}\right)^i \Delta f_{t+i}$. So if η is large, as Engel and West (2005) argue, and log GDP is correlated with fundamentals f_t , then variations in the ex ante forecasts, $\mathbb{E} \left[\Delta^Q \ln GDP_{Q(3)} \middle| \Omega_{w(1)} \right]$, should track changes in $s_t - \mathbb{E}_t^D f_t = \mathbb{E}_t^D \sum_{i=1}^{\infty} \left(\frac{\eta}{1+\eta}\right)^i \Delta f_{t+i}$. This is the element in the real-time forecasts picked up by the spot rate and lagged fundamentals. Table 4 showed that the estimated variance contributions from these variables are small yet statistically significant – exactly what we should expect to find if there is little variation in the ex-ante forecasts.

The forecasting power of the order flows for the real-time estimates works through a different mechanism. Recall that our order flows are not public information, so their forecasting power for the change in the real-time estimates cannot come via changes in the ex-ante forecasts. Instead, the order flows must be correlated with the flow of public information concerning the fundamental over the forecast horizon. For the case of GDP, this is the second term in the decomposition above. The only difference between $\ln GDP_{Q(2)|w(13)}$ and $\ln GDP_{Q(2)|w(1)}$ is that the former estimate incorporates the information in public data releases between week 1 and 13. With this perspective, our results in Tables 3 and 4 imply that the end-user flows convey information about future fundamentals that is subsequently revealed by the arrival of public data releases. Clearly, these releases represent information that is incremental to the information embedded in spot rates at the beginning of the forecast period. Our empirical findings therefore provide strong corroboration for Proposition 5.

4.3 Exchange Rate Dynamics and Information Flow

One notable feature of the results in Tables 3 and 4 is that the forecasting power of our end-user flows for fundamentals appears stronger at longer forecasting horizons. We interpret this finding as evidence that some of the information conveyed by the order flows only shows up in public news

releases many months later. In this section we investigate two implications of this interpretation. First, we examine whether end-user flows have forecasting power for changes in the exchange rate. Second, we consider whether the forecasting power of flows is consistent with their ability to forecast the future flow of information concerning exchange rate fundamentals.

To understand how our forecasting results for fundamentals relate to the forecastability of the exchange rate, we return to the model. In particular, we consider the implications of Proposition 1 for the change in the log spot rate.

Proposition 6 *When dealer quotes for the price of foreign currency satisfy (11), the change in the log spot rate between the start of period t and $t + \tau$ is*

$$\Delta^\tau s_{t+\tau} \equiv s_{t+\tau} - s_t = \varphi^\tau (s_t - \mathbb{E}_t^\mathbb{D} f_{t,\tau}) + \frac{1}{1+\eta} \sum_{i=0}^{\infty} \left(\frac{\eta}{1+\eta}\right)^i (\mathbb{E}_{t+\tau}^\mathbb{D} - \mathbb{E}_t^\mathbb{D}) f_{t+\tau+i} \quad (33)$$

where $\varphi^\tau \equiv \left(\frac{1+\eta}{\eta}\right)^\tau - 1 > 0$ and $f_{t,\tau} \equiv \frac{\varphi^\tau + 1}{\varphi^\tau(1+\eta)} \sum_{i=0}^{\tau-1} \left(\frac{\eta}{1+\eta}\right)^i f_{t+i}$.

Equation (33) shows us that the change in spot rate comprises two components. The first term on the right identifies the expected depreciation rate $\mathbb{E}_t^\mathbb{D}[s_{t+\tau} - s_t]$, which is proportional to the gap between the current spot rate and expected “near-term” fundamentals, $f_{t,\tau}$. The second term identifies the impact of new information regarding future fundamentals received by dealers between the start of periods t and $t + \tau$, $(\mathbb{E}_{t+\tau}^\mathbb{D} - \mathbb{E}_t^\mathbb{D})f_{t+\tau+i}$. This will be the only term making a significant contribution to the change in spot rates over short and medium horizons. The reason is that reasonable estimates of the semi-interest elasticity, η , fall between 20 to 60 (Engel and West 2005), so φ^τ will be close to zero until τ becomes very large. Any variation in $s_t - \mathbb{E}_t^\mathbb{D} f_{t,\tau}$ will therefore have little impact on the realized change in spot rates. Consequently, we should expect short- and medium-term changes in spot rates to be mainly driven by the arrival of new information concerning the future course of fundamentals.

The implications of our findings in Tables 3 and 4 for forecasting returns should now be clear. If our end-user flows forecast changes in the real-time estimates of variable \varkappa because they contain information about the future flow of public information concerning \varkappa , the flows should also have forecasting power for future changes in spot rates if \varkappa is correlated with exchange rate fundamentals. In other words, our results in Tables 3 and 4 suggest that end-user flows ought to predict $(\mathbb{E}_{t+\tau}^\mathbb{D} -$

$\mathbb{E}_t^D)f_{t+\tau+i}$ if the macro variables we examined are correlated with fundamentals.

To examine this hypothesis, we estimate the following forecasting regression:

$$\Delta^\tau s_{t+\tau} = a_0 + a_1(r_t - \hat{r}_t) + \sum_{j=1}^6 \beta_j x_{j,t}^\tau + \omega_{t+\tau}, \quad (34)$$

where $r_t - \hat{r}_t$ is the interest differential between one month Eurodollar and Euromark deposits and $x_{j,t}^\tau$ is the order flow from segment j in weeks $t-\tau$ to t . We include the interest differential to control for any variations in expected depreciation (i.e., $\varphi^\tau(s_t - \mathbb{E}_t^D f_{t,\tau})$ in equation 33). The regression errors $\omega_{t+\tau}$ pick up news concerning future fundamentals that is not correlated with the end-user flows.

Table 5: Forecasting Exchange Rate Changes									
Weeks τ	$\hat{r} - r$	Corporate		Traders		Investors		R^2	χ^2 (p-value)
		US	NUS	US	NUS	US	NUS		
1	0.102 (0.409)	0.482 (0.317)	-0.033 (0.136)	0.089 (0.102)	-0.153 (0.198)	-0.346 (0.238)	0.142 (0.141)	0.027	8.056 (0.234)
2	0.147 (0.324)	0.509*** (0.263)	-0.037 (0.104)	0.088 (0.082)	-0.09 (0.148)	-0.449*** (0.188)	0.163** (0.096)	0.074	17.239 (0.008)
3	0.176 (0.305)	0.615*** (0.215)	-0.034 (0.090)	0.095* (0.073)	-0.084 (0.137)	-0.432*** (0.177)	0.145** (0.082)	0.121	24.500 (0.001)
4	0.202 (0.302)	0.544*** (0.177)	-0.042 (0.084)	0.094* (0.068)	-0.097 (0.125)	-0.517*** (0.158)	0.137** (0.072)	0.163	30.738 (<0.001)
Notes: The table reports coefficient and standard errors from regressions of future returns measured over horizons τ of one to four weeks, on an (unreported constant), the current interest differential and order flows cumulated over the last 4 weeks. The left hand column reports Wald statistics for the null that all the coefficients on order flow are zero. Estimates are calculated at the weekly frequency. The standard errors correct for heteroskedasticity and the moving average error process induced by overlapping forecasts (2 - 4 week results). *, **, and *** denote significance at the 10%, 5% and 1% levels.									

Table 5 reports the results of estimating (34) for horizons τ of one to four weeks. Two features of the table are striking. First, many of the β_i coefficients on the end-user flows are highly statistically significant, particularly the US corporate and long-term investor flows. The right hand column

shows Wald statistics for the joint significance of all six flow segments that are highly significant beyond the one week horizon. By contrast, none of the coefficients on the interest differential are statistically significant (although they do have the correct positive sign). The second striking feature concerns the degree of forecastability as measured by the R^2 statistics. The forecasting power rises with the horizon, reaching 16 percent at four weeks. By comparison, the R^2 statistics from Fama-type regressions (where the rate of depreciation is regressed on the interest differential) are generally in the 2-4 percent range. Here all the forecasting power comes from the order flows. If we omit the interest differentials and re-estimate the regressions, the estimated coefficients on the flows and the R^2 statistics are essentially unchanged.

The results in Table 5 point to a remarkably strong within-sample relation between order flows and future exchange rate changes. However, there is a long tradition in the exchange rate literature of considering out-of sample forecasting performance. In Evans and Lyons (2005) we examined the out-of-sample forecasting performance of the six order flows for $\Delta^\tau s_{t+\tau}$ with the restrictions $a_0 = 0$ and $a_1 = 1$. At the four week horizon the out-of-sample forecasts accounted for a highly significant 15.7 percent of the variation in excess returns. This degree of forecastability closely matches the in-sample R^2 statistic in Table 5. So even though the statistics in the table relate to the within-sample relation between order flows and changes in the exchange rate, they are representative of the true out-of-sample forecasting power of order flows.

In view of these results, it is now natural to ask whether the predictive power of order flows for exchange rate changes is consistent with their ability to forecast the future flow of information concerning fundamentals. To address this question, we need to take a stand on the relation between true fundamentals, f_t , and our real-time estimates. We consider 6 different measures based on the variables for which we have real-time estimates, f_t^M : (i) the difference between log GDP in the US and German, $y - \hat{y}$, (ii) the US-German log price ratio, $p - \hat{p}$, (iii) the US-German log M1 ratio, $m - \hat{m}$, (iv) the US log M1 to GDP ratio, $m - y$, (v) the German log M1 to GDP ratio, $\hat{m} - \hat{y}$, and (vi) the log M1-GDP differential between the US and Germany, $(m - y) - (\hat{m} - \hat{y})$. For each measure of fundamentals, f_t^M , we first calculate the projection of the quarterly change in f_t^M on the six order flows, $Proj(\Delta^Q f_{t+Q}^M | \{x_{j,t}\})$, as the fitted value from the regression

$$\Delta^Q f_{t+Q}^M \equiv f_{t+13}^M - f_t^M = \mu_0 + \sum_{j=1}^6 \mu_j x_{j,t}^A + \epsilon_{t+13}.$$

We then estimate

$$\Delta^\tau s_{t+\tau} = b_0 + b_1(r_t - \hat{r}_t) + b_2 Proj(\Delta^Q f_{t+Q}^M | \{x_{j,t}\}) + \zeta_{t+\tau}. \quad (35)$$

If the predictive power of order flows for future changes in the exchange rate are due to their ability to forecast the future flow of information concerning measured fundamentals, i.e. $\mathbb{E}_{t+\tau}^D f_{t+13}^M - \mathbb{E}_t^D f_{t+13}^M$, then the estimates of b_2 in (35) should be positive and significant. Moreover, if our measure of fundamentals, f_t^M , is closely correlated with actual fundamentals, f_t , none of the individual order flows $x_{j,t}$, should have incremental predictive power for $\Delta^\tau s_{t+\tau}$ beyond their role in the projection $Proj(\Delta^Q f_{t+Q}^M | \{x_{j,t}\})$.

Table 6 reports the estimates of the forecasting regression (35) for horizons τ of one to four weeks, using the projections of the six different fundamentals measures. Once again, we find the results rather striking. First, the coefficient estimates display a similar pattern across all four forecast horizons. The coefficients on the projections involving the log GDP and price ratios are small and statistically insignificant. By contrast, the coefficients on projections for the log M1 ratios, M1 to income ratios, and the M1-GDP differentials are all highly significant. This constitutes direct empirical evidence that the end-user flows are conveying information about the future course of fundamentals, and it is this information that gives flows their forecasting power for future changes in spot rates. The second noteworthy feature concerns the R^2 statistics. A comparison of the R^2 statistics in Table 5 with the statistics in the lower three rows of each panel in Table 6 shows that the forecasting power of the projections is almost as high as that of the underlying order flows. For example, at the four week horizon the R^2 statistic from (35) using the projection of the quarterly change in $(m - y) - (\hat{m} - \hat{y})$ is 13.3 percent, while the R^2 from estimating (34) is 16.3 percent. The use of the projection places restrictions on the way that the six flows enter (35), but these restrictions do little to impair the forecasting ability of flows for future exchange rates. The right hand column of Table 6 provides more formal evidence on this idea. Here we report LM statistics for the null hypothesis that the residuals from (35) are unrelated to the six flows. If order flows have forecasting power for exchange rates for reasons that are unrelated to the role in conveying information about fundamentals, or the fundamentals measures used in the projections are only weakly correlated with true fundamentals, we should find that order flows have

some residual forecast power, so the null ought to be rejected. However, as the table shows, we fail to reject the null in all the cases where the projection coefficients appear significant.

Table 6: Information Flow Regressions									
Horizon	$\hat{r} - r$	$y - \hat{y}$	$p - \hat{p}$	$m - \hat{m}$	$m - y$	$\hat{m} - \hat{y}$	$(m - y) - (\hat{m} - \hat{y})$	R^2	LM (p-value)
1 week	-0.229	-0.229						<0.001	7.859
	(0.369)	(0.949)							(0.249)
	-0.194		-0.290					0.001	8.525
	(0.367)		(0.507)						(0.202)
	0.161			0.589**				0.023	1.115
	(0.387)			(0.218)					(0.981)
	0.04				0.436**	-0.700**		0.025	0.783
	(0.398)				(0.280)	(0.281)			(0.993)
2 weeks	0.110						0.585**	0.023	1.158
	(0.381)						(0.219)		(0.979)
	-0.215	0.013						<0.001	21.882
	(0.307)	(0.704)							(0.001)
	-0.208		-0.129					0.001	23.977
	(0.312)		(0.420)						(0.001)
	0.199			0.656**				0.060	na
	(0.317)			(0.185)					
3 weeks	0.055				0.467*	-0.771**		0.063	3.741
	(0.317)				(0.227)	(0.231)			(0.712)
	0.136						0.639**	0.059	4.787
	(0.310)						(0.184)		(0.571)
	-0.208	-0.061						<0.001	32.498
	(0.308)	(0.659)							(<0.001)
	-0.198		-0.093					<0.001	39.596
	(0.317)		(0.399)						(<0.001)
	0.245			0.708**				0.104	4.475
	(0.310)			(0.169)					(0.613)
	0.101				0.524**	-0.825**		0.109	4.364
	(0.301)				(0.212)	(0.199)			(0.628)
	0.182						0.694**	0.102	5.926
	(0.302)						(0.172)		(0.432)

Table 6: Information Flow Regressions (cont.)

Horizon	$\hat{r} - r$	$y - \hat{y}$	$p - \hat{p}$	$m - \hat{m}$	$m - y$	$\hat{m} - \hat{y}$	$(m - y) - (\hat{m} - \hat{y})$	R^2	LM (p-value)
4 weeks	-0.214 (0.315)	-0.094 (0.651)						<0.001	52.375 (<0.001)
	-0.200 (0.327)		-0.106 (0.383)					<0.001	53.402 (<0.001)
	0.248 (0.316)			0.709** (0.156)				0.135	8.033 (0.236)
	0.122 (0.302)				0.564** (0.193)	-0.799** (0.186)		0.138	9.129 (0.166)
	0.186 (0.307)						0.697** (0.162)	0.133	10.109 (0.120)

Notes: The table reports coefficients and asymptotic standard errors from regressions of future returns measured over horizons of one to four weeks on the current interest differential, and the projection of the future quarterly change macro fundamentals on current order flows from the six user-user segments. Fundamentals are listed at the head of each column. The left hand column report LM statistics for the null that the regression residuals are unrelated to order flows. (The LM statistic could not be computed for the case labelled “na” because the projection was perfectly correlated with one or more of the order flows.) Standard errors correct for heteroskedasticity and the moving average error process induced by overlapping forecasts (2 - 4 week results). *, **, and *** denote significance at the 10%, 5% and 1% levels.

5 Conclusion

The aim of this paper has been to analyze the links between macro fundamentals, transaction flows and exchange rate dynamics. First, we presented a micro-founded general equilibrium model that provides the theoretical foundation for understanding how dispersed information concerning macro fundamentals is conveyed to spot rates via transaction flows. We then examine the empirical implications of the model. We found that transaction flows have significant forecasting power for macro fundamentals – incremental forecasting power beyond that contained in exchange rates and other variables. We also showed that proprietary transaction flows have significant forecasting power for future exchange returns and that this forecasting power reflects their ability to predict how “the market” will react to the flow of information concerning macro fundamentals. In sum, we find strong support for the idea that exchange rates vary as “the market” assimilates dispersed information regarding macro fundamentals from transaction flows.

Let us conclude with some perspective. Our results provide a qualitatively different view of

why macroeconomic variables perform so poorly in accounting for exchange rates at horizons of one year or less. This view is different from both the traditional macro and the emerging “micro” perspectives. Unlike the macro perspective, we do not view all new information concerning macro fundamentals as being immediately embedded into the exchange rate. Much information about macro fundamentals is dispersed and it takes time for “the market” to fully assimilate its implications for the spot exchange rate. It is this assimilation process that accounts for (much of) the disturbances in exchange rate equations. Our approach also differs from the extant micro perspective because models offered thus far (e.g., Evans and Lyons 2002a,b) have interpreted the information conveyed by transaction flows as orthogonal to macro fundamentals. This information is viewed, instead, as relating to the other driver within the broader asset pricing literature, termed stochastic discount factors, expected returns or portfolio balance effects. Most readers of this micro literature have adopted the same view: transaction flow effects on exchange rates are about pricing errors, not about fundamentals. Our findings, by contrast, suggest that transaction flows are central to the process by which expectations of future macro variables are impounded into exchange rates.

References

- Bacchetta, P., and E. van Wincoop (2006), Can information dispersion explain the exchange rate disconnect puzzle? *American Economic Review*, 93, pp. 552-576.
- Campbell, J. Y., and L. Viceira (2002), *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*, Clarendon Lectures in Economics, Oxford University Press.
- Cheung, Yin-Wong, M. Chinn, and A. Pascual (2005), Empirical exchange rate models of the nineties: Are any fit to survive? *Journal of International Money and Finance*, 24, pp. 1150-1175.
- Croushore, D. and T. Stark (2001), A real-time data set for macroeconomists, *Journal of Econometrics* 105, pp. 111-130.
- Engel, C., and K. West (2005), Exchange rates and fundamentals, *Journal of Political Economy*, 113: 485-517.
- Evans, M. (2005), Where are we now? Real-time estimates of the macro economy, *The International Journal of Central Banking*, 1, pp. 127-175.
- Evans, M., and R. Lyons (2002a), Order flow and exchange rate dynamics, *Journal of Political Economy*, 110, pp. 170-180.
- Evans, M., and R. Lyons (2002b), Informational integration and FX trading, *Journal of International Money and Finance*, 21, pp. 807-831.
- Evans, M., and R. Lyons (2004), A New Micro Model of Exchange Rate Dynamics, NBER Working Paper 10379, March.
- Evans, M., and R. Lyons (2005), Meese-Rogoff Redux: Micro-based Exchange Rate Forecasting, *American Economic Review Papers and Proceedings*, pp. 405-414.
- Evans, M., and R. Lyons (2006), Understanding Order Flow, *International Journal of Finance and Economics*, 11, pp. 1-24.
- Faust, J., J. Rogers, and J. Wright (2003), Exchange rate forecasting: The errors we've really made, *Journal of International Economics*, 60, pp. 35-59.

- 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 (2005), Currency returns, institutional investor flows, and exchange rate fundamentals, *Journal of Finance* 60, pp. 1535-1566.
- Hayek, F. (1945), The use of knowledge in society, *American Economic Review*, 35: 519-530.
- Lyons, R. (1997), A Simultaneous Trade Model of the Foreign Exchange Hot Potato, *Journal of International Economics*, 42, pp. 275-298.
- Mark, N. (1997), Exchange Rates and Fundamentals: Evidence on Long-horizon Predictability, *American Economic Review*, 85 pp. 201-218.
- Meese, R., and K. Rogoff (1983), Empirical exchange rate models of the seventies, *Journal of International Economics*, 14, pp. 3-24.
- Osler, C. (2003), Currency Orders and Exchange-Rate Dynamics: Explaining the Success of Technical Analysis, *Journal of Finance* 58, pp. 1791-1819

A Appendix

This Appendix has three sections. Appendix A.1 provides a detailed description of the model outlined in Section 1 of the paper. Appendix A.2 derives the results presented in Propositions 1 - 6. Finally, Appendix A.3 contains a brief description of how the real-time estimates are computed.

A.1 Model Details

A.1.1 Households

Under assumptions A1 and A2 in the text, the consumption and portfolio decisions facing the representative US household can be written as the following dynamic programming problem:

$$\begin{aligned} \mathbb{J}(W_{\text{H},t}) &= \max_{\alpha_t, \alpha_t^\Lambda, \alpha_t^M, C_{\text{H},t}} \left\{ \frac{1}{1-\gamma} C_{\text{H},t}^{1-\gamma} + \frac{\chi}{1-\gamma} (\alpha_t^M W_{\text{H},t}/P_t)^{1-\gamma} + \delta \mathbb{E}_t^{\text{H}} \mathbb{J}(W_{\text{H},t+1}) \right\}, \\ \text{s.t.} \quad W_{\text{H},t+1} &= R_t (ER_{\text{H},t+1}^M W_{\text{H},t} - P_t C_{\text{H},t}), \\ ER_{\text{H},t+1}^M &= 1 + \left(\frac{S_{t+1} \hat{R}_t}{S_t R_t} - 1 \right) \alpha_t + \left(\frac{R_{t+1}^Q}{R_t} - 1 \right) \alpha_t^\Lambda - \left(\frac{R_t - 1}{R_t} \right) \alpha_t^M \end{aligned}$$

where $R_{t+1}^Q \equiv (Q_{t+1} + D_{t+1})/Q_t$ is the return on US equity and

$$W_{\text{H},t} = R_{t-1} B_{\text{H},t-1} + S_t \hat{R}_{t-1} \hat{B}_{\text{H},t-1} + R_t A_{\text{H},t-1} + M_{\text{H},t-1}$$

is the dollar value of wealth at the beginning of period t . $ER_{\text{H},t+1}^M$ is the excess return on wealth between periods t and $t+1$ that depends on the share of wealth held in euro bonds $\alpha_t \equiv S_t \hat{B}_{\text{H},t}/(\hat{R}_t W_{\text{H},t})$, US equity $\alpha_t^\Lambda \equiv Q_t A_{\text{H},t}/W_{\text{H},t}$, and real balances $\alpha_t^M \equiv M_{\text{H},t}/W_{\text{H},t}$. Solving this problem gives the following first-order conditions:

$$C_{\text{H},t} : \mathbb{E}_t^{\text{H}} \left[\delta \left(\frac{C_{\text{H},t+1}}{C_{\text{H},t}} \right)^{-\gamma} \frac{P_t}{P_{t+1}} R_t \right] = 1, \quad (\text{A1a})$$

$$\alpha_t^M : \left(\frac{M_{\text{H},t}}{P_t C_{\text{H},t}} \right)^{-\gamma} = \frac{1}{\chi} \frac{R_t - 1}{R_t}, \quad (\text{A1b})$$

$$\alpha_t^\Lambda : \mathbb{E}_t^{\text{H}} \left[\delta \left(\frac{C_{\text{H},t+1}}{C_{\text{H},t}} \right)^{-\gamma} \frac{R_{t+1}^Q}{R_t} \right] = 1, \quad (\text{A1c})$$

$$\alpha_t : \mathbb{E}_t^{\text{H}} \left[\delta \left(\frac{C_{\text{H},t+1}}{C_{\text{H},t}} \right)^{-\gamma} \frac{S_{t+1} P_t}{S_t P_{t+1}} \hat{R}_t \right] = 1. \quad (\text{A1d})$$

To characterize optimal household decisions, we work with log normal approximations to the first-order conditions and a log linearization of the budget constraint. We first combine the identity $\alpha_t^M \equiv M_{H,t}/W_{H,t}$ with the first-order condition for real balances and the definition of $ER_{H,t+1}^M$. The budget constraint can then be rewritten as:

$$\frac{W_{H,t+1}}{W_{H,t}} = R_t \left(ER_{H,t+1} - (1 + \Gamma(R_t)) \frac{P_t C_{H,t}}{W_{H,t}} \right),$$

where $\Gamma(R) \equiv \chi^{1/\gamma} \left(\frac{R-1}{R} \right)^{1-\frac{1}{\gamma}}$ and

$$ER_{H,t+1} \equiv 1 + \left(\frac{S_{t+1} \hat{R}_t}{S_t R_t} - 1 \right) \alpha_t + \left(\frac{R_{t+1}^Q}{R_t} - 1 \right) \alpha_t^A.$$

Notice that the coefficient on the consumption-wealth ratio includes the $\Gamma(R_t)$ 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:

$$\Delta w_{H,t+1} = r_t + k + \frac{1}{\rho} (er_{H,t+1} - \varsigma r_t) - \frac{1-\rho}{\rho} (p_t + c_t - w_{H,t}), \quad (\text{A2})$$

where $\rho \equiv 1 - \mu(1 + \Gamma(R))$, $\varsigma \equiv \frac{\mu(\gamma-1)}{\gamma(R-1)R} \Gamma(R)$, $k \equiv \ln \rho + (1 - \frac{1}{\rho}) \ln \mu + \varsigma/\rho$, and μ is the steady state value of the consumption-wealth ratio, $P_t C_{H,t}/W_{H,t}$. Using the definition of $ER_{H,t+1}$ above, we follow Campbell and Viceira (2002) in approximating the log excess return on wealth by:

$$\begin{aligned} er_{H,t+1} &= \alpha_t^A (r_{t+1}^Q - r_t) + \alpha_t (\Delta s_{t+1} + \hat{r}_t - r_t) + \frac{1}{2} \alpha_t^A (1 - \alpha_t^A) \mathbb{V}_t^H (r_{t+1}^Q) \\ &\quad + \frac{1}{2} \alpha_t (1 - \alpha_t) \mathbb{V}_t^H (\Delta s_{t+1}) - \alpha_t \alpha_t^A \mathbb{C}\mathbb{V}_t^H (r_{t+1}^Q, \Delta s_{t+1}), \end{aligned} \quad (\text{A3})$$

where $\mathbb{V}_t^H(\cdot)$ and $\mathbb{C}\mathbb{V}_t^H(\cdot, \cdot)$ denote the variance and covariance conditioned on information Ω_t^H . This second-order approximation holds exactly in the continuous-time limit when the spot exchange rate and the price of other assets follow diffusion processes.

We can now use (A2), (A3) and the log linearized first-order conditions to characterize the optimal choice of consumption, real balances and the portfolio shares. Combining the log linearized

versions of (A1c) and (A1d) with (A2) and (A3) we obtain:

$$\begin{bmatrix} \alpha_t \\ \alpha_t^Q \end{bmatrix} = \frac{\rho}{\gamma} (\Xi_t^H)^{-1} \begin{bmatrix} \mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t + \frac{1}{2} \mathbb{V}_t^H(s_{t+1}) - \psi_{H,t}^s \\ \mathbb{E}_t^H r_{t+1}^Q - r_t + \frac{1}{2} \mathbb{V}_t^H(r_{t+1}^Q) - \psi_{H,t}^Q \end{bmatrix}, \quad (\text{A4})$$

where

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

for $v = \{s, r^Q\}$. The matrix Ξ_t^H is the conditional covariance of the vector $(\Delta s_{t+1}, r_{t+1}^Q)'$. $\mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t - \psi_{H,t}^s$ and $\mathbb{E}_t^H r_{t+1}^Q - r_t - \psi_{H,t}^Q$ are the risk-adjusted expected excess returns on euro bonds and US equities. The variance terms arise because we are working with log excess returns. $\psi_{H,t}^v$ identifies the hedging factor associated with euro bonds ($v = s$) and US equities ($v = r^Q$). All that now remains is to characterize the demand for real balances and the consumption-wealth ratio. The former is found by log linearizing (A1b):

$$m_{H,t} - p_t = \varpi + c_{H,t} - \eta r_t, \quad (\text{A5})$$

where $\varpi \equiv \frac{1}{\gamma} \ln \chi + rR\eta$ and $\eta \equiv 1/\gamma(R-1) > 0$. An approximation to the log consumption wealth-ratio is found by combining (A2) with the linearized version of (A1a):

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

We can characterize the behavior of the representative European household in a similar way. Specifically, the linearized budget constraint is:

$$\Delta \hat{w}_{\widehat{H},t+1} \cong \hat{r}_t - \Delta \hat{p}_{t+1} + k + \frac{1}{\rho} (er_{\widehat{H},t+1} - \varsigma \hat{r}_t) - \frac{1-\rho}{\rho} (\hat{p}_t + \hat{c}_{\widehat{H},t} - \hat{w}_{\widehat{H},t}), \quad (\text{A6})$$

where the log excess return is approximated by:

$$\begin{aligned} er_{\widehat{H},t+1} &\cong \hat{\alpha}_t^A (r_{t+1}^{\widehat{Q}} - \hat{r}_t) + \tilde{\alpha}_t (i_t - \Delta s_{t+1} - i_t^*) + \frac{1}{2} \hat{\alpha}_t^A (1 - \hat{\alpha}_t^A) \mathbb{V}_t^{\widehat{H}}(r_{t+1}^{\widehat{Q}}) \\ &\quad + \frac{1}{2} \tilde{\alpha}_t (1 - \tilde{\alpha}_t) \mathbb{V}_t^{\widehat{H}}(\Delta s_{t+1}) + \alpha_t^A \tilde{\alpha}_t \mathbb{C} \mathbb{V}_t^{\widehat{H}}(r_{t+1}^{\widehat{Q}}, \Delta s_{t+1}). \end{aligned} \quad (\text{A7})$$

with $\tilde{\alpha}_t \equiv 1 - \hat{\alpha}_t^A - \hat{\alpha}_t - \hat{\alpha}_t^M$ as the share in dollar bonds in household wealth, and $r_{t+1}^{\widehat{Q}}$ denoting

the log return on European equity. The optimal portfolio shares are:

$$\begin{bmatrix} \tilde{\alpha}_t \\ \hat{\alpha}_t^Q \end{bmatrix} = \frac{\rho}{\gamma} \left(\hat{\Xi}_t^{\hat{\Pi}} \right)^{-1} \begin{bmatrix} r_t - \mathbb{E}_t^{\hat{\Pi}} \Delta s_{t+1} - \hat{r}_t + \frac{1}{2} \mathbb{V}_t^{\hat{\Pi}} (\Delta s_{t+1}) - \psi_{\hat{\Pi},t}^{-s} \\ \mathbb{E}_t^{\hat{\Pi}} r_{t+1}^Q - \hat{r}_t + \frac{1}{2} \mathbb{V}_t^{\hat{\Pi}} (r_{t+1}^Q) - \psi_{\hat{\Pi},t}^Q \end{bmatrix} \quad (\text{A8})$$

where

$$\psi_{\hat{\Pi},t}^{\omega} = \gamma \mathbb{C}\mathbb{V}_t^{\hat{\Pi}} (\hat{c}_{\hat{\Pi},t+1} + \hat{p}_t - \hat{w}_{\hat{\Pi},t+1}, v_{t+1}) + (1 - \gamma) \mathbb{C}\mathbb{V}_t^{\hat{\Pi}} (\Delta \hat{p}_{t+1}, v_{t+1}),$$

for $v = \{-s, r^Q\}$ and $\hat{\Xi}_t^{\hat{\Pi}}$ is the conditional covariance matrix for the vector $(-\Delta s_{t+1}, r_{t+1}^Q)'$. The demand for log real balances is given by:

$$\hat{m}_{\hat{\Pi},t} - \hat{p}_t = \varpi + \hat{c}_{\hat{\Pi},t} - \eta \hat{r}_t, \quad (\text{A9})$$

and the log consumption wealth ratio by:

$$\hat{c}_{\hat{\Pi},t} + \hat{p}_t - \hat{w}_{\hat{\Pi},t} = \frac{\rho k}{1-\rho} + \left(1 - \frac{1}{\gamma}\right) \mathbb{E}_t^{\hat{\Pi}} \sum_{i=0}^{\infty} \rho^{i+1} (\hat{r}_{t+i} - \Delta \hat{p}_{t+1+i}) + \mathbb{E}_t^{\hat{\Pi}} \sum_{i=1}^{\infty} \rho^{i-1} (er_{\hat{\Pi},t+i} - \varsigma \hat{r}_{t+i-1}). \quad (\text{A10})$$

A.1.2 Firms

The pricing problem facing the US firm can be written as the following dynamic programming problem:

$$\begin{aligned} \mathbb{Q}(K_t) &= \max_{P_t^{\text{US}}, \hat{P}_t^{\text{US}}} \mathbb{E}_t^{\text{US}} \{D_t/P_t + \delta \Lambda_{t+1} \mathbb{Q}(K_{t+1})\}, \\ \text{s.t.} \quad \frac{D_t}{P_t} &= \left(\frac{P_t^{\text{US}}}{P_t}\right)^{1-\theta} C_t + \frac{S_t \hat{P}_t}{P_t} \left(\frac{\hat{P}_t^{\text{US}}}{\hat{P}_t}\right)^{1-\theta} \hat{C}_t, \\ K_{t+1} &= (1 - \varrho) K_t + \xi_t K_t^\nu - \left(\frac{P_t^{\text{US}}}{P_t}\right)^{-\theta} C_t - \left(\frac{\hat{P}_t^{\text{US}}}{\hat{P}_t}\right)^{-\theta} \hat{C}_t, \end{aligned}$$

where $\Lambda_{t+1} \equiv \Lambda_{t+1,t}$ is the stochastic discount factor between periods t and $t+1$. Recall that \mathbb{E}_t^{US} denotes expectations conditioned on the information available to US firms at the start of period t , Ω_t^{US} . The first-order conditions for the US firm's problem are

$$\begin{aligned} P_t^{\text{US}} &: 0 = \mathbb{E}_t^{\text{US}} \left\{ \frac{(1-\theta)}{P_t} \left(\frac{P_t^{\text{US}}}{P_t}\right)^{-\theta} C_t + \delta \theta \frac{\Lambda_{t+1} \mathbb{Q}'(K_{t+1})}{P_t} \left(\frac{P_t^{\text{US}}}{P_t}\right)^{-\theta-1} C_t \right\}, \quad \text{and} \\ \hat{P}_t^{\text{US}} &: 0 = \mathbb{E}_t^{\text{US}} \left\{ \frac{(1-\theta)}{\hat{P}_t} \frac{S_t \hat{P}_t}{P_t} \left(\frac{\hat{P}_t^{\text{US}}}{\hat{P}_t}\right)^{-\theta} \hat{C}_t + \delta \theta \frac{\Lambda_{t+1} \mathbb{Q}'(K_{t+1})}{\hat{P}_t} \left(\frac{\hat{P}_t^{\text{US}}}{\hat{P}_t}\right)^{-\theta-1} \hat{C}_t \right\}. \end{aligned}$$

$\mathbb{Q}'(K_t)$ is the marginal value of capital that satisfies the envelope equation $\mathbb{Q}'(K_t) = \delta \mathbb{E}_t^{\text{US}}[\Lambda_{t+1} \mathbb{Q}'(K_{t+1}) R_t^k]$, where $R_t^k \equiv (1 - \varrho) + \nu \xi_t K_t^{\nu-1}$. Simplifying these equations and log-linearizing gives the following expressions for the log prices set by US firms

$$p_t^{\text{US}} = \mathbb{E}_t^{\text{US}} p_t + \mathbf{m}_t^{\text{US}}, \quad \text{and} \quad \hat{p}_t^{\text{US}} = \mathbb{E}_t^{\text{US}} p_t + \hat{\mathbf{m}}_t^{\text{US}} - \mathbb{E}_t^{\text{US}} s_t,$$

where \mathbf{m}_t^{US} and $\hat{\mathbf{m}}_t^{\text{US}}$ are the percentage markups in the price of US goods over the expected log US price-level, $\mathbb{E}_t^{\text{US}} p_t$. Let $\mathbf{n}_t \equiv \lambda_t + \ln \mathbb{Q}'(K_t)$ where $\lambda_t = \ln \Lambda_t$. We can now identify the markups by

$$\begin{aligned} \mathbf{m}_t^{\text{US}} &= \mathbb{E}_t^{\text{US}} \mathbf{n}_{t+1} - \ln \left(\frac{\theta-1}{\theta\delta} \right) + \frac{1}{2} \mathbb{V}_t^{\text{US}} (c_t + \theta p_t + \mathbf{n}_{t+1}) - \frac{1}{2} \mathbb{V}_t^{\text{US}} ((\theta-1) p_t + c_t), \quad \text{and} \\ \hat{\mathbf{m}}_t^{\text{US}} &= \mathbb{E}_t^{\text{US}} \mathbf{n}_{t+1} - \ln \left(\frac{\theta-1}{\theta\delta} \right) + \frac{1}{2} \mathbb{V}_t^{\text{US}} (\hat{c}_t + \theta \hat{p}_t + \mathbf{n}_{t+1}) - \frac{1}{2} \mathbb{V}_t^{\text{US}} ((\theta-1) p_t + \hat{c}_t + \varepsilon_t), \end{aligned}$$

where $\mathbf{n}_t = \ln \delta + \mathbb{E}_t^{\text{US}} \mathbf{n}_{t+1} + r_t^k + \lambda_t + \frac{1}{2} \mathbb{V}_t^{\text{US}} (\mathbf{n}_{t+1})$.

The pricing problem facing the EU firm is analogous and produces the following approximations for the log prices of EU goods:

$$\hat{p}_t^{\text{EU}} = \mathbb{E}_t^{\text{EU}} \hat{p}_t + \hat{\mathbf{m}}_t^{\text{EU}}, \quad \text{and} \quad p_t^{\text{EU}} = \mathbb{E}_t^{\text{EU}} \hat{p}_t + \mathbb{E}_t^{\text{EU}} s_t + \mathbf{m}_t^{\text{EU}},$$

with markups

$$\begin{aligned} \hat{\mathbf{m}}_t^{\text{EU}} &= -\ln \left(\frac{\theta-1}{\theta\delta} \right) + \mathbb{E}_t^{\text{EU}} \hat{\mathbf{n}}_{t+1} + \frac{1}{2} \mathbb{V}_t^{\text{EU}} (\hat{c}_t + \theta \hat{p}_t + \hat{\mathbf{n}}_{t+1}) - \frac{1}{2} \mathbb{V}_t^{\text{EU}} ((\theta-1) p_t + \hat{c}_t) \quad \text{and} \\ \mathbf{m}_t^{\text{EU}} &= -\ln \left(\frac{\theta-1}{\theta\delta} \right) + \mathbb{E}_t^{\text{EU}} \hat{\mathbf{n}}_{t+1} + \frac{1}{2} \mathbb{V}_t^{\text{EU}} (c_t + \theta p_t + \hat{\mathbf{n}}_{t+1}) - \frac{1}{2} \mathbb{V}_t^{\text{EU}} ((\theta-1) p_t + c_t - \varepsilon_t), \end{aligned}$$

where $\hat{\mathbf{n}}_t \equiv \hat{\lambda}_t + \ln \mathbb{Q}'(\hat{K}_t)$ and $\hat{\mathbf{n}}_t = \ln \delta + \mathbb{E}_t^{\text{EU}} \hat{\mathbf{n}}_{t+1} + \hat{r}_t^k + \hat{\lambda}_t + \frac{1}{2} \mathbb{V}_t^{\text{EU}} (\hat{\mathbf{n}}_{t+1})$ with $\hat{r}_t^k \equiv \ln[(1 - \varrho) + \nu \xi_t \hat{K}_t^{\nu-1}]$.

We can now relate the real exchange rate to the pricing decisions of firms. In particular, if we first write the real exchange rate as

$$\varepsilon_t \equiv \frac{S_t \hat{P}_t}{P_t} = \left\{ \frac{P_t^{\text{US}1-\theta} \left(S_t \hat{P}_t^{\text{US}} / P_t^{\text{US}} \right)^{1-\theta} + P_t^{\text{EU}1-\theta} \left(S_t \hat{P}_t^{\text{EU}} / P_t^{\text{EU}} \right)^{1-\theta}}{P_t^{\text{US}1-\theta} + P_t^{\text{EU}1-\theta}} \right\}^{1/(1-\theta)}$$

and then take a log-linear approximation around the symmetric steady state of $\mathcal{E} = 1$, we obtain

$$\begin{aligned}\varepsilon_t &= \frac{1}{2}(s_t + \hat{p}_t^{\text{US}} - p_t^{\text{US}}) + \frac{1}{2}(s_t + \hat{p}_t^{\text{EU}} - p_t^{\text{EU}}), \\ &= \frac{1}{2}(s_t - \mathbb{E}_t^{\text{US}} s_t) + \frac{1}{2}(s_t - \mathbb{E}_t^{\text{EU}} s_t) + \frac{1}{2}(\hat{\mathbf{m}}_t^{\text{US}} - \mathbf{m}_t^{\text{US}}) + \frac{1}{2}(\hat{\mathbf{m}}_t^{\text{EU}} - \mathbf{m}_t^{\text{EU}}).\end{aligned}$$

Hence, the log real exchange rate varies with the expectational errors of firms concerning the current spot rate, and the differential in the price markups between the European and US markets.

A.1.3 Dealers

Dealers make four decisions each period. They choose euro price quotes $S_{d,t}^{\text{I}}$ and $S_{d,t}^{\text{II}}$ at the start of trading rounds I and II. They initiate trades, $T_{d,t}$, against other dealers' quotes in round II trading, and they choose consumption, $C_{d,t}$, after trading is complete. We refer below to this consumption decision as the round III decision. It proves useful to first consider the optimal choices of $T_{d,t}$ and $C_{d,t}$ before examining how $S_{d,t}^{\text{I}}$ and $S_{d,t}^{\text{II}}$ are determined.

Define $W_{d,t}^i = B_{d,t}^i + S_t^i \hat{B}_{d,t}^i$ as the dollar wealth of dealer d at the start of round- i in period t , where S_t^i is the price quoted by other dealers. The trading problem facing dealer d in round II can be written as

$$\begin{aligned}\mathbb{V}(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}}) &= \max_{T_{d,t}} \mathbb{E}[\mathbb{V}(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}}) | \Omega_{d,t}^{\text{II}}] \\ \text{s.t.} \quad W_{d,t}^{\text{III}} &= W_{d,t}^{\text{II}} + (S_{d,t}^{\text{II}} - S_t^{\text{II}}) \mathcal{T}_{d,t}^{\text{II}}, \quad \text{and} \quad \hat{B}_{d,t}^{\text{III}} = \hat{B}_{d,t}^{\text{II}} + T_{d,t} - \mathcal{T}_{d,t}^{\text{II}}.\end{aligned} \tag{A11}$$

where $\mathbb{V}(W_{d,t}^i, \hat{B}_{d,t}^i)$ denotes the value function for the dealer defined over wealth, $W_{d,t}^i$, and euro bond holdings, $\hat{B}_{d,t}^i$. Notice that $\mathcal{T}_{d,t}^{\text{II}} \notin \Omega_{d,t}^{\text{II}}$ so the dealer's choice of trade, $T_{d,t}$, cannot be conditioned on incoming orders from other dealers, $\mathcal{T}_{d,t}^{\text{II}}$, (i.e., interdealer trade takes place simultaneously). In round III dealers choose consumption so solve

$$\begin{aligned}\mathbb{V}(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}}) &= \max_{C_{d,t}} \left\{ \frac{1}{1-\gamma} C_{d,t}^{1-\gamma} + \delta \mathbb{E}[\mathbb{V}(W_{d,t+1}^{\text{I}}, \hat{B}_{d,t+1}^{\text{I}}) | \Omega_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}}] \right\} \\ \text{s.t.} \quad W_{d,t+1}^{\text{I}} &= R_t(W_{d,t}^{\text{III}} - P_t C_{d,t}) + (S_{t+1}^{\text{I}} \hat{R}_t - R_t S_t^{\text{II}}) \hat{B}_{d,t}^{\text{III}}, \quad \text{and} \quad \hat{B}_{d,t+1}^{\text{I}} = \hat{R}_t \hat{B}_{d,t}^{\text{III}}.\end{aligned} \tag{A12}$$

The first-order conditions associated with the problems in (A11) and (A12) are

$$0 = \mathbb{E}[\mathbb{V}_2(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}], \quad \text{and} \quad (\text{A13})$$

$$C_{d,t}^{-\gamma} = \delta R_t P_t \mathbb{E}[\mathbb{V}_1(W_{d,t+1}^{\text{I}}, \hat{B}_{d,t+1}^{\text{I}})|\Omega_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{III}}], \quad (\text{A14})$$

where $\mathbb{V}_i(\cdot, \cdot)$ denotes the i 'th. partial derivative of the dealer's value function.

Next, we consider the quote problems facing the dealer at the start of round I and II. The round I problem can be written as

$$\begin{aligned} \mathbb{V}(W_{d,t}^{\text{I}}, \hat{B}_{d,t}^{\text{I}}) &= \max_{S_{d,t}^{\text{I}}} \mathbb{E}[\mathbb{V}(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})|\Omega_{d,t}^{\text{I}}] \\ \text{s.t. } W_{d,t}^{\text{II}} &= W_{d,t}^{\text{I}} + (S_t^{\text{II}} - S_t^{\text{I}})(\hat{B}_{d,t}^{\text{I}} - \mathcal{T}_{d,t}^{\text{I}}) + (S_{d,t}^{\text{I}} - S_t^{\text{I}})\mathcal{T}_{d,t}^{\text{I}}, \quad \text{and} \quad \hat{B}_{d,t}^{\text{II}} = \hat{B}_{d,t}^{\text{I}} - \mathcal{T}_{d,t}^{\text{I}}, \end{aligned} \quad (\text{A15})$$

and the round II problem as

$$\begin{aligned} \mathbb{V}(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}}) &= \max_{S_{d,t}^{\text{II}}} \mathbb{E}[\mathbb{V}(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}] \\ \text{s.t. } W_{d,t}^{\text{III}} &= W_{d,t}^{\text{II}} + (S_{d,t}^{\text{II}} - S_t^{\text{II}})\mathcal{T}_{d,t}^{\text{II}}, \quad \text{and} \quad \hat{B}_{d,t}^{\text{III}} = \hat{B}_{d,t}^{\text{II}} + \mathcal{T}_{d,t}^{\text{II}} - \mathcal{T}_{d,t}^{\text{I}}. \end{aligned} \quad (\text{A16})$$

Recall that all dealers choose quotes simultaneously, so the choice of $S_{d,t}^i$ cannot be conditioned on the quotes of other dealers, i.e., $S_{n,t}^i$ for $n \neq d$ and $i = \{\text{I, II}\}$. Furthermore, quotes are good for any amount and are available to all households in round I, and all dealers in round II. Consequently, $(S_{d,t}^i - S_t^i)\mathcal{T}_{d,t}^i$ will have a limiting value of $-\infty$ if $S_{d,t}^i$ differs from S_t^i in trading round $i = \{\text{I, II}\}$. We establish below that $\mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})|\Omega_{d,t}^{\text{I}}]$ and $\mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}]$ are positive so dealers must quote a common price in each trading round, i.e., $S_{d,t}^{\text{I}} = S_t^{\text{I}}$ and $S_{d,t}^{\text{II}} = S_t^{\text{II}}$.

Now we turn to the determination of S_t^{I} and S_t^{II} . A dealer will only be willing to quote at the beginning of each trading round if doing so does not reduce his expected utility. In round II trading, the marginal utility associated with incoming orders, $\mathcal{T}_{d,t}^{\text{II}}$, is

$$\mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})(S_{d,t}^{\text{II}} - S_t^{\text{II}}) - \mathbb{V}_2(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}].$$

This term equals zero when $S_{d,t}^{\text{II}} = S_t^{\text{II}}$ and round II trades are chosen optimally satisfying (A13). Thus, incoming orders from other dealers have no effect on the dealers' expected utility at the

margin provided he has the opportunity to initiate trades and quotes a common price to avoid arbitrage.

In round I, the marginal utility associated with incoming orders, $\mathcal{T}_{d,t}^I$, is zero when

$$0 = \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})[(S_{d,t}^{\text{I}} - S_t^{\text{I}}) - (S_t^{\text{II}} - S_t^{\text{I}})] - \mathbb{V}_2(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})|\Omega_{d,t}^{\text{II}}].$$

Applying the no-arbitrage restriction of common round I quotes, and substituting the envelope condition $\mathbb{V}_2(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}}) = \mathbb{E}[\mathbb{V}_2(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}]$ from problem (A11) gives

$$\begin{aligned} 0 &= \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})(S_t^{\text{II}} - S_t^{\text{I}}) + \mathbb{V}_2(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}], \\ &= (S_t^{\text{II}} - S_t^{\text{I}})\mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})|\Omega_{d,t}^{\text{II}}], \end{aligned}$$

where the second line follows from (A13) and the fact that both S_t^{II} and S_t^{I} are a function of common information $\Omega_t^{\text{D}} \subset \Omega_{d,t}^{\text{II}}$. Thus dealers will not be made worse off at the margin by incoming orders during round I trading if the common quote is the same in each round: $S_t^{\text{I}} = S_t^{\text{II}} = S_t$.

Finally, we determine the value of S_t . With $S_{d,t}^i = S_t^i = S_t$ for $i = \{\text{I, II}\}$, it is straightforward to establish that

$$\begin{aligned} \mathbb{V}_1(W_{d,t}^{\text{I}}, \hat{B}_{d,t}^{\text{I}}) &= \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}})|\Omega_{d,t}^{\text{I}}] = \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{I}}], \\ \mathbb{V}_1(W_{d,t}^{\text{II}}, \hat{B}_{d,t}^{\text{II}}) &= \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}], \quad \text{and} \\ \mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}}) &= \delta R_t \mathbb{E}[\mathbb{V}_1(W_{d,t+1}^{\text{I}}, \hat{B}_{d,t+1}^{\text{I}})|\Omega_{d,t}^{\text{II}}, \hat{B}_{d,t}^*]. \end{aligned}$$

Hence, the first order condition in (A12) implies that

$$C_{d,t}^{-\gamma} = \mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})P_t. \tag{A17}$$

Consequently,

$$\mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{I}}] = \mathbb{E}[C_{d,t}^{-\gamma}/P_t|\Omega_{d,t}^{\text{I}}] > 0, \quad \text{and} \quad \mathbb{E}[\mathbb{V}_1(W_{d,t}^{\text{III}}, \hat{B}_{d,t}^{\text{III}})|\Omega_{d,t}^{\text{II}}] = \mathbb{E}[C_{d,t}^{-\gamma}/P_t|\Omega_{d,t}^{\text{II}}] > 0,$$

as noted above. We also have

$$\mathbb{V}_2(W_{d,t}^I, \hat{B}_{d,t}^I) = \mathbb{E} \left[\mathbb{E}[\mathbb{V}_2(W_{d,t}^{III}, \hat{B}_{d,t}^{III}) | \Omega_{d,t}^{II}] \middle| \Omega_{d,t}^I \right] + (S_t^{II} - S_t^I) \mathbb{E}[\mathbb{V}_1(W_{d,t}^{II}, \hat{B}_{d,t}^{II}) | \Omega_{d,t}^I], \quad (\text{A18})$$

and

$$\begin{aligned} \mathbb{V}_2(W_{d,t}^{III}, \hat{B}_{d,t}^{III}) &= \delta \mathbb{E} \left[\mathbb{V}_1(W_{d,t+1}^I, \hat{B}_{d,t+1}^I) (S_{t+1}^I \hat{R}_t - R_t S_t^{II}) \middle| \Omega_{d,t}^{II}, \hat{B}_{d,t}^* \right] \\ &\quad + \delta \hat{R}_t \mathbb{E}[\mathbb{V}_2(W_{d,t+1}^I, \hat{B}_{d,t+1}^I) | \Omega_{d,t}^{II}, \hat{B}_{d,t}^*]. \end{aligned} \quad (\text{A19})$$

When $S_t^{II} = S_t^I$, (A18) and (A13) imply that $\mathbb{V}_2(W_{d,t}^I, \hat{B}_{d,t}^I) = 0$. Using this result and the fact that $\mathbb{V}_1(W_{d,t+1}^I, \hat{B}_{d,t+1}^I) = \mathbb{E}[\mathbb{V}_1(W_{d,t+1}^{III}, \hat{B}_{d,t+1}^{III}) | \Omega_{d,t+1}^I]$, (A19) becomes

$$\mathbb{V}_2(W_{d,t}^{III}, \hat{B}_{d,t}^{III}) = \delta \mathbb{E} \left[\mathbb{V}_1(W_{d,t+1}^{III}, \hat{B}_{d,t+1}^{III}) (S_{t+1} \hat{R}_t - R_t S_t) | \Omega_{d,t}^{II}, \hat{B}_{d,t}^* \right]. \quad (\text{A20})$$

Taking expectations conditioned on Ω_t^D , using (A13) and the law of iterated expectations, we get

$$S_t = \frac{\hat{R}_t}{R_t} \frac{\mathbb{E}[\mathbb{V}_1(W_{d,t+1}^{III}, \hat{B}_{d,t+1}^{III}) S_{t+1} | \Omega_t^D]}{\mathbb{E}[\mathbb{V}_1(W_{d,t+1}^{III}, \hat{B}_{d,t+1}^{III}) | \Omega_t^D]}. \quad (\text{A21})$$

Equation (A21) identifies the price at which dealers are willing to fill incoming orders for euros in rounds I and II based on common information, Ω_t^D , and the trading environment of our model. To gain more perspective on its implications, we take a log-normal approximation to (A21):

$$\mathbb{E}[s_{t+1} - s_t | \Omega_t^D] + \hat{r}_t - r_t = \psi, \quad (\text{A22})$$

where $\psi \equiv -\mathbb{V}_t^D(s_{t+1}) - \mathbb{C}\mathbb{V}_t^D(\ln V_1(W_{d,t+1}^{III}, \hat{B}_{d,t+1}^{III}), s_{t+1})$. This is the form of equation (13) in the text. It says that log spot rate, s_t , implied by the common dealer quotes must be such that the expected log excess return based on Ω_t^D compensates the dealers for filling incoming euro orders from households and other dealers.

Finally, we consider the consumption and trading decisions of each dealer. Combining the first-order condition in (A14) with our results on the marginal utility of wealth gives

$$C_{d,t}^{-\gamma} = \delta R_t \mathbb{E}[C_{d,t+1}^{-\gamma} (P_t/P_{t+1}) | \Omega_{d,t}^{II}, \hat{B}_{d,t}^{III}]. \quad (\text{A23})$$

This is the standard consumption-Euler equation. Notice, however, that dealers can condition their period- t choices on their holdings of euro bonds after round II is complete, i.e., $\hat{B}_{d,t}^{\text{II}}$. The optimal choice of round II trade, $T_{d,t}$, is governed by the first-order condition in (A13). Combining this expression with (A20) and (A17) gives

$$0 = \delta \mathbb{E} \left[\left(\frac{C_{d,t+1}}{C_{d,t}} \right)^{-\gamma} \frac{P_t}{P_{t+1}} \left(S_{t+1} \hat{R}_t - R_t S_t \right) \middle| \Omega_{d,t}^{\text{II}} \right]. \quad (\text{A24})$$

This equation takes the form of a standard first-order condition governing the portfolio choice between dollar and euro bonds.

A.1.4 Equilibrium

An equilibrium in this model is described by: (i) the consumption and portfolio decisions of households, (ii) the price-setting decisions of firms, (iii) the interest rate decisions of central banks, (iv) the quote, trade and consumption decisions of dealers, consistent with market clearing in the equity, bonds, money and goods markets. Assumptions A1 and A2 imply that all the equities issued by US and European firms are held by the domestic representative household. Thus market clearing implies that the ex-dividend prices of US and EU equity, Q_t and \hat{Q}_t satisfy

$$\begin{aligned} \text{US equity} & : \quad Q_t + D_t = P_t \mathbb{Q}_t^{\text{US}}, & \text{and} \\ \text{EU equity} & : \quad \hat{Q}_t + \hat{D}_t = \hat{P}_t \mathbb{Q}_t^{\text{EU}}, \end{aligned}$$

where \mathbb{Q}_t^i is the value of the real dividend stream of firm $i = \{\text{US}, \text{EU}\}$ to domestic households under an optimal price-setting policy. Market clearing also implies that the optimal share of equities in households' portfolios satisfy

$$\alpha_t^{\text{Q}} W_{\text{H},t} = Q_t, \quad \text{and} \quad \hat{\alpha}_t^{\text{Q}} W_{\text{H},t} = \hat{Q}_t,$$

because the number of outstanding shares issued by each firm is normalized to one.

Market clearing in the euro bond market requires that the dollar value of aggregate euro orders

received by dealers in round I trading equal aggregate household order flow:

$$\sum_{d=1}^D S_t T_{d,t}^I = x_t.$$

In round II all trading takes place between dealers. Hence, the dollar value of incoming orders for euros received by all dealers must equal the aggregate dealer order flow:

$$\sum_{d=1}^D S_t T_{d,t}^{II} = \sum_{d=1}^D S_t T_{d,t}.$$

At the end of each period, the aggregate holdings of euro bonds by US households, $\hat{B}_{H,t}$, EU households, $\hat{B}_{\hat{H},t}$, and non-households (i.e. dealers and central banks), \hat{B}_t , must sum to zero:

$$\hat{B}_{H,t} + \hat{B}_{\hat{H},t} + \hat{B}_t = 0.$$

In the money markets, central banks accommodate households' demand for currency at a chosen nominal interest rate. These interest rates are set as

$$\begin{aligned} \text{FED} &: r_t = \frac{1}{\eta} \{ \varpi + \mathbb{E}_t^{\text{FED}} [p_t + c_t - m_t] \} = \frac{1}{\eta} \{ \varpi + p_t + \mathbb{E}_t^{\text{FED}} c_t - m_t^* \}, \quad \text{and} \\ \text{ECB} &: \hat{r}_t = \frac{1}{\eta} \{ \varpi + \mathbb{E}_t^{\text{ECB}} [\hat{p}_t + \hat{c}_t - \hat{m}_t] \} = \frac{1}{\eta} \{ \varpi + \hat{p}_t + \mathbb{E}_t^{\text{ECB}} \hat{c}_t - \hat{m}_t^* \}. \end{aligned}$$

Market clearing in goods markets is demand-determined in each national market given the prices chosen by firms. Aggregate US consumption comprises the consumption of US households and dealers: $C_t = C_{H,t} + \sum_{d=1}^D C_{d,t}$ where $C_{H,t} \equiv \int_0^{1/2} C_{h,t} dh$. Aggregate EU consumption comprises the consumption of EU households $\hat{C}_t = \hat{C}_{\hat{H},t} \equiv \int_{1/2}^1 \hat{C}_{h,t} dh$. The implications of price-setting for dividends via their effect on consumer demand are incorporated into the firms' decision-making problems.

A.2 Proofs of Propositions

Proposition 1 Consider the expected demand for money by US and EU households implied by (A5) and (A9) given prices and interest rates, conditioned on dealers' common information, Ω_t^D :

$$\mathbb{E}_t^D m_{H,t} - p_t = \varpi + \mathbb{E}_t^D c_{H,t} - \eta r_t, \quad (\text{A25a})$$

$$\mathbb{E}_t^D \hat{m}_{\hat{H},t} - \hat{p}_t = \varpi + \mathbb{E}_t^D \hat{c}_{\hat{H},t} - \eta \hat{r}_t, \quad (\text{A25b})$$

If Ω_t^D is a subset of the period- t information available to the FED and the ECB, $\mathbb{E}_t^D m_{H,t} = \mathbb{E}_t^D \mathbb{E}_t^{\text{FED}} m_{H,t}$ and $\mathbb{E}_t^D \hat{m}_{\hat{H},t} = \mathbb{E}_t^D \mathbb{E}_t^{\text{ECB}} \hat{m}_{\hat{H},t}$ by iterated expectations. Combining (A25) with these restrictions and the central banks' policy rules gives

$$\mathbb{E}_t^D m_t^* - p_t = \varpi + \mathbb{E}_t^D c_{H,t} - \eta r_t, \quad (\text{A26a})$$

$$\mathbb{E}_t^D \hat{m}_t^* - \hat{p}_t = \varpi + \mathbb{E}_t^D \hat{c}_{\hat{H},t} - \eta \hat{r}_t. \quad (\text{A26b})$$

To derive equation (11), we first use (A26) to substitute for r_t and \hat{r}_t in (A22). This gives

$$s_t = \frac{\eta}{1+\eta} \mathbb{E}_t^D s_{t+1} + \frac{1}{1+\eta} \mathbb{E}_t^D f_t, \quad (\text{A27})$$

where fundamentals, f_t , are defined in (12). Solving this equation forward and applying the law of iterated expectations gives (11). Notice that if $\mathbb{E}_t^D m_{H,t} \neq \mathbb{E}_t^D \mathbb{E}_t^{\text{FED}} m_{H,t}$ and $\mathbb{E}_t^D \hat{m}_{\hat{H},t} \neq \mathbb{E}_t^D \mathbb{E}_t^{\text{ECB}} \hat{m}_{\hat{H},t}$, because central banks have less information than dealers, we still get (A27) from (A25) and (A22), but fundamentals depend on $m_{H,t} - \hat{m}_{\hat{H},t}$ rather than $m_t^* - \hat{m}_t^*$. The present value expression for the log spot rate is therefore robust to different information assumptions regarding dealers and central banks provided we adjust the definition of fundamentals accordingly.

Proposition 2 Let $er_{t+1}^Q \equiv r_{t+1}^Q - r_t + \frac{1}{2} \mathbb{V}_t^H(r_{t+1}^Q) - \psi_{H,t}^Q$ be the risk adjusted log excess return on US equities. We may now rewrite the portfolio allocation equation in (A4) as

$$\alpha_t = \Theta_t^H \left(\mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t + \frac{1}{2} \mathbb{V}_t^H(s_{t+1}) - \psi_{H,t}^s \right) - \Psi_t^H \mathbb{E}_t^H er_{t+1}^Q, \quad (\text{A28})$$

where $\Theta_t^H \equiv \frac{\rho}{\gamma} \mathbb{V}_t^H(r_{t+1}^Q)/|\Xi_t^H|$ and $\Psi_t^H \equiv \frac{\rho}{\gamma} \mathbb{C}\mathbb{V}_t^H(r_{t+1}^Q, s_{t+1})/|\Xi_t^H|$. Households know that dealers quote spot rates in accordance with (11). So the expected excess return on euro bonds can be written as

$$\mathbb{E}_t^H \Delta s_{t+1} + \hat{r}_t - r_t = \mathbb{E}_t^D \Delta s_{t+1} + \hat{r}_t - r_t + \nabla \mathbb{E}_t^H s_{t+1} = \nabla \mathbb{E}_t^H s_{t+1} + \psi.$$

Combining this expression with (A28) gives us

$$\alpha_t = \Theta_t^H \nabla \mathbb{E}_t^H s_{t+1} - \Psi_t^H \mathbb{E}_t^H e r_{t+1}^Q + \Theta_t^H \left(\frac{1}{2} \mathbb{V}_t^H(s_{t+1}) + \psi - \psi_{\hat{H},t}^s \right). \quad (\text{A29})$$

In the case of European households, their desired share of wealth held in euro bonds, $\hat{\alpha}_t$, is by definition equal to $1 - \tilde{\alpha}_t - \hat{\alpha}_t^Q - (\hat{P}_t \hat{C}_{\hat{H},t} + \hat{M}_{\hat{H},t})/\hat{W}_{\hat{H},t}$. Substituting for $\tilde{\alpha}_t$ and $\hat{\alpha}_t^A$ from (A8), $\hat{P}_t \hat{C}_{\hat{H},t}/\hat{W}_{\hat{H},t}$ from (A10) and $\hat{M}_{\hat{H},t}/\hat{W}_{\hat{H},t}$ from (A9) in this definition gives

$$\hat{\alpha}_t = 1 + \Theta_t^{\hat{H}} (\mathbb{E}_t^{\hat{H}} \Delta s_{t+1} + \hat{r}_t - r_t - \frac{1}{2} \mathbb{V}_t^{\hat{H}}(\Delta s_{t+1}) + \psi_{\hat{H},t}^{-s}) - \Psi_t^{\hat{H}} \mathbb{E}_t^{\hat{H}} e r_{t+1}^{\hat{Q}} - \exp(\hat{c}_{\hat{H},t} + \hat{p}_t - \hat{w}_{\hat{H},t}) (1 + \exp(\varpi - \eta \hat{r}_t)) \quad (\text{A30})$$

where $\Theta_t^{\hat{H}} \equiv \frac{\rho}{\gamma} (\mathbb{V}_t^{\hat{H}}(r_{t+1}^{\hat{Q}}) + \beta_t^{\hat{H}} \mathbb{V}_t^{\hat{H}}(r_{t+1}^{\hat{Q}}))/|\hat{\Xi}_t^{\hat{H}}|$ and $\Psi_t^{\hat{H}} \equiv \frac{\rho}{\gamma} (\mathbb{V}_t^{\hat{H}}(s_{t+1}) + \beta_t^{\hat{H}} \mathbb{V}_t^{\hat{H}}(r_{t+1}^{\hat{Q}}))/|\hat{\Xi}_t^{\hat{H}}|$. Proceeding as above, we obtain

$$\hat{\alpha}_t = 1 + \Theta_t^{\hat{H}} \nabla \mathbb{E}_t^{\hat{H}} s_{t+1} - \Psi_t^{\hat{H}} \mathbb{E}_t^{\hat{H}} e r_{t+1}^{\hat{Q}} + \Theta_t^{\hat{H}} \left(\psi_{\hat{H},t}^{-s} + \psi - \frac{1}{2} \mathbb{V}_t^{\hat{H}}(s_{t+1}) \right) - \exp(\hat{c}_{\hat{H},t} + \hat{p}_t - \hat{w}_{\hat{H},t}) (1 + \exp(\varpi - \eta \hat{r}_t)). \quad (\text{A31})$$

Equations (A29) and (A31) show that the desired portfolio shares for euro 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 equities, (iii) risk premia; and in the case of European households; the consumption wealth and money wealth ratios. Substituting the expressions for α_t and $\hat{\alpha}_t$ in the order flow equation (16), and linearizing around a symmetric steady state where expectations of dealers and households are the same gives (17).

Proposition 3 Let $\Omega_t^H = \{\Omega_t^D, v_t\}$ for some vector of variables v_t so that $\Omega_t^D \subset \Omega_t^H$. From Bayesian updating we know that

$$\begin{aligned} \mathbb{E}[\mathcal{X}_{t+1} | \Omega_t^\omega, v_t] &= \mathbb{E}[\mathcal{X}_{t+1} | \Omega_t^\omega] + \mathbb{B}_{\mathcal{X},v}(v_t - \mathbb{E}[v_t | \Omega_t^\omega]), \\ \mathbb{B}_{\mathcal{X},v} &= \mathbb{V}_t^\omega(v_t)^{-1} \mathbb{C}\mathbb{V}_t^\omega(\mathcal{X}_{t+1}, v_t). \end{aligned} \quad (\text{A32})$$

for some random variable \varkappa_{t+1} and information set Ω_t^ω . Applying this equation in the case where $\varkappa_{t+1} = \mathbb{E}[\mathbf{y}_{t+1}|\Omega_{t+1}^D]$, $\Omega_t^\omega = \Omega_t^D$, 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 $\varkappa_{t+1} = \mathbf{y}_{t+1}$, $\Omega_t^\omega = \Omega_t^D$, 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:

$$\mathbb{E}_t^H \mathbb{E}_{t+1}^D \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1} = \kappa (\mathbb{E}_t^H \mathbf{y}_{t+1} - \mathbb{E}_t^D \mathbf{y}_{t+1}), \quad (\text{A33})$$

where

$$\kappa \equiv \mathbb{B}_{\mathbb{E}_{t+1}^D \mathbf{y}_{t+1}, v_t} (\mathbb{B}'_{\mathbf{y}_{t+1}, v_t} \mathbb{B}_{\mathbf{y}_{t+1}, v_t})^{-1} \mathbb{B}'_{\mathbf{y}_{t+1}, v_t}.$$

Now we combine (20) and (A33) to give $\nabla \mathbb{E}_t^H s_{t+1} = \pi \kappa \nabla \mathbb{E}_t^H \mathbf{y}_{t+1}$ which is (21a). Applying the same technique to the foreign forecast differential gives $\nabla \mathbb{E}_t^{\hat{H}} s_{t+1} = \pi \hat{\kappa} \nabla \mathbb{E}_t^{\hat{H}} \mathbf{y}_{t+1}$ where $\hat{\kappa}$ is the foreign counterpart of κ . This is equation (21b). Substitution for $\nabla \mathbb{E}_t^H s_{t+1}$ and $\nabla \mathbb{E}_t^{\hat{H}} s_{t+1}$ in (17) with these expressions gives (22).

Proposition 4 First we use (A32) with $\mathbf{y}_{t+1} = \varkappa_{t+1}$, $\Omega_t^\omega = \Omega_t^D$, and $\Omega_{t+1}^D = \{\Omega_t^D, v_t\}$ to give

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

Next we combine this expression with (26):

$$\Delta s_{t+1} = r_t - \hat{r}_t + \psi + \pi \mathbb{B}_{\mathbf{y}_{t+1}, v_t} (v_t - \mathbb{E}[v_t|\Omega_t^D]).$$

Now note that the vector v_t denotes the new information available to dealers between the start of periods t and $t+1$. Thus, period t order flow x_t is an element of v_t . We can therefore write:

$$\Delta s_{t+1} = r_t - \hat{r}_t + \psi + b(x_t - \mathbb{E}_t^D x_t) + \zeta_{t+1}$$

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. To see how the correlation between order flow and spot rates depends on the degree of information aggregation, we simply use (22) to substitute for x_t in the definition of $\mathbb{B}_{\mathbf{y}_{t+1}, x_t}$. In particular, we first write

$$\pi \mathbb{B}_{\mathbf{y}_{t+1}, x_t} \mathbb{V}_t^{\mathbb{D}}(x_t) = \phi \pi \mathbb{C} \mathbb{V}_t^{\mathbb{D}}(\mathbf{y}_{t+1}, \nabla \mathbb{E}_t^{\mathbb{H}} \mathbf{y}'_{t+1}) \kappa' \pi' + \hat{\phi} \pi \mathbb{C} \mathbb{V}_t^{\mathbb{D}}(\mathbf{y}_{t+1}, \nabla \mathbb{E}_t^{\widehat{\mathbb{H}}} \mathbf{y}'_{t+1}) \hat{\kappa}' \pi' + \pi \mathbb{C} \mathbb{V}_t^{\mathbb{D}}(\mathbf{y}_{t+1}, o_t),$$

and use the identity $\mathbf{y}_{t+1} \equiv \mathbb{E}_t^{\mathbb{D}} \mathbf{y}_{t+1} + \mathbb{E}_t^{\omega} \mathbf{y}_{t+1} - \mathbb{E}_t^{\mathbb{D}} \mathbf{y}_{t+1} + (\mathbf{y}_{t+1} - \mathbb{E}_t^{\omega} \mathbf{y}_{t+1})$ for $\omega = \{\mathbb{H}, \widehat{\mathbb{H}}\}$ to give

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

Proposition 5 Consider the projection of $\Delta f_{t+\tau}$ on $s_t - \mathbb{E}_t^{\mathbb{D}} f_t$ and the unexpected component of order flow $x_t - \mathbb{E}_t^{\mathbb{D}} x_t$:

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

Order flow has incremental forecasting power when β_x differs from zero. To show that this is indeed the case, we first note that $\beta_x (x_t - \mathbb{E}_t^{\mathbb{D}} x_t) + \epsilon_{t+\tau}$ must equal the projection error in (30), $\epsilon_{t+\tau}$, because $x_t - \mathbb{E}_t^{\mathbb{D}} x_t$ is uncorrelated with $s_t - \mathbb{E}_t^{\mathbb{D}} f_t$. Consequently, β_s takes the same value as it did in (30) and:

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

Using the identity $\Delta f_{t+\tau} \equiv \nabla \mathbb{E}_t^{\omega} \Delta f_{t+\tau} + \mathbb{E}_t^{\mathbb{D}} \Delta f_{t+\tau} + (\Delta f_{t+\tau} - \mathbb{E}_t^{\omega} \Delta f_{t+\tau})$ for $\omega = \{\mathbb{H}, \widehat{\mathbb{H}}\}$ to substitute for $\Delta f_{t+\tau}$, and (22) to substitute for order flow, we find that

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

The final step is to substitute for $\Delta f_{t+\tau}$ using the fact that $f_t = C y_t$.

Proposition 6 First we iterate (A27) forward τ periods to get

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

Subtracting $(\frac{\eta}{1+\eta})^\tau s_t$ from both sides and re-arranging gives

$$\mathbb{E}_t^\mathbb{D} \Delta^\tau s_{t+\tau} = \left(\left(\frac{1+\eta}{\eta} \right)^\tau - 1 \right) s_t - \left(\frac{1+\eta}{\eta} \right)^\tau \frac{1}{1+\eta} \mathbb{E}_t^\mathbb{D} \sum_{i=0}^{\tau-1} \left(\frac{\eta}{1+\eta} \right)^i f_{t+i}.$$

Combining this equation with the identity $\Delta^\tau s_{t+\tau} = \mathbb{E}_t^\mathbb{D} \Delta^\tau s_{t+\tau} + s_{t+\tau} - \mathbb{E}_t^\mathbb{D} s_{t+\tau}$, we find that

$$\begin{aligned} \Delta^\tau s_{t+\tau} &= \left(\left(\frac{1+\eta}{\eta} \right)^\tau - 1 \right) s_t - \left(\frac{1+\eta}{\eta} \right)^\tau \frac{1}{1+\eta} \mathbb{E}_t^\mathbb{D} \sum_{i=0}^{\tau-1} \left(\frac{\eta}{1+\eta} \right)^i f_{t+i} + s_{t+\tau} - \mathbb{E}_t^\mathbb{D} s_{t+\tau}, \\ &= \varphi^\tau (s_t - \mathbb{E}_t^\mathbb{D} f_{t,\tau}) + s_{t+\tau} - \mathbb{E}_t^\mathbb{D} s_{t+\tau}, \end{aligned}$$

where $\varphi^\tau \equiv (\frac{1+\eta}{\eta})^\tau - 1 > 0$ and $f_{t,\tau} \equiv \frac{\varphi^\tau + 1}{\varphi^\tau (1+\eta)} \sum_{i=0}^{\tau-1} (\frac{\eta}{1+\eta})^i f_{t+i}$. Using (11) to substitute for $s_{t+\tau} - \mathbb{E}_t^\mathbb{D} s_{t+\tau}$ gives equation (33).

A.3 Real-Time Estimates

We provide a brief description of how we computed the real-time estimates of a monthly log series \varkappa . Computing the real-time estimates for a quarterly series like GDP follows analogously and is described in detail by Evans (2005). Let $\Delta \varkappa_t$ denote the increment to the monthly value for $\varkappa_{M(\tau)}$, where $M(\tau)$ is the last date of month τ . Next, define the partial sum

$$\widetilde{\Delta \varkappa}_t \equiv \sum_{i=M(\tau-1)+1}^{\min\{M(\tau), t\}} \Delta \varkappa_i$$

as the cumulative daily contribution to $\varkappa_{M(\tau)}$ in month τ . Notice that when $t = M(\tau)$, the monthly change in $\varkappa_{M(\tau)}$, $\Delta^M \varkappa_{Q(\tau)} = \widetilde{\Delta \varkappa}_{M(\tau)}$. The daily dynamics of $\widetilde{\Delta \varkappa}_t$ are described by

$$\widetilde{\Delta \varkappa}_t = (1 - dum_t) \widetilde{\Delta \varkappa}_{t-1} + \Delta \varkappa_t, \tag{A34}$$

where dum_t is a dummy variable equal to one on the first day of each month, and zero otherwise. To accommodate the presence of variable reporting lags, let $\Delta^{M(j)} \varkappa_t$ denote the monthly growth in \varkappa ending on day $M(\tau - j)$ where $M(\tau)$ denotes the last day of the most recently completed month

and $t \geq M(\tau)$. Monthly growth in the last (completed) month is given by

$$\Delta^{M(1)}\varkappa_t = (1 - dum_t) \Delta^{M(1)}\varkappa_{t-1} + dum_t \widetilde{\Delta}\varkappa_{t-1}. \quad (\text{A35})$$

When t is the first day of a new month, $dum_t = 1$, so $\Delta^{M(1)}\varkappa_{M(\tau)+1} = \widetilde{\Delta}\varkappa_{M(\tau)} = \Delta^M\varkappa_{M(\tau)}$. On all other days, $\Delta^{M(1)}\varkappa_t = \Delta^{M(1)}\varkappa_{t-1}$. To accommodate occasions where the reporting lag is more than a month, we track monthly growth two months back via the recursion:

$$\Delta^{M(2)}\varkappa_t = (1 - dum_t) \Delta^{M(2)}\varkappa_{t-1} + dum_t \Delta^{M(1)}\varkappa_{t-1}. \quad (\text{A36})$$

Equations (A34), (A35) and (A36) enable us to define the link between the daily contributions, $\Delta\varkappa_t$, and data releases. Suppose the reporting lag for the release on day t is less than one month. Then if $\widehat{\Delta}\varkappa_t$ is the released value for the growth in \varkappa during the last month on day t ,

$$\widehat{\Delta}\varkappa_t = \Delta^{M(1)}\varkappa_t. \quad (\text{A37})$$

If the reporting lag is longer than a month (but less than two),

$$\widehat{\Delta}\varkappa_t = \Delta^{M(2)}\varkappa_t. \quad (\text{A38})$$

We incorporate the information contained in the monthly data releases on other variables in a similar manner. (Incorporating information from quarterly data releases is more complex, see Evans 2005 for details.) Specifically, let z_t^i denote the value of another series, released on day t , that relates to activity in the last completed month. We assume that

$$z_t^i = \beta_i \Delta^{M(1)}\varkappa_t + u_t^i. \quad (\text{A39})$$

where u_t^i is an i.i.d. $N(0, \sigma_i^2)$ shock. In cases where the reporting lag is two months,

$$z_t^i = \beta_i \Delta^{M(2)}\varkappa_t + u_t^i. \quad (\text{A40})$$

It is important to recognize that (A37) - (A40) allows for variations in the reporting lag from data

release to data release.

To complete the model, we specify the dynamics for the daily increments, $\Delta \varkappa_t$. We assume that

$$\Delta \varkappa_t = \sum_{i=1}^k \phi_i \Delta^{M(i)} \varkappa_t + \zeta_t, \quad (\text{A41})$$

where ζ_t is an i.i.d. $N(0, \sigma_\zeta^2)$ shock.

Finding the real time estimates of \varkappa requires a solution to two related problems. First, there is a pure inference problem of how to compute $\mathbb{E}[\varkappa_{M(\tau)}|\Omega_t]$ using the signalling equations (A37) - (A40), and the $\Delta \varkappa_t$ process in (A41), given values for all the parameters in these equations. Second, we need to estimate these parameters. The Kalman Filtering algorithm provides a solution to both problems. In particular, given a set of parameter values, the algorithm provides the means to compute the real-time estimates $\mathbb{E}[\varkappa_{M(\tau)}|\Omega_t]$. The algorithm also allows us to construct a sample likelihood function from the data series, so that the model's parameters can be computed by maximum likelihood.

To use the algorithm, we write the model in state space form. For the case where $k = 1$, the dynamics described by equations (A34) - (A36) and (A41) can be represented by the matrix equation:

$$\begin{bmatrix} \widetilde{\Delta \varkappa_t} \\ \Delta^{M(1)} \varkappa_t \\ \Delta^{M(2)} \varkappa_t \\ \Delta \varkappa_t \end{bmatrix} = \begin{bmatrix} 1 - dum_t & 0 & 0 & 1 \\ dum_t & 1 - dum_t & 0 & 0 \\ 0 & dum_t & 1 - dum_t & 0 \\ 0 & \phi_1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \widetilde{\Delta \varkappa_{t-1}} \\ \Delta^{M(1)} \varkappa_{t-1} \\ \Delta^{M(2)} \varkappa_{t-1} \\ \Delta \varkappa_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \zeta_t \end{bmatrix},$$

or, more compactly

$$\mathbb{Z}_t = \mathbb{A}_t \mathbb{Z}_{t-1} + \mathbb{V}_t. \quad (\text{A42})$$

This is the state equation of the state space form.

The link between the data releases on \varkappa and the elements of \mathbb{Z}_t are described by (A37) and (A38):

$$\widehat{\Delta \varkappa_t} = \begin{bmatrix} 0 & \text{ML}_t^1(\varkappa) & \text{ML}_t^2(\varkappa) & 0 \end{bmatrix} \mathbb{Z}_t, \quad (\text{A43})$$

where $\text{ML}_t^i(\varkappa)$ denotes a dummy variable that takes the value of one when the reporting lag for

series \varkappa lies between $i - 1$ and i months, and zero otherwise. The link between the releases for the other series and elements of \mathbb{Z}_t are described by (A39) and (A40):

$$z_t^i = \begin{bmatrix} 0 & \beta_i \text{ML}_t^1(z^i) & \beta_i \text{ML}_t^2(z^i) & 0 \end{bmatrix} \mathbb{Z}_t + u_t^i. \quad (\text{A44})$$

Stacking (A43) and (A44) for series $i = 1, 2, \dots, g$ gives

$$\begin{bmatrix} \widehat{\Delta \varkappa}_t \\ z_t^1 \\ \vdots \\ z_t^g \end{bmatrix} = \begin{bmatrix} 0 & \text{ML}_t^1(\hat{\varkappa}) & \text{ML}_t^2(\hat{\varkappa}) & 0 \\ 0 & \beta_1 \text{ML}_t^1(z^1) & \beta_1 \text{ML}_t^2(z^1) & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \beta_g \text{ML}_t^1(z^g) & \beta_g \text{ML}_t^2(z^g) & 0 \end{bmatrix} \mathbb{Z}_t + \begin{bmatrix} 0 \\ u_t^1 \\ \vdots \\ u_t^g \end{bmatrix},$$

or

$$\mathbb{X}_t = \mathbb{C}_t \mathbb{Z}_t + \mathbb{U}_t. \quad (\text{A45})$$

This equation links the vector of *potential* data releases for day t , \mathbb{X}_t , to elements of \mathbb{Z}_t . The vector of actual data releases for day t , \mathbb{Y}_t , is related to the vector of *potential* releases by

$$\mathbb{Y}_t = \mathbb{B}_t \mathbb{X}_t,$$

where \mathbb{B}_t is a $n \times (g + 1)$ selection matrix that “picks out” the $n \geq 1$ data releases for day t . Combining this expression with (A45) gives us the observation equation:

$$\mathbb{Y}_t = \mathbb{B}_t \mathbb{C}_t \mathbb{Z}_t + \mathbb{B}_t \mathbb{U}_t. \quad (\text{A46})$$

Equations (A42) and (A46) describe a state space form which can be used to find real-time estimates of variable \varkappa in two steps. In the first, we obtain the maximum likelihood estimates of the model’s parameters. For this purpose the sample likelihood function is built up recursively by applying the Kalman Filter to (A42) and (A46). The second step applies the Kalman Filter to (A42) and (A46) to calculate the real-time estimates of \varkappa using the maximum likelihood parameter estimates.

The real-time estimates for US variables use data releases on quarterly GDP and 18 monthly releases: Nonfarm Payroll, Employment, Retail Sales, Industrial Production, Capacity Utiliza-

tion, Personal Income, Consumer Credit, Personal Consumption Expenditures, New Home Sales, Durable Goods Orders, Construction Spending, Factory Orders, Business Inventories, the Government Budget Deficit, the Trade Balance, NAPM index, Housing Starts, the Index of Leading Indicators, Consumer Prices and M1. The real-time estimates for German variables use data releases on quarterly GDP and 8 monthly releases: Employment, Retail Sales, Industrial Production, Manufacturing Output, Manufacturing Orders, the Trade Balance, Consumer Prices and M1. We allow for 10 lags in the daily increment process when estimating real-time GDP, and 7 lags for the other variables. These specifications appear to capture all the time-series variation in the data. In particular, we are unable to reject the null hypothesis of no serial correlation in the Kalman Filter innovations evaluated at the maximum likelihood estimates for any of our models.