



Research Division
Federal Reserve Bank of St. Louis
Working Paper Series



**The Dynamic Interaction of Trading Flows, Macroeconomic
Announcements and the CAD/USD Exchange Rate:
Evidence from Disaggregated Data**

**Nikola Gradojevic
and
Christopher J. Neely**

Working Paper 2008-006C
<http://research.stlouisfed.org/wp/2008/2008-006.pdf>

January 2008
Revised August 2009

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

The Dynamic Interaction of Trading Flows, Macroeconomic Announcements and the CAD/USD Exchange Rate: Evidence from Disaggregated Data

Nikola Gradojevic*

Christopher J. Neely†

August 14, 2009

Abstract: We explore the relationship between disaggregated trading flows, the Canada/U.S. dollar (CAD/USD) market and U.S. macroeconomic announcements with a novel data set of unprecedented breadth and length. Foreign financial trading flows appear to demand liquidity, contemporaneously driving the CAD/USD while commercial trading flows seem to be price sensitive, providing liquidity in response to exchange rate movements. Despite strong contemporaneous correlations with trading flows, exchange rate returns are generally not predictable, except for some intriguing success at long horizons. This failure contrasts with much, but not all, previous research on the topic. While two types of CAD trading flows and the CAD/USD appear to be cointegrated, such structure is probably spurious. There appear to be structural breaks in the order-flow-exchange rate VECM systems in 1994-1996 and 1998-1999.

Keywords: Foreign exchange; Order flow; Market microstructure; Macroeconomic announcements.
JEL Classification: F31; F37; C32.

* Faculty of Business Administration,
Lakehead University
955 Oliver Road, Thunder Bay,
Ontario P7B 5E1, Canada;
Tel: 1 807 3438419; Fax: 1 807 3438443;
nikola.gradojevic@lakeheadu.ca

† Assistant Vice President
Research Department
Federal Reserve Bank of St. Louis
Tel.: 1 314 4448731; Fax: 1 314 4448731;
neely@stls.frb.org

The authors thank Chris D'Souza, Jasmina Hasanhodzic, Michael King, Brian Lucey, Carol Osler, Mark Taylor and seminar participants at the Missouri Economics Conference, the Bank of Canada, the Applied Economics Ruby Anniversary Editors' Conference, the International Federation of Technical Analysts and the Midwest Finance Association 2009 Meetings for helpful comments. Nikola Gradojevic gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada. The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

1. Introduction

Lyons' (1995) seminal work has inspired a large literature on the relationship of order flow and asset prices. Much of this literature has studied exchange rates and some has focused on three issues: the identity of demanders/providers of liquidity, the ability of order flow to predict exchange rates and the joint response of asset prices and order flow to macroeconomic announcements. In addition to those issues, this paper investigates long-run relations and the stability of those relations with a novel data set pertaining to the CAD/USD market.

Factors, such as equity market developments or need for commercial transactions, create order flow that can be loosely considered an "exogenous" driver of the exchange rate (Lyons (2001)). While market-making banks provide immediate liquidity to these customers, they have limited appetite to accumulate inventory and so liquidity must ultimately come from other, price-sensitive customers, who respond strongly to exchange rate fluctuations.

The first issue is to identify liquidity buyers/sellers in order to understand their motivations and the forces driving exchange rates.¹ A handful of papers have studied this issue: Bjørnnes, Rime and Solheim (2005), for example, find that non-financial customers passively provide liquidity in the SEK/EUR market while financial customers take an active role. Likewise, Marsh and O'Rourke (2005) argue that commercial order flow is price sensitive. Boyer and van Norden (2006) note that the price responsiveness of commercial order flow contrasts with the usual predictions of the microstructure literature.

The second issue is whether order flow predicts exchange rates. This question is important both for scientific and practical reasons. From a scientific point of view, the microstructure literature holds that order flow reveals private information, including expectations of long-term macro fundamentals, which should forecast exchange rate movements. Therefore the forecasting performance of types of order flow are important to validate or challenge existing microstructure theories. From a practical point of view, the

¹ The literature describes exogenous order flow as buying liquidity while the price-sensitive, endogenous order flow is said to sell it. Osler (2008) very nicely reviews the literature on foreign exchange microstructure.

Bank of Canada might wish to forecast exchange rate movements for policy purposes, or private traders might wish to use their own trading/order flow data to inform trading rules or quoting behavior.

The predominant view of the literature is that order flow does forecast exchange rates. Evans and Lyons (2002a, 2002b, 2005b) and Gradojevic (2007a) have documented short-term predictive relationships while Killeen, Lyons, and Moore (2006) find order flows to be informative at medium (up to six months) and long horizons. Other researchers, such as Danielsson, Payne and Luo (2002), Berger et al. (2008), and Sager and Taylor (2008), have dissented from the view that order flows predict exchange rates, at least for the data sets they studied.

The third issue is to evaluate the joint response of order flow and exchange rates to macroeconomic surprises. Do exchange rates simply react directly to such surprises or do such events also prompt agents to revise their expectations and reallocate their portfolios in a way that reveals private information? Many researchers, including Simpson, Ramchander, and Chaudhry (2005) and Faust et al. (2003), have studied how exchange rates respond to the unexpected component of macroeconomic announcements. Several groups of authors—Han and Kling (1999), Christie-David and Chaudhry (2000), and Doukas and Switzer (2004)—have specifically studied the effect of macro announcements on the CAD/USD. Similarly, Hayo and Neuenkirch (2009) study the effects of monetary policy communication and macro news on financial markets, including the CAD/USD.

But there has been much less study of the joint response of order flow and exchange rates. Evans and Lyons (2002c) find that DEM/USD order flow has a greater influence after announcements while Evans and Lyons (2005a) discover sustained effects of announcements with DEM/EUR/USD data. Love and Payne (2008) and Evans and Lyons (2008) find that order flow largely mediates the effect of announcements on exchange rates. Using USD/EUR, GBP/EUR and USD/GBP data, Love and Payne (2008) find additional order flow and a greater price impact after news releases while Evans and Lyons (2008) use an innovative heteroskedasticity-based identification scheme to identify structural effects in DE/USD data. Savaser (2006) shows that price contingent orders in the USD/GBP market help explain

price reaction to announcements. These papers have studied the impact of news announcements on order flows and exchange rates in the DEM/EUR/USD and USD/GBP markets, using data sets that contain a single bank's order flow over a period of four months to six years. But we are unaware of any previous examinations of how CAD trading flows react to announcements.

This paper considers the effect of macro surprises, the identity of liquidity providers, and the predictive content of Canadian trading flow data— net transactions of banks by customer type—with a unique data set from the Bank of Canada, collected from six major dealing banks in Canada, that spans 15 years of daily observations (1990-2004). Unlike order flow, trading flow data are not specifically signed according to which party initiates the transaction. Nevertheless, to the extent that banks are passive market makers and do not initiate transactions—at least with certain types of customers—trading flows should behave similarly to order flows. This disaggregated data set is much longer and contains a much higher proportion of transactions volume than previously used order flow data sets.² These advantages provide an unparalleled view of the Canadian market, much greater power to confirm or reject hypotheses and the ability to study the stability of relations over time.

To presage our results, trading flow-exchange rate correlations are consistent with the view that foreign financial traders demand liquidity, which commercial and interbank trading flow supplies. These relationships are fairly stable over time and are accompanied by fairly good in-sample predictability. Neither ex ante long-horizon or vector error correction (VECM) methods reliably forecast exchange rates out-of-sample by statistical criteria, however. Although long-horizon regressions have some intriguing successes, this might be fortuitous. In contrast to the inconsistent evidence of exchange rate predictability, the VECM can reliably forecast trading flows at short-horizons. We have seen very little previous investigation of the predictability of order/trading flows.

To supplement our statistical forecast metrics, we follow Gradojevic (2007a) and Rime, Sarno, and

² We calculate that our six banks had 41 to 61 percent of the market, using the trading volumes from the BIS Triennial Surveys for 1998, 2001 and 2004. In contrast, for example, Evans and Lyons (2005) rely on 10-15% of the total volume in USD/EUR (USD/DEM) over 1993-1999.

Sojli (2007) in evaluating the profitability of trading signals from lagged trading flows and exchange rates. This exercise complements the innovative work of Schulmeister (2006) and Rime, Sarno, and Sojli (2007), who studied the interaction of order flow and trading rules. Using the Evans and Lyons (2002a) order flow data, Schulmeister found that technical trading rule signals explain 30% of the variance of DEM/USD and JPY/USD order flow changes. Rime, Sarno, and Sojli (2007) use EUR/USD, GBP/USD and JPY/USD tick and order flow over a one-year period to construct a profitable trading rule. In contrast to the Rime, Sarno, and Sojli (2007) results, we find that trading rules based on trading flows could not have produced profits in real time; our results are consistent with a semi-strong form of market efficiency.

The length of our data set permits us to investigate long-run relationships with far greater power than previous investigations and to study their structural stability. In summary, the VECM study shows cointegrating relationships between cumulative foreign financial, Canadian financial and commercial trading flows and cumulative returns of the CAD/USD using fifteen years (1990-2004) of daily order flow data. In contrast to Killen, Lyons and Moore (2006), who find that order flow is strongly exogenous, no variable is even weakly exogenous in our preferred three-variable VECM. The long-run structure appears to be unstable, however. There is strong evidence of structural breaks in the VECM system for most of the data set but especially around 1994 and late 1998-1999. Further work—omitted for brevity—shows that the trading flow equations are the source of the instability.

The next section of the paper describes the data. Section 3 studies the contemporaneous relations between the trading flows and the exchange rate while Section 4 describes the forecasting exercises. Section 5 examines the structural stability of the systems and Section 6 concludes.

2. Data description

The Bank of Canada coordinates and the Canadian Foreign Exchange Committee compiles the trading flow dataset, which accounts for approximately 41- to 61-percent of all CAD/USD transactions over the period 1990-2004 (1994-2004 for Canadian-domiciled investment transactions). Trading flows are the net purchases-less-sales of USD by customers of the top six Canadian commercial banks. Other

things equal, positive trading flow should raise the CAD/USD spot closing rates from the Bank of Canada, appreciating the USD. Foreign exchange rates also influence trading flows, however. If bank customers are price sensitive, then a rising CAD/USD exchange rate—which makes U.S. goods and services relatively more expensive—discourages purchases of USD and reduces net trading flow.³

The trading flows are disaggregated by type of customer as follows: Commercial client transactions (denoted CC) include all transactions with resident and non-resident non-financial customers; Canadian-domiciled investment transactions (denoted CD) include all transactions with non-dealer financial institutions located in Canada; Foreign institution transactions (denoted FD) include all transactions with foreign financial institutions, such as dealers, pension funds, mutual funds and hedge funds; Interbank transactions denoted (IB) pertain to other Canadian-domiciled financial institutions, such as chartered banks, credit unions, investment dealers, and trust companies.

The commercial (CC) transactions are motivated by trades in real goods and services, while the foreign financial (FD) and Canadian financial (CD) transactions are motivated by international portfolio considerations. Transactions between the top six banks would be netted out of our data and therefore the interbank (IB) trading flows represent smaller Canadian depository institutions' excess purchases from the top six banks, which probably reflect the smaller banks' commercial clients' net orders for USD.

As discussed previously, the trading flows are equivalent to order flows to the extent that the top six Canadian commercial banks are market makers, passively responding to orders by altering quotes rather than by initiating transactions. This reasonably describes the way that large banks usually transact with non-dealer financial, commercial and smaller bank clients. The possibility of some bank-initiated transactions in our data, at least for some client categories, means that trading flows only approximate order flows. Of course, all order flow data sets only approximate true order flow because all such data sets only contain a small fraction of any category of total order flow. For example, Evans and Lyons

³ In the interests of brevity, we omit graphs and summary statistics on the data in this paper but we describe them more fully online via http://research.stlouisfed.org/econ/cneely/Data_Appendix_The_Dynamic_Interaction.pdf.

(2005b) rely on 10-15% of the total transactional volume over the 1993-1999 period. We believe that the very broad coverage and length of the Bank of Canada trading flow data make it a very good approximation to true order flow.

We examine how exchange rate and trading flows react to the surprise component of 26 U.S. macro announcements: business inventories, capacity utilization, consumer confidence, construction spending, CPI, consumer credit, advance durables, new orders at factories, fed funds target, advance real GDP, preliminary real GDP, final real GDP, housing starts, initial claims, industrial production, leading indicators, manufacturing composite index, non-farm payrolls, new home sales, PCE, personal income, PPI, retail sales, retail sales ex vehicles, trade balance, and the U.S. government fiscal surplus/deficit. The surprise component for the i^{th} shock ($i=1, \dots, 26$) at time t , $n_{i,t}$, is the standardized difference between the actual release value and the Money Market Services (MMS) expectations of the announcements:

$$(1) \quad n_{i,t} = (\text{Actual}_{i,t} - \text{Forecasted}_{i,t})/\sigma_i,$$

where σ_i is the standard deviation of the difference between the actual and forecasted value for the i^{th} announcement.⁴ This standardization allows us to directly compare regression coefficients on the shocks.

3. Correlations between the exchange rate and trading flows

A number of authors have examined the contemporaneous relationship between exchange rate changes and order flows—e.g., Evans and Lyons (2002a, 2002b), Danielsson, Payne, and Luo (2002), Payne (2003), Fisher and Hillman (2003), Marsh and O’Rourke (2005), Froot and Ramadorai (2005), Bjønnes, Rime and Solheim (2005).⁵ Researchers often simply regress exchange rate returns on contemporaneous order flow—as Sager and Taylor (2008) clearly explain—often finding a strong relation

⁴ While one could add Canadian announcements to the list, Ito and Roley (1987), Laakkonen (2004) and Ehrmann and Fratzscher (2005) provide evidence that U.S. announcements tend to have larger effects on dollar rates than their foreign counterparts.

⁵ Carpenter and Wang (2003) and Gradojevic and Yang (2006) consider nonlinear contemporaneous relationships. Other researchers look at determinants of correlations, such as time of day (Osler (2002)) or the permanence of returns (Froot and Ramadori (2005)), or the type of order flow (Bjønnes, Rime, and Solheim (2003), Osler (2002, 2003)) or the impact of news (Love and Payne (2008)).

between the variables. Evans and Lyons (2002a) claim that contemporaneous order flows explain 64 percent of daily variation in DEM returns. Evans (2007) makes strong claims about these correlations:

“This contemporaneous relationship between depreciation rates and interdealer order flows appears robust. It holds for many different currencies, and for different currency-order flow combinations (e.g., Evans and Lyons 2002b, Payne 2003 and Froot and Ramadorai 2005). It is also worth emphasizing that order flow’s impact on spot rates is very persistent.” —Evans (2007), page 6.

Generally speaking, both the effect of trading flows on the exchange rate and the feedback from exchange rates to trading flows determine their contemporaneous correlation. These correlations can be interpreted as structural if one is willing to exclude two-way intraday feedback between the variables. Even in the absence of such a strong assumption, clearly signed correlations indicate the predominant causal direction. Assuming downward sloping demand curves, positive (negative) correlations indicate that the trading flow is predominantly driving (reacting to) the exchange rate.

To investigate the contemporaneous relationship, we estimate unconditional and rolling correlations of the CAD/USD returns with the four types of trading flows.⁶ Figure 1 shows these 500-day backward-looking rolling and unconditional correlations with their 2-standard error bands, over 1990 to 2004:12. The foreign domiciled investment (FD) trading flows have a positive, sizeable (0.47) and fairly stable correlation with cumulative CAD/USD returns. This suggests that FD flows predominantly drive the exchange rate, rather than reacting to it. Such trading flows probably stem from portfolio rebalancing between equity and bond markets in the United States and Canada. Thus, they might well be insensitive to high frequency exchange rate movements.⁷ In contrast to the strong positive correlation of the CAD/USD with FD trading flows, Canadian domiciled financial (CD) trading flows have slightly

⁶ In all the exercises presented in this paper, the CAD/USD returns include the overnight interest differential implied by euro-market rates supplied by the BIS. Computing correlations—rather than regression coefficients—generalizes the Evan-Lyons procedure in a way that emphasizes the endogeneity of both variables.

⁷ Alternatively, the positive FD correlation with exchange rate returns might reflect trend-following behavior—buying USD following USD appreciation—by foreign financial firms. This interpretation does not make clear what, if any, trading flows drive the exchange rate. It is, however, consistent with the fact—to be shown in Table 1—that positive exchange returns forecast higher future FD trading flows.

negative (-0.03), statistically insignificant correlation with cumulative CAD/USD returns.

Commercial client transactions (CC) clearly have a negative, sizeable (-0.28) and fairly stable correlation with cumulative CAD/USD returns. The negative correlation suggests that CC trading flows predominantly react to exchange rate movements. This is sensible: Commercial trading flows tend to buy (sell) more USD when the USD is less (more) expensive. Until the end of 1998, interbank trading flows behaved much like commercial trading flows, showing a large negative correlation with exchange rate returns that averaged about -0.3 (Figure 1).⁸ At about the turn of 1998-1999, however, the correlation rose dramatically and the correlation became much closer to zero, at about -0.05.

Although these correlations are not structural, they suggest some plausible economic interpretations. First, the positive FD-return and negative CC-return correlations are stable in value and significantly different from zero. The signs of these correlations are consistent with the idea that foreign financial firms trade foreign exchange for quasi-exogenous reasons such as equity rebalancing—demanding liquidity—and commercial clients respond to price changes by buying more of the relatively cheaper currency—supplying liquidity. The correlations among the trading flows themselves also support this interpretation that FD trading flows buy liquidity from the others, FD trading flows are negatively correlated with other trading flows: FD trading flows have correlations of -0.11, -0.38 and -0.17 with the CD, CC and IB, respectively. Correlations between CD, CC and IB are small. The idea that financial firms buy liquidity and commercial firms sell liquidity is consistent with the work of Lyons (2001), Evans and Lyons (2004), Marsh and O'Rourke (2005), Bjønnes, Rime, and Solheim (2005), and Gradojevic (2007b).

Second, the high positive value for the FD-return correlation versus the almost zero CD-return correlation is consistent with Evans and Lyon's (2005b) finding that non-U.S. order flow was less informative for predicting the USD/EUR. In the present case, foreign (largely U.S.) financial trading

⁸ The 500-day (or two-year) rolling correlations in Figure 1 are backward looking, so the sharp increase near the end of 2000 indicates that the data changed drastically near the end of 1998.

flows (FD) might be much more informative and have a much larger price impact than Canadian financial (CD) trading flows. If Canadian financial transactions are uninformative, this would contrast with the results in Marsh and O'Rourke (2005), Bjønnes, Rime and Solheim (2005) and Lyons (2001), who find that financial order flow is very informative. Alternatively, the CD trading flows might be equally or more informative but much more price sensitive than the FD trading flows.

Third, in contrast to the relative stability of the FD, CC, and CD correlations, the IB trading flow correlation with CAD/USD returns is unstable over time, rising sharply at the beginning of 1999.⁹ The reason for this instability is not clear, but the Russian default in August 1998 and the November collapse of Long-Term Capital Management sharply reduced risk tolerance.¹⁰ Gradojevic (2007b) shows that FD selloffs drove a large CAD depreciation in August 1998. Subsequently, the Bank of Canada intervened heavily in the forex market and raised interest rates by a record 1 percentage point on August 27, 1998.

4. Forecasting the exchange rate and trading flows

The difficulty in forecasting exchange rates has plagued international financial research since Meese and Rogoff (1983) showed that monetary models' forecasts could not outperform a simple no-change forecast of the exchange rate. Even with 25 years of study, there is scant evidence that monetary models can consistently and significantly outperform a naïve random walk (e.g., see Faust, Rogers, and Wright (2003)). More recently, Engel and West (2004, 2005) have proposed a solution to the puzzle: If fundamentals are sufficiently persistent, then exchange rates will be nearly unforecastable because they will reflect the discounted sum of expectations of these very persistent fundamentals. The order flow literature has offered a chance to forecast exchange rates by monitoring the release of private information.

“When we compare the true, ex ante forecasting performance of a micro-based model against both a standard macro model and a random walk, we find that the micro-based model consistently

⁹ The negative IB-exchange rate relationship contrasts with Lyons' (2001) conclusion that interdealer US dollar purchases lead to the U.S. dollar appreciation.

¹⁰ In mid-August of 1998, the Russian Central Bank announced a moratorium on all repayment of foreign debt owed by Russian banks and private borrowers.

outperforms both: microbased forecasts account for roughly 16 percent of the variance in monthly spot rate changes.” — Evans and Lyons (2005b, p. 413)

While Evans and Lyons (2005b) concentrate on statistical forecast ability, papers such as Rime, Sarno, and Sojli (2007) have found economic value: Order flow variables can inform profitable trading rules.

This paper considers whether trading flows forecast cumulative returns to the CAD/USD. In doing so, we consider three types of forecasts: 1) in-sample Granger causality; 2) long-horizon regressions, as used by Mark (1995); and 3) VECM forecasts to exploit potential long-term relations in the data, such as those found by Killen, Lyons, and Moore (2006). Granger causality tests are the least demanding forecasting technique as they require only 1-step ahead, in-sample predictability. Long-horizon regressions can be useful if the data generating process is a threshold model, as Kilian and Taylor (2003) argued. Unfortunately, long-horizon regressions do not impose possible permanent relations among the cumulated returns and trading flows or imply a complete data generating process. VECMs remedy these deficiencies. We consider whether the VECM provides either statistical forecasting value, or economic forecasting value (a profitable trading rule).

Because the CD-CAD/USD contemporaneous relation is very weak and the IB-CAD/USD relation appears to be unstable, the forecasting exercises will be reported with only FD (financial) and CC (non-financial or commercial) trading flow data. Results from forecasting exercises with all four trading flows in the systems and also with bivariate VARs/VECM were performed but will be omitted for brevity. The inference between the different formulations is broadly consistent; discrepancies with the reported results will be noted where appropriate.

4.1 Granger causality

To test whether trading flows Granger cause foreign exchange returns, we regressed the foreign exchange returns on 5 lags of itself and up to 5 lags of trading flows. The Granger causality regressions for exchange rate returns are expressed as follows:

$$(2) \quad r_t = \beta_0 + \sum_{j=1}^5 \alpha_j r_{t-j} + \sum_{j=1}^P \sum_{i=1}^2 \beta_{i,j} TF_{i,t-j} + \sum_{i=1}^N a_i n_{i,t} + \varepsilon_t$$

where r_t denotes the return on the CAD/USD exchange rate, $TF_{i,t}$ is the t^{th} observation on the i^{th} trading flow (financial or non-financial), and $n_{i,t}$ is the standardized i^{th} news shock at time t , defined in equation (1). The trading flow Granger causality regressions are expressed as follows:

$$(3) \quad TF_{i,t} = \beta_0 + \sum_{j=1}^P \alpha_j r_{t-j} + \sum_{j=1}^5 \sum_{i=1}^2 \beta_{i,j} TF_{i,t-j} + \sum_{i=1}^N a_i n_{i,t} + \varepsilon_t.$$

General-to-specific lag length tests chose 2 trading flow lags for the exchange rate equation, 3 exchange rate lags for predicting financial trading flows and 1 lag for predicting non-financial trading flows.

Table 1 shows that the test for Granger causality—a likelihood ratio test that the coefficients on two lags of the FD and CC trading flows (4 coefficients) were jointly zero—had a p-value of 0.096, indicating that one cannot reject that trading flows Granger cause foreign exchange returns at the ten percent level. Further tests clearly reject the null that exchange rates do not Granger cause financial trading flows at any reasonable significance level. The positive CAD/USD coefficients (0.066, 0.087 and 0.044) indicate that financial trading flows are trend following. Higher CAD/USD prices predict excess FD purchases of USD. In contrast to their ability to predict financial trading flows, exchange rates fail to Granger cause non-financial trading flows. The p-value for the test of no forecastability is 0.912.

4.2 Long-horizon regressions

Mark (1995) introduced long-horizon regressions—the regression of h-period returns on a set of independent variables—to the exchange rate literature. While their efficacy in predicting exchange rates with monetary fundamentals is questionable, they might well forecast exchange rates with microstructure variables (Evans and Lyons (2005b), Sager and Taylor (2008)).

There are at least two important issues in specifying such long-horizon regressions. The first is whether to use the whole sample, or, more realistically, rolling or expanding samples, to estimate the coefficients in the regression. Whole sample coefficient estimates maximize efficiency if the system is

stable, but also incorporate a look-ahead bias. In contrast, expanding and rolling procedures create true ex ante forecasts. These latter procedures start with a 500-day (2-year) initial estimation period, forecast out-of-sample over the forecast horizons, update the data set by one day, reestimate the coefficients, forecast out-of-sample, etc. The expanding window gets larger each day while the rolling procedure maintains a 500-day estimation window.

The second important issue is what information set to use in forecasting exchange rate returns from time t to time $t+h$: Does one use only independent variables available at time t or does one also permit variables from time t to time $t+h$? Only the first method is an ex ante forecast. Meese and Rogoff (1983) introduced the second method—using future values of regressors—to emphasize the extreme difficulty of forecasting exchange rates with macroeconomic fundamentals. As Evans and Lyons (2005b) note, this is really a test of a significant in-sample, contemporaneous relation. The regressions for these two assumptions about forcing variables can be written as follows:

$$(4) \quad \sum_{k=1}^K r_{t+k} = \beta_0 + \sum_{i=1}^2 \beta_i TF_{i,t} + \sum_{i=1}^N a_i \sum_{k=0}^K n_{i,t+k} + \varepsilon_t$$

$$(5) \quad \sum_{k=0}^K r_{t+k} = \beta_0 + \sum_{i=1}^2 \beta_i \sum_{k=0}^K TF_{i,t+k} + \sum_{i=1}^N a_i \sum_{k=0}^K n_{i,t+k} + \varepsilon_t$$

where $n_{i,t+k}$ is the i^{th} news surprise on day $t+k$. Note that the macro announcement regressors are the sums of the shocks in each category over the forecast horizon.

To evaluate the performance of the long-horizon regressions, we calculate MSE ratios as the MSE produced by the long-horizon forecasts over the MSE produced by a no-change forecast. To measure the probability of obtaining the actual MSE ratio if there were no forecastability, we bootstrapped 400 new samples of exchange rate returns with no predictability and then estimated the long-horizon regressions on the bootstrapped data with the various assumptions about the information set and the use of full, expanding or rolling samples. The p-value is the proportion of simulated MSE ratios that are less than the

actual MSE ratios. The bootstrapping method avoids difficulties with small sample biases created by persistent regressors and the complications of overlapping forecast horizons.

Table 2 shows the results of long-horizon regressions to predict exchange rates with trading flows over 1-, 5-, 20-, 60- and 120-day horizons, using full, expanding and rolling samples, with and without future values of the trading flows. We can draw three conclusions from Table 2.

First, when the information set includes future values of the regressors—see panels 1, 3 and 5—the “forecasts” do very well, clearly outperforming the no-change forecasts and doing progressively better as the forecast horizon increases. For example, the ratios of the MSEs of the full sample, long-horizon forecast to the no-change forecast decline from 0.763 at a 1-day horizon to only 0.434 at the 120-day horizon. The p-values in the final column indicate that long-horizon MSEs at 1-, 5- and 20-day horizons are significantly better than what one would expect from such a long-horizon regression if there were no relation between trading flows and exchange rates. The strong contemporaneous correlation between the trading flows and exchange returns drives this significance.

Second, panels 2, 4, and 6 show that when one does not use future values of the trading flows in the regressions, the no-change benchmark outperforms the long-horizon regressions. Even with the full sample coefficients, the long-horizon MSEs are only marginally better than the no-change forecast. The MSE ratios are about 0.98-0.99 at all horizons in this case (see panel 2). The forecasts from expanding samples with no future values of the trading flows (panel 4) produce MSE ratios slightly greater than one—from 1.007 to 1.031—indicating that the no-change forecast is superior to the long-horizon forecast. Finally, the rolling samples without future values of the independent variables produce some evidence of predictability in panel 6. For 60- and 120-day horizons, the rolling sample MSEs dominate those from the no-change forecasts or expanding samples with p-values of 0.065 and 0.138, respectively. We must admit that these intriguing improvements might arise by chance, rather than reflecting real predictability.

Third, the p-values from the expanding and rolling exercises indicate that these procedures do significantly better than one would expect if there were no information in the independent variables. The

p-values indicate that these forecasts do significantly better at the shorter horizons than one would expect if there were no forecastability in the data.¹¹

In summary, realistic constraints on the coefficient estimates and the information set substantially reduce or eliminate the ability of trading flows to forecast exchange rates. Trading flow forecasts might modestly improve on the benchmark at 60- and 120-day horizons. The fact that rolling forecasts outperform the whole sample and expanding sample forecasts at long horizons indicates that the data generating process is probably unstable.

4.3 VECM forecasts with a statistical evaluation metric

Granger causality tests focused narrowly on 1-day forecastability while the long-horizon regressions forced a stark choice between forecasting with current data or implausibly using future values of the regressors. Neither method appropriately models long-run relations between cumulative returns and cumulative trading flows, which may be important for determining the channel through which trading flows affect exchange rates. While liquidity and inventory effects should be transitory, information effects should be permanent, implying a cointegrating relation, and such information effects should differ by type of trading flow (Froot and Ramadori (2005), Killen, Lyons, and Moore (2006)).

To determine the structure of the VAR/VECM, we first tested for unit roots in CAD/USD cumulative returns, cumulative FD trading flows, and cumulative CC trading flows. Augmented Dickey-Fuller (Dickey and Fuller (1981)) and Phillips and Perron (1988) tests failed to reject the unit root null for all three cumulative variables but clearly reject it for their first differences.¹² Therefore we estimated a VAR in levels to permit long-run cointegrating relations between the endogenous variables. Although likelihood ratio tests reject restrictions to fewer than 13 lags, we follow the Schwarz criterion for

¹¹ Because the simulated data are unforecastable and the long-horizon regression is misspecified, the useless predictors in the simulated long-horizon regressions slightly raise the mean VECM MSE above that of the stable no-change forecast. Therefore, the 1-day ahead mean simulated MSE ratio is actually somewhat greater than one in all cases. At very long horizons, the MSE ratio distribution is so wide that there is little power to reject the null.

¹² The Johansen test also chose 2 cointegrating vectors in a 5-variable VAR system, which included all 4 trading flows, but we were unable to reject the exclusion of interbank order flow from the cointegrating relations.

parsimony's sake and use 2 lags. Experimentation indicates that the results are robust to much longer lag lengths. The Johansen (1988) test indicated two cointegrating relations with time trends—implying a VECM structure. Tests rejected the null of weak exogeneity for all three variables. This contrasts with the results of Killen, Lyons and Moore (2006), who find that order flow is strongly exogenous. We omit further discussion of the cointegrating relations and the VECM structure in the interests of brevity and because later work will show that the VECM is almost certainly unstable and the structure is spurious.

The VECM can be written as follows:

$$(6) \quad \Delta Y_t = B_0 + a\beta' Y_{t-1} + B_1 \Delta Y_{t-1} + B_2 \Delta Y_{t-2} + CX_t + \varepsilon_t$$

where Y_t is a vector of the endogenous variables: cumulative exchange rate returns, cumulative financial and non-financial trading flows and a time trend. ΔY_t contains the differences of the endogenous variables but not the difference of the time trend and X_t is the vector of macro surprises and a constant. The macro surprises enter as simple exogenous variables having only a contemporaneous effect on trading flows and exchange rates.

The VECM procedures are similar to those from the long-horizon regressions. We forecast the three endogenous variables using full, expanding and rolling samples over 1-, 5-, 20-, 60- and 120-day horizons. The expanding and rolling procedures start forecasting at about 1993:01, updating coefficients every 250 days for ease of estimation but always making ex ante forecasts. The rolling window is 750 observations or about 3 years of data.¹³ To evaluate forecast quality, we compare the VECM forecast MSEs to those from two naïve benchmarks: 1) A no-change forecast for exchange rates; and 2) a random walk with drift for the two trading flows. The benchmark trading flow forecasts estimate the drift over the whole sample, which is some advantage to the benchmark. Bootstrapping the benchmark models provides the probabilities (i.e., p-values) that the VECM forecasts outperform what one would expect under the null of no predictability.

¹³ A rolling window of 500 observations produces similar inference.

The full-sample VECM forecasts outperform the benchmarks of no-change for the exchange rate returns and the in-sample mean for the trading flows. The top, leftmost panel of Table 3 shows that the MSE ratios decline from 0.996 at a 1-day horizon to 0.761 at a 120-day horizon for exchange rate returns. The bootstrapped p-values indicate that these MSE ratios are not significantly less than one would expect in the absence of predictability, however. That is, the advantage of the whole sample forecast permits the VECM to appear to outperform the naïve benchmarks, in sample, even without real predictability. The MSE ratios for the financial (FD) trading flows decline from 0.938 at 1-day to 0.457 at 120-days and the non-financial (CC) trading flows decline from 0.945 at 1-day to 0.411 at 120-days. The MSE ratios for the trading flow forecasts are always significantly less than one would expect, in the absence of predictability. That is, the p-values in the top-center and top-right panels are essentially zero.

Using the full sample to estimate the VECM coefficients clearly gives an unrealistic advantage to the VECM procedure. With expanding or rolling sample estimation, the VECM no longer outperforms the no-change benchmark for exchange rate prediction at any horizon.

The VECM does appear to predict trading flows at short horizons with expanding sample coefficients, however. The VECM forecasts beat the random-walk-with-drift benchmark at 1- and 5-day horizons, creating MSE ratios of 0.960 and 0.979, for the financial trading flows. Likewise, the VECM forecasts beat the in-sample drift benchmark for non-financial trading flows at the 1-day horizon with expanding samples, creating an MSE ratio of 0.982. The p-values for these trading flow forecasts generally indicate that such forecasts do significantly better than one would expect if trading flows were unforecastable.

Unfortunately, the VECM does not duplicate the promising long-horizon, long-run regression forecasting results for the CAD/USD.¹⁴ The VECM with rolling sample coefficients generally underperforms both the expanding sample VECM and the random walk benchmark. The 750-day rolling window appears to be too short to effectively estimate the VECM coefficients. 500-day rolling sample

¹⁴ The long-horizon and VECM exchange rate forecast statistics are comparable because both methods forecast the sum of returns over the forecast horizon.

results—omitted for brevity—are even poorer. This failure of short-term exchange rate forecasting is consistent with a semi-strong form of market efficiency, the literature on exchange rate forecasting with fundamentals and Sager and Taylor’s (2008) conclusions on order flow’s forecasting power.

4.4 VECM forecasts with an economic evaluation metric

The previous exercises show that lagged exchange rates, lagged trading flows and announcements do not forecast the CAD/USD, at least at short horizons, by statistical criteria. Several papers—e.g., Clarida et al. (2003) and Dueker and Neely (2007)—have shown that a model can generate economic predictability (i.e., a profitable trading rule) while failing to exhibit statistical predictability. Market participants might be very interested in whether one can profitably trade on such information.

To evaluate the economic value of lagged exchange rates and lagged trading flows, we construct trading rules that switch between long/short positions in the foreign currency, according to the VECM forecasts. A long position in the USD at date t means that the rule borrows CAD, converts them to USD at the closing rate for date t , and earns the USD overnight rate and the overnight exchange rate return while paying the CAD overnight rate. A short position borrows USD to invest in CAD.

To reduce transactions costs that are entailed by trading on small expected changes in the exchange rate, we consider forecasts over multiperiod forecast horizons (1-, 5-, 20-, 60- and 120-days) and we introduce a band-of-inaction, which has proved useful in papers like Dueker and Neely (2007). The rule will only switch position from long to short or vice versa if the expected average exchange rate change from t to $t + h$ exceeds the size of the filter, f_h , which depends on the forecast horizon.

$$(7) \quad \text{If } z_{t-1} = -1, \quad z_t = +1, \quad \text{if } E \left[\sum_{i=1}^h r_{t+i} \mid I_t \right] \geq f_h$$

$$\quad \quad \quad \quad \quad = -1, \quad \text{otherwise.}$$

$$\text{If } z_{t-1} = +1, \quad z_t = -1, \quad \text{if } E \left[\sum_{i=1}^h r_{t+i} \mid I_t \right] < -f_h$$

$$\quad \quad \quad \quad \quad = +1, \quad \text{otherwise.}$$

For the full sample estimates, the optimal filter was chosen ex post to maximize the net returns over the

whole sample, conditional on the full sample coefficients. For the expanding and rolling samples, the optimal filter sizes were chosen ex ante each day, to maximize the net return to the previous day.

The cumulative excess return r for a trading rule giving signal z_t at time t , over the period from time zero to time T , conducting n trades, with transaction cost c_t , is as follows:

$$(8) \quad r = \sum_{t=1}^T \left(z_t r_{t+1} - I(z_t \neq z_{t-1}) c_t \right),$$

where r_{t+l} is the daily excess return to a long position in the USD from t to $t+l$. As in Neely, Weller and Ulrich (2009), we assume that the transactions cost (c_t) declines linearly over time, from 5.75 basis points on January 2, 1990 to almost 2 basis points on December 30, 2004. This decline approximates the surely uneven decline in real-world transactions costs over the sample.

Table 4 shows the results of trading rules based on h -period forecasts from the VECM system. The results are again similar to those for the statistical evaluation of the long-horizon regression and VECM forecasts. Using the whole sample, with ex post filters, one obtains reasonable annual net returns, ranging from 0.85 with the 1-day forecasts to 4.34 percent with the 60-day forecasts. Realistically constrained procedures—expanding or rolling samples with ex ante filters—are unprofitable, however. Only the 60-day forecast trading rules using rolling sample coefficients produce a positive excess return and that 46 basis point annual gain is not statistically significant. The marginal profitability of the 60-day rolling forecasts is consistent with the success of the 60-day long-horizon regressions. Unfortunately, the 120-day forecast trading rules are not profitable, even before transactions costs.

The 1-day forecasts are the poorest for all estimation methods. This is consistent with the lack of out-of-sample statistical short-horizon exchange rate predictability and with the low technical trading profitability of the CAD (Sweeney (1986), Neely, Weller and Ulrich (2009)). The negative short-term forecasting results, however, contrasts with Rime, Sarno, and Sojli (2007), who find that order flows do forecast exchange rates.

4.5 The Effect of macro announcements on the CAD/USD and trading flows

We estimate the effects of macro announcement surprises on the exchange rate and trading flows in a reduced form VECM, which means that the surprise coefficients are not isolated structural responses but reduced form coefficients, i.e., total responses of the three endogenous variables after they interact. For example, the exchange rate response to GDP announcement surprises is not only the exchange rate's direct response to the GDP announcement but includes the impact on the CAD/USD from trading flow responses to the GDP announcement and the initial impact of the release on the exchange rate. Because these VECM coefficients describe total responses that include interactions among the endogenous variables they do not necessarily accord with our intuition about the marginal response of these variables to news. Most papers in the announcement literature suffer from this problem, to one degree or another. Further confusing the issue, the correlation in U.S. and Canadian business cycles means that a U.S. surprise might change the relative value of the U.S./Canadian variables much less than one might think.

Table 5 shows strong patterns in how the CAD/USD, financial and non-financial trading flows react to several types of macro surprises. The effect of macroeconomic news on trading flows definitely depends on the type of trading flow, which should not be surprising as the purposes of financial and non-financial trading flows differ. Specifically, non-financial trading flows tend to respond to macro surprises in opposite ways than do financial trading flows and CAD/USD returns. The correlation between the FX return coefficients (column 2) and the financial trading flow coefficients (column 4) is 0.734. The coefficients on the non-financial trading flows are negatively correlated with the coefficients on both the exchange rate (-0.507) and the financial trading flows (-0.626). One might conjecture that CAD/USD responds similarly to the financial trading flows because the financial trading flows are driving the exchange rate. The non-financial trading flows might be price-sensitive, responding to the exchange rate, providing liquidity.

The CAD/USD and financial trading flows generally respond negatively—at least the significant responses—to positive surprises about future U.S. output. That is, financial trading flows and the

exchange rate decline in response to a fall in U.S. business inventories, higher GDP numbers, higher new home sales, and a rise in the U.S. trade surplus. Trade surplus responses are particularly strong, perhaps because higher trade surpluses reduce pressure for monetary or intervention policies to depreciate the USD. Or, a net increase in foreign demand for domestic goods might cause the home currency to appreciate, as in a Mundell-Fleming-type Large Open Economy Model.

This negative response of the CAD/USD to positive news about U.S. output is counterintuitive as exchange rates typically respond positively to output shocks in monetary models (Meese and Rogoff (1983)). Vlaar (2007) potentially resolves this puzzle by showing that the response of an exchange rate to an output shock depends on the origin of the output shocks. The important non-farm payroll announcement surprises elicit a positive reaction from the CAD/USD and financial trading flows, in contrast to the negative response to other types of real shocks.

Non-financial trading flows do not usually react significantly to news about the real economy, except for a negative reaction to improvement in the U.S. trade balance. The point estimates for the coefficients on non-financial trading flows, while usually insignificant, tend to be opposite in sign to those for the CAD/USD and financial trading flows.

Counterintuitively, the CAD/USD tends to increase when the PCE or CPI indices come in higher than expected. Purchasing power parity would suggest that higher U.S. prices would lead to a lower CAD/USD exchange rate. On the other hand, because U.S. inflation expectations were well anchored over the period and the Federal Reserve closely watched the PCE for much of the sample, a positive PCE shock might indicate higher future U.S. interest rates but would not change the expected price level. Non-financial trading flows respond negatively to U.S. PCE and CPI shocks. The negative response of commercial clients to a price shock is in line with the short-run PPP hypothesis — higher cost USD decreases CC trading flow for the USD. Conditional on the positive exchange rate reaction to a U.S. CPI shock, commercial demand for USD might decline because the USD is more expensive.

5. Structural stability

The order/trading flow literature has not previously considered whether the order flow/exchange rate relations are stable but the poor performance of the forecasts suggests this possibility. Theory does not suggest an obvious candidate for a cause or date of a structural break in the relationships between the trading flows and exchange rates. Although the unusual asymptotic properties of the VECM parameters complicate some tests for parameter constancy, one can examine the log likelihood of the recursively estimated system to check for stability, without prespecifying a break date.

Figure 2 shows the time series of the normalized test statistic of the constancy of the log likelihood. The horizontal line denotes the 5 percent critical value. The test statistics, which measure the distance between the subsample and full sample covariance matrices, begin roughly at the beginning of 1993, after a 3-year initial subsample, and go through 2004. The test is done in 2 ways: 1) by reestimating the full system at each date, which produces the statistic $X(t)$; 2) by concentrating out the short run parameters (coefficients on ΔY_{t-1} to ΔY_{t-p} in (6)) and reestimating only the long run parameters in α and β , which produces the statistic $RI(t)$. The very high test statistics provide strong evidence of breaks throughout the sample. As the subsample length approaches the full sample, the power of the test declines, of course, and one cannot reject the null of equal covariance matrices.

What is the source of the instability? The long-run parameters seem to exhibit roughly the same instability as the whole system, which suggests that the long-run parameters are the source of the problem. The rising test statistics in the early part of the sample are troublesome because one would expect the test statistics to decline as the subsample became closer to the whole sample. Thus, the greatest evidence of instability is in 1994-1996. This is the period in which the Bank of Canada began collecting CD trading flow data, which might have affected the categorization of other trading flows. In addition, other events might have affected the CAD in 1994. Gradojevic (2007b) noted heavy selling of the CAD by foreign financial institutions in April 1994, which might have been driven by rising U.S. interest rates. Likewise, the Mexican peso crisis drew attention to large Canadian fiscal and current

account deficits. Late 1998-1999 appears to be another period of instability, where the test statistic does not decline as one would expect. During this period the correlation between IB trading flows and the exchange rate shifted strongly toward zero in Figure 1. This period closely followed the August 1998 Russian default and the September 1998 collapse of Long-Term Capital Management, both of which greatly reduced financial market tolerance for risk.

6. Conclusions

Using a unique data set from the Bank of Canada with unprecedented coverage and length, this paper has studied four questions about CAD trading flows and the CAD/USD exchange rate: 1) Who provides/demands liquidity? 2) Do trading flows forecast exchange rates? 3) How do trading flows and exchange rates react to U.S. macroeconomic surprises? 4) Are the relations between trading flows and the exchange rate stable over time? The length and breadth of our data set provides great power to investigate these hypotheses and determine if previous results are robust.

Foreign financial (FD) trading flows are positively correlated with CAD/USD returns, while commercial (CC) trading flows are negatively correlated with CAD/USD returns. These correlations are sizeable and stable over time. This is consistent with an economic structure in which financial trading flows predominantly drive the exchange rate while commercial trading flows mostly respond to lower prices and provide liquidity. Alternatively, this structure would be consistent with financial trading flows exhibiting trend following behavior, perhaps because of the influence of technical analysis. In contrast to the strong correlations with FD and CC, the CAD/USD returns have low and insignificant correlation with Canadian domiciled financial trading flows (CD) and an unstable correlation with interbank (IB) trading flows.

To investigate whether trading flows forecast exchange rate returns, we employ Granger causality tests, long-horizon regressions and a VECM model. While financial trading flows do Granger cause exchange rates, there is only weak evidence that trading flows can forecast exchange rates out-of-sample. Specifically, long-horizon regressions with rolling samples display some modest ability to predict

exchange rates at the longest (60- and 120-day) horizons. This evidence might arise by chance, however. Consistent with a semi-strong form of market efficiency, lagged trading flows do not predict exchange rate returns out-of-sample by either statistical or economic criteria with a VECM. This negative finding contrasts with that of Rime, Sarno, and Sojli (2007), who study EUR/USD, GBP/USD and JPY/USD data.

In contrast to the mostly negative forecasting results for the CAD/USD, VECMs can forecast both foreign financial (FD) and commercial (CC) trading flows at short horizons. Curiously, FD is trend following: Higher CAD/USD prices tend to predict excess FD purchases of USD. Technical trading might produce this trend-following behavior.

The length of the sample permits us to investigate the structure and stability of the VECM system. In particular, Johansen (1988) tests indicated that the VECM system with cumulated exchange rate returns, foreign financial and commercial trading flows has two cointegrating relations. These permanent relations are consistent with trading flows affecting exchange rates through information channels. None of the variables were weakly exogenous, in contrast to Killen, Lyons and Moore (2006), who found that order flow is strongly exogenous. Tests strongly rejected a constant VECM log likelihood casting doubt on the relevance of the information channels, however. The strongest evidence for a structural break in the sample was fairly early, around 1994. There was also evidence of further instability in late 1998, around the time of the Russian default and the failure of Long-Term Capital Management.

Several types of U.S. macroeconomic announcements—GDP, housing starts, PCE, CPI, and trade balance—influence the CAD/USD exchange rate or trading flows to a statistically significant degree. There are strong patterns in the reduced form responses to macro surprises. Surprises that raise foreign financial trading flows also tend to raise the CAD/USD but reduce commercial trading flow. This pattern might arise because announcement surprises substantially drive exchange rate responses through their effect on foreign financial trading flows and elicit a liquidity provision response from commercial order flow.

References

- Berger, D. W., A. P. Chaboud, S.V. Chernenko, E. Howorka, and J.H. Wright. (2008). Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data, *Journal of International Economics* 75(1), 93-109.
- Bjønnes, G.H., D. Rime, D, and H.O.Aa. Solheim (2003). Volume and Volatility in the FX-Market: Does It Matter Who You Are? In *Exchange Rate Modelling: Where Do We Stand?* edited by P. De Grauwe, pp. 39–52. Cambridge, MA: MIT Press, CESifo Seminar Series.
- Bjønnes, G.H., D. Rime, and H.O.Aa. Solheim (2005). Liquidity Provision in the Overnight Foreign Exchange Market, *Journal of International Money and Finance* 24 (2), 175-96.
- Boyer, M. M., and S. van Norden (2006). Exchange Rates and Order Flow in the Long Run. *Finance Research Letters* 3(4), 235–243.
- Carpenter, A., and J. Wang (2003). Sources of Private Information in FX Trading. Unpublished manuscript, University of New South Wales.
- Clarida, R.H., L. Sarno, M.P. Taylor, and G. Valente (2003). The Out-of-sample Success of Term Structure Models as Exchange Rate Predictors: A step beyond. *Journal of International Economics* 60(1), 61–83.
- Christie-David, R., and M. Chaudhry (2000). Currency Futures, News Releases, and Uncertainty Resolution. *Global Finance Journal* 11(1-2), 109–27.
- Danielsson, J., R.G. Payne, and J. Luo. (2002). Exchange Rate Determination and Inter-Market Order Flow Effects. Mimeo, London School of Economics Financial Markets Group.
- Dickey, D.A., and W.A. Fuller (1981). Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root. *Econometrica* 49(4), 1057–1072.
- Doukas, J., and L.N. Switzer (2004). Bi-National News Effects and Exchange Rate Futures: The Case of Canadian Dollar Futures Contracts. Unpublished manuscript.

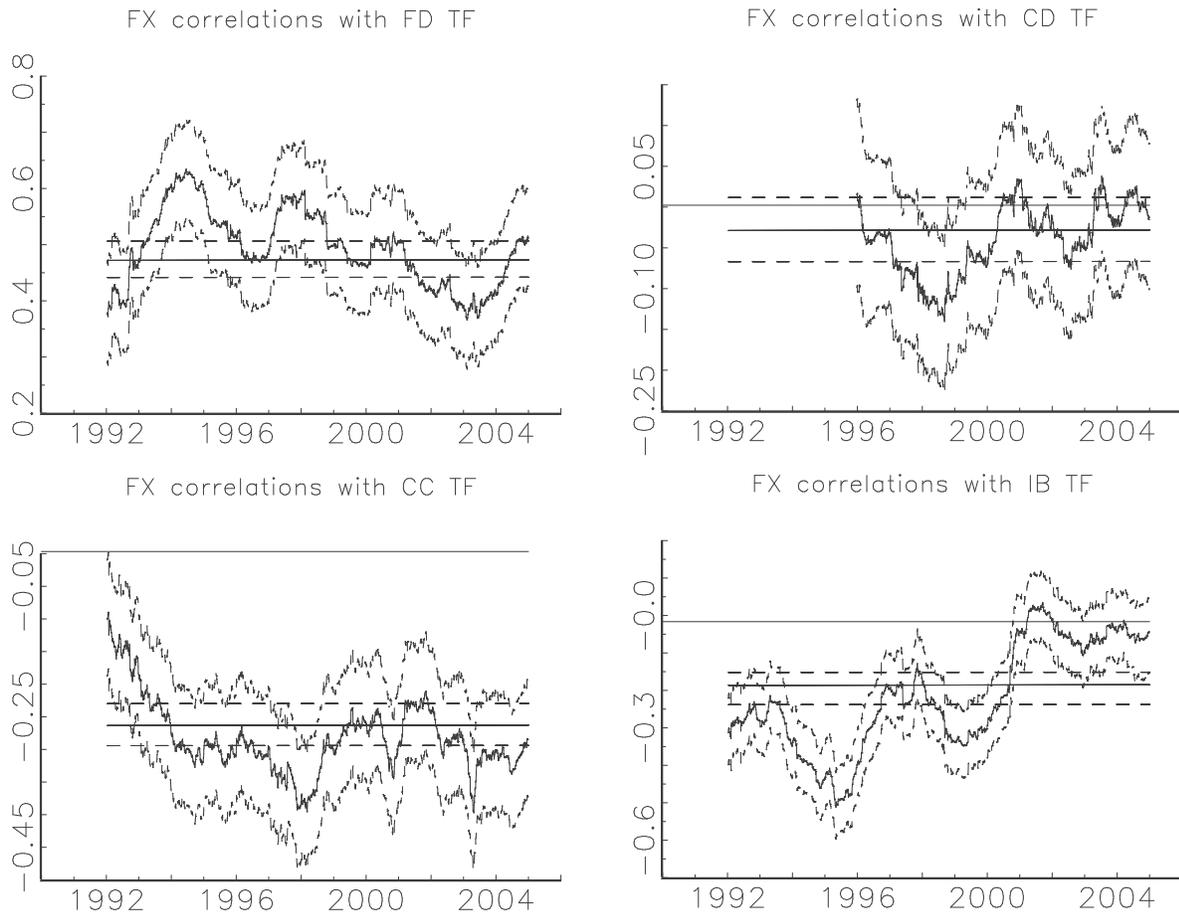
- Dueker, M., and C.J. Neely (2007). Can Markov Switching Models Predict Excess Foreign Exchange Returns? *Journal of Banking and Finance* 31(2), 279–296.
- Ehrmann, M., and M. Fratzscher (2005). Exchange Rates and Fundamentals: New Evidence from Real-Time Data. *Journal of International Money and Finance* 24(2), 317–41.
- Engel, C., and K. West (2004). Accounting for Exchange Rate Variability in Present Value Models when the Discount Factor is Near 1. *American Economic Review* 94(2), 119–25.
- Engel, C., and K. West (2005). Exchange Rates and Fundamentals. *Journal of Political Economy* 113(3), 485–517.
- Evans, M.D.D.. (2007). Foreign Exchange Market Microstructure. Unpublished manuscript, Georgetown University.
- Evans, M.D.D., and R.K. Lyons (2002a) Order Flow and Exchange Rate Dynamics. *Journal of Political Economy* 110(1), 170-80.
- Evans, M.D.D., and R.K. Lyons (2002b). Informational Integration and FX Trading. *Journal of International Money and Finance* 21(6), 807–31.
- Evans, M.D.D., and R.K. Lyons (2002c). Time-Varying Liquidity in Foreign Exchange. *Journal of Monetary Economics* 49(5), 1025-51.
- Evans, M.D.D., and R.K. Lyons (2004). A New Micro Model of Exchange Rate Dynamics. NBER Working Paper 10379.
- Evans, M.D.D., and R.K. Lyons (2005a). Do currency markets absorb news quickly? *Journal of International Money and Finance* 24(2), 197-217.
- Evans, M. D. D., and R.K. Lyons (2005b). Meese-Rogoff Redux: Micro-based Exchange-Rate Forecasting. *American Economic Review* 95(2), 405–414.
- Evans, M.D.D., and R.K. Lyons (2008). How is Macro News Transmitted to Exchange Rates? *Journal of Financial Economics* 88(1), 26-50

- Faust, J., J.H.Rogers, and J.H. Wright (2003). "Exchange Rate Forecasting: The Errors We've Really Made." *Journal of International Economics* 60(1), 35-59.
- Faust, J., J.H.Rogers, E. Swanson, and J.H. Wright (2003). Identifying Effects of Monetary Policy Shocks on Exchange Rates using High Frequency Data. *Journal of European Economic Association* 1(5), 1031–1057.
- Fisher, P., and R. Hillman (2003). Comments on Richard K. Lyons, Explaining and Forecasting Exchange Rates with Order Flows. In Economic Policy Forum: Explaining and Forecasting Exchange Rates with Order Flows.
<http://www.cepii.fr/anglaisgraph/publications/economieinter/rev96/rev96fischer.pdf>.
- Froot, K. A., and T. Ramadorai (2005). Currency Returns, Intrinsic Value and Institutional Order Flows. *Journal of Finance*, 60, 1535–66.
- Gradojevic, N. (2007a). The Microstructure of the Canada/U.S. Dollar Exchange Rate: A Robustness Rest. *Economics Letters* 94(3), 426–432.
- Gradojevic, N. (2007b) A Market Microstructure Analysis of the Canadian Dollar Depreciation Episodes in the 1990s, *Applied Financial Economics* 17(17), 1377–1387.
- Gradojevic, N., and J. Yang (2006). Non-Linear, Non-Parametric, Non-Fundamental Exchange Rate Forecasting. *Journal of Forecasting* 25(4), 227–45.
- Han, L. and J.L. King (1999). Foreign Exchange Futures Volatility: Day-of-the-Week, Intraday, and Maturity Patterns in the Presence of Macroeconomic Announcements. *Journal of Futures Markets* 19(6), 665–93.
- Hayo, B., and M. Neuenkirch (2009). Domestic or U.S. News: What Drives Canadian Financial Markets? MAGKS Papers on Economics 200908, Philipps-Universität Marburg.
- Ito, T. And V. Roley (1987). News from the U.S. and Japan: Which Moves the Yen/Dollar Exchange Rate? *Journal of Monetary Economics* 19(2), 255-77.

- Johansen, S. (1988). Statistical Analysis of Cointegration Vectors. *Journal of Economic Dynamics and Control* 12(2-3), 213–254.
- Kilian, L., and M.P. Taylor (2003). Why Is It So Difficult to Beat the Random Walk Forecast of Exchange Rates? *Journal of International Economics* 60(1), 85–107.
- Killeen, W., R.K. Lyons, and M. Moore (2006). Fixed versus Flexible: Lessons from EMS Order Flow. *Journal of International Money and Finance* 25(4), 551–579.
- Laakkonen, H. (2004). The Impact of Macroeconomic News on Exchange Rate Volatility. Working paper, Bank of Finland.
- Love, R., and R. Payne (2008). Macroeconomic News, Order Flow and Exchange Rates. *Journal of Financial and Quantitative Analysis* 43(2), 467-88.
- Lyons, R.K. (1995). Tests of Microstructural Hypotheses in the Foreign Exchange Market. *Journal of Financial Economics* 39(2-3), 321–51.
- Lyons, R.K. (2001). The Microstructure Approach to Exchange Rates. Cambridge, MA: MIT Press.
- Mark, N.C. (1995). Exchange Rates and Fundamentals: Evidence on Long Horizon Predictability. *American Economic Review* 85(1), 210–18.
- Marsh, I.W., and C. O'Rourke, (2005). “Customer Order Flow and Exchange Rate Movements: Is There Really Information Content?” Working Paper, Cass Business School.
- Meese, R.A., and K. Rogoff (1983). Empirical Exchange Rate Models of the Seventies: Do They Fit out of Sample? *Journal of International Economics* 14(1-2), 3–24.
- Neely, C.J., P.A. Weller, and J.M. Ulrich (2009). The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market. *The Journal of Financial and Quantitative Analysis* 44(2), 467-488.
- Payne, R.G. (2003). Informed Trade in Spot Foreign Exchange Markets: An Empirical Investigation. *Journal of International Economics* 61(2), 307–29.
- Osler, C.L. (2002). Stop-Loss Orders and Price Cascades in Currency Markets. Federal Reserve Bank of New York Staff Report, 150.

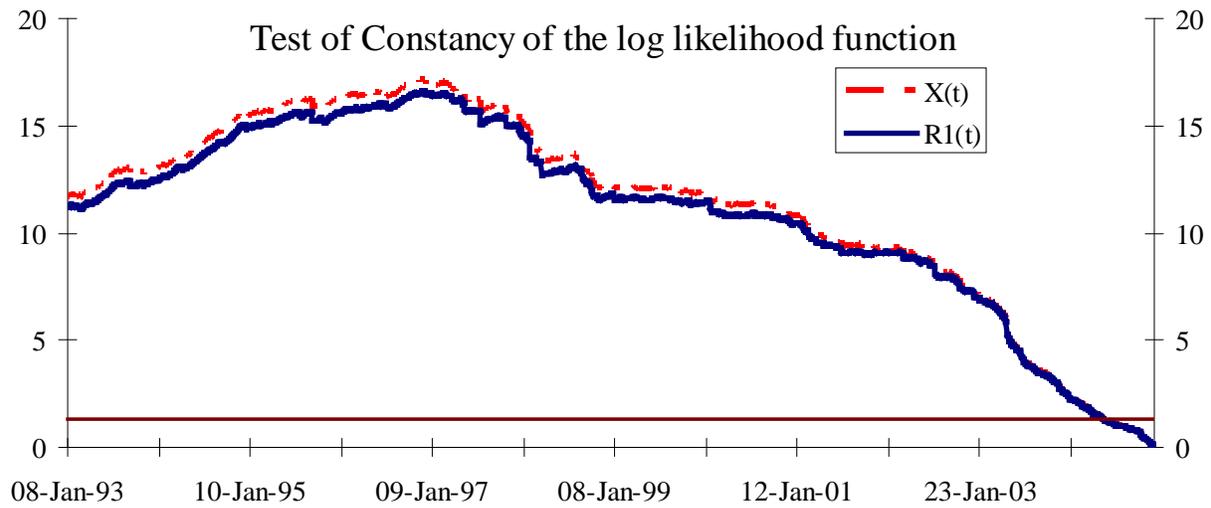
- Osler, C.L. (2003). Currency Orders and Exchange Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis. *Journal of Finance* 58(5), 1791–1820.
- Osler, C.L. (2008). Foreign Exchange Microstructure: A Survey of the Literature, Forthcoming, Springer Encyclopedia of Complexity and System Science.
- Phillips, P. C. B. and P. Perron (1988). Testing for a Unit Root in Time Series Regression. *Biometrika* 75(2), 335–346.
- Rime, D., L. Sarno, and E. Sojli (2007). Exchange Rate Forecasting, Order Flow and Macroeconomic Information. Norges Bank Research Department Working Paper 2007/2.
- Sager, M., and M.P. Taylor (2008). Commercially Available Order Flow Data and Exchange Rate Movements: Caveat Emptor. *Journal of Money, Credit, and Banking* 40(4), 583-625.
- Savaser, T. (2006). Exchange Rate Response to Macro News: Through the Lens of Microstructure, mimeo.
- Schulmeister, S. (2006). The Interaction between Technical Currency Trading and Exchange Rate Fluctuations. *Finance Research Letters* 3(3), 212–233.
- Simpson, M.W., S. Ramchander, and M. Chaudhry (2005). The Impact of Macroeconomic Surprises on Spot and Forward Foreign Exchange Markets. *Journal of International Money and Finance* 24(5), 693–718.
- Sweeney, R. (1986). Beating the Foreign Exchange Market. *Journal of Finance* 41(1), 163–182.
- Vlaar, P.J.G. (2007). GDP Growth and Currency Valuation: The Case of the Dollar. *Journal of International Money and Finance* 26(8), 1424-49.

Figure 1: Rolling correlations of the CAD/USD with trading flows



Notes: The figure shows rolling correlations and a 2-standard error band, using 500-day windows, of the CAD/USD with the 4 types of order flow over the period 1990 through 2004. The figure also shows the correlations over the same sample and a 2-standard error band for those figures.

Figure 2: Test for the constancy of the log-likelihood



Notes: The figure illustrates normalized test statistics of the null that the VECM covariance matrix from the subsample is equal to the VECM covariance matrix from the full sample. Test statistics greater than one reject stability of the system at the 5 percent level. The test is done in 2 ways: 1) reestimating the full system at each date, for which the statistics are denoted as $X(t)$; 2) concentrating out the short run parameters (coefficients on dY) and reestimating only the long run parameters in α and β , for which the statistics are denoted as $RI(t)$. The sample is 1990 through 2004.

Table 1: Granger causality results

Dependent variable	FX Return		Financial Flows		Non-financial Flows	
	coefficients	(s.e.)	coefficient	(s.e.)	coefficient	(s.e.)
Constant	-0.002	(0.006)	0.055	(0.007)	-0.035	(0.004)
FX Return{1}	-0.012	(0.021)	0.066	(0.020)	0.001	(0.014)
FX Return{2}	0.013	(0.020)	0.087	(0.022)		
FX Return{3}	0.040	(0.022)	0.044	(0.022)		
FX Return{4}	-0.012	(0.021)				
FX Return{5}	-0.005	(0.022)				
Fin TF{1}	-0.017	(0.018)	0.133	(0.021)	-0.013	(0.014)
NonFin TF{1}	-0.058	(0.030)	-0.037	(0.031)	0.153	(0.020)
Fin TF{2}			0.021	(0.021)	-0.007	(0.012)
NonFin TF{2}			-0.007	(0.035)	0.061	(0.021)
Fin TF{3}			0.033	(0.023)	0.014	(0.011)
NonFin TF{3}			-0.055	(0.033)	0.045	(0.020)
Fin TF{4}			0.037	(0.020)	-0.006	(0.011)
NonFin TF{4}			0.021	(0.032)	0.021	(0.025)
Fin TF{5}			0.057	(0.019)	0.004	(0.011)
NonFin TF{5}			-0.034	(0.033)	0.021	(0.020)
RSQ	0.02		0.08		0.05	
LogL	-1385.5		-1717.7		255.7	
T	3710		3710		3710	
# pars	34		40		38	
LR test for no GC at lag	4.686		30.252		0.012	
DF for above LR	2		3		1	
P-value for above LR	0.096		0.000		0.912	

Notes: The table shows the results of Granger causality tests that 1) FD and CC trading flows cause CAD/USD returns; 2) CAD/USD returns cause financial (FD) trading flows; and 3) CAD/USD returns cause non-financial (CC) trading flows. The sample is 1990 through 2004.

Table 2: Long-horizon forecast results

Window Size	Future values of TF regressors?	Forecast horizon	MSE Regression	MSE No change	MSE ratio	No change p- value
Whole sample	Yes	1	0.096	0.125	0.763	0.000
Whole sample	Yes	5	0.448	0.632	0.708	0.000
Whole sample	Yes	20	1.684	2.533	0.665	0.000
Whole sample	Yes	60	4.538	8.133	0.558	0.310
Whole sample	Yes	120	7.446	17.138	0.434	0.885
Whole sample	No	1	0.124	0.125	0.989	0.630
Whole sample	No	5	0.626	0.633	0.990	0.880
Whole sample	No	20	2.506	2.533	0.989	0.910
Whole sample	No	60	8.036	8.135	0.988	0.913
Whole sample	No	120	17.015	17.143	0.993	0.995
Expanding	Yes	1	0.105	0.136	0.773	0.000
Expanding	Yes	5	0.492	0.680	0.723	0.000
Expanding	Yes	20	1.914	2.729	0.701	0.000
Expanding	Yes	60	5.285	8.743	0.604	0.110
Expanding	Yes	120	8.723	18.304	0.477	0.678
Expanding	No	1	0.137	0.136	1.007	0.000
Expanding	No	5	0.689	0.680	1.013	0.013
Expanding	No	20	2.770	2.727	1.016	0.160
Expanding	No	60	8.933	8.742	1.022	0.398
Expanding	No	120	18.869	18.306	1.031	0.560
Rolling	Yes	1	0.110	0.136	0.807	0.000
Rolling	Yes	5	0.479	0.680	0.704	0.000
Rolling	Yes	20	1.377	2.729	0.505	0.000
Rolling	Yes	60	1.501	8.743	0.172	0.000
Rolling	Yes	120	1.363	18.304	0.074	0.005
Rolling	No	1	0.145	0.136	1.068	0.023
Rolling	No	5	0.721	0.680	1.059	0.015
Rolling	No	20	2.808	2.727	1.029	0.030
Rolling	No	60	8.321	8.742	0.952	0.065
Rolling	No	120	16.370	18.306	0.894	0.138

Notes: The table shows the results of the long-horizon regressions in which exchange rate returns over horizons of 1-, 5-, 20-, 60- and 120-days are regressed on lagged exchange rate returns, trading flows and macro announcements. Column 1 describes the sample over which coefficients were estimated for each specification. Column 2 notes whether contemporaneous values of the trading flow variables are used to predict the exchange rate return, or whether it was an ex ante prediction. Column 3 is the forecast horizon. Columns 4 and 5 are the mean squared error for the long horizon regression and the no-change forecast, respectively. Column 6 is the ratio of those MSEs. Column 7 is the proportion of bootstrapped MSE ratios that are smaller than the actual number in column 6. The sample is 1990 through 2004.

Table 3: VECM forecasting results with statistical evaluation criteria

	Whole sample				Whole sample				Whole sample		
	Horizon	Mtgle MSE	p- values		Horizon	Drift MSE	p- values		Horizon	Drift MSE	p- values
	ratios				ratios				ratios		
FX returns	1	0.996	0.45	Financial TF	1	0.938	0.00	Non-financial TF	1	0.945	0.00
FX returns	5	0.987	0.47	Financial TF	5	0.902	0.00	Non-financial TF	5	0.891	0.00
FX returns	20	0.956	0.54	Financial TF	20	0.811	0.00	Non-financial TF	20	0.726	0.00
FX returns	60	0.878	0.56	Financial TF	60	0.604	0.00	Non-financial TF	60	0.515	0.00
FX returns	120	0.761	0.47	Financial TF	120	0.457	0.00	Non-financial TF	120	0.411	0.00
	Expanding sample				Expanding sample				Expanding sample		
FX returns	1	1.031	0.59	Financial TF	1	0.960	0.00	Non-financial TF	1	0.982	0.00
FX returns	5	1.127	0.56	Financial TF	5	0.979	0.00	Non-financial TF	5	1.025	0.03
FX returns	20	1.489	0.72	Financial TF	20	1.030	0.01	Non-financial TF	20	1.125	0.10
FX returns	60	2.114	0.86	Financial TF	60	1.111	0.03	Non-financial TF	60	1.302	0.18
FX returns	120	2.418	0.92	Financial TF	120	1.364	0.12	Non-financial TF	120	1.545	0.29
	Rolling samples				Rolling samples				Rolling samples		
FX returns	1	1.273	0.99	Financial TF	1	1.063	0.07	Non-financial TF	1	1.251	0.95
FX returns	5	2.180	0.99	Financial TF	5	1.347	0.24	Non-financial TF	5	2.032	0.93
FX returns	20	4.286	0.99	Financial TF	20	1.896	0.26	Non-financial TF	20	3.516	0.91
FX returns	60	4.899	0.98	Financial TF	60	2.378	0.40	Non-financial TF	60	4.063	0.88
FX returns	120	3.882	0.92	Financial TF	120	2.463	0.44	Non-financial TF	120	3.834	0.89

Notes: The table presents results from VECM forecasts of a three-variable system using CAD/USD returns, financial (FD) trading flows, and commercial (CC) trading flows with 2 cointegrating relations, a trend in the cointegrating relation and 2 lags in the VECM representation. The top, middle and bottom panels display results estimating the coefficients with the whole sample, an expanding sample and a rolling sample, respectively. The MSE ratio p-values are the proportion of bootstrapped MSE ratios that are less than the actual MSE ratio, using a null model assuming no-change forecasts for the exchange rate and random walks with drift for the trading flow variables. The random walk with drift is always estimated over the whole sample, for both the actual MSE ratios and the bootstrapped data. The sample is 1990 through 2004.

Table 4: VECM trading rule results

Whole sample, ex post filter					
Horizon	1	5	20	60	120
Gross AR	2.85	4.89	3.26	4.48	2.86
Net AR	0.85	4.23	3.11	4.34	2.77
t statistic	0.57	2.86	2.10	2.94	1.87
Sharpe	0.15	0.77	0.57	0.79	0.50
Sharpe SE	0.27	0.31	0.29	0.31	0.29
TradesPeryear	51.48	18.68	4.37	4.08	2.85
Filter Size	0.010	0.030	0.020	0.000	0.010
Expanding sample, ex ante filter					
Horizon	1	5	20	60	120
Gross AR	0.12	-0.10	-1.65	-0.39	-0.73
Net AR	-1.22	-0.49	-1.84	-0.56	-0.83
t statistic	-0.71	-0.29	-1.08	-0.33	-0.48
Sharpe	-0.21	-0.08	-0.32	-0.10	-0.14
Sharpe SE	0.30	0.30	0.30	0.30	0.30
TradesPeryear	36.95	12.43	6.11	5.05	3.12
Filter Size	0.010	0.008	0.001	0.061	0.076
Rolling samples, ex ante filter					
Horizon	1	5	20	60	120
Gross AR	-1.23	-1.12	0.03	0.66	-1.68
Net AR	-2.39	-1.40	-0.16	0.46	-1.77
t statistic	-1.40	-0.82	-0.09	0.27	-1.04
Sharpe	-0.41	-0.24	-0.03	0.08	-0.31
Sharpe SE	0.31	0.30	0.30	0.30	0.30
TradesPeryear	30.98	8.48	5.14	5.32	2.42
Filter Size	0.006	0.126	0.000	0.098	0.015

Notes: The table shows the results of trading rules based on h -period forecasts from the VECM system. Three subpanels from top to bottom show results from forecasts with coefficients from full samples, expanding samples and rolling samples. The rows show gross annual return, net annual return (net of transactions costs), the t statistic for the null that the net annual return equals zero, the Sharpe ratio, its standard error, trades per year and the mean filter size in basis points. The sample is 1990 through 2004.

Table 5: The effect of macroeconomic shocks on the CAD/USD and trading flows

	FX returns		Financial TF		Non-financial TF	
	Coefficient	t-stat	Coefficient	t-stat	Coefficient	t-stat
Business Inventories	0.0327	(1.2195)	0.0064	(2.1571)	0.0010	(0.5760)
Capacity Utilization	0.0285	(0.7935)	0.0030	(0.7569)	-0.0019	-(0.8347)
Consumer Confidence	0.0165	(0.5955)	-0.0049	-(1.6182)	0.0019	(1.0516)
Construction Spending	0.0148	(0.5611)	0.0002	(0.0690)	0.0000	(0.0062)
CPI	0.0414	(1.3476)	0.0065	(1.9142)	-0.0061	-(3.1059)
Consumer Credit	-0.0040	-(0.1494)	0.0023	(0.7920)	-0.0019	-(1.0983)
Advance Durables	0.0324	(1.1876)	0.0043	(1.4347)	-0.0028	-(1.6274)
new orders at factories	0.0197	(0.6457)	-0.0009	-(0.2816)	-0.0004	-(0.1981)
fed funds target	-0.0036	-(0.1147)	0.0012	(0.3352)	0.0015	(0.7586)
Real GDP Advance	-0.0031	-(0.0718)	0.0011	(0.2325)	0.0030	(1.0847)
Real GDP Preliminary	-0.0822	-(1.8222)	-0.0092	-(1.8566)	0.0018	(0.6166)
Real GDP Final	-0.0948	-(1.7714)	-0.0028	-(0.4695)	0.0000	(0.0044)
Housing Starts	-0.0018	-(0.0518)	-0.0062	-(1.6576)	0.0022	(1.0221)
Initial Claims	0.0039	(0.2790)	-0.0011	-(0.7301)	-0.0001	-(0.1275)
Industrial Production	-0.0213	-(0.5165)	-0.0002	-(0.0338)	0.0031	(1.1913)
Leading Indicators	0.0368	(0.7480)	0.0096	(1.7706)	-0.0038	-(1.1975)
Mfg Comp Index	-0.0037	-(0.1286)	0.0058	(1.8653)	-0.0031	-(1.7149)
NFP	0.0562	(2.1782)	0.0016	(0.5706)	-0.0011	-(0.6367)
New home sales	-0.0531	-(1.9156)	-0.0111	-(3.6331)	-0.0003	-(0.1918)
PCE	0.0679	(2.2360)	0.0054	(1.6102)	-0.0017	-(0.8797)
Personal Income	-0.0286	-(0.8316)	-0.0022	-(0.5862)	-0.0001	-(0.0274)
PPI	-0.0309	-(1.0966)	-0.0028	-(0.9092)	-0.0008	-(0.4172)
Retail Sales	-0.0081	-(0.2056)	0.0002	(0.0521)	0.0021	(0.8502)
Retail Sales ex Vehicles	0.0479	(1.4069)	0.0021	(0.5593)	-0.0005	-(0.2393)
Trade Balance	0.1001	(3.6779)	0.0091	(3.0481)	-0.0049	-(2.7909)
Govt fiscal surplus/de ficit	0.0530	(1.6392)	0.0028	(0.7802)	-0.0014	-(0.6948)
Corr($B_{FX\ ret}$, $B_{Fin\ TF}$)	0.734					
Corr($B_{FX\ ret}$, $B_{Non-fin\ TF}$)	-0.507					
Corr($B_{Fin\ TF}$, $B_{Non-fin\ TF}$)	-0.626					

Notes: The table displays coefficients on macro surprises from the VECMs. From left to right, the panel displays VECM coefficients and standard errors for the CAD/USD return equation, the financial trading flow equation and the non-financial trading flow equation. Shaded cells indicate statistically significant coefficients. The sample is 1990 through 2004.