

Currency Returns, Intrinsic Value, and Institutional-Investor Flows

KENNETH A. FROOT and TARUN RAMADORAI*

ABSTRACT

We decompose currency returns into (permanent) intrinsic-value shocks and (transitory) expected-return shocks. We explore interactions between these shocks, currency returns, and institutional-investor currency flows. Intrinsic-value shocks are: dwarfed by expected-return shocks (yet currency returns overreact to them); unrelated to flows (although expected-return shocks correlate with flows); and related positively to forecasted cumulated-interest differentials. These results suggest flows are related to short-term currency returns, while fundamentals better explain long-term returns and values. They also rationalize the long-observed poor performance of exchange-rate models: by ignoring the distinction between permanent and transitory exchange-rate changes, prior tests obscure the connection between currencies and fundamentals.

WHAT DRIVES EXCHANGE-RATE VALUES? It is an age-old question, with a heritage of rich modeling and important theoretical insights. There has been frustration with the data, however, which generally have not held up their part of the bargain. A generation after Meese and Rogoff (1983), traditional exchange-rate fundamentals specified by theory continue to forecast more poorly than a random walk. There are some important exceptions—such as the evidence at long horizons presented by Mark (1995)—but by and large there remains no well-accepted “traditional” model of exchange-rate determination.

This frustration has led to a search for alternatives that better explain exchange-rate changes. Evans and Lyons (2002) propose net foreign exchange order flow as a candidate.¹ They find that daily interdealer order flow explains about 60% of daily exchange-rate changes and argue that flows are the

*Froot is André R. Jakurski Professor of Business Administration at Harvard University and is a research fellow and Chairman of the Insurance Program at the NBER; Ramadorai is at the University of Oxford and CEPR. We thank an anonymous referee, Mark Aguiar, Atindra Barua, John Campbell, Randy Cohen, Frank Diebold, Bernard Dumas, Rich Lyons, Colin Mayer, Paul O’Connell, Andrei Shleifer, Jeremy Stein, Tuomo Vuolteenaho, Joshua White, and especially Jeff Frankel for comments; seminar participants at the Harvard Finance and International Economics Seminars and the NBER Spring IFM Program Meetings for discussions; the Department of Research at Harvard Business School for research support; and State Street Corporation and State Street Associates for data. All errors remain our responsibility. This is a substantially revised version of NBER Working Paper no. 9101.

¹ In addition, Rime (2001) finds that weekly U.S. Treasury data on flows help explain exchange-rate movements. Wei and Kim (1997) and Cai et al. (2001) find that large-trader positions explain currency volatility better than do news announcements or traditional fundamentals. See Cheung and Chinn (2001) for survey evidence from practitioners on order flow information.

proximate cause of exchange-rate movements. This and related evidence has given rise to the view that the interaction between flows and foreign exchange microstructure might hold the elusive key—at least empirically and perhaps even theoretically—to exchange-rate determination.

How might flows affect currency values? We distinguish three generic ways. One is that flow is correlated with news about long-run or intrinsic currency value. We term this the “strong flow-centric” view. It holds that flows can cause current and future exchange-rate changes through private information about value that, when released, permanently impacts exchange rates (see Kyle (1985) for such a model). Alternatively, causality might run the other way, so that changes in intrinsic value might cause flow. Brennan and Cao (1997), for example, argue that, conditional on positive news, investors who update their priors by more will concurrently buy the currency. Either way, under the strong flow-centric view, flows should be positively correlated with changes in a currency’s intrinsic value.²

Second is a weaker version of this, which we refer to as the “weak flow-centric view.” It says that institutional flows are related to *deviations* from intrinsic value, but not to intrinsic value itself. Flows therefore have only transitory price effects, associated with things such as liquidity changes, price pressure, and temporary preference and other demand shocks. In equity markets, for example, a growing list of papers (e.g., Barberis and Shleifer (2003), Hong and Stein (1999), and Daniel, Hirshleifer, and Subrahmanyam (1998)) appeal to the flows of certain groups of traders to explain apparent patterns of transitory price behavior; that is, short-term momentum and longer term reversion. In these papers, transitory demand shocks cause both flow and price deviations away from intrinsic value. However, the weak flow-centric view is also agnostic about the direction of causality. Its prediction is simply that flows and returns will be positively correlated at short horizons, but uncorrelated at long horizons, once the reversion of the currency to its intrinsic value fully factors in.

A third and final possibility is that the institutional-investor flows we measure are unrelated to currency values, even over the short run. We call this the “fundamentals-only view” and it serves throughout as a convenient null hypothesis: flows explain no portion of unexpected returns, that is, either changes in intrinsic value or deviations from intrinsic value. Exchange-rate fundamentals, not flows, explain currency values.

In equities, there is already considerable evidence rejecting the fundamentals-only view in favor of the flow-centric views (Froot, O’Connell, and Seasholes (2001), Froot and Ramadorai (2001), Cohen, Gompers, and Vuolteenaho (2002), and Choe, Kho, and Stulz (2001) are recent examples). Like interbank currency flows, institutional-investor flows in equities appear contemporaneously correlated with returns. The correlation is repeatedly positive, suggesting that, as a group, these investors (not those who trade with them) take liquidity, rather than provide it. However, there is as yet

² Throughout the paper, we define “intrinsic-value” shocks as changes in the real exchange rate, holding constant expected excess currency returns. See Section I.A.

no evidence to distinguish between the strong and weak flow-centric views, since no studies have controlled for intrinsic value. Previous rejections of the fundamentals-only view in favor of a flow-centric view therefore cannot distinguish between strong and weak versions.

The goal of this paper is to assess these three views, controlling, in particular, for shocks to intrinsic value. We do this using an extremely general exchange-rate model whose only assumption is that purchasing power parity (PPP) holds in expectation in the long run. We then show how unexpected returns can be decomposed into two components: permanent “intrinsic-value” shocks and transitory deviations from intrinsic value, which we call “expected-return” shocks. This decomposition for currencies is similar to the widely used approximate decomposition of equity returns derived by Campbell and Shiller (1988) and Campbell (1991). An important difference, however, is that our derivation is exact, not approximate. This follows because currency returns are log linear in continuously compounded interest rates, whereas stock returns are not log linear in dividends.

Both components of unexpected returns are innovations, and therefore are unpredictable. However, expected-return shocks have, by definition, transitory impacts on value: Unexpected increases (decreases) in the current exchange rate are associated with simultaneously offsetting decreases (increases) in cumulated future expected returns. Intrinsic-value shocks, by contrast, have permanent effects on value: They do not result in a change—offsetting or otherwise—in future expected returns. Indeed, we demonstrate that if long-run PPP holds and expected returns are constant, then intrinsic-value news is identically equal to innovations in cumulated expected future real interest differentials.

As with Campbell and Shiller, our return decomposition can be readily implemented empirically using a vector autoregression (VAR). Intrinsic-value news is computed as a residual, after subtracting from realized unexpected returns the cumulated forecast of changes in expected returns (i.e., subtracting estimated expected-return news). The VAR variables are motivated directly by our model, which equates by identity changes in expected returns to changes in cumulated real interest differentials and the real exchange rate. Along with these variables, we include institutional-investor flows, allowing us to distinguish among our three hypotheses.

Our flow panel data set—high-frequency daily data covering 7 years and 18 exchange rates—provides us with a unique opportunity to measure long-horizon as well as short-horizon effects. Indeed, we find that the cross-sectional power of the panel allows us to make useful statements about long-run and permanent components, even though our time series extends only 7 years. This gives us the ability to better fuse an interest in exchange-rate determination at the macroeconomic level with issues that arise in the currency microstructure literature. Thus, the scope of our data, as well as our methodology, distinguishes our empirical results from other flow-related papers.

When we perform our decomposition into intrinsic-value and expected-return shocks, the data tell us that intrinsic-value news is relatively small—its

volatility is roughly one-half that of expected-return news. In addition, we find that intrinsic-value news and expected-return news are negatively correlated, and we show that this implies currency returns *overreact* to intrinsic-value news. Initial overreaction is substantial; we measure it to be on average between 24% and 56% of the intrinsic-value shock. We also find that 30 trading days (approximately 1.5 months) *after* the shock, the amount of overreaction has actually grown farther, so that the average overreaction grows to between 27% and 89%. Viewed from the currency response 30 days out, this implies a pattern of short-term *underreaction* (i.e., momentum) and longer term overreaction (i.e., reversal) in currency returns. Such patterns are qualitatively similar to what emerges from equity index return decompositions, reinforcing the view that common mechanisms lie behind short-run momentum and long-run reversal (see, e.g., Campbell (1991), Campbell and Ammer (1993), and Vuolteenaho (2002)).³

This decomposition is a powerful tool for distinguishing our three views. For example, like Evans and Lyons (2002), we find a strong positive contemporaneous correlation between daily currency returns and flows, consistently averaging about 25% in the major currencies.⁴ However, the decomposition shows that this correlation is entirely driven by long horizon expected-return shocks; that is, an appreciation associated with a current inflow is, over time, expected to reverse because of lower expected future returns. Flows are unrelated to the permanent component of exchange-rate changes, intrinsic-value shocks.

We also turn our decomposition toward “trend chasing”—the idea that current flows respond extrapolatively to past returns. Numerous behavioral asset-pricing models have at least one class of investors who behave this way, and it is a common empirical result as well.⁵ We find that, conditional on a surprise appreciation, institutional investors do indeed tend to buy if an intrinsic-value shock was responsible. However, they *sell* if a short-term expected-return shock was responsible for the appreciation. This finding—selling subsequent to transitory appreciations and buying otherwise—suggests that investors distinguish between intrinsic-value and expected-return shocks. It also suggests that mechanical views of institutional investors as “trend chasers” are too simplistic to describe institutional-investor behavior.

³ Our underreaction finding is similar to, but smaller in magnitude than, momentum effects in equity markets (see Jegadeesh and Titman (1993) and Rouwenhorst (1998)). For example, in individual stocks Vuolteenaho (2002) finds the underreaction component to be 40% of the permanent shock. The overreaction finding for equities is in DeBondt and Thaler (1985) and Jegadeesh and Titman (2001), among others.

⁴ The correlation of returns with institutional investor flows is positive for every currency, but lower than that computed with the interbank transaction count used by Evans and Lyons (2002). This should be expected, in that purchases recorded in the interbank data represent actions taken by the “active” counterparty (i.e., the one taking liquidity). Our institutional investors may sometimes “purchase” through a limit order at the bid, so that they may provide, rather than take, liquidity. As a result, we would expect the daily correlation between returns and flows to be lower in our data.

⁵ See Griffin, Harris, and Topaloglu (2003) for a review of other literature and empirical results on trend chasing.

We also provide evidence that flows anticipate returns. In our data, flow surprises positively forecast excess returns over relatively short horizons (up to 30 trading days). However, flow surprises *negatively* forecast excess returns at longer horizons (i.e., between 30 trading days and an infinite horizon). This too suggests momentum and reversal, and helps corroborate the finding that flows are positively related to contemporaneous and future returns but unrelated to long-horizon returns.

Similarly, when we run naïve regressions of flows and returns at different horizons, we find that at horizons of 1 or 2 months the flow/return correlation rises to almost 45%, reliably higher than at daily horizons. The positive noncontemporaneous correlations implied by trend-chasing and anticipation effects are behind this increase. However, at longer horizons, noncontemporaneous correlations become negative, so much so that at horizons of 3 or more years the total correlation between flows and returns is negative (although indistinguishable from zero). At these horizons, trend-chasing and anticipation effects appear fully reversed. These noncontemporaneous effects are not examined in previous work on foreign exchange and are the underlying driver for the conclusion that transitory components dominate the relationship between exchange rates and flows.

A key prediction of our model is that the permanent component of returns (i.e., intrinsic-value news) will be positively correlated with innovations in real interest differentials. We test this and find it holds significantly in the data, in contrast to the insignificant relationship between permanent return components and flows. This suggests that low power and/or noise in permanent returns does not explain our inability to detect a relationship between intrinsic-value news and flows.

Overall, our results hint at two major conclusions. First, our decomposition provides considerable evidence rejecting the strong flow-centric view in favor of the weak flow-centric view. Flows appear to be bound up with transitory currency under- and overreactions, but unrelated to the permanent component of exchange-rate surprises. Yet, these exchange-rate surprises are strongly related to important fundamental variables, as predicted by theory. The results therefore provide strong support for models of investor-driven transitory deviations from intrinsic value (such as Barberis and Shleifer (2003), Hong and Stein (1999), and Daniel et al. (1998)), and against models that suggest investor-driven movements in intrinsic value itself (such as that of Kyle (1985)).

Second, the results suggest to us a possible rationale for the failure of traditional exchange-rate models. We find that the permanent component of currency returns is the only portion of currency returns related to important fundamentals. Yet the variance of the permanent component is small—approximately 15–20% of the total variance of returns. The large transitory exchange-rate changes we measure act to obscure the relationship between the permanent component of returns and important fundamentals. Thus, much as in Mark (1995) and Clarida and Gali (1994), we find evidence contrary to Meese and Rogoff (1983, 1988): longer term

returns both forecast and respond contemporaneously to changes in observed fundamentals.⁶

The rest of this paper is organized as follows. Section I derives our return decomposition and introduces the associated VAR. Section II discusses our data and methodology for pricing forward contracts. A naïve approach to estimating the relationship between flows and returns at multiple horizons is presented and estimated in Section III. Section IV then discusses the main VAR results and their interpretation. Section V concludes. The appendices that follow explain the construction of our flow data, the estimation of simulated standard errors, and the analytic connection between our VAR and naïve estimators.

I. Return Decomposition and the VAR

A. A Foreign Exchange Return Decomposition

Our first step is to decompose currency excess returns into components due to innovations in intrinsic value and innovations in expected excess returns. By definition, the log excess return on foreign exchange is equal to the change in the log real exchange rate plus the log real interest rate differential

$$r_{t+1} = (\delta_{t+1} - \delta_t) + (i_t^* - \pi_{t+1}^*) - (i_t - \pi_{t+1}), \quad (1)$$

where δ is the value of the log real exchange rate defined in terms of U.S. dollars per unit of foreign exchange; i and i^* are, respectively, the continuously compounded one-period U.S. dollar and foreign currency riskless interest rates; and π and π^* are the associated continuously compounded rates of inflation.

Solving this equation forward, subject to expectations based on time- t information, we have that the log real exchange rate is the sum of expected future real interest differentials, less cumulated log expected excess returns:

$$\delta_t = \sum_{i=1}^{\infty} E_t(d_{t+i} - r_{t+i}), \quad (2)$$

where d_t is the real interest differential between times $t - 1$ and t , $d_t = (i_{t-1}^* - \pi_t^*) - (i_{t-1} - \pi_t)$. In solving equation (1) forward, we impose the terminal condition $\lim_{i \rightarrow \infty} E_t \delta_{t+i} = 0$, assuring convergence. This implies that, in the long run, PPP holds in expectation. Equation (2) therefore says that any appreciation away from PPP occurs because future real interest differentials are temporarily higher than required returns. The real appreciation decays over time, as the real interest differential less expected excess returns reverts to its unconditional mean.

⁶ The presence of large transitory components also reconciles our results with those of Evans and Lyons (2003), who find that at daily horizons, news announcements are transmitted to exchange rates partly through flows. At such horizons, our results suggest that the flow channel is associated only with the transitory component of returns.

Using equations (1) and (2), we can write the unexpected one-period return as the sum of the innovations in expected future real interest differentials and excess currency returns

$$r_{t+1} - E_t(r_{t+1}) = (E_{t+1} - E_t) \sum_{i=1}^{\infty} (d_{t+i} - r_{t+i+1}). \tag{3}$$

This decomposition equation differs from Campbell’s in that it is exactly, not approximately, linear. The approximation in the case of equities occurs because dividends are additive to cashflow, so that log returns can be written in terms of log prices and log dividends only by approximation. In the exchange-rate example, however, cash flow accretion (interest) occurs proportionally, and therefore the log return is exactly additive in log real interest differentials and log real exchange-rate changes. An alternative interpretation for equation (3) is that it represents a decomposition of excess currency returns into a permanent and transitory component as in Beveridge and Nelson (1981). Changes in future expected excess returns generate temporary fluctuations, as the current impact of a future change creates an equal and opposite movement in the currency.

Equation (3) expresses unexpected excess currency returns as the difference between “intrinsic-value” news and “expected-return” news. Intrinsic-value news is defined to be the innovation in the expected present value of future interest differentials, $(E_{t+1} - E_t) \sum_{i=1}^{\infty} d_{t+i}$. It can be thought of as the excess currency return that would prevail at a given time if expected future currency returns were held constant. Expected-return news, which is $(E_{t+1} - E_t) \sum_{i=1}^{\infty} r_{t+i+1}$, is the innovation in the exchange rate that is attributable to a change in required excess returns, holding intrinsic value constant. Naturally, an increase in future expected returns, given intrinsic value, results in a current depreciation. Defining

$$v_{iv,t} = (E_{t+1} - E_t) \sum_{i=1}^{\infty} d_{t+i}$$

and

$$v_{er,t} = (E_{t+1} - E_t) \sum_{i=1}^{\infty} r_{t+i+1},$$

equation (3) can therefore be written as

$$r_{t+1} - E_t(r_{t+1}) = v_{iv,t} - v_{er,t}. \tag{4}$$

A surprise appreciation of a currency is associated either with an increase in expected future real interest differentials given required returns, or a decrease in required future returns given real interest differentials.

Note that equation (2) holds for any conditioning information used efficiently. It holds with perfect foresight and under expectations formed rationally subject

to any information set. By the law of iterated expectations therefore equation (4) is an unpredictable innovation. Hence, we have considerable leeway in positing variables that affect expected returns and/or intrinsic values.

The decomposition in equation (4) is useful for several purposes. First, it allows us to compare the magnitudes of intrinsic-value and expected-return shocks in currencies. In equities, this decomposition has been the subject of considerable research. The results suggest that, for equity indexes, expected-return news dominates intrinsic-value news, whereas, for individual equities, the reverse seems to be true.⁷

Second, the decomposition permits us to understand how changes in flows and intrinsic value interact with currency returns. To investigate this interaction, we need to define the VAR setup.

B. VAR Specification

The specification begins with a standard VAR built to allow for proper forecasting of the terms in equation (2). It includes equations determining returns (r_t), flows normalized by own-flow *SD* (standard deviation) (f_t), interest differentials (d_t), and the real exchange rate (δ_t):

$$z_t = \Gamma z_{t-1} + u_t, \quad (5)$$

where $z_t' = [r_t \ f_t \ d_t \ \delta_t]$. We include these variables in order to forecast expected future intrinsic values, flows, and returns. Real exchange rates and, especially, interest differentials, are widely noted as informative about excess returns (see Cumby and Huizinga (1991), Baxter (1994), and Engel and West (2003)).⁸ The variable z_t and the companion matrix Γ are defined to allow for a general number of lags (see Campbell and Shiller (1987)). We assume that Γ is constant across currencies within a given subsample and that the covariance matrix, $E[u_t u_t'] = \Sigma$, allows for contemporaneous correlation of the residuals across currencies.

The VAR impulse response allows us to identify how shocks affect cumulated expected innovations. Specifically, the innovation in cumulative expected

⁷ Campbell and Shiller (1988) and Campbell (1991) find that expected return news dominates cash flow news in the postwar period for equity indexes. Vuolteenaho (2002) and Cohen et al. (2002) perform a similar exercise for individual stocks and find that expected return news is less than half as large as cashflow news. Campbell and Clarida (1987) decompose excess currency returns into changes in the long-run real exchange rate and long-run real interest differentials. However, they assume that changes in expected returns are perfectly correlated with changes in real interest differentials.

⁸ We also considered breaking up the real interest differential equation into two equations (nominal interest differentials and inflation differentials), in order to allow the dynamics of the processes to differ. This had no qualitative effect on the results, other than the expected reduction in statistical precision. (The number of free parameters is proportional to the number of equations squared.) We therefore present the four-equation version here.

future changes $k \geq 1$ periods forward, $(E_t - E_{t-1}) \sum_{j=1}^{\infty} [z_{t+j} - z_{t+k+j}]$, is given by $\Phi(k)u_t$, where⁹

$$\Phi(k)u_t = (\Gamma - \Gamma^{k+1})(I - \Gamma)^{-1}u_t. \tag{6}$$

We pick out cumulated expected changes in any VAR variable by premultiplying by the appropriate selection vector. For example, the innovation in cumulated expected excess returns is given by $e1'\Phi(k)u_t$, where

$$e1' = [1 \quad 0 \quad \dots \quad 0]. \tag{7}$$

Analogously, the innovation in cumulated expectations of the second variable, currency flows, is given by $e2'\Phi(k)u_t$, where $e2' = [0 \quad 1 \quad \dots \quad 0]$.

The total impulse response from a shock to exchange-rate returns is the sum of the innovation in cumulative expected future return changes, $e1'\Phi(k)u_t$ plus the shock itself, $e1'u_t$, or

$$e1'\Psi(k)u_t = e1'(\Phi(k) + I)u_t, \tag{8}$$

where $\Psi(k) = (\Phi(k) + I)$. For accumulations beginning at time t , we denote the infinite-horizon cumulative innovation matrices as $\Phi = \Phi(\infty) = \Gamma(I - \Gamma)^{-1}$ and $\Psi = \Psi(\infty) = \Phi(\infty) + I$. For accumulations beginning at time $t + k$, infinite horizon cumulative innovations are $\Psi - \Psi(k)$ and $\Phi - \Phi(k)$, respectively.

The impulse responses from the VAR map directly back to the decomposition derived in Section I.A. According to equation (4) above, the permanent component of currency surprises is $v_{iv,t} = e1'u_t + v_{er,t}$. From the VAR, expected-return news is just $v_{er,t} = e1'\Phi u_t$ and intrinsic-value news is the total return impulse response, $v_{iv,t} = e1'\Psi u_t$. In other words, intrinsic-value news is the exchange-rate innovation that occurs when expected returns are held constant; it is observed by adding back all future expected-return innovations to the current return shock.

This methodology provides some assurance that omitted forecasting variables are not a problem. The permanent component of returns is just the residual from the return equation. If there were left-out variables in our VAR, their omission would bias the results toward too little transitory variation and too much permanent variation. However, as we report below, this bias does not seem to be a problem empirically: we find that the variation in the permanent component—intrinsic-value news—is relatively small in comparison with that of the transitory component—expected-return news.¹⁰

⁹To derive equation (6), note that $(E_t - E_{t-1})[z_{t+j}] = \Gamma^j u_t$, and therefore by the perpetuity formula $(E_t - E_{t-1}) \sum_{j=1}^{\infty} z_{t+j} = \Gamma(I - \Gamma)^{-1}u_t$. Applying this same logic to z_{t+k+j} and subtracting yields equation (6).

¹⁰Other papers similarly argue that parsimony is appropriate in comparable settings, especially given the specific nature of the left-out variable bias; see, for example, Cohen, Gompers, and Vuolteenaho (2002), Vuolteenaho (2002), Campbell and Ammer (1993), Campbell (1991), and Campbell and Shiller (1988).

C. VAR Relationships

Using the VAR and the return decomposition, we address a number of questions about the relationship between currencies, flows, and intrinsic values. The framework is well suited to address multiple horizons, which turn out to be especially important given the evidence below.

First, how big are expected-return and intrinsic-value shocks relative to currency-return shocks? We can answer this by computing the variance of excess returns, which from equation (4), is

$$\sigma_{fx}^2 = \sigma_{iv}^2 + \sigma_{er}^2 - 2\rho_{iv,er}\sigma_{iv}\sigma_{er}, \quad (9)$$

or, equivalently, in the notation from Section I.B,

$$e1'\Sigma e1 = e1'\Psi\Sigma\Psi'e1 + e1'\Phi\Sigma\Phi'e1 - 2e1'\Psi\Sigma\Phi'e1. \quad (10)$$

When useful, we can further decompose the right-hand side by distinguishing between cumulated innovations in expected returns over short horizons $\Phi(k)$ (up to k days) and long horizons, $\Phi - \Phi(k)$ (beginning with $k + 1$ days).

Second, the variance decomposition in equation (9) can be used to answer the question: How much does the currency move with a 1% shock to intrinsic value? Vuolteenaho (2002) interprets this coefficient as a measure of overreaction. That is, if the currency appreciates by more (less) than 1%, there is "overreaction" (underreaction). Obviously underreaction can occur only if the correlation between intrinsic-value and excess return shocks is sufficiently positive. From equation (10), the overreaction coefficient is given by

$$\beta_{\text{over}} = 1 - \frac{e1'\Phi(k)\Sigma\Psi'e1}{e1'\Psi\Sigma\Psi'e1} - \frac{e1'(\Phi - \Phi(k))\Sigma\Psi'e1}{e1'\Psi\Sigma\Psi'e1}. \quad (11)$$

Given that intrinsic-value shocks are the residual from our VAR, any measurement error in that model biases β_{over} toward 1, making it difficult to find evidence of over/underreaction.

Third, once we have determined the extent to which currencies under- and/or overreact, we measure the subsequent flow response. For example, if the currency initially underreacts to intrinsic-value news, do institutional investors exploit this by buying immediately thereafter?

D. VAR Estimation Issues

Several issues arise in the implementation of our VAR with daily data. First, since we are interested in lower-frequency dynamics, we want to add many lags. We use the Schwartz Bayes criterion to determine that the optimal lag length is 65 trading days (3 months). To ensure tractability and better out-of-sample behavior, we impose continuity restrictions on the coefficients. Specifically, we aggregate lags of days 2–5, 6–10, 11–21, 22–43, and 44–65, forcing the coefficients within each subperiod to be identical. In this way, we hope to detect

predictability over months, in a manner similar to studies that might run VARs using monthly data.

Second, inflation data are monthly. Given that there is no good fix for this, we simply assume that monthly inflation occurs smoothly through the month. Since inflation shocks are highly persistent, this is not a terrible assumption. However, if, for example, there is a once-and-for-all mid-month surprise increase in the domestic price level, we will not correctly match the timing of the inflation shock with the daily data on exchange rates, flows, and interest rates. With respect to inflation shocks, we will therefore blur cause and effect in the sense of Granger. In order to better align the CPI's with their announcement dates, we lag their entry into the information set by 2 weeks.

Third, we restrict the coefficients on lags past the first of the real exchange rate to be zero in the VAR. Given the high persistence of the real exchange rate, there is no need to include these extra terms. Eliminating them has little effect on the results.

Finally, we calculate all standard errors using the delete-one jackknife methodology of Shao and Wu (1989) and Shao (1989). The jackknife estimator does not require normality and is heteroskedasticity consistent. Figures in the paper, unless labeled otherwise, present ± 2 standard error bounds computed this way.¹¹

II. Data: Foreign Exchange Transactions of Institutional Investors

Our FX transactions data come from State Street Corporation. State Street is one of the world's largest global custodians, responsible for approximately 7 trillion U.S. dollars of institutional-investor portfolios. State Street records all transactions in these portfolios, including foreign exchange, underlying securities, and derivatives.

The data cover foreign-exchange transactions conducted in 111 currencies by 13,230 funds over 1,735 trading days from June 20, 1994 through February 9, 2001. We filter out currencies classified by the IMF as being pegged or fixed, as well as currencies whose transaction flows are relatively sparse on a daily basis. This leaves 18 countries, plus Euroland.¹² See Appendix A for additional data information, including the details of our methodology for valuing forward contracts. Valuation of individual forward contracts is necessary in order to properly aggregate contemporaneous flow values.

Even after removing all but 18 of the 111 currencies, some differences across exchange-rate regimes inevitably remain. However, parameter values across currencies are actually very similar. To provide perspective on this, we

¹¹ Standard errors were also computed using the delta method, with comparable results (not presented).

¹² The currencies (measured against the U.S. dollar) are for Australia, Canada, Euroland, Japan, New Zealand, Norway, Sweden, Switzerland, United Kingdom, Mexico, Indonesia, Korea, Philippines, Singapore, Taiwan, Poland, India, and South Africa. Pre-euro, Euroland flows represent aggregates across the 11 Euroland countries and are paired with the Deutsche mark prior to the introduction of the euro.

consistently report estimates for a subset of six “major” currencies, in addition to the full sample (Australia, Canada, Euroland, Japan, Switzerland, and the United Kingdom). We tend to err on the side of larger panels for three reasons. First, and most importantly, our methodology holds out the prospect of estimating long-horizon effects, and these require larger cross sections to achieve sufficient power. Second, the experiment we have in mind here is one that draws randomly across time and currency to summarize how, on average, currencies behave in response to an unexpected flow, exchange-rate change, or change in interest differential. In future research, it may be possible to implement a more general model with currency-specific coefficients. Finally, it helps to note that in any case our estimates are often not reliably different across the major countries (see below).¹³

We measure returns against the U.S. dollar. However, the U.S. dollar return against, say, the Swiss franc, is not really the appropriate exchange-rate change to compare with net U.S. dollar purchases of Swiss francs. This is because flows into Swiss francs are on average funded—directly or indirectly—with a basket of euros, pounds, yen, and so forth. Unfortunately, it is impossible to observe the basket weights directly, since a single Swiss euro transaction may be effected through several transactions, with the U.S. dollar serving in each of them as vehicle currency. As an alternative approach, we estimate the weights from the slope coefficients in a regression of excess Swiss franc returns on excess returns (against the U.S. dollar) of the pound, euro, yen, and Australian dollar. We then use the residuals from this regression as our measure of excess returns on the Swiss franc. We employ this methodology for all currencies except those of Canada, the United Kingdom, Japan, Australia, and Euroland.

Table I presents some descriptive statistics for the flow data. It shows the mean and *SD* of absolute flows expressed in U.S. dollars. Transaction volumes are, as expected, heavily weighted toward the major currencies. Within the major countries, the euro and yen flows run at about three times the flows for Australia, Canada, and Switzerland. Flows are considerably smaller for other countries, with the possible exception of Sweden and New Zealand.

Table I also reports flow and excess currency return autocorrelations. Flows are strongly autocorrelated, with coefficients averaging approximately 17% for the major currencies and even higher for some of the other Asian currencies.¹⁴ The autocorrelations are quite similar across countries, particularly within the majors. They are smaller than those observed in daily equity flows, perhaps

¹³ This group of countries seemed to us the most natural definition of “major” currencies. On the suggestion of the referee, we considered narrowing the group down to the four largest currencies. However, we found that with such a limited cross section we were unable to detect reversion back toward PPP. Such a finding is consistent with the literature on purchasing power parity, which shows that substantial cross-sectional information is required to detect mean reversion in the real exchange rate (see Froot and Rogoff (1995)).

¹⁴ These autocorrelations are smaller than those observed in daily equity flows, perhaps because foreign exchange markets are more liquid than markets for individual equities. Froot et al. (2001) find that the daily autocorrelation of cross-border equity transactions aggregated to a country level are usually 30% and higher.

Table I
Descriptive Statistics: Foreign Exchange Flows and Currency Returns

The sample period is from June 20, 1994 to February 9, 2001. The flow data are from State Street Corporation (SSC), and US\$ currency excess returns (spot US\$/foreign currency price change less the interest differential) are computed using Datastream data. The first two columns report the mean of absolute daily aggregate flows (a), and the SD of daily net flows in hundreds of millions of US\$. The third column reports the daily partial autocorrelation of aggregate flows. The fourth column reports the second partial autocorrelation under the restriction that days 2–5 have the same coefficient. The fifth column reports the same partial autocorrelation for days 6–10. Columns 6, 7, and 8 report the partial autocorrelations for currency excess returns, with the same coefficient restrictions imposed. Column 9 reports the daily contemporaneous correlation between flows and currency excess returns.

	Flow Moments		Flow Autocorrelations			Return Autocorrelations			Flow–Return Correlation
	$\hat{\mu}(f^a)$	$\hat{\sigma}(f^a)$	$\hat{\rho}_1(f^a)$	$\hat{\rho}_2(f^a)$	$\hat{\rho}_3(f^a)$	$\hat{\rho}_1(r)$	$\hat{\rho}_2(r)$	$\hat{\rho}_3(r)$	$\rho(f^a, r)$
	US\$ 100M	US\$ 100M							
Majors	2.748	3.785	0.170	0.095	0.021	0.042	−0.020	−0.001	0.278
All	2.871	3.935	0.133	0.107	0.057	0.047	−0.029	0.010	0.096
Euroland	1.650	2.260	0.170	0.040	0.008	0.013	0.002	−0.001	0.328
Japan	1.437	2.132	0.188	0.053	0.000	0.047	−0.004	0.007	0.306
United Kingdom	0.886	1.254	0.220	0.009	0.007	0.040	−0.002	−0.010	0.383
Switzerland	0.684	1.038	0.189	0.032	0.001	0.092	−0.033	−0.001	0.258
Canada	0.541	0.912	0.140	0.057	0.001	0.060	−0.022	−0.002	0.145
Australia	0.463	0.657	0.095	0.035	0.014	−0.002	−0.005	−0.002	0.281
Sweden	0.237	0.382	0.130	0.007	0.012	0.049	−0.016	−0.018	0.121
New Zealand	0.141	0.307	0.030	0.018	0.005	−0.068	−0.035	0.005	0.071
Korea	0.103	0.213	0.219	0.049	0.035	0.123	−0.092	0.084	0.024
Singapore	0.091	0.193	0.035	0.037	0.017	−0.019	0.003	−0.027	0.050
Norway	0.070	0.156	0.091	0.006	0.010	0.029	−0.025	−0.013	0.124
Mexico	0.047	0.101	0.119	0.079	−0.003	−0.094	0.015	0.015	0.055
South Africa	0.051	0.092	0.147	0.029	0.039	0.050	−0.002	−0.001	0.091
Taiwan	0.050	0.109	0.033	0.059	0.047	0.163	−0.005	0.007	0.062
India	0.028	0.080	0.004	0.085	0.030	0.113	0.022	−0.005	0.024
Indonesia	0.026	0.044	0.206	0.058	0.010	0.114	−0.003	0.012	0.038
Poland	0.023	0.074	0.066	−0.016	−0.016	−0.023	−0.019	−0.002	0.119
Philippines	0.021	0.035	0.126	0.072	0.042	0.062	−0.054	−0.019	0.102

because foreign exchange markets are more liquid than markets for individual equities (see Froot et al. (2001) for evidence on equities).

Finally, Table I shows that the correlation of daily flows with excess returns is in the 15–30% range for the major currencies. Again, these correlations are similar across countries, particular within the majors that are almost identical, except for Canada. Canada still shows a strong positive correlation, at one-half the size of the other majors. We also find extremely large Canadian flows near the end of the sample.

III. Naïve Estimation and Results

Before we turn to the full multivariate relationship between flows, returns, and intrinsic values, it is useful to provide some statistics that describe the simple bivariate flow/return relationship. Indeed, as it turns out, a considerable portion of what we learn from the bivariate case will apply to the multivariate case. Hence, we estimate a simple correlation between log foreign exchange excess returns and flows

$$\rho_j(K) = \text{cor}(r_{t,j}(K), F_{t,j}(K)), \quad (12)$$

where $r_{t,j}(K) = \sum_{k=1}^K r_{t+1-k,j}$, is the k -period cumulated log excess return on currency j , $r_{t,j}$ is the excess return on the j th currency measured against an appropriate basket of currencies, $F_{t,j}(K) = \sum_{k=1}^K f_{t+1-k,j}$ is the cumulated normalized U.S.-dollar net purchases into currency j , and $f_{t,j} = \text{flow}_{t,j} / \sigma(\text{flow}_j)$ is the normalized currency- j flow, where the normalization is based on own-currency flow *SD*. The normalization helps ensure comparability of flows across countries in the following sections. This also matters when we estimate correlations for currency panels.

We can measure the flow/return correlation, ρ , for any horizon across the panel of currencies. Clearly, this simple correlation coefficient is related to the coefficients derived in Section I. We clarify that relationship in Appendix C.

A. Results from the Naïve Approach

Currency panel estimates of flow–return correlations from equation (12) are reported in Figure 1. We exploit panels throughout the paper, as they give us an enormous advantage: the 18 exchange rates in the full panel allow us to reach maximum horizons of about 7 years. With time series information only, we could make statements of comparable statistical precision for horizons of only a few months. Given that our goal is to distinguish between permanent and transitory returns and to relate these components to flows and intrinsic values, power at long horizons is crucial.

In Figure 1, the x -axis is in log terms, running from 1 day horizons (10^0 days) to 1,735-day horizons ($>10^3$ days). The figure shows 90% confidence bounds derived from the Monte Carlo procedure described in Appendix B. The bounds




Figure 1. Contemporaneous correlation of flows and returns. This figure shows the evolution of the contemporaneous correlation coefficient between flows and returns (vertical axis) at different return horizons (horizontal axis, log scale in days). Results for majors are presented first, followed by those for all countries. The dashed red line is the in sample correlation coefficient computed using overlapping return windows, while the blue lines represent 90th percentile Monte Carlo confidence intervals, estimated as specified in Appendix B. A dotted black line indicates the position of 0 on the vertical axis.

are computed under the null hypothesis that flows and returns are each i.i.d., with a daily contemporaneous correlation given by the actual observed value.

In order to provide a sense for cross-country differences, we employ two separate panels, for major and all 18 currencies, respectively. The results are strongly similar for the panels and, subject to power considerations, for individual currencies as well. There are two important features of the results to emphasize.

First, the correlations start off positive at the 1 day horizon and increase from there. This implies positive noncontemporaneous correlations, which are statistically significant at horizons from approximately 2 to 70 trading days for the majors and 3 to 40 trading days for all currencies, with both peaking at around 20 trading days (about 1 month). Thus, at medium-term horizons, flows must either positively anticipate returns or positively follow recent returns.¹⁵

Second, beyond these medium-term horizons, the overall correlation falls, crossing 0 at about 300 trading days for both panels. This is the point where the positive daily correlation is just offset by *negative* noncontemporaneous correlations. Put differently, at a horizon of about 1.2 years there is approximately no correlation between currency flows and returns. At these horizons, noncontemporaneous correlations are statistically negative, since the dashed line has already crossed the lower confidence bound for the majors (and has nearly done so for all currencies). At still longer horizons, where the cross-sectional effects dominate, the point estimate for the noncontemporaneous correlations is negative. The implication is striking: those currencies that rose in the medium term amidst inflows declined over the longer term back to where they began and perhaps even further.

To supplement the figure, Table II provides estimates of (12) for both panels and the currencies individually. The results for individual countries are essentially the same as those discussed in the panels using Figure 1. This suggests that our panel aggregation does little to obscure important features of the data.

At this stage, one can interpret the facts as suggesting that any impact of flows on currencies is transitory. If this holds up, any information contained in flows is not about intrinsic value *per se*. These naive conclusions are tentative, in that we have not yet separated returns into permanent and transitory components and determined their relationship with flows and intrinsic values. We turn to that next.

IV. Results from the VAR

A. Exchange-Rate Variance Decomposition and Overreaction

Our first result concerns the variance decomposition of currency excess returns, given by equation (9) and estimated in Table III. For both the major

¹⁵ The standard error bands in Figure 1 are calculated based on the observed contemporaneous correlation of 1-day flows and returns, at approximately 28% for the majors and 10% for all currencies. Transcendence of these bands therefore indicates statistically significant noncontemporaneous correlation.

Table II
Country Univariate Correlations, Aggregate Flows
and Excess Returns

This table presents estimates of the contemporaneous correlation coefficient $\hat{\rho}^h$ between flows $F_t(h)$ and currency excess returns $r_t(h)$ at different return horizons h . Flows are summed in an overlapping fashion at each successive return horizon and matched against currency excess returns for the specified period, over the same horizon. Flows are measured as a percentage of own country *SD*, and returns are reported in basis points. The columns are arranged in ascending order of return horizon $h = 1, 5, 20, 60, 120, 240$, and 400 days.

	$\hat{\rho}^1$	$\hat{\rho}^5$	$\hat{\rho}^{20}$	$\hat{\rho}^{60}$	$\hat{\rho}^{120}$	$\hat{\rho}^{240}$	$\hat{\rho}^{400}$
Majors	0.278	0.389	0.428	0.338	0.219	0.081	-0.027
All	0.096	0.150	0.180	0.156	0.070	0.014	-0.055
Euroland	0.328	0.472	0.518	0.401	0.210	-0.036	-0.073
Japan	0.306	0.470	0.563	0.567	0.538	0.396	0.343
United Kingdom	0.383	0.486	0.509	0.356	0.144	0.275	0.467
Switzerland	0.258	0.387	0.359	0.302	0.420	0.438	0.129
Canada	0.145	0.180	0.261	0.307	0.232	0.150	-0.074
Australia	0.281	0.373	0.419	0.373	0.303	0.352	0.094
Sweden	0.121	0.244	0.319	0.402	0.358	0.161	0.005
New Zealand	0.071	0.116	0.146	-0.155	-0.427	-0.583	-0.647
Korea	0.024	0.059	0.121	0.194	0.176	0.206	0.335
Singapore	0.050	0.119	0.279	0.248	0.075	0.083	0.385
Norway	0.124	0.150	0.061	0.133	0.008	-0.238	-0.275
Mexico	0.055	0.205	0.393	0.426	0.277	0.190	0.044
South Africa	0.091	0.102	0.206	0.325	0.158	0.239	0.138
Taiwan	0.062	0.171	0.284	0.418	0.466	0.597	0.756
India	0.024	0.042	0.050	0.005	-0.041	0.037	0.232
Indonesia	0.038	0.073	0.087	0.101	0.023	0.029	-0.038
Poland	0.119	0.232	0.295	0.237	-0.065	-0.296	-0.187
Philippines	0.102	0.112	0.153	0.152	0.122	0.109	0.124

and all countries, we find expected-return shocks are considerably larger than intrinsic-value shocks. Naturally, this is dominated by longer term expected-return shocks (cumulated from $k = 30$ trading days onward), whose *SD* alone is 200% of that of intrinsic-value shocks for the majors. For short-term expected-return shocks (cumulated up to $k = 30$ trading days), the corresponding number is about 9%. Expected-return shocks become even more important once we include emerging markets in the panel. For all countries together, the *SDs* of long-term and short-term expected-return shocks are 247% and 52%, respectively, of the *SD* of intrinsic-value shocks. In both samples, expected-return shocks at long horizons appear to be the dominant source of currency excess returns.

Second, using Table III, we estimate the under/overreaction coefficient of exchange rates to intrinsic-value shocks given in equation (11). For all countries, $\beta_{\text{over}} = 1 - 0.34 + 0.89 = 1.55$. This says that over short horizons, the exchange rate *underreacts* by 34% (compared to the level it reaches in 30 days), whereas over longer horizons, it *overreacts* by 89% to an

Table III
Variance Decomposition

This table shows the components of the variance of excess currency returns. These are estimated using the intrinsic value and expected-return decomposition obtained from our vector autoregression (VAR) estimates. The columns present, in order, the total variance of currency excess returns; the variance of the intrinsic value component of excess returns; the variance of the expected-return component of excess returns; the covariance between the two components, expected return, and intrinsic value; the variance of short horizon expected returns (k signifies 30 trading days); the variance of long horizon expected returns (from $k + 1$ onward); and the covariance of short and long horizon expected returns. These estimates are presented for the major countries first, followed by the estimates for all countries. Standard errors are presented below coefficients in parentheses, and are estimated using the delete-1 jackknife method.

	σ_{fx}^2	σ_{iv}^2	σ_{er}^2	$\sigma_{er,iv}$	$\sigma_{er(1,k)}^2$	$\sigma_{er(k+1,\infty)}^2$	$\sigma_{er(1,k),er(k+1,\infty)}$
Majors	2,804.94 (88.29)	537.27 (109.28)	2,022.12 (1,647.53)	-122.79 (768.75)	4.33 (11.86)	2,079.72 (1,698.51)	-30.96 (177.75)
All	6,704.69 (500.30)	1,047.3 (418.19)	4,514.48 (1,879.72)	-571.47 (756.26)	277.51 (640.22)	6,289.72 (3,185.19)	-1,026.41 (1,843.65)

intrinsic-value shock. The net initial overreaction is therefore 55%. The corresponding numbers are qualitatively similar, though smaller, for the major currencies, where $\beta_{\text{over}} = 1 - 0.04 + 0.27 = 1.23$.

Figure 2 depicts the VAR impulse response of cumulative excess returns of majors and all countries to a 50 basis-point intrinsic-value (permanent) shock. Both graphs show that the exchange rate appreciates immediately in response to the good intrinsic-value news. However, the peak in cumulated excess returns occurs later by about 70 days for the major countries and 12 days for all countries. This implies an additional 3 and 17 basis points in return for majors and all currencies, respectively. After the follow-on increase, there is a noticeable return decline in both groups at horizons of 100 days and more. These estimates suggest that currencies tend initially to overreact to intrinsic-value changes. However, this overreaction would be even greater if there was not a small amount of *underreaction* to 30-day cumulated excess returns.

B. The Flow/Return Relationship: Impulse Responses

Next, we explore the flow response to return shocks. We make use of the return decomposition by examining the flow response to both intrinsic-value (permanent) shocks and expected-return (temporary) shocks. Figure 3 shows the flow response to an appreciation caused by a positive intrinsic-value shock. In the major currencies, there are strong subsequent inflows over short horizons and then, after that at longer horizons, outflows. This pattern lines up with the under/overreaction results in the currency itself. In the full set of currencies, the short-term inflows persist at virtually all horizons. This is consistent with a trend-chasing story—institutional investors buy currencies following a positive permanent shock.




Figure 2. Excess return impulse response to an intrinsic-value shock. This figure shows the response of cumulated excess currency returns (in basis points) to a 50 basis point shock to intrinsic-value news. The response is shown at different return horizons (horizontal axis, log scale in days). Results for majors are presented first, followed by those for all countries. The dashed red line is the response of cumulated excess returns, while the blue lines display ± 2 standard error bounds estimated using the delete-1 jackknife method.




Figure 3. Currency flow impulse response to an intrinsic-value shock. This figure shows the response of cumulated currency flows (in units of own country *SD*) to a 50 basis point appreciation in the currency resulting from a shock to intrinsic-value news. The response is shown at different return horizons (horizontal axis, log scale in days). Results for majors are presented first, followed by those for all countries. The dashed red line is the response of cumulated currency flows, while the blue lines display ± 2 standard error bounds estimated using the delete-1 jackknife method.

Figure 4 shows the response to an appreciation caused by a transitory short-term expected-return shock. (Note the sign of the effect: a *decline* in expected return generates an appreciation.) If investors trend chase, we should see inflows subsequent to the appreciation, even though expected returns have fallen. The figure, however, shows the reverse: when the currency appreciates due to a decline in short-term expected returns, there is a strong *outflow* in the short run. This suggests that institutional investors respond to a short-term transitory appreciation by selling and to a permanent appreciation by buying.

C. Unpacking VAR Relationships

Now we turn to the estimates of price impact, anticipation, trend-chasing, and related effects measured in Table IV. First, the price impact effect in cell (1, 1) shows the correlation of current flow and return surprises to be significant at 28% for the major currencies and 9% for all countries, corroborating the naïve correlations reported above. Second, cell (2, 1) is our estimate of short-term trend chasing—return surprises predicting changes in expected short-run future flows. At 29% and 15%, for majors and all countries, respectively, estimated short-term trend chasing is somewhat larger in magnitude than price impact. Both are highly statistically positive. Third, cells (3, 1) and (4, 1) contain estimates of longer term trend chasing. These estimates are not statistically significant for either the majors or all currencies. Today's surprise appreciations have no apparent relation to changes in long-term expected flows.

Next, we turn to the anticipation effect—the comovement of *expected* future returns with flows. Cell (1, 2) shows that at short horizons, there is positive anticipation (reliably positive for all countries), so that flows positively predict returns over the next 30 trading days intervals. However, cell (1, 3) shows that over longer horizons, the anticipation effect turns reliably *negative*—positive flow surprises result in negative updates of expected returns at horizons greater than 30 trading days. This is in line with the previously identified overreaction effects and negative long-horizon naïve correlations. Note that, for both panels, the sum of cells (1, 2) and (1, 3) is negative. After the initial “impact effect” of a contemporaneous positive flow surprise, the flow surprise actually *reduces* total infinite-horizon expected returns.

Cells (2, 2), (2, 3) and (3, 2), as well as (3, 3) measure the comovement of expected returns and expected flows. Cell (2, 2), for example, reports the comovement between short-term expected flows and short-term expected returns. This is positive for both samples and marginally statistically significant for all countries. Cell (3, 3) reports that the covariance for longer term expected flows with expected returns is not statistically significant.

Finally, we examine the cells in column 4 (which are sums across the preceding three columns). These contain the comovements between intrinsic-value news, $v_{iv,t}$, and flow innovations. None of the estimates in column 4 are statistically different from zero. This says that flow surprises and innovations in expected future flows are not reliably correlated with intrinsic-value shocks. We can interpret cell (4, 4) as a kind of infinite-horizon correlation between




Figure 4. Currency flow impulse response to a short run expected-return shock. This figure shows the response of cumulated currency flows (in units of own country *SD*) to a 50 basis point appreciation in the currency resulting from a shock to short run expected returns. The response is shown at different return horizons (horizontal axis, log scale). Results for majors are presented first, followed by those for all countries. The dashed red line is the response of cumulated currency flows, while the blue lines display ± 2 standard error bounds estimated using the delete-1 jackknife method.

Table IV
Relationships between Flows and Returns

These tables summarize the relationships between flows and returns using the intrinsic value and expected-return decomposition obtained from our vector autoregression (VAR) estimates. Estimates for the major countries are presented first, followed by those for all countries. Here $k = 30$ trading days, for computing short horizon cumulated expectations of the innovations from our model. Each cell contains estimates of the scaled covariances between the elements indicated in the row and column headings. To scale these covariances, we divide each entry in the cells by $(e1' \Sigma e1 e2' \Sigma e2)^{1/2}$, the product of the SDs of the innovations to flows ($e2'u$) and returns ($e1'u$). Standard errors are below coefficients in parentheses, and are estimated using the delete-1 jackknife method.

Flows ↓	Returns →	Unexpected Innovation $e1'u$	Expected Short-Term Innovation $e1'\Phi(k)u$	Expected Long-Term Innovation $e1'(\Phi - \Phi(k))u$	Intrinsic Value Innovation $e1'\Psi u = v_{iv,t}$
Majors					
Unexpected Innovation $e2'u$		0.282 (0.013)	0.037 (0.054)	-0.278 (0.127)	0.04 (0.122)
Expected Short-term Innovation $e2'\Phi(k)u$		0.285 (0.079)	0.02 (0.033)	-0.248 (0.125)	0.057 (0.103)
Expected Long-term Innovation $e2'(\Phi - \Phi(k))u$		-1.213 (0.866)	-0.015 (0.098)	1.103 (0.953)	-0.125 (0.448)
Total Flow Innovation $e2'\Psi u$		-0.647 (0.865)	0.042 (0.099)	0.577 (0.863)	-0.028 (0.242)
All					
Unexpected Innovation $e2'u$		0.094 (0.007)	0.105 (0.047)	-0.168 (0.057)	0.032 (0.036)
Expected Short-term Innovation $e2'\Phi(k)u$		0.151 (0.053)	0.095 (0.049)	-0.203 (0.080)	0.043 (0.042)
Expected Long-term Innovation $e2'(\Phi - \Phi(k))u$		0.673 (0.607)	0.147 (0.230)	-0.723 (0.642)	0.096 (0.197)
Total Flow Innovation $e2'\Psi u$		0.919 (0.620)	0.347 (0.280)	-1.094 (0.673)	0.171 (0.259)

flows and returns. It is, in fact, closely related to the naïve k -period-horizon flow/return correlations calculated above. (Appendix C derives the exact relationship between the two.)

Overall, the results are very similar to the naïve case: they suggest that the long-horizon comovement of flows and returns is essentially zero. At short horizons, flows and returns appear positively correlated: there is evidence of same-day price impact, of short-term trend chasing, and of short-term anticipation. However, at longer horizons, we are not able to detect any relationship between returns and flows. Moreover, intrinsic-value shocks—the permanent components of exchange-rate surprises—are essentially uncorrelated with flow shocks. Flows are correlated only with the transitory component of exchange-rate surprises. This provides strong evidence in support of the weak flow-centric view and against both the strong flow-centric and fundamentals-only views.

Table V
Relationships between Interest Differentials and Returns

These tables summarize the relationships between interest differentials and returns using the intrinsic value and expected-return decomposition obtained from our vector autoregression (VAR) estimates. Estimates for the major countries are presented first, followed by those for all countries. Here $k = 30$ trading days, for computing short horizon cumulated expectations of the innovations from our model. Each cell contains estimates of the scaled covariances between the elements indicated in the row and column headings. To scale these covariances, we divide each entry in the cells by $(e1'\Sigma e1e3'\Sigma e3)^{1/2}$, the product of the *SDs* of the innovations to interest differentials ($e3'u$) and returns ($e1'u$). Standard errors are below coefficients in parentheses, and are estimated using the delete-1 jackknife method.

Returns → Interest Differentials ↓	Unexpected Innovation $e1'u$	Expected Short-Term Innovation $e1'\Phi(k)u$	Expected Long-Term Innovation $e1'(\Phi - \Phi(k))u$	Intrinsic Value Innovation $e1'\Psi u = v_{iv,t}$
Majors				
Unexpected Innovation $e3'u$	0.004 (0.011)	0.001 (0.069)	0.048 (0.062)	0.052 (0.033)
Expected Short-term Innovation $e3'\Phi(k)u$	1.575 (0.799)	0.085 (1.111)	-0.586 (1.296)	1.074 (0.890)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$	20.805 (16.556)	0.264 (1.722)	-19.892 (18.004)	1.177 (8.041)
Total Int. Differential Innovation $e3'\Psi u$	22.384 (16.516)	0.35 (1.965)	-20.43 (18.315)	2.304 (8.782)
All				
Unexpected Innovation $e3'u$	0.003 (0.029)	0.078 (0.095)	0.037 (0.089)	0.118 (0.048)
Expected Short-term Innovation $e3'\Phi(k)u$	1.287 (1.395)	1.645 (1.845)	-0.515 (1.994)	2.418 (1.130)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$	6.202 (10.020)	1.301 (2.865)	-4.758 (9.736)	2.745 (2.880)
Total Int. Differential Innovation $e3'\Psi u$	7.493 (10.401)	3.023 (4.245)	-5.235 (10.124)	5.281 (3.810)

D. The Relationship between Returns and Interest Differentials

The next set of questions pertains to the relationship between returns and interest differentials. How strongly are their innovations related? Which tends to anticipate the other? Are innovations to returns and interest differentials highly correlated? To address these questions, we perform a similar decomposition to Table IV, only for returns and interest differentials instead of for returns and flows.

Table V summarizes the interest-differentials/returns interaction. Cell (1, 1) shows positive, but insignificant, contemporaneous correlation between shocks to interest differentials and shocks to returns. Similarly, cells (2, 1) and (3, 1) show that return shocks are positively correlated with *future* innovations in expected interest differentials, reliably at short horizons, but not at long horizons. Surprise currency appreciations positively

Table VI
Relationships between Interest Differentials and Returns
in a Three-Equation VAR without Flows

These tables summarize the relationships between interest differentials and returns using the intrinsic value and expected-return decomposition obtained from a vector autoregression (VAR) containing only excess returns, interest differentials, and real exchange rates. Estimates for the major countries are presented first, followed by those for all countries. Here $k = 30$ trading days, for computing short horizon cumulated expectations of the innovations from our model. Each cell contains estimates of the scaled covariances between the elements indicated in the row and column headings. To scale these covariances, we divide each entry in the cells by $(e1'\Sigma e1e3'\Sigma e3)^{1/2}$, the product of the SDs of the innovations to interest differentials ($e3'u$) and returns ($e1'u$). Standard errors are below coefficients in parentheses, and are estimated using the delete-1 jackknife method.

Interest Differentials ↓	Returns →	Unexpected Innovation $e1'u$	Expected Short-Term Innovation $e1'\Phi(k)u$	Expected Long-Term Innovation $e1'(\Phi - \Phi(k))u$	Intrinsic Value Innovation $e1'\Psi u = v_{iv,t}$
Majors					
Unexpected Innovation $e3'u$		0.003 (0.011)	0.001 (0.069)	0.049 (0.062)	0.054 (0.033)
Expected Short-term Innovation $e3'\Phi(k)u$		1.591 (0.805)	0.045 (1.119)	-0.486 (1.284)	1.15 (0.873)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$		21.221 (17.070)	0.208 (1.762)	-20.341 (18.561)	1.088 (8.404)
Total Int. Differential Innovation $e3'\Psi u$		22.815 (17.030)	0.255 (1.999)	-20.778 (18.867)	2.292 (9.121)
All					
Unexpected Innovation $e3'u$		0.003 (0.029)	0.079 (0.096)	0.036 (0.089)	0.118 (0.049)
Expected Short-term Innovation $e3'\Phi(k)u$		1.272 (1.393)	1.639 (1.863)	-0.486 (2.010)	2.425 (1.145)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$		5.974 (10.008)	1.437 (2.777)	-4.679 (9.659)	2.732 (2.874)
Total Int. Differential Innovation $e3'\Psi u$		7.249 (10.383)	3.156 (4.174)	-5.129 (10.038)	5.276 (3.812)

anticipate future increases in real interest differentials, especially over short horizons.

Second, innovations in interest differentials are positively, though not reliably, correlated with expected future returns (cells (1, 2) and (1, 3)). Moreover, cells (2, 2) and (3, 2) suggest that changes in expected future interest differentials comove positively with changes in expected future returns. In other words, innovations in real interest differentials have some forecasting power for short-term currency movements. These results are individually statistically insignificant, but collectively are significant when cumulated (see column 4). There we report that intrinsic-value shocks are positively correlated with current, short-term expected, and long-term expected innovations in interest differentials. All the estimates are positive, and most are statistically significant for the full panel of all countries, where we have greater precision.

Table VII
Relationships between Interest Differentials and Flows

These tables summarize the relationships between interest differentials and flows using the intrinsic value and expected-return decomposition obtained from our vector autoregression (VAR) estimates. Estimates for the major countries are presented first, followed by those for all countries. Here $k = 30$ trading days, for computing short horizon cumulated expectations of the innovations from our model. Each cell contains estimates of the scaled covariances between the elements indicated in the row and column headings. To scale these covariances, we divide each entry in the cells by $(e2'\Sigma e2e3'\Sigma e3)^{1/2}$, the product of the *SDs* of the innovations to interest differentials ($e3'u$) and flows ($e2'u$). Standard errors are below coefficients in parentheses, and are estimated using the delete-1 jackknife method.

Flows → Interest Differentials ↓	Unexpected Innovation $e2'u$	Expected Short-Term Innovation $e2'\Phi(k)u$	Expected Long-Term Innovation $e2'(\Phi - \Phi(k))u$	Total Flow Innovation $e2'\Psi u$
Majors				
Unexpected Innovation $e3'u$	0.009 (0.008)	-0.056 (0.090)	-0.056 (0.140)	-0.103 (0.146)
Expected Short-term Innovation $e3'\Phi(k)u$	1.94 (0.900)	0.164 (1.500)	-2.814 (3.324)	-0.71 (3.583)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$	6.728 (6.750)	6.334 (5.757)	-29.508 (40.116)	-16.447 (30.424)
Total Int. Differential Innovation $e3'\Psi u$	8.677 (6.754)	6.441 (5.881)	-32.377 (42.365)	-17.26 (32.955)
All				
Unexpected Innovation $e3'u$	0.002 (0.006)	0.012 (0.044)	-0.005 (0.152)	0.009 (0.166)
Expected Short-term Innovation $e3'\Phi(k)u$	0.639 (0.424)	0.777 (0.873)	0.956 (3.081)	2.373 (3.445)
Expected Long-term Innovation $e3'(\Phi - \Phi(k))u$	-0.996 (2.175)	-0.238 (2.362)	2.902 (9.086)	1.668 (12.892)
Total Int. Differential Innovation $e3'\Psi u$	-0.354 (2.312)	0.551 (2.655)	3.853 (10.760)	4.05 (14.441)

Third, the cells in column 4 show covariations between innovations in interest differentials and permanent exchange-rate shocks. All the covariances are positive, and, for the most part, statistically positive for all countries. This provides evidence that at least one measure of exchange-rate fundamentals—real interest differentials—seems positively related to the *permanent* component of exchange changes.

Table VI contains analogous estimates of the interest-differential/return relation, but in a three-equation VAR excluding flows. The pattern of the estimates is extremely similar to that in Table V. This provides further evidence that there is little in the relationship between interest differentials and returns in Table V that is driven by flows.

E. The Relationship between Interest Differentials and Flows

Table VII uses the same template to summarize the interaction between interest differentials and flows. We make two important points. First, cell

(2, 1) shows that current flow shocks positively predict future changes in interest differentials, in both panels at short horizons. Short-run changes in interest differentials are therefore anticipated by both exchange-rate changes (Tables V and VI) and flows (Table VII). Second, no other cell is statistically significant, including the infinite horizon covariance in cell (4, 4). Thus, at long horizons, innovations in flows and interest differentials seem unrelated, even while innovations in exchange rates and interest differentials are positively linked. This provides further evidence against the strong-flow-centric view.

V. Conclusion

Our goal in this paper is to understand the interaction between the currency flows of institutional investors and currency returns, controlling for the exchange rate's intrinsic value. To do this, we decompose returns into intrinsic-value and expected-return components, which represent, respectively, the permanent and transitory components of excess returns. We then proceed to determine whether flows comove predominantly with the intrinsic-value or expected-return portions.

We find that flows are strongly positively related to the expected-return component at short horizons, but negatively related at longer horizons. Specifically, a positive shock to today's flows results in positive revisions of expected 30-day currency returns, but in even larger negative revisions of longer horizon expected returns. Altogether, inflows have no lasting positive impact on expected returns, and may even have a negative impact. Indeed, we find that a flow surprise is essentially unrelated to the permanent component of a contemporaneous total return shock.

Furthermore, when we look specifically for comovement at long horizons between the main driver of intrinsic value—real interest differentials—and flows, we find none. This negative result does not seem to be driven by lack of power: we find that real interest differentials are in fact positively related to the permanent component of returns. In other words, at least one important exchange-rate fundamental is linked to the permanent component of currency returns, as theory would suggest. This same important fundamental is essentially unrelated to flows. Taken together, these facts are consistent with what we call the “weak flow-centric view”: flows have short-term impacts on returns, but these do not persist, and are even reversed over the long term.

The conclusions are very much the same using a far simpler approach that measures the simple flow/return correlation, but does not attempt to control for fundamentals. That is, the simple correlation between flows and returns is positive at short horizons of 1 day or 1 month (where it peaks), or even 1 quarter. However, when the return horizon is extended to 9 months or 1 year and beyond, we find that the return/flow correlation drops rapidly to zero. Flows seem to result in momentum effects at shorter horizons and roughly offsetting reversal effects at longer horizons.

Within these broad-based results, we also find interesting details. For example, there is some evidence of anticipation—flows being positively correlated

with future excess returns—and considerable evidence of trend chasing—excess returns being positively correlated with future flows. However, when we divide excess returns into permanent and transitory components, we find that trend chasing breaks into two opposing effects. That is, investors strongly purchase a currency subsequent to a permanent positive return, but sell it subsequent to a transitory positive return.

In addition, our return decomposition provides evidence that transitory shocks are large relative to permanent shocks. This is similar to what has been found for equity index returns. We also find that exchange-rate shocks generally exceed their permanent component initially—that is, they overreact—and go on to exceed them by even more over the following 30 trading days. Overreaction also reconciles our results with the finding in Evans and Lyons (2003) that at very short horizons (1 day), news announcements affect the exchange rate through flows. Our results suggest that, at such horizons, fundamental news *does* affect the exchange rate through flow, but only through the relatively large transitory component of return (which we find is related to flow), not the permanent component (which we find is unrelated to flow).

Lastly, we find that the permanent component of currency returns covaries positively with cumulated shocks to real interest differentials. Yet, these same shocks appear almost uncorrelated with the total return (i.e., the sum of permanent and transitory return components). Our decomposition therefore filters out a large amount of transitory noise, making it possible to detect a positive relationship between measures of intrinsic value and returns. These findings would seem to help explain the long-standing difficulty of detecting a relationship between currency values and common fundamentals, except at long horizons (see, e.g., Meese and Rogoff (1983) and Mark (1995)).

Putting these pieces together, we find no support for a strong flow-centric view. Current and future flow shocks do not have much information (except over very short horizons) about current and future innovations in real exchange rates, interest differentials, or intrinsic-value shocks. Flows seem best in explaining transitory excess returns, such as short-run underreaction and long-run overreaction. As a result, the evidence seems to best fit the weak version of the flow-centric view. At long horizons, we find support for the fundamentals-only view: shocks to interest differentials matter and flows do not.

Appendix A: Computing Transaction Present Values and Flows

In the foreign exchange transactions data, each currency transaction reported indicates future value amounts bought and sold (these are forward as well as spot FX contracts). To aggregate across trades on a given trade date, we compute the present value of currency bought and sold. To do so, we discount both sides (currency bought and currency sold) of each forward transaction by the interest rate corresponding to both the time interval between trade and settlement dates and the currency.

We use the following formula for computing present values (where the operator $\lfloor x \rfloor$ rounds x down to the nearest integer):

$$PV_t^c = \tau_t^c A_{t,m}^c, \quad \text{where } \tau = \left((1 + i_{m,t}^c)^{\lfloor m/T^c \rfloor} (1 + (i_{m,t}^c)(m - \lfloor m/T^c \rfloor)) \right)^{-1}. \quad (\text{A1})$$

The interest rate for currency c and maturity m is given by $i_{m,t}^c$. The discount factor, τ_t^c , is applied to the amount transacted forward, $A_{t,m}^c$, in currency c at time t , in local currency, settling at time $t + m$. The value $A_{t,m}^c$ can be either positive (currency bought) or negative (currency sold).

In transactions for which interest rates in one currency are unavailable, we convert the PV of the other currency in the transaction (assuming interest rates are available for that currency) at that day's 11:00 a.m. daily spot rate. For transactions that do not match the maturities of available interest rates, we linearly interpolate or extrapolate.

We also filter out transactions that are likely to have important data entry or interest-rate-measurement errors. Specifically, we remove transactions where the difference between the PV amounts bought and sold is greater than \$1,000,000 or where the percentage difference is greater than 30%. We exclude transactions in the sparsely traded currencies of Brazil, Kenya, Luxembourg, Peru, Russia, Turkey, and Zimbabwe. After applying these filters, we are left with a total of 6,402,392 transactions across all maturities.

To get net flows on date t , we then aggregate transaction present values by currency, summing all signed present values of trades from date t .

A. 1. Interest Rate, Exchange Rate, and Inflation Data

Interest rate data (daily interbank interest rates for various horizons from 1 week to 1 year) are from Datastream. When interbank rates are unavailable, we use corporate and/or deposit rates, and, if necessary, country treasury rates. We use swap or sovereign bond rates at various maturities from 1 to 25 years.

Spot currency rates are 11:00 a.m. EST rates from WMR/Reuters as far back as possible. Prior to that, we use Datastream or Barclays Bank data.

Inflation is constructed from differences in monthly log Consumer Price Index (CPI) from the international financial statistics of the IMF. We construct a daily inflation series from the monthly series by assuming inflation occurs smoothly within the month.

Appendix B: Constructing Standard Errors for Naïve Correlation Estimates

Standard errors for naïve correlation estimates are by Monte Carlo. These correlation estimates are computed between flows and returns for the panel of N countries and T days, $\rho(K) = \text{cor}(r_{j,t}(K), F_{j,t}(K))$, where $j = 1, \dots, N$ and $t = 1, \dots, T$ at different return horizons K . We compute correlations for all horizons K between 1 and T , so that for given K , we use (NT/K) observations,

where each observation is the sum of K time series elements in each country. When $K = T$, the panel correlation is an $N \times 1$ cross section, with each observation of flows representing the total net inflow over the entire sample period for a country and each observation of returns representing the total excess currency return over the same period for the corresponding country. When $K = 1$, $\rho(K) = \rho(1) = \text{cor}(r_{j,t}, F_{j,t})$, which is the daily panel correlation, computed using NT observations for both flows and returns.

To compute Monte Carlo standard errors, we assume flows and returns are multivariate normal, with moments derived from our daily panel of flows and returns. Define $\mu_f = \frac{1}{NT} \sum_{j=1}^N \sum_{t=1}^T f_{j,t}$, $\mu_r = \frac{1}{NT} \sum_{j=1}^N \sum_{t=1}^T r_{j,t}$, and $\sigma_f^2, \sigma_r^2, \sigma_{fr}$, respectively, as the variance of daily flows computed across the panel, the variance of returns computed similarly, and the covariance between the daily panels of flows and returns. We draw ω_f and ω_r as simulated flows and returns, respectively, each an $NT \times 1$ vector, from the multivariate normal distribution:

$$\begin{pmatrix} \omega_f \\ \omega_r \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_f \\ \mu_r \end{pmatrix}, \begin{pmatrix} \sigma_f^2 & \sigma_{fr} \\ \sigma_{fr} & \sigma_r^2 \end{pmatrix} \right). \quad (\text{B1})$$

We treat these draws as our flow and return data, and compute correlations at different return horizons K . We draw 10,000 such samples from the distribution and generate panel correlations for each return horizon using these draws. The results are then sorted to generate point-wise confidence intervals for our sample estimates at each return horizon.

Appendix C: Connecting the VAR with the Naïve Approach

Using the VAR, it is possible to derive the R^2 statistics from the naïve approach. In addition, we can use the VAR to be more precise about low-frequency comovements.

We can approximate the K -period return and flow, respectively, as

$$r_{t,j}(K) = \sum_{k=1}^K r_{t+1-k,j} \approx \sum_{k=1}^K e1' \Psi(k) u_{t+1-k}, \quad (\text{C1})$$

$$F_{t,j}(K) = \sum_{k=1}^K f_{t+1-k,j} \approx \sum_{k=1}^K e2' \Psi(k) u_{t+1-k}. \quad (\text{C2})$$

The approximation occurs because we assume that before the start of the sample, we have $u_t = 0$, for all $t \leq 0$. As a result, the conditional expectation of future returns and flows is built up entirely from innovations that occur during the sample period.

With these assumptions, and noting that the u 's are uncorrelated at all leads and lags, we can write the correlation coefficient between $r(K)$ and $F(K)$ as

$$\rho(K) = \frac{\sum_{k=1}^K e1'\Psi(k)\Sigma\Psi(k)'e2}{\left(\sum_{k=1}^K e2'\Psi(k)\Sigma\Psi(k)'e2\right)^{1/2} \left(\sum_{k=1}^K e1'\Psi(k)\Sigma\Psi(k)'e1\right)^{1/2}}. \quad (C3)$$

This expression reminds us that multiperiod returns and flows combine news from various periods. In a three-period return, for example, first-period news and its impact on second-period and third-period expected outcomes is combined with later news, including, for example, third-period news, where any impact on later period expectations is expressed only in subsequent third-period observations. To isolate the multiperiod comovement attributable to current-period news, it is more appropriate to use the VAR impulse response. This produces a total K -period correlation between returns and flows from current-period news described by

$$\rho(K) = \frac{e1'\Psi(K)\Sigma\Psi(K)'e2}{(e1'\Psi(K)\Sigma\Psi(K)'e1)^{1/2} (e2'\Psi(K)\Sigma\Psi(K)'e2)^{1/2}}. \quad (C4)$$

REFERENCES

- Barberis, Nicholas, and Andrei Shleifer, 2003, Style investing, *Journal of Financial Economics* 68, 161–199.
- Baxter, Marianne, 1994, Real exchange rates and real interest differentials: Have we missed the business-cycle relationship?, *Journal of Monetary Economics* 33, 5–37.
- Beveridge, Stephen, and Charles R. Nelson, 1981, A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to the measurement of the Business Cycle, *Journal of Monetary Economics* 7, 151–174.
- Brennan, Michael, and H. Henry Cao, 1997, International portfolio investment flows, *Journal of Finance* 52, 1851–1880.
- Cai, Jun, Yan-Leung Cheung, Raymond S. K. Lee, and Michael Melvin, 2001, Once-in-a-generation yen volatility in 1998: Fundamentals, intervention, and order flow, *Journal of International Money and Finance* 20, 327–347.
- Campbell, John Y., 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3–37.
- Campbell, John Y., and Richard Clarida, 1987, The dollar and real interest rates, *Carnegie-Rochester Conference Series on Public Policy* 27, 103–139.
- Campbell, John Y., and Robert Shiller, 1987, Cointegration and tests of present value models, *Journal of Political Economy* 95, 1062–1088.
- Campbell, John Y., and Robert Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Cheung, Yin-Wong, and Menzie D. Chinn, 2001, Currency traders and exchange rate dynamics: A survey of the U.S. market, *Journal of International Money and Finance* 20, 439–471.

- Choe, Hyuk, Bong-Chan Kho, and Rene M. Stulz, 2001, Do domestic investors have more valuable information about individual stocks than foreign investors?, Working paper 8073, NBER.
- Clarida, Richard, and Jordi Gali, 1994, Sources of real exchange rate fluctuations: How important are nominal shocks?, *Carnegie-Rochester Conference Series on Public Policy* 41, 1–56.
- Cohen, Randy, Paul Gompers, and Tuomo O. Vuolteenaho, 2002, Who underreacts to cashflow news? Evidence from trading between individuals and institutions, *Journal of Financial Economics* 66, 409–462.
- Cumby, Robert E., and John Huizinga, 1991, The predictability of real exchange rate changes in the short and long run, *Japan and the World Economy* 3, 17–38.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 1998, Investor psychology and security market under- and over-reactions, *Journal of Finance* 53, 1839–1886.
- DeBondt, Werner F. M., and Richard H. Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.
- Engel, Charles, and Kenneth D. West 2003, Exchange rates and fundamentals, European Central Bank working paper no. 248.
- Evans, Martin D. D., and Richard K. Lyons, 2002, Order flow and exchange rate dynamics, *Journal of Political Economy* 110, 170–180.
- Evans, Martin D. D., and Richard K. Lyons, 2003, How is macro news transmitted to exchange rates?, Working paper no. 9433, NBER.
- Froot, Kenneth A., Paul O'Connell, and Mark S. Seasholes, 2001, The portfolio flows of international investors, *Journal of Financial Economics* 59, 151–193.
- Froot, Kenneth A., and Tarun Ramadorai, 2001, The information content of international portfolio flows, Working paper no. 8472, NBER.
- Froot, Kenneth A., and Kenneth Rogoff, 1995, Perspectives on PPP and long-run real exchange rates, in Gene Grossman and Kenneth Rogoff, eds.: *Handbook of International Economics*, Vol. 3 (Elsevier Science Publishers B.V., Amsterdam).
- Griffin, John M., Jeffrey H. Harris, and Selim Topaloglu, 2003, The dynamics of institutional and individual trading, *Journal of Finance* 58, 2285–2320.
- Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.
- Jegadeesh, Narasimhan, and Sheridan Titman, 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *Journal of Finance* 56, 699–720.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1335–1355.
- Mark, Nelson, 1995, Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review* 85, 201–218.
- Meese, Richard, and Kenneth Rogoff, 1983, Empirical exchange rate models of the seventies, *Journal of International Economics* 14, 3–24.
- Meese, Richard, and Kenneth Rogoff, 1988, Was it real? The exchange rate-interest differential relation over the modern floating rate period, *Journal of Finance* 43, 933–948.
- Rime, Dagfinn, 2001, U.S. exchange rates and currency flows, Working paper no. 4, Stockholm Institute for Financial Research.
- Rouwenhorst, K. Geert, 1998, International momentum strategies, *Journal of Finance* 53, 267–284.
- Shao, Jun, 1989, The efficiency and consistency of approximations to the jackknife variance estimators, *Journal of the American Statistical Association* 84, 114–119.
- Shao, Jun, and C. F. Jeff Wu, 1989, A general theory for jackknife variance estimation, *Annals of Statistics* 17, 1176–1197.
- Vuolteenaho, Tuomo O., 2002, What drives firm-level stock returns?, *Journal of Finance* 57, 233–264.
- Wei, Shang-Jin, and Jungshik Kim, 1997, The big players in the foreign exchange market: do they trade on information or noise?, Working paper no. 6256, NBER.