Multiscale Analysis of Foreign Exchange Order Flows and Technical Trading Profitability

This version: September, 2013

Nikola Gradojevic*

Camillo Lento†

Abstract

This paper investigates the multiscale (frequency-dependent) relationship between technical trading profitability and feedback trading effects in the Canada/U.S. dollar foreign exchange market. The results suggest weak evidence that technical trading activities of financial and non-financial customers drive frequent violations of the FX market microstructure assumption that exchange rate movements are driven by order flow. After controlling for transaction costs, we find that the contribution of financial customers in feedback trading dominates the contribution of non-financial customers at lower frequencies, while the opposite holds at higher frequencies. An additional, novel contribution is that technical indicators constructed from order flows can be profitable.

Keywords: Foreign Exchange Markets; Order Flows; Technical Trading; Frequency Domain.

JEL Classification: F31; G14; C53.

* Corresponding author. IÉSEG School of Management (LEM-CNRS), Lille Catholic University, 3, rue de la Digue, 59000 Lille, France. Phone: +33 (0) 320 545 892. Email: n.gradojevic@ieseg.fr; The Rimini Center for Economic Analysis, Italy.

† Faculty of Business Administration, Lakehead University, 955 Oliver Road, Thunder Bay, Ontario P7B 5E1, Canada. Phone: 1 807 3438387. Fax: 1 807 3438023. E-mail: clento@lakeheadu.ca. School of Accounting, Economics and Finance, University of Southern Queensland, West Street, Toowoomba, Queensland, Australia.

1. Introduction

In technical analysis, technical trading rules are constructed from historical price (and volume) information. Specifically, traders place their orders by mechanically applying mathematical transformations and rules based on past and present prices. However, this notion contrasts the major assumption underlying foreign exchange (FX) market microstructure theory that exchange rate movements are driven by order flow and not vice-versa (Lyons, 2001). The argument in favor of this assumption stems from the classical equity microstructure literature such as Hasbrouck (1991), Glosten and Milgrom (1985) and O'Hara (1995): in rational markets, *aggregate* order flow should reflect innovations in dispersed information, rather than being the result of "momentum" (or "feedback") trading strategies followed by *some* FX traders.[1] Notwithstanding the theoretical validity of this argument, Schulmeister (2006) finds that currency order flows can in fact be driven by technical trading signals. He uncovers strong feedback effects where a rising exchange rate triggers buy technical trading signals and thereby strengthens the appreciation trend.[2] Although estimating price effects is not within the scope of this paper, it is worth noting that in case of any evidence of reverse causality, the linear estimate of the size of price effects would be biased at a given time scale.

Only a select few papers have directly tested the causality assumption in FX markets and they focused on a particular data frequency (typically daily). Killeen et al. (2006) find that Granger causality runs from interdealer order flow to price, and not vice versa, for the DM/FRF exchange rate. However, several papers documented statistically significant reverse

[1] Positive (negative) feedback trading is systematic buying (selling) in response to price increases, and selling (buying) in response to price decreases (Evans and Lyons, 2002).

[2] Alternatively, feedback effects in the FX market can also be caused by liquidity provision. Bjønnes, Rime and Solheim (2005), for example, find that non-financial customers are passive liquidity providers in the SEK/EUR market. In the same vein, D'Souza (2008) documents that, in addition to commercial clients, dealers are key participants in the provision of liquidity in the Canada/U.S. dollar market.

causality effects. For instance, Sager and Taylor (2008) perform Granger causality tests on the data from Evans and Lyons (2002) and reveal that causality runs from the DM/USD and JPY/USD exchange rate returns to corresponding interdealer order flows. They also present evidence against the causality assumption for customer order flows. This evidence corroborates Marsh and O'Rourke (2005) who argue that commercial order flow is price sensitive. Similarly, Boyer and van Norden (2006) conclude that interdealer order flow responds to the FRF/USD spot rate innovations. They note that the price responsiveness of commercial order flow contrasts with the usual predictions of the microstructure literature. Gradojevic and Neely (2008) demonstrate the ability of the Canada/U.S. dollar returns to predict financial order flows, but not non-financial order flows. Lyons (2001) finds some evidence that falling prices induce additional selling in the JPY market and refers to that phenomenon as "distressed selling." Recently, Gradojevic (2012) showed that reverse causality in FX market microstructure is frequency-dependent. This evidence is generally in accord with arguments that, in financial markets, the data generating process (DGP) is a complex network of layers with each layer corresponding to a particular frequency. Thus, a successful characterization of such DGP should be estimated with techniques that account for intra- and inter-frequency dynamics (Dacorogna et al., 2001).[3]

The main goal of this paper is to test the strength of the relationship between multiscale feedback effects (i.e., reverse causality in the frequency domain) and technical trading profitability in the Canada/U.S. dollar market. The motivation for this research is to study the behavior of technical traders with different time horizons such as daily, weekly, bi-weekly and monthly. More importantly, this exploration seeks to understand whether the dominance of technical trading is the kind of "irrational" behavior that governs feedback

[3] The idea that the causality relationship between two variables may have different characteristics at different time-scales can also be found in Gençay et al. (2001). They use wavelet multiresolution analysis of money growth and inflation, and show that for Argentina, Brazil, Chile, Israel, Mexico and Turkey the nature of the causality changes with wavelet scales (periods between two and 32 months).

trading. The profitability of moving average and trading range breakout technical indicators is tested on the Canada/U.S. dollar exchange rate, and cumulative financial and non-financial order flows. This paper is novel to the literature because it is the first study, to the authors' best knowledge, to use currency order flows (i.e., proxies for trading volume) for technical trading rule calculations. Also, it is important to emphasize that the analysis presented in this paper differs from the existing literature (e.g., Schulmeister, 2006, Rime et al., 2010 and Gradojevic and Neely, 2008) in that it accounts for the wide range of FX trading horizons in the frequency domain.

In the first part of the paper, the frequency domain causality tests are performed on Canada/U.S. dollar returns, and financial and non-financial order flows. It is confirmed that, in general, there is very little evidence of a stable causal relationship between order flows and returns running in either direction. The null hypothesis of no predictability of FX returns by spot non-financial order flows is not rejected at weekly and shorter horizons, while it is rejected for financial order flows at horizons between 3 and 8 days. In terms of reverse causality from price to non-financial order flows, we document that daily and weekly frequencies exhibit feedback effects. On the contrary, financial order flows are found to be driven by price changes at longer horizons. Next, for both order flow types, we test the hypothesis that technical trading rules constructed at the 'reverse causality frequencies' are more profitable than the trading rules generated from data sampled at other frequencies. Our evidence of very short run technical trading profitability indicates that non-financial customers were engaged in feedback trading primarily at the daily horizon. Also, we present the findings of medium (bi-weekly and monthly) horizon profitability that can be attributed to the technical trading of financial customers.

In general, even after accounting for transactions costs, substantial technical trading excess returns are found for all three time series whereas the profitability in general increases with the time horizon.[4] Furthermore, the results indicate that technical trading rules are more profitable with non-financial order flows than with financial order flows. This suggests that there is more technical trading information content in the trades of non-financial customers (dominated by Canadian corporations) than in the trades of financial customers (dominated by foreign dealers). Alternatively stated, Canadian-domiciled and corporate customer transactions may provide insight regarding the fundamental trends in the Canadian economy, which act as a leading indicator for the exchange rate.

Considering the evidence of stable increasing trend in technical trading profitability with the trading horizon, it can be concluded that reverse causality effects at a particular frequency do not imply technical trading profitability at such frequency. In addition, the trading intensity, measured as the number of trades triggered by technical indicators, is not related to reverse causality. These findings suggest that feedback trading effects cannot be explained by the predominant activity of technical traders and that future research efforts should turn to "liquidity provision" as the more likely prevalent form of feedback behavior in the Canada/U.S. dollar market (see Gradojevic and Neely, 2008).

In the next section, the methodology for causality in the frequency domain is reviewed. The data and the construction of technical trading strategies are presented in Section 3. Section 4 discusses the findings. The final section concludes the paper.

2. Causality in the frequency domain

The test for causality in the frequency domain by Breitung and Candelon (2006) originates from Geweke (1982) and Hosoya (1991). Let $z_t = [x_t, y_t]'$ be a two-dimensional time series vector with t = 1,...,T. It is assumed that z_t has a finite-order VAR representation

[4] The findings that technical trading strategies can be profitable at medium horizons is consistent with Neely and Weller (2003) and Harris and Yilmaz (2009).

$$\Theta(L)z_t = \varepsilon_t,$$

(1)

where $\Theta(L) = I - \Theta_1 L - \ \dots \ - \Theta_p L^p$ is a 2×2 lag polynomial with $L^k z_t = z_{t-k}$. It is assumed that the vector ε_t is white noise with $E(\varepsilon_t) = 0$ and $E(\varepsilon_t \varepsilon_t') = \Sigma$, where Σ is a positive definite matrix. Next, let G be the lower triangular matrix of the Cholesky decomposition $G'G = \Sigma^{-1}$, such that $E(\eta_t \eta_t') = I$ and $\eta_t = G\varepsilon_t$. The system is assumed to be stationary, implying the following MA representation:

$$z_t = \Phi(L)\varepsilon_t = \begin{bmatrix} \Phi_{11}(L) & \Phi_{12}(L) \\ \Phi_{21}(L) & \Phi_{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix}$$

$$= \Psi(L)\eta_t = \begin{bmatrix} \Psi_{11}(L) & \Psi_{12}(L) \\ \Psi_{21}(L) & \Psi_{22}(L) \end{bmatrix} \begin{bmatrix} \eta_{1t} \\ \eta_{2t} \end{bmatrix}$$

(2)

where $\Phi(L) = \Theta(L)^{-1}$ and $\Psi(L) = \Phi(L)G^{-1}$. Using this representation, the spectral density of x_t can be expressed as

$$f_x(\omega) = \frac{1}{2\pi}\left\{ \left|\Psi_{11}(e^{-i\omega})\right|^2 + \left|\Psi_{12}(e^{-i\omega})\right|^2 \right\}$$

(3)

The measure of causality suggested by Geweke (1982) and Hosoya (1991) is defined as

$$M_{y\to x}(\omega) = \log \frac{2\pi f_x(\omega)}{\left|\Psi_{11}(e^{-i\omega})\right|^2}$$

(4)

$$= \log\left[1 + \frac{\left|\Psi_{12}(e^{-i\omega})\right|^2}{\left|\Psi_{11}(e^{-i\omega})\right|^2}\right]$$

(5)

This measure is zero if $|\Psi_{12}(e^{-i\omega})| = 0$ in which case it is said that *y* does not cause *x* at frequency ω. To test the hypothesis that *y* does not cause *x* at frequency ω the following null hypothesis is used:

$$M_{y \to x}(\omega) = 0 \tag{6}$$

Yao and Hosoya (2000) estimate $M_{y \to x}(\omega) = 0$ by replacing $|\Psi_{11}(e^{-i\omega})|$ and $|\Psi_{12}(e^{-i\omega})|$ from Equation (5) with estimates obtained from the fitted VAR. However, this approach is not appropriate since $|\Psi_{12}(e^{-i\omega})|$ is a complicated nonlinear function of the VAR parameters.[5] Breitung and Candelon (2006) resolve this problem by showing that the null hypothesis $M_{y \to x}(\omega) = 0$ is equivalent to a linear restriction on the VAR coefficients. First, they use $\Psi(L) = \Theta(L)^{-1} G^{-1}$ and $\Psi_{12}(L) = -\dfrac{g^{22}\Theta_{12}(L)}{|\Theta(L)|}$ (where g^{22} is the lower diagonal element of G^{-1} and $|\Theta(L)|$ is the determinant of $\Theta(L)$) to express the null hypothesis as

$$|\Theta_{12}(e^{-i\omega})| = \left| \sum_{k=1}^{p} \theta_{12,k} \cos(k\omega) - \sum_{k=1}^{p} \theta_{12,k} \sin(k\omega) i \right| = 0, \tag{7}$$

where $\theta_{12,k}$ is the (1,2)-element of Θ_k. Thus, a necessary and sufficient set of conditions for $|\Theta_{12}(e^{-i\omega})| = 0$ is

$$\sum_{k=1}^{p} \theta_{12,k} \cos(k\omega) = 0, \tag{8}$$

[5] Yao and Hosoya (2000) propose a numerical estimation procedure instead of an exact analytical expression.

$$\sum_{k=1}^{p} \theta_{12,k} \sin(k\omega) = 0, \tag{9}$$

The notation can be simplified by letting $a_j = \theta_{11,j}$ and $\beta_j = \theta_{12,j}$. Then, the VAR equation for x_t can be written as

$$x_t = a_1 x_{t-1} + ... + a_p x_{t-p} + \beta_1 y_{t-1} + ... + \beta_p y_{t-p} + \varepsilon_{1t}. \tag{10}$$

The hypothesis $M_{y \to x}(\omega) = 0$ is equivalent to the linear restriction

$$H_0 : R(\omega)\beta = 0, \tag{11}$$

where $\beta = \left[\beta_1, ..., \beta_p\right]'$ and

$$R(\omega) = \begin{bmatrix} \cos(\omega) & \cos(2\omega) & ... & \cos(p\omega) \\ \sin(\omega) & \sin(2\omega) & ... & \sin(p\omega) \end{bmatrix}$$

The ordinary F statistic for (11) is approximately distributed as $F(2, T-2p)$ for $\omega \in (0, \pi)$. As in Breitung and Candelon (2006), to assess the statistical significance of the causal relationship between exchange rate returns and order flows, the causality measure for $\omega \in (0, \pi)$ is compared to the 5% critical value of a χ^2-distribution with 2 degrees of freedom (5.99).[6]

3. Data Description and Technical Trading Rules

3.1. Data description

The data are sampled at the daily frequency and are obtained from the Bank of Canada. They span the period between October 10, 1994 and September 30, 2005. This represents a total of

[6] Breitung and Candelon (2006) study the local power of the test when the frequency being tested converges to the true frequency and show that the Wald statistic is asymptotically distributed as noncentral χ^2.

2,798 observations of daily returns (r_t) and order flows. If S_t denotes an exchange rate at time t, then $r_t = \log(S_t) - \log(S_{t-1})$.

The order flow data are aggregate daily trading flows (in Canadian dollars) for eight major Canadian commercial banks:

- Commercial client transactions (CC) include all transactions with resident and non-resident non-financial customers.

- Foreign institution transactions (FD) include all transactions with foreign financial institutions, such as FX dealers.

The CC transactions are motivated by trades in real goods and services, while the FD transactions are motivated by international portfolio considerations. These order flows represent approximately 40-60% of all Canada/U.S. dollar transactions. Using the definition from Lyons (2001), order flows are measured as the difference between the number of currency purchases (buyer-initiated trades) and sales (seller-initiated trades). *Ceteris paribus*, positive (negative) order flow should raise (lower) the Canada/U.S. dollar spot closing rates S_t, appreciating (depreciating) the USD. In the remainder of the paper the CC transactions will be referred to as non-financial order flows, while the FD transactions will be referred to as financial order flows.

3.2. Technical Trading Strategies

We calculate two simple momentum-based technical trading rules to proxy for trading activity. Specifically, the moving average cross-over rule (MACO) and the trading range break-out rule (TRBO) are calculated with the Canada/U.S. dollar exchange rate, the non-financial order flows, and the financial order flows.[7] We have selected these two trading rules because they are commonly employed in the prior literature (e.g., Brock, Lakonishok and LeBaron, 1992 and Park et al., 2007). In addition, we have computed each trading rule with three different sets of parameters in order to reduce the sensitivity of our results to a single trading rule parameter.

The MACO generates a buy (sell) signal whenever the short moving average is above (below) the long moving average. The following combinations of short and long periods are employed: (1, 50), (5, 50), and (10, 50). The TRBO rule generates a buy signal when the price breaks-out above the resistance level and a sell signal when the price breaks below the support level. The three periods employed are 5, 10, and 20.

The profitability of the trading rules is determined by comparing the returns generated by the trading signals to the buy-and-hold strategy returns. Similar to Gencay (1998), the returns generated from the trading rules are adjusted for transaction costs (i.e., both the bid-ask spread and brokerage trading costs). Our methodology relies on this relatively simple technique for analyzing the profitability of the trading rules because of the possible problems related to non-linear trading models such as computational expensiveness, overfitting, data snooping and difficulties of interpreting the results.[8]

The returns from the buy-and-hold strategy are calculated by investing in the security at the beginning of the data set, given the trading rule parameters, and holding the security until the end of the data set. To minimize the measurement error due to non-synchronous trading made evident by Scholes and Williams (1977), the investor will be long in the market one day after the trading signal is generated. Therefore, once a buy signal is generated, the investor will be long on the following day, and returns will be calculated based on the market

---

[7] Filter rules were not feasible for this data set due to an insufficient number of generated trading signals at lower frequencies.
[8] See White (2005) for a thorough discussion of these issues.

returns. Finally, the position is carried forward if the investor is long (short) and a buy (sell) signal is generated. Risk-adjusted return analysis in the form of Sharpe ratios is also performed as a robustness check.

4. Results

4.1. Multiscale causality

This subsection is adapted from Gradojevic (2012) and it reports the results of causality tests in the frequency domain for two bivariate systems: one for each order flow and exchange rate returns. Both Dickey Fuller and Phillips-Perron tests reject the null hypothesis of a unit root in all time series at the 1% significance level (p-value=0.000). According to the AIC criterion, the likelihood ratio and the final prediction error criteria a VAR(6) model was selected for both systems.[9]

Before we test the causality in the frequency domain, we carefully check whether VAR(6) is an appropriate representation of the data. In particular, we test for possible conditional heteroskedasticity and autocorrelation in the residuals. The evidence based on the Ljung-Box Q test statistic confirms that the residuals of our VAR(6) systems do not suffer from serial correlation. In the system with non-financial order flows, the Q statistics for the two equations are as follows: 20.83 (p-value=0.40) and 11.44 (p-value=0.93); i.e., the null hypothesis of no serial correlation in the residuals is not rejected. In the system with financial order flows, the Q statistics for the two equations are as follows: 18.67 (p-value=0.54) and 25.26 (p-value= 0.19); i.e., the null hypothesis of no serial correlation in the residuals is not rejected. The same can be shown by using the Lagrange-multiplier test statistic and the null

[9] The Schwarz's Bayesian information criterion and the Hannan and Quinn information criterion suggest one lag for the bivariate system with non-financial order flow and two lags for the bivariate system with financial order flow. We find autocorrelation in the residuals of the lower order VAR systems and opt for a more conservative approach that utilizes six lags.

hypothesis of no autocorrelation at lag order L. For non-financial order flows: L=1 (p-value=0.90), L=2 (p-value=0.32), L=3 (p-value=0.10), L=4 (p-value=0.07), L=5 (p-value=0.22) and L=6 (p-value=0.08). For financial order flows: L=1 (p-value= 0.11), L=2 (p-value=0.10), L=3 (p-value=0.37), L=4 (p-value=0.23), L=5 (p-value=0.11) and L=6 (p-value=0.09).

Next, we use the Engle-Granger test in order to test for conditional heteroskedasticity in the VAR systems. The null hypothesis of no ARCH effects is not rejected in the non-financial order flow equation (p-value=1.00) and also in the financial order flow equation (p-value=0.99). However, we detect conditional heteroskedasticity in the exchange rate returns equations for both order flow types (p-value=0.00). The presence of ARCH effects in exchange rate returns may appear problematic, but Bodart and Candelon (2009) perform a simulation study and conclude that the causality test in the frequency domain from Breitung and Candelon (2006) is robust to such data features. In light of these findings, we conjecture that the results of our multiscale causality tests are unaffected by the observed conditional heteroskedasticity effects. We use this feature of the methodology as one of the arguments for working in the frequency domain.[10]

Figure 1 presents the causality measure between non-financial order flows and exchange rate returns for all frequencies ($\omega \in (0, \pi)$) along with the 5% critical value (5.99) that is represented with a horizontal dashed line. The top panel indicates that the null hypothesis of no causality is rejected when $\omega < 0.35$ which corresponds to frequencies with a wavelength of roughly more than four weeks (20 days = $2\pi/\omega$). Thus, FX microstructure theory appears to be correct only at such horizons. In other words, it takes at least four weeks

[10] A time-domain approach may be appropriate as well. In fact, in a related paper, the VECM/VAR systems based on these variables have already been covered (see Gradojevic and Neely, 2009), but this contribution does not empirically test the links between the observed "reverse causality" effects and technical trading activity.

for customer order flows to become informative for the price. This shows that the relevance of non-financial order flow strongly depends on the horizon length.

The bottom panel of Figure 1 reveals evidence of reverse causality in the short run for $\omega \in [0.6, 1.4]$ (between 4.5 and 10.5 days, i.e., the weekly and bi-weekly lag horizons) and $\omega > 2.2$ (shorter than three days, i.e., the daily lag horizon). These findings explain the inability of non-financial order flows to predict returns for frequencies higher than 0.35 and stresses that forecasting exchange rates with order flows requires care. It can be concluded that linear exchange rate models that employ non-financial order flows produce unbiased estimates of the size of price effects at medium to long horizons.

[Insert Figure 1 about here]

The results for financial order flows are displayed in Figure 2 and they are in stark contrast to the ones for non-financial transactions. Financial order flows are informative in the range $\omega \in [0.8, 2]$ corresponding to a cycle length between 3 and 8 days (top panel of Figure 2). However, the bottom panel of Figure 2 does not reject the null hypothesis of no predictability for $\omega < 1.6$ thereby indicating bi-directional causality in the range $\omega \in [0.8, 1.6]$. Hence, for $\omega \in [1.6, 2]$ (lag horizon 3-4 days) the estimates support the FX microstructure causality assumption, while the reverse causality problems are present at lag horizons longer than roughly 4 days. Hence, over the range when non-financial order flow conforms with the causality assumption, financial order flows exhibit feedback trading effects. The findings that financial order flow is a poor predictor of exchange rate returns at longer horizons is consistent with Berger et al. (2008) who find that the association between interdealer order flow and exchange rate returns weakens at longer horizons.

[Insert Figure 2 about here]

Next, the paper considers the robustness of the results with respect to time period. The data set is divided into three subsets according to Gradojevic and Neely (2008) as follows: 1994-1997 (821 observations), 1998-2001 (1012 observations), and 2002-2005 (965 observations).[11] First, the relationship between non-financial order flow and exchange rate returns will be examined. The causal impact of non-financial order flows on exchange rate returns appears to be very sensitive to the time period (upper right panel of Figure 3). Although the 1994-1997 period resembles the results for the entire sample (non-financial order flow is informative when $\omega < 0.8$, i.e., for the eight-day and longer horizons), causality is not observed for 1998-2001 and only in the very short run for 2002-2005 (when $\omega > 2.6$). Furthermore, according to Figure 3, reverse causality is not present in 1994-1997, but the figures for the other two periods are similar to the bottom panel of Figure 1. In all, forecasting exchange rate returns with non-financial order flows at longer horizons does not violate the causality assumption of microstructure theory.

The robustness analysis for financial order flows identifies the 2002-2005 period as the one when transactions were not informative. The figures for the other two periods mirror the top panel of Figure 2. Reverse causality for longer horizons is present in 1998-2001 and 2002-2005, but not in the 1994-1997 period (solid line in Figure 3). Therefore, it is important to note that the informativeness of financial order flows is sensitive to time period and frequency. Moreover, as the relationship between financial order flows and exchange rate returns tends to occasionally reflect feedback trading, special attention should be paid to forecasting at longer horizons.

[Insert Figure 3 about here]

[11] They utilized the test for the constancy of the log-likelihood on the same data set and found structural instability in 1998 and 2001.

It can also be noted that the 1994-1997 period was in line with the causality assumption for both order flow types, i.e., reverse causality effects were not observed and order flows were mostly informative for the price. This period was characterized by stable exchange rates when the Canadian dollar traded in a relatively narrow range around 0.73 U.S. dollars. The period of stability lasted until the heavy fluctuations in 1998 caused mainly by the economic crisis in emerging markets in Asia, Russian default, and the collapse of Long-Term Capital Management.

4.2. Technical Trading Strategies

Recall that the main purpose of this paper is to explore the potential links between reverse causality and technical trading activity at various frequencies. Table 1 presents the profits generated from the technical trading strategies calculated with the Canada/U.S. dollar return data, CC data (non-financial order flows), and FD data (financial order flows).

[Insert Table 1 about here]

Subsection 4.1 identified reverse causality for the daily and weekly frequencies in the case of non-financial customers and for the bi-weekly and monthly frequencies in the case of financial customers. Therefore, if there is a relationship between reverse causality and the profits from technical trading rules in regards to the activity of non-financials (financials), one would expect technical trading profits from daily and weekly frequencies to be larger (smaller) than the profits from the bi-weekly and monthly frequencies.

Table 1 provides evidence that the bi-weekly and monthly frequency generated the largest profits from technical analysis, as opposed to the daily and weekly frequencies. For example, the average profits from the six trading rules were negative at the daily frequency for the return data, CC data, and FD data. The return and FD data also generated negative profits at the weekly frequency. However, the average profits from the trading rules are positive for the return, CC and FD data at both bi-weekly and monthly frequencies.

Therefore, Table 1 finds no evidence of a robust relationship between the average technical trading profitability and reverse causality across time scales. However, different results emerge when analyzing the results for individual rules as opposed to the average results. In particular, for the MACO(5,50), MACO(10,50) and TRBO(20) rules, profits are positive for all frequencies when using non-financial order flow, and lowest at the monthly horizon. This is consistent with the findings on reverse causality reported in Panel B of Figure 1. Using financial order flow, the profits from these same trading rules are negative at the daily horizon and positive at the weekly, bi-weekly and monthly horizons, which is consistent with Panel B of Figure 2. Thus it would appear that the relationship between trading rule profitability and reverse causality may depend on the lag length employed in the trading rule.

Interesting results also emerge when analyzing the relationship between technical trading profitability and reverse causality on the individual data sets. For example, Table 1 also reveals that the technical trading rules were the most profitable when calculated with the non-financial order flow data. The CC data results in larger average profits than the return and FD data for all four frequencies examined. In addition, the trading rules generated positive profits on three of the four frequencies examined, whereas the returns and FD data generated positive profits on only two of the four frequencies. Recall that CC transactions are with resident and non-resident non-financial customers, whereas the FD transactions are with foreign financial institutions. The fact that the trading rules were more profitable with CC data than either the FD data or return data suggests that there is more momentum based technical trading information content in the trades of non-financial customer (e.g., Canadian corporations) than in the trades of financial customers (e.g., foreign dealers).

Table 1 also indicates a stable, increasing trend in technical trading profitability with the trading horizon. That is, the profits from the technical trading rules increase as the frequency of the data decreases (i.e., from daily, to bi-weekly, to monthly). The profits from the technical trading rules are the largest on all three data sets when calculated with the monthly frequency. With both the CC and FD data, the profits from the technical trading rules increased from daily to bi-weekly to monthly. The return data resulted in the largest profits with the monthly data frequency, followed by the daily and then bi-weekly frequencies

As an additional sensitivity test, we present the results before transaction costs (raw returns). The raw returns can rule out the possibility that the increase in trading rule profit when moving from daily to weekly frequency is caused simply by a reduction in the number of trades (i.e., the higher frequency trades will generate more trades, and increase transaction costs). Table 2 presents the raw returns from the technical trading strategies.

[Insert Table 2 about here]

The results for the technical trading returns before transaction costs from Table 2 suggest that it is possible that the inverse relationship between profits and the data frequency (Table 1) could arise due to less trades being triggered (i.e., less transaction costs). The results from the price series suggest a U-shape of returns across frequencies (Figure 4) with higher profits at the lowest and highest frequencies.

[Insert Figure 4 about here]

Recall that CC order flows show strong reverse causality effects at the daily frequency. This is in line with Table 2 (and Figure 4), which suggests that non-financial customers engage in technical trading at the daily frequency (evidenced by the large returns at the daily level on the price data). This could be the result of Canadian multinational corporations entering the foreign exchange market on a daily basis to hedge/settle transactions denominated in foreign currencies. The FD order flow displays reverse causality at the weekly and lower frequencies. Again, the returns without transaction costs suggest that financial customers engage in technical trading at lower frequencies (evidence by the larger returns at the monthly level on the price data). This could be the result of foreign institutions being a market maker for long term foreign currency forward contracts.

Overall, Table 2 and Figure 4 can be interpreted to provide support for the notion that non-financial customers (CC order flow) drive technical trading at the daily frequency, while financial customers (FD order flow) drive technical trading at the bi-weekly and monthly horizons. Overall, we find some support for the multiscale dependency of reverse causality on the technical trading profitability when analyzing the returns without transaction costs and the individual trading rules. However, the relationship is not robust as it is not supported by the average trading rules profits.

Table 3 presents the number of trades generated from the technical trading strategies calculated with the Canada/U.S. dollar return data, CC data (non-financial order flows), and FD data (financial order flows). It reveals that the average number of trades generated with the return, CC, and FD data all decrease as the frequency of the data decreased. For example, the daily frequency with the return data generated 532 trading signals over the time period analyzed. The number of trades decreased to 85, 35, and 17 with weekly, bi-weekly and monthly data frequency, respectively. The same pattern is clearly evident with the CC and FD data.

[Insert Table 3 about here]

An additional sensitivity test is conducted to determine if the number of trades declines as the frequency of the data decreases. For example, one may argue that it is intuitively apparent that for a given sample period, a trading rule based on monthly returns will generate fewer trades than one based on daily data. Accordingly, Table 4 presents the average number of trades generated against the expected number of trades at the weekly, bi-weekly, and monthly frequencies. The expected number of trades is calculated based on the number of trades actually generated at the daily level. For example, with 250 trading days per year, and 12 months per year, one would expect 26 trades with monthly return data (i.e., 532 / [250/12]). The expected number of trades at the weekly and bi-weekly level is calculated in the same manner, but, with 52 weekly and 26 bi-weekly periods, respectively. Table 4 presents the results.

[Insert Table 4 about here]

The results from Table 4 support the results in Table 3 as the number of trades generated in for the weekly, bi-weekly and monthly data are all below their expectations. Table 4 provides further support to the results that the number of trades decrease as the frequency of the data decreases.

An additional sensitivity test is conducted to assess whether the riskiness of the technical trading profits has any impact on the overall results. For example, it may be difficult to assess which trading strategy is optimal based on profits alone as risk is also an important factor in assessing an optimal strategy.

Accordingly, the Sharpe ratio (SR) is calculated to measure the riskiness of the technical trading strategies. The SR is calculated as the ratio of the mean return of the trading strategy to its standard deviation. Higher SRs are desirable because they indicate either higher mean returns or less volatility. A negative SR indicates that a buy-and hold strategy would perform better than the technical trading rule. The SRs for the profits and profits before transaction costs are presented in Table 5 and Table 6, respectively.

[Insert Table 5 and Table 6 about here]

Recall that Table 1 provided evidence that the bi-weekly and monthly frequency generated the largest profits from technical analysis, as opposed to the daily and weekly frequencies. The SR analysis (Table 5) is consistent with the profit analysis as the SRs are larger at the bi-weekly and monthly frequency. For example, the SR for the return time series is -0.018, -0.007, 0.008, and 0.042 for the daily, weekly, bi-weekly, and monthly frequencies. Similar patterns are also identified, on average, for the CC and FD data.

The SR analysis with profits before transaction costs (Table 6) is consistent with main results as the SRs for the return and CC increase as the frequency decreases from daily to monthly. Therefore, the SR analysis supports the conclusion that a robust relationship does not exists between the average risk-adjusted technical trading profitability and reverse causality across time scales.

The SR analysis also supports the conclusion that using the non-financial order flow data results in the optimal trading strategy as the risk-adjusted profits for three of the four frequencies (daily, weekly, and bi-weekly), and outperformed the financial order flow data at all four frequencies (Table 5). Before adjusting for transaction costs, the non-financial order flow data results in the optimal risk-adjusted trading strategy for all four frequencies (Table 6).

Taken together, the results show no robust evidence of a relationship between technical trading profitability and reverse causality across time scales in terms of either technical trading profitability or trading intensity.

5. Conclusions

The first goal of this study is to critically investigate the causality relationship between spot exchange rates and currency order flows in the Canada/U.S. dollar market. The recent popularity of FX market microstructure models stems from the argument that order flow is the key variable in the process by which dispersed private and public information is aggregated until it becomes embedded in exchange rates. The central hypothesis underlying this process is that order flows drive exchange rates and not vice versa.

Next, we investigate the potential links between the reverse causality and technical trading activities across time scales. It is important to understand whether feedback effects are driven by momentum trading strategies that can be tracked using the daily Canada/U.S. dollar exchange rates, and non-financial and financial order flows over the 1994-2005 time period. This investigation represents our second major research question.

In contrast to the microstructure theory, the evidence shows that both the existence and direction of the causal relationship depends on the customer (order flow) type, frequency, and time period. In general, non-financial order flows are informative in the medium to long run, while financial order flows are good predictors of exchange rates over a narrow range of frequencies with wave-lengths between 3 and 8 days. The usefulness of financial order flows is plagued by feedback trading effects that are dominant at a lag horizon longer than 4 days. Some short-run reverse causality problems are also evident for non-financial order flows. Finally, the causality relationship is influenced by the FX market regime and it appears to follow the theoretical predictions in relatively stable markets (e.g., the 1994-1997 period). All of the preceding findings complement the current state of the microstructure literature and represent valuable practical guidelines for exchange rate modeling using order flow variables.

In regards to the second research question, we find some support for the multiscale dependency of reverse causality on the technical trading profitability when analyzing the returns without transaction costs and the individual trading rules. However, the relationship is not robust as it is not supported by the average trading rules profits. These findings suggest that feedback trading effects cannot be explained by the predominant activity of technical traders and that "liquidity provision" might be responsible for the presented evidence of reverse causality in the Canada/U.S. dollar market (Gradojevic and Neely, 2008). It is worthwhile to note the striking result that technical trading rules that employ both order flow types (i.e., proxies for volume) can be profitable, and that the profits from the technical trading rules increase as the frequency of the data decreases. Furthermore, we demonstrate the superiority in profitability of non-financial order flows. This suggests that the trades of Canadian-domiciled and corporate customers are more informative in the technical trading sense with respect to the Canadian dollar than those of the foreign financial institutions.

In closing, it is worth noting that the messages of this paper neither invalidate nor validate the causality assumption of market microstructure literature. It is possible that the analysis of intra-frequency (high frequency) dynamics or data for other exchange rates and dealing banks might shed more light on the issue, and have more robustness in rejecting the null hypothesis. However, the evidence is consistent with Sager and Taylor (2008) and warns that predicting exchange rates with microstructure variables is a more complex task than has been previously thought.

Acknowledgements

Nikola Gradojevic gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada. Also, the authors wish to thank Ramo Gençay for valuable advice on an earlier version of the paper, seminar participants at the Bank of Canada and Ryerson University (International Monetary and Financial Economics Workshop), and Scott Hendry (Bank of Canada) for providing the data. Giovanni Scalzo contributed to the manuscript with some excellent research assistance.

References

Berger, D., Chaboud, A., Chernenko, S., Howorka, E., Wright, J., 2008. Order flow and exchange rate dynamics in electronic brokerage system data. J. Int. Econ. 75, 93–109.

Bjønnes, G.H., Rime, D., Solheim, H.O.A., 2005. Liquidity provision in the overnight foreign exchange market. J. Int. Money Finance 24, 175–196.

Bodart, V., Candelon, B., 2009. Evidence of interdependence and contagion using a frequency domain framework. Emerg. Mark. Rev. 10, 140–150.

Boyer, M., van Norden, S., 2006. Exchange rates and order flow in the long run. Financ. Res. Lett. 3, 235–243.

Breitung, J., Candelon, B., 2006. Testing for short- and long-run causality: A frequency-domain approach. J. Econometrics 132, 363–378.

Brock, W., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. J. Finance 47, 1931–1764.

Dacorogna, M., Gençay, R., Muller, U., Olsen, R., Pictet, O. 2001. An Introduction to High-Frequency Finance. Academic Press, San Diego.

D'Souza, C., 2008. The role of foreign exchange dealers in providing overnight liquidity. Bank of Canada Working Paper 2008-44.

Evans, M., Lyons, R., 2002. Order flow and exchange rate dynamics. J. Pol. Econ. 110, 170–180.

Evans, M., Lyons, R.. 2005. Meese-Rogoff redux: Micro-based exchange rate forecasting. Amer. Econ. Rev. 95, 405–414.

Gençay, R., 1998. Optimization of technical trading strategies and the profitability in security markets. Econ. Lett. 59, 249–254.

Gençay, R., Selçuk, F., Whitcher, B., 2001. An Introduction to Wavelets and Other Filtering Methods in Finance and Economics. Academic Press, San Diego.

Geweke, J., 1982. Measurement of linear dependence and feedback between multiple time series. J. Amer. Stat. Assoc. 77, 304–324.

Glosten, L., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. J. Financ. Econ. 14, 71–100.

Gradojevic, N., 2007. The microstructure of the Canada/U.S. dollar exchange rate: A robustness test. Econ. Lett. 94, 426–432.

Gradojevic, N., 2012. Frequency domain analysis of foreign exchange order flows. Econ. Lett. 115 (1), 73–76.

Gradojevic, N., Neely, C., 2008. The dynamic interaction of trading flows, macroeconomic announcements and the CAD/USD exchange rate: Evidence from disaggregated data. Working Paper 2008-006C.

Harris, R.D.F., Yilmaz, F., 2009. A momentum trading strategy based on the low frequency component of the exchange rate. J. Banking Finance 33, 1575–1585.

Hasbrouck, J., 1991. Measuring the information content of stock trades. J. Finance 46, 179–207.

Hosoya, Y., 1991. The decomposition and measurement of the interdependence between second-order stationary process. Probab. Theory Relat. Fields 88, 429–444.

Killeen, W., Lyons, R., Moore, M., 2006. Fixed versus flexible: Lessons from EMS order flow. J. Int. Money Finance 25, 551–579.

Lyons, R., 2001. The Microstructure Approach to Exchange Rates. The MIT Press, Cambridge, USA.

Marsh, I. W., O'Rourke, C. 2005. Customer order flow and exchange rate movements: Is there really information content? Working Paper.

Neely, C. J., Weller, P. A., 2003. Intraday technical trading in the foreign exchange market. J. Int. Money Finance 22, 223–237.

O'Hara, M., 1995. Market Microstructure Theory. Blackwell Publishers. Cambridge, Mass.

Osler, C., Vandrovych, V., 2009. Hedge funds and the origins of private information in currency markets. Working Paper No. 1484711.

Park, C.H., Irwin, S.H., 2007. What do we know about the profitability of technical analysis? J. Econ. Surveys 21, 786–826.

Rime, D., Sarno, L., Sojli, E., 2010. Exchange rate forecasting, order flow and macroeconomic information. J. Int. Econ. 80, 72–88.

Sager, M., Taylor, M., 2008. Commercially available order flow data and exchange rate movements: Caveat emptor. J. Money Credit Bank. 40, 583–625.

Scholes, M., Williams, J., 1977. Estimating betas from non-synchronous data. J. Financ. Econ. 5, 309–327.

Schulmeister, S., 2006. The interaction between technical currency trading and exchange rate fluctuations. Financ. Res. Lett. 3, 212–233.

Yao, F., Hosoya, Y., 2000. Inference on one-way effect and evidence in Japanese macroeconomic data. J. Econometrics 98, 225–255.

White, H., 2005. Approximate nonlinear forecasting methods, in: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), Handbook of Economics Forecasting. Elsevier, New York, pp. 460–512.

Figures

Figure 1. Causality tests for non-financial (CC) order flows.

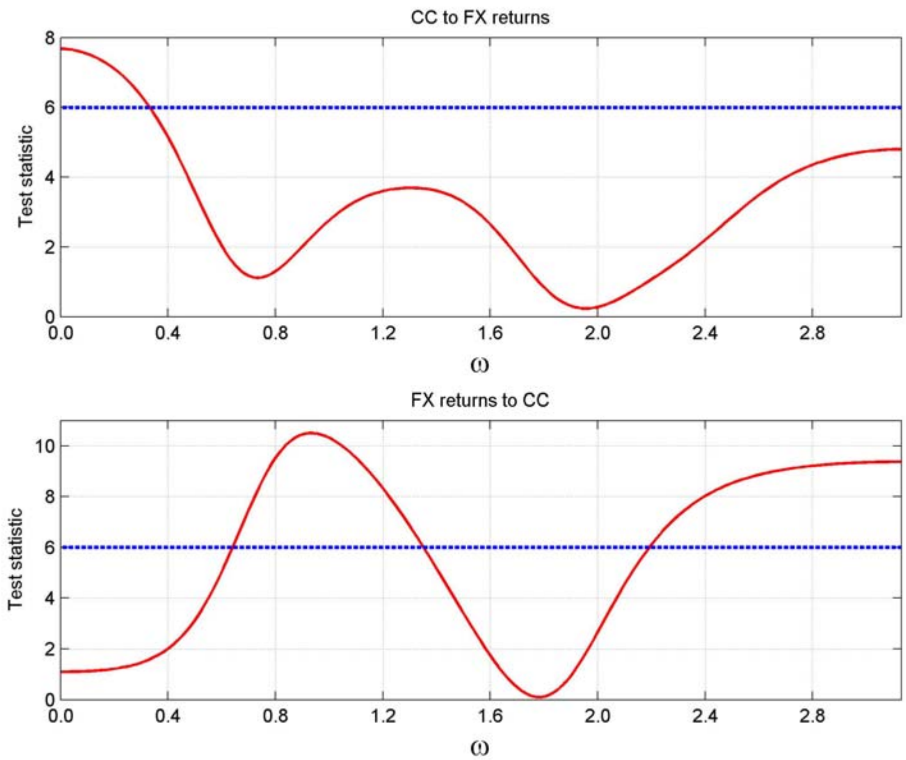

Top panel: Causality tests (non-financial order flows to FX returns). Bottom panel: Causality tests (FX returns to non-financial order flows). The values of the χ^2 test statistic are given by a solid line. The 5% critical value (5.99) that is given by a horizontal dashed line. The null hypotheses are 1) that non-financial order flow does not cause FX returns at frequency ω (top) and 2) that FX returns do not cause non-financial order flow at frequency ω (bottom).

Figure 2. Causality tests for financial (FD) order flows.

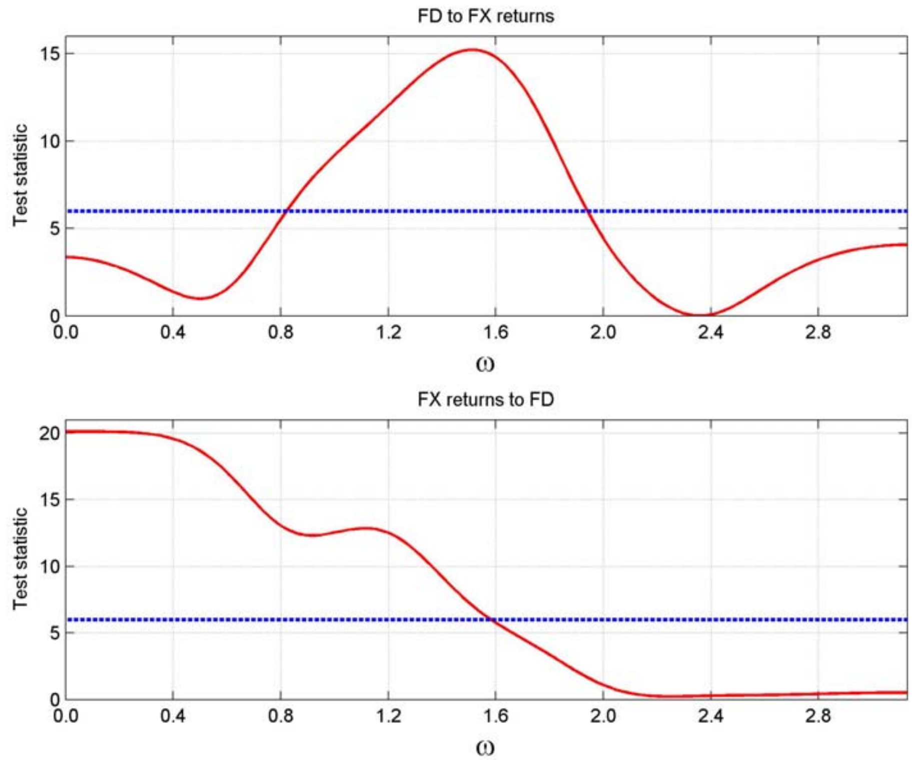

Top panel: Causality tests (financial order flows to FX returns). Bottom panel: Causality tests (FX returns to financial order flows). The values of the χ^2 test statistic are given by a solid line. The 5% critical value (5.99) that is given by a horizontal dashed line. The null hypotheses are 1) that financial order flow does not cause FX returns at frequency ω (top) and 2) that FX returns do not cause financial order flow at frequency ω (bottom).

Figure 3. Causality tests for order flows across subsamples.

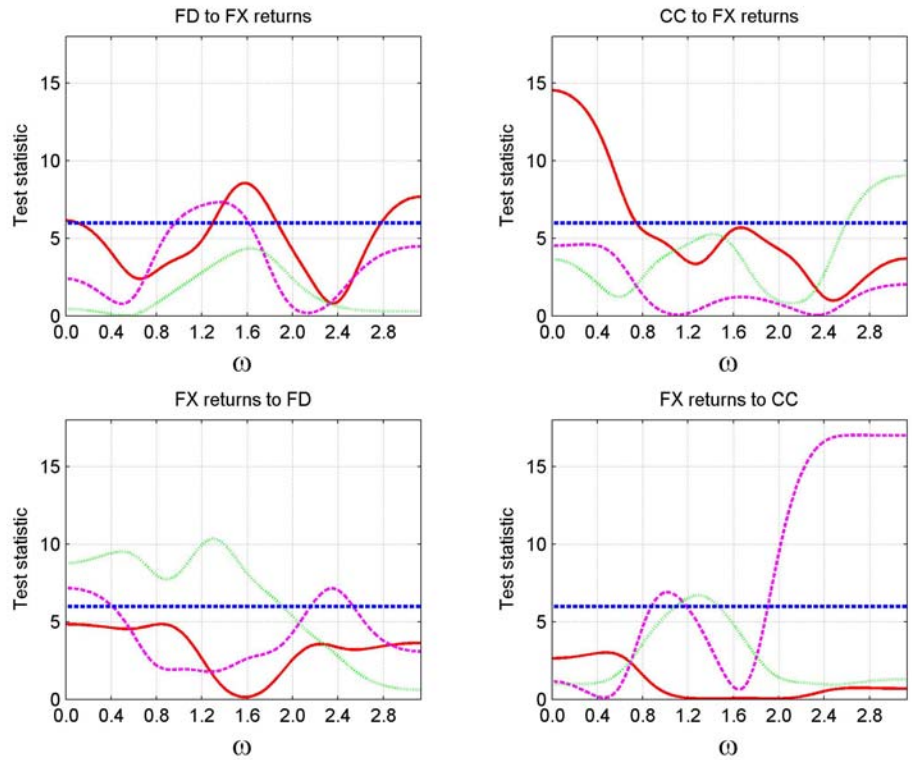

Top left panel: Causality tests (financial order flows to FX returns). Bottom left panel: Causality tests (FX returns to financial order flows). Top right panel: Causality tests (non-financial order flows to FX returns). Bottom right panel: Causality tests (FX returns to non-financial order flows).The values of the χ^2 test statistic are given by a solid line for the 1994-1997 subsample, dashed line for the 1998-2001 subsample and dotted line for the 2002-2005 subsample. The 5% critical value (5.99) that is given by a horizontal dashed line. The null hypotheses are 1) that order flows does not cause FX returns at frequency ω (top panels) and 2) that FX returns do not cause order flows at frequency ω (bottom panels).

Figure 4.

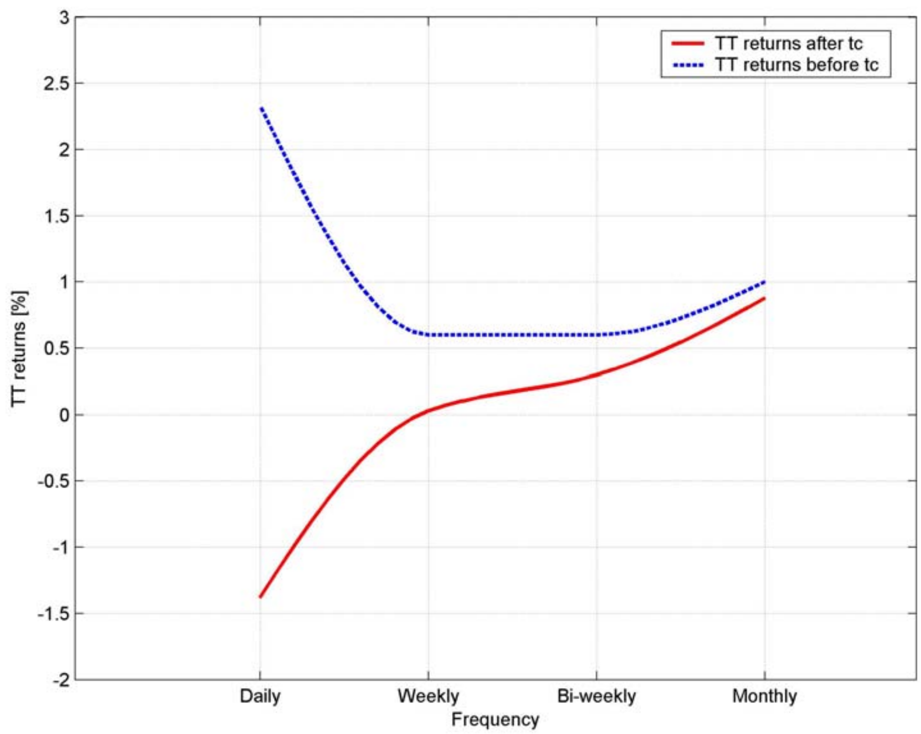

Notes: The dashed line represents multiscale (from the daily to the monthly horizon) average technical trading returns using the Canada/U.S. dollar returns before transactions costs are taken into account. The solid line represents multiscale (from the daily to the monthly horizon) average technical trading returns using the Canada/U.S. dollar returns after transactions costs.

Tables

Table 1 – Profits generated by technical trading rules

	MACO (1,50)	MACO (10,50)	MACO (5,50)	TRBO (5)	TRBO (10)	TRBO (20)	Average
Daily Data							
Return	-4.31%	-0.40%	-0.49%	-3.04%	-1.11%	1.09%	-1.38%
CC	-4.42%	1.21%	1.59%	-3.57%	-0.50%	1.56%	-0.69%
FD	-5.76%	-1.32%	-0.58%	-4.68%	-1.61%	-0.36%	-2.38%
Weekly							
Return	-1.30%	0.53%	0.62%	-0.88%	-0.07%	1.26%	0.03%
CC	0.90%	1.76%	1.83%	1.48%	1.50%	1.90%	1.56%
FD	-7.00%	5.44%	1.11%	-5.84%	1.33%	2.88%	-0.35%
Bi-weekly							
Return	-0.95%	-0.13%	1.05%	-0.35%	0.69%	1.51%	0.30%
CC	1.38%	1.66%	1.85%	1.64%	1.74%	1.21%	1.58%
FD	0.03%	-0.03%	-0.01%	0.54%	0.67%	0.16%	0.23%
Monthly							
Return	0.25%	1.13%	1.52%	0.72%	1.03%	0.64%	0.88%
CC	1.14%	1.55%	0.95%	1.37%	1.25%	1.02%	1.21%
FD	0.15%	-0.01%	-0.01%	0.49%	0.08%	0.14%	0.14%

Table 2 – Profits before transaction costs generated by technical trading rules

	MACO (1,50)	MACO (10,50)	MACO (5,50)	TRBO (5)	TRBO (10)	TRBO (20)	Average
Daily Data							
Return	2.98%	2.28%	1.36%	2.83%	1.89%	2.62%	2.33%
CC	0.97%	3.06%	2.52%	0.90%	1.56%	2.61%	1.94%
FD	0.21%	0.78%	0.61%	0.29%	0.99%	1.05%	0.66%
Weekly							
Return	0.03%	0.91%	0.85%	0.05%	0.38%	1.40%	0.60%
CC	1.56%	1.89%	1.95%	2.03%	1.76%	1.99%	1.86%
FD	-2.12%	6.78%	1.71%	-1.36%	3.50%	3.47%	2.00%
Bi-weekly							
Return	-0.19%	0.07%	1.17%	0.12%	0.88%	1.58%	0.60%
CC	1.64%	1.73%	1.93%	1.86%	1.83%	1.26%	1.71%
FD	0.33%	0.03%	0.00%	0.81%	0.75%	0.17%	0.35%
Monthly							
Return	0.53%	1.21%	1.57%	0.87%	1.10%	0.70%	1.00%
CC	1.28%	1.61%	1.00%	1.45%	1.32%	1.06%	1.29%
FD	0.24%	0.00%	0.00%	0.57%	0.09%	0.15%	0.17%

Table 3 – Number of trades generated by technical trading rules

	MACO (1,50)	MACO (10,50)	MACO (5,50)	TRBO (5)	TRBO (10)	TRBO (20)	Average
Return data							
Daily	584	204	145	1,003	727	529	532
Weekly	36	35	24	203	133	81	85
bi-weekly	9	8	11	91	53	38	35
Monthly	1	0	1	44	30	25	17
CC data							
Daily	435	141	71	869	665	528	452
Weekly	12	11	11	165	132	113	74
bi-weekly	6	5	7	74	59	51	34
Monthly	1	0	2	35	31	23	15
FD data							
Daily	482	161	91	867	650	506	460
Weekly	12	11	4	162	123	99	69
bi-weekly	3	2	0	80	61	52	33
Monthly	1	0	0	40	33	33	18

Table 4 – Number of trades generated versus expected number of trades

	Average	Expected Trades	Difference
Return data			
Daily	532		
Weekly	85	111	(25)
bi-weekly	35	55	(20)
Monthly	17	26	(9)
CC data			
Daily	452		
Weekly	74	94	(20)
bi-weekly	34	47	(13)
Monthly	15	22	(6)
FD data			
Daily	460		
Weekly	69	96	(27)
bi-weekly	33	48	(15)
Monthly	18	22	(4)

Table 5 – Sharp ratios of profits generated by technical trading rules

	MACO (1,50)	MACO (10,50)	MACO (5,50)	TRBO (5)	TRBO (10)	TRBO (20)	Average
Daily Data							
Return	(0.048)	(0.009)	(0.009)	(0.034)	(0.014)	0.007	(0.018)
CC	(0.065)	0.012	0.020	(0.004)	(0.011)	0.017	(0.005)
FD	(0.055)	(0.016)	(0.008)	(0.044)	(0.017)	(0.007)	(0.024)
Weekly							
Return	(0.051)	0.010	0.017	(0.044)	(0.010)	0.035	(0.007)
CC	0.031	0.078	0.083	0.048	0.060	0.083	0.064
FD	(0.033)	0.013	0.001	(0.034)	(0.001)	0.004	(0.008)
Bi-weekly							
Return	(0.055)	(0.016)	0.051	(0.027)	0.026	0.072	0.008
CC	0.086	0.112	0.136	0.109	0.121	0.083	0.108
FD	(0.006)	(0.008)	(0.001)	0.013	0.018	(0.001)	0.003
Monthly							
Return	0.010	0.075	0.020	0.042	0.068	0.040	0.042
CC	0.008	0.008	0.007	0.008	0.008	0.007	0.008
FD	0.001	(0.003)	(0.059)	0.016	(0.002)	0.005	(0.007)

Table 6 – Sharp ratios of profits before transaction costs generated by technical trading rules

	MACO (1,50)	MACO (10,50)	MACO (5,50)	TRBO (5)	TRBO (10)	TRBO (20)	Average
Daily Data							
Return	0.028	0.020	0.011	0.027	0.017	0.023	0.021
CC	0.010	0.040	0.036	0.001	0.019	0.034	0.023
FD	(0.000)	0.003	0.002	0.001	0.007	0.005	0.003
Weekly							
Return	(0.007)	0.023	0.025	(0.013)	0.005	0.040	0.012
CC	0.062	0.084	0.089	0.073	0.072	0.088	0.078
FD	(0.014)	0.018	0.003	(0.017)	0.007	0.006	0.001
Bi-weekly							
Return	(0.018)	(0.006)	0.057	(0.004)	0.035	0.076	0.023
CC	0.106	0.118	0.143	0.125	0.128	0.088	0.118
FD	0.005	(0.005)	(0.000)	0.023	0.021	(0.001)	0.007
Monthly							
Return	0.031	0.080	0.023	0.054	0.074	0.043	0.051
CC	0.122	0.155	(0.025)	0.136	0.130	0.112	0.105
FD	0.006	(0.002)	(0.058)	0.020	(0.002)	0.005	(0.005)