High-Frequency Exchange Rate Forecasting:
An Autoregressive Conditional Multinomial–Autoregressive Conditional Duration Approach

Charlie X. Cai * [a]
University of Leeds

Qi Zhang [b]
University of Leeds

This draft: 03 March 2012

Abstract

We examine a practical application of the autoregressive conditional multinomial–autoregressive conditional duration (ACM–ACD) model (Russell and Engle, 2005) to high-frequency exchange rate forecasting. We have two prime objectives: to examine the model's forecasting performance and to assess the feasibility of constructing intra-day trading rules. Applying the model to three pairs of currencies, we find strong predictability in the high-frequency quote change data, with the rate of correct predictions varying from 54 to 70 per cent. We demonstrate that filtering the data, by increasing the threshold of mid-price change for evaluating price movement events, can improve forecasting performance. However, over-filtering the intra-day data reduces its information content and lowers forecasting performance. Further, we find that a trading strategy following the predictions of the model can generate an average 5 to 15 per cent daily return. However, this quickly diminishes once transaction costs as low as one basis point are considered, since such a strategy requires on average eight transactions per minute.

JEL classification:

Keywords: Foreign exchange, High-Frequency Data, Forecasting, Duration Model .

[a] *Senior Lecturer in Accounting and Finance and Deputy Director of the International Institute of Banking and Financial Services (IIBFS) Leeds University Business School, Leeds, UK*

[b] *Lecturer in Accounting and Finance, Leeds University Business School, Leeds, UK.*

* Corresponding Author: Charlie Cai, x.cai@lubs.leeds.ac.uk

High-Frequency Exchange Rate Forecasting: An Autoregressive Conditional Multinomial–Autoregressive Conditional Duration Approach

1 INTRODUCTION

Foreign exchange rate forecasting is an ongoing challenge to academics, practitioners and policy makers. A large number of studies have been dedicated to the examination of exchange rate dynamics. These studies have produced at least two main conclusions among other insights. First, although macroeconomic fundamentals should in theory dictate the purchasing power of a country's currency, they are less useful in predicting exchange rates, given the nature of low-frequency announcements[1]. Second, better information sets or modelling techniques may help improve short-term forecasting of exchange rate movements; for example, the use of customer order-flow information[2]. While the random walk model is difficult to surpass in forecasting the mean, Hong, Li and Zhao (2007) show that sophisticated time series models are preferable for out-of-sample density forecasting. However, there are some difficulties in the direct application of such models[3]. In this regard, a successful model for the mean forecasting is much more useful practically, since it can be directly applied to constructing quotation strategy for dealers and trading strategy for traders.

Given the challenge of foreign exchange rate forecasting as described above, the current paper aims to explore the use of the autoregressive conditional multinomial–autoregressive

[1] It is widely documented that traditional exchange rate determination models perform poorly in explaining and forecasting exchange rates (Meese and Rogoff, 1983; Cheung et al., 2005; Engel et al., 2008).

[2] Motivated by the market microstructure literature, Evans and Lyons (2002, 2005) show empirically that order flow is an important determinant of daily exchange rates. Several studies in this area try to bring macroeconomic and microstructure approaches together to study exchange rates (Bacchetta and van Wincoop, 2004; Evans and Lyons, 2007; Rime et al., 2010). These studies show that the short-term determinants of exchange rate movements are more likely driven by the short-term supply and demand fluctuations of the exchange rate market.

[3] Hong et al., (2007) suggest that a better density forecast can be used in many ways (risk management, option pricing etc). However, they also point out that direct application of adopting these models and evaluation of the performance in these applications requires further exploration in terms of implementation.

conditional duration (ACM–ACD) model, as introduced by Russell and Engle (2005), in high-frequency foreign exchange rate forecasting. The motivation for this is threefold.

First, following the technological advances of the last two decades, high-frequency trading at the intra-day level has become popular in financial markets. Technical trading rules embedded in so called 'algorithm trading' play an important role in financial markets. Indeed, high-frequency algorithmic trading has been judged one of the major drivers behind the growth of foreign exchange market turnover in the Bank for International Settlements (BIS, 2010) report on global foreign exchange market activity in 2010. However, most of the previous literature on intra-day exchange rate forecasting has focused on regular time intervals such as 30 minutes or one hour. None of the existing studies consider tick-by-tick frequency with irregular time spacing[4]. Our study fills this void by modelling the exchange rate movement tick by tick and evaluating the out-of-sample forecasting performance.

Second, there is empirical evidence suggesting that the ACM–ACD model performs well in high-frequency forecasting. Zhang, Cai and Keasey (2009) demonstrate, in the context of the US equity market, that it provides better forecasting performance than either the asymmetric ACD model (AACD) – an alternative asymmetric extension of the ACD model proposed by Bauwens and Giot (2003) – or the random walk model. They also show that the forecast performance of the model generally decreases as the liquidity of the stock increases, with the exception of the most liquid stocks. This suggests that there is a potential gain in forecasting performance by applying this model to the exchange rate market, the most liquid of financial markets.

Finally, the ACM–ACD model is designed for modelling irregularly spaced high-frequency dynamics. It incorporates an important set of information, the durations between price changes, into modelling the price dynamics. Theoretical studies (Easley and O'Hara, 1987) in market microstructure show that the time between intra-day market events, like trades and

[4] One exception is Engle and Russell's (1997) study, which examines high-frequency quote changes in the foreign exchange rate market using an ACD model. However, their study does not produce quote change forecasts, since it models only duration and not the direction of quote updates.

quote updates, may contain information about market states. Dealers receive information from their customer orders and they also learn about market-wide order flow from brokered inter-dealer trades (Loyns, 2001: Chapter 3). Given that order-flow information is valuable and private to the dealers, in this study we explore whether a sophisticated model utilizing both the duration and direction of quote change information, both of which are publicly available, would help to abstract private information and produce better out-of-sample forecasting of quote updates.

Specifically, we examine the following two research questions. (1) Does the ACM–ACD model provide better forecasting of the high-frequency exchange rate than a random walk model? (2) Can the forecasting provided by the ACM–ACD model be used to generate profitable trading rules?

We examine three currencies against the dollar (USD) – the euro (EUR), the British pound (GBP) and the Japanese yen (JPY) – for the year 2010. The high-frequency bid and ask quotes for spot foreign exchange rates are obtained from GAIN Capital. Using weekly and daily estimation windows, with a daily rolling frequency, we assess the out-of-sample forecasting by the ACM–ACD model at the intra-day tick-by-tick level.

Our analysis yields the following key findings. First, the ACM–ACD model provides much improved forecasting over that of the benchmark random walk model and others reported in the existing literature. The forecast accuracy of the models ranges from 55 to 63 for the three currency pairs when no data filter is applied. The correlation analyses and Wald test on the relationship between the predicted and realized quotation movements also confirm the positive and significant prediction power of the models. We further show that when a modest filter is applied to the raw data (by increasing the tick size threshold for counting the number of quote changes), the prediction accuracy increases to between 59 and 70%. However, when a higher threshold of price movement is imposed, the prediction accuracy drops to a range between 53 and 54%. This suggests that over-filtering may reduce the useful information within the data.

Second, we find that a switching strategy based on the directional forecasts can earn substantial returns before transaction costs. Specifically, we simulate a portfolio switching the position from longing one currency to longing the other when the directional forecasts change, while holding the position when same-signed forecasts are produced. This strategy can generate an average daily return, accumulated from intra-day transactions, ranging from 5 to 15%, from prediction models before filtering. These returns are much better than the buy-and-hold return, which is of the order of approximately 0.03% daily. While applying filters improves directional prediction accuracy, it does not always produce greater economic benefit. Only the return in the GBP/USD series is improved after applying data filters. Our risk adjustment tests based on the certainty-equivalent return produce similar conclusions to the portfolio return measures.

Third, we document the significant impact of transaction costs on the trading strategy. The documented before-cost return requires high-frequency automatic trading; specifically, of about eight trades per minute, according to our findings. When transaction costs of half a basis point are applied to each intra-day transaction, the profit quickly turns to loss for most of the strategies, except for the two models for the JPY using filtered data. If transaction costs of one basis point are considered, none of the models generate positive returns. Thus, although the models deliver high accuracy in predicting quotation updates, the predictions cannot be translated into abnormal profit due to the large impact of transaction costs on high-frequency trading. This is consistent with the efficient market hypothesis, which suggests that the reason why these prediction patterns exist is because it is costly to take these arbitrage opportunities.

However, this finding does not entirely undermine the usefulness of the model. A model with better prediction accuracy than those on offer in the existing literature may be used in two possible ways. First, a profitable trading strategy could be implemented if one could obtain relatively low transaction costs; this is more feasible for large institutional traders. Second, it could be used to help dealers abstract information from other dealers' quotations in the market. They would thereby be acting as front-runners, but in a legal and legitimate way; in anticipating

other dealers' quote movements, they could make better-informed decisions regarding their own quotations.

The rest of the paper is organized as follows. In Section 2 we briefly review the literature on exchange rate forecasting and introduce the ACM–ACD model in the context of high-frequency forecasting. In Section 3 we describe the research design and data. In Section 4 we present the main results, and in Section 5 we conclude.

2 THE ACM–ACD MODEL FOR HIGH-FREQUENCY FORECASTING

Due to data limitations, most of the early studies on exchange rate forecasting focus on lower frequency data (Dooley and Shafer, 1983; Sweeney, 1986; Levich and Thomas, 1993; Neely et al., 1997). In the past two decades, advances in computer technology have improved the infrastructure of the trading platform. Due to the resultant increase in the practical importance of high-frequency trading, there is a growing literature on the predictability of price changes both at intra-day level (Osler, 2000, 2001; Curcio et al., 1997; Neely and Weller, 2003) and at ultra-high frequency in milliseconds (see Brogaard (2010) for a useful survey of these ultra-high-frequency trading rules).

One of the main findings to emerge from these studies is the importance of order-flow information to short-term exchange rate movements. Many studies document that order flow is useful in predicting exchange rate movements in both short and long horizons (e.g., Evans, 2002; Evans and Lyons, 2002; Payne, 2003). Furthermore, order flows are key media for transmitting macroeconomic announcements into exchange rate movements (Evans, 2002; Evans and Lyons, 2002). However, using order flow for predicting future price presents two challenges. First, order flow information is private to the dealers; it contains information regarding customers' expectations about future exchange rate movements. Only dealers and brokers are privileged to observe some of these order flows. Clients expect their trading interests to be treated as highly confidential and, in general, there are rules that prohibit dealers and brokers from profiting

directly from the misuse of this confidential information[5]. Nevertheless, the privilege of being able to observe order flow can indirectly benefit dealers. It is shown that dealers do consider order flow to be informative and that they set prices accordingly (Lyons, 1995). Second, and more important, order flows are not all equally informative. Identifying which orders are the most informative requires further information regarding the identity behind the order. In other words, private information is required, in order to be able to interpret the observed order flow. Therefore, direct access to order-flow information without the contextual knowledge of a dealer to interpret it will not necessarily lead to more informed trading decisions.

Given that there is no publicly available data for order flow, and given the challenge of abstracting information from such data, we explore forecasting in the foreign exchange (FX) market based on dealers' quotes. In contrast with order flow, FX quotations – like those in other financial markets – are publicly available, since advertising their quotations can attract order flow for dealers. We study whether, from the dynamic adjustments in a quotation, we can abstract the private information of the dealers. With advances in the development of financial econometrics, we are able to explore this dynamic in two important dimensions: duration and quotation changes. That is to say, we are able to jointly model when the next quotation will be posted and, conditional on the duration and past quotation changes, the direction in which it will change.

Quotation updates by dealers are inherently in irregular time intervals. This irregularly spaced data presents a challenge to econometric modelling, since the majority of time series techniques are based on equally spaced intervals between data arrivals. As Engle and Russell (1997) point out, the choice of analysis interval is itself a very important part of the analysis, since the interval must yield a balance between noise (heteroscedasticity) and information. In this regard, Engle and Russell's (1998) autoregressive conditional duration (ACD) model is advantageous in modelling such data. They propose a point process model for inter-temporally correlated event arrival times to model data arriving at unequal time intervals.

[5] See Harris (2003: Chapter 11) for a good discussion on legal and illegal front-runner activity in the financial trading industry.

Another characteristic found in financial data is the discreteness of price change due to the minimum tick size and the human tendency to use rounding. Russell and Engle (2005) propose a new modelling framework to jointly model price changes and durations. They decompose the joint distribution into the product of the conditional distribution of the price change and the marginal distribution of the time interval. Specifically, they model the former by means of the autoregressive conditional multinomial (ACM) model and the latter by means of the ACD model (Engle and Russell, 1998).

Following Engle and Russell (1997, 1998), we use the thinning algorithm to model the intensity of quote arrivals signifying price change. This model facilitates prediction of the length of time before prices change. Furthermore, the market microstructure model of Easley and O'Hara (1992) clearly suggest links between quotation update and the arrival of new information. Therefore, the ACM–ACD model of quotation changes can also be seen as a model that captures the rate of information arrival.

We lay out the details of our empirical design in the next section.

3 EMPIRICAL DESIGN

In this section we first explain the procedure for applying the ACM–ACD model to the out-of-sample forecast. We then propose measurements and tests for evaluating forecast performance both statistically and economically. Finally, we present the exchange quote data used in this study.

3.1 MODELLING AND FORECASTING DESIGN

We present a detailed description of the ACM–ACD model setup in the Appendix. In this section, we focus our discussion on the application of the model to out-of-sample forecasting. In essence, our objective is to use the ACM–ACD approach to model the dynamic structure of the quote change process and use the estimated dynamics to generate forecasts. We take the following considerations into account for our empirical design.

First, we consider the options for the length of the estimation window. Zhang et al., (2009) demonstrate that the ACM–ACD model captures the characteristics of the micro data very well with a sample length of 20 days, in the context of the US equity market. However, in comparison to the equity market, the foreign exchange rate market is much more liquid and fast moving. For example, there are more than 8000 quotations in the most active period of a FX trading day in our sample. Furthermore, the trading week is a natural frequency cycle in the FX market. Therefore we choose to examine daily and weekly estimation windows.

Second, we consider how often the estimation of the model should be updated for out-of-sample forecasting. In low-frequency contexts, modelling often rolls by taking one observation at a time. Such a rolling strategy is not practical in a high-frequency context, given that the time per estimation is much longer than the duration between observations. Typically, for a model using one-day data it takes about 5 seconds to estimate the parameters, while the median duration of quote change is only about one second; thus, by the time one obtains the forecast for the 'next' step price change, the realized price change has already moved several steps ahead. We therefore opt for a daily rolling window. The out-of-sample forecast on each day is generated using the model calibrated with data available from the day before.

Specifically, at the beginning of every trading day, we use the past day's (or five days') intra-day data to fit the ACM–ACD model and obtain the model parameters. From the second quotation change of the day after 2am Eastern Time, we make predictions about directional change in the next quotation, by applying the model parameters with the lag duration and directional information. We record all predictions and update the model estimates at the end of the trading day for the next day's forecasts. Detail regarding the definition of a trading day is given in Section 3.3, where we discuss the data.

Finally, we consider the criteria used for recording a quote price change. The default threshold used in our analysis is zero change; in other words, any quotation update leading to a change in the mid-quote level would be recognized as an event. A threshold as low as zero would encompass many events containing very small price changes; potentially, these may contain less

information. Furthermore, when such a model is used for forecasting, a price change signal may be correctly generated but the realized price change may yet be very small, due to the information used to build the model. Such forecasting may not produce economic significance. There are therefore potential gains to be had from increasing the threshold for recording events before fitting them to the models.

By contrast, there is a danger of over-filtering, where information is greatly reduced as the result of a high threshold being applied. This effect is documented in Zhang et al. (2009), who conclude that the optimal threshold to apply to ACM–ACD forecasting ofto apply in NYSE stocks is 1/8; whereas higher thresholds, such as ¼, reduce forecasting performance. We therefore explore the effect of different thresholds (filters) on the model's performance.

3.2 EVALUATING FORECASTING PERFORMANCE

We evaluate the performance of the ACM–ACD model from two angles: statistical and economic significance. First, we examine the statistical significance of the model's forecasting accuracy. To this end, we construct a variable (*correct*) to measure the percentage of forecasts that are in the same direction as the realized quote change. Following Dacorogna et al. (2001), we judge whether a forecasting model is better than the random walk by examining whether the value of *correct* is significantly higher than 50%. We approximate the confidence level of *correct* by the 99.99% confidence level of the random walk model.

(1) $$\sigma_{0.01\%} \approx \frac{3.719016}{2\sqrt{n}}$$

where n is the number of predictions. The factor 2 comes from the assumption of an equal probability of having positive or negative signals; a setup analogous to that of evaluating the fairness of a coin by comparing coin-tossing outcomes to these confidence levels. When the value of (*correct* minus 50%) is larger than $\sigma_{0.01\%}$, we conclude that the forecasting model is significantly better than a random walk.

In addition, we calculate the correlations between the prediction signals and the realized directions in order to evaluate the capacity of the model to predict the realized directions. We

also estimate logistic models with realized direction as dependent variable and predicted direction as independent variable. Wald test is used to test the significance of the coefficient on the predicted direction variable.

Second, it is suggested that a model which has apparently performed well according to statistical evaluation criteria does not necessarily generate economic benefit (Brooks et al., 2001). We perform our final evaluation by measuring the profitability of the trading strategy based on the model's forecasts. Specifically, we construct a portfolio following the buy–sell signals generated by the model.

For each day, we obtain the model parameters and apply them with the duration and direction data to form a one-step forecast. To illustrate the construction of the portfolio, we might consider the EUR/USD quote series. If the model predicts an appreciation in the EUR against the USD, we will long the EUR. Since the model parameter is estimated overnight and fixed during the day, the signals generated take very little time on receiving a quotation change. We assume that one would be able to transact at the new quotation midpoint. (Transaction costs are discussed later). We hold the currency until the next signal changes. The predictions are updated whenever a quotation change event occurs. If a new predicted direction is different from the previous signal, we reverse our position; otherwise, we continue holding the existing position.

We then collect the return for each transaction. We report the average return of each transaction, the average interval of trades, and the average accumulated daily return. As a benchmark, we implement a buy-and-hold strategy whereby a currency is bought and held throughout a trading day. We show the return for holding either currency in a given pair. We calculate the returns and variance for each strategy. We compare the utility value of the portfolios by calculating the certainty-equivalent return (CER) as follows:

(2) $$CER_{S,\gamma} = \bar{\mu}_S - \frac{\gamma}{2}\sigma_S^2 ,$$

where $\bar{\mu}_S$ and σ_S^2 are the mean and variance of the daily returns from each strategy, *S*. γ is the risk aversion coefficient, which takes on a value of 1, 3 or 5.

3.3 DATA

We use high-frequency bid and ask quotes from GAIN Capital for spot foreign exchange rates between our three currency pairings against the dollar (USD) – the euro (EUR), the British pound (GBP) and the Japanese yen (JPY) – during the year 2010.

Preliminary examination of the data shows that quotation activity is mainly concentrated between 2am to 1pm Eastern time. Similar to previous studies (Neely and Weller, 2003; Rime et al.,2010), we choose to model the intra-day dynamic over the period when the market is widely active. This period covers the end of Asian trading hours, the whole trading session in Europe and a substantial part of US trading hours. The quotation file contains timestamps in seconds. For multiple quote changes within a second, we obtain an average mid-quote price to capture the overall information during that second.

****

Table 1 about here

****

Table 1 reports the summary statistics of the price and return of the exchange rate. Panel A shows the range and distribution of the exchange rate during the sample period of 2010. Panel B shows the mid-quote percentage changes in basis points. N represents the number of quote changes in the sample.

Panel A of Table 1 shows that 2010 covers a wide range of exchange rate movement for the three paired currencies against the USD. We present the price level here mainly as a reference point for our discussion of the threshold for filtering quote changes.

Panel B contains more information about the dynamics. Given the quotation style, the quote changes are calculated based on the quotation measuring the return of longing the EUR and GBP, and shorting the JPY. Overall, both the EUR and GBP depreciate against the USD during the sample period, while the JPY appreciates. The average price changes between quotes are very small even when expressed in basis points. The largest jumps for the EUR and GBP are

in the order of 20 to 30 basis points, while quote changes at the extremes are much larger for the JPY.

4 RESULTS

4.1 FORECASTING ACCURACY

Before reporting the estimation results, we present a summary of the durations for the mid-quote movements in our sample. Table 2 summarizes the statistics under different model specifications. Durations are measured in seconds except for the maximum value, which is measured in hours. The Filter column indicates the threshold parameters used to filter the quotation changes. Only movements greater than the filter parameter are recorded as quotation change events. The filters in Table 2 are expressed as 10,000 times the quotation unit; for example, the threshold of 1 for the EUR implies that only when a quotation change exceeds 0.0001 USD will it be recorded as a quotation change event. The Est. Win. column shows the number of days used to estimate the models. The N column indicates the total number of quotation change events recorded according to the filtering rules.

Table 2 shows that the mean duration is skewed by large maximum values. The median duration shows that quotes are updated typically every two to three seconds and most often every second. Filtering the price change increases durations and reduces the total number of events, as expected. When a relatively large threshold is applied, the median durations become significantly higher and the number of events drops dramatically, as shown in the last row of each currency pair. The drop in the number of events is also due to the fact that we have to retain two more hours in each day to use for estimating the initial value of the day. In contrast, more modest threshold values produce small increases in duration, while the falls in the number of events are still substantial. These variations in duration and number of events provide a wide range of conditions for testing the sensitivity of the model to the information-to-noise ratio of the data.

****

Tables 2 and 3 about here

****

Table 3 reports the percentage of correct directional forecasts of the models. A forecasting signal is generated only when there is a change in predicted direction; that is, only predictions requiring a change of position would be recorded as a signal. For example, when a predicted quote movement changes from negative to positive, a signal of positive quote change is generated. The variable *correct* is assigned a value of 1 when the predicted signal bears the same sign as the realized quotation change, and zero otherwise.

Comparing the number of signals in this table to the number of quote changes in Table 2, we can see that there is a substantial reduction in the number of observations in Table 3; this suggests positive autocorrelation in the predictions (that is, many consecutive same-signed predictions).

The Mean column shows the mean of the variable *correct* for all the signals generated during the sample period, which is a measure equivalent to the percentage of correct sign predictions by a given model. It shows that all the models produce a correct ratio of more than 50%, the benchmark ratio for a random walk sequence. Given the large values of N, we apply a much stricter significance level of 0.01% for testing the statistical significance of these predictions. Comparing the mean value with the confidence interval suggests that these predictions are significantly better than those of a random walk. Correlation analysis confirms that these predictions are capable of anticipating the same direction as the realized quote changes: all correlations are positive and significant, and the best model has a correlation as high as 35%. Furthermore, the Wald test suggests strong predictability of the models, with all results significant at 0.01%.

These findings provide strong support for the use of the ACM–ACD approach for high-frequency exchange rate modelling. The strength of these predictions is much better than those reported in the existing literature. For example, Dacorogna et al. (2001: Chapter 9) report

correct ratios, for their out-of-sample predictions relating to ten USD currency pairs, ranging from 50.5 to 54.2% and having a maximum value of 6% for the correlation between predictions and realized values. Furthermore, our model performance is comparable with the best results reported in Zhang et al. (2009), who study 50 NYSE equities out of sample forecast using the ACM–ACD model. They report a mean forecast accuracy of 54% and a maximum of 69%. They further show that prediction accuracy generally decreases as the quotation update frequency increases, except for the group of stock with the highest quotation update frequency, which has a mean number of daily update of 780. In comparison to equity trading, FX trading is much higher frequency and faster moving. Our findings thus complement their study, showing that the ACM–ACD model performs equally well with very high-frequency data.

When estimation windows are extended from one day to five days, the model performances are generally improved but only marginally. This suggests that a five-day window, covering a full week's pattern of trading, contains more information than a single trading-day window. We therefore maintain a five-day estimation window for most of the models.

We find that the largest performance improvement comes from applying filters to the raw quote changes. Table 3 shows that filtering the raw data increases the information-to-noise ratio, and that there is an especially large improvement when the first step of filtering is applied. Both the mean measure of accuracy and the correlations show substantial improvement in all currency pairs. Furthermore, since the JPY's quotation unit is different, it allows us to explore more steps in the thresholds. The model performance is at its best when a filter of 0.01 is applied in the JPY series, producing a 70% accuracy rate. However, when larger filters are applied, performances decline; indeed, those models using the largest threshold performed worse than the initial model without filtering.

Overall, the out-of-sample forecast performance of the ACM–ACD models is statistically significant and markedly superior to those in the existing literature. Forecast performance is further enhanced when modest filters are applied.

In the next section we further evaluate the economic significance of the forecasts.

4.2 ECONOMIC SIGNIFICANCE

We examine the economic significance of the forecast model by considering the possibility of profiting from these predictions. We first review the before-cost portfolio return and then examine the effect of transaction costs.

4.2.1 BEFORE TRANSACTION COSTS

Table 4 reports the performance of the long–short strategies based on the model forecast. For reference we also record the benchmark buy-and-hold (B&H) strategy for each currency of the pairs. Table 4 shows that the mean daily returns of the strategies are very high; for example, for the first three models for the EUR, the average daily accumulated returns are all higher than 10%. This is in strong contrast to the B&H strategies, where a daily return of less than 0.03% is recorded. Generally, the CERs of the strategies based on the model predictions are positive for low risk aversion, with the exception of the models that applied large thresholds; whereas all of the B&H strategies result in a negative CER. Similar contrasts in performance between our model forecast and the B&H benchmarks are found in the other two pairs. However, we should remember that these returns do not take into account transaction costs. Before we examine the impact of transaction costs in detail, in the next subsection, we note that the frequency of trades is very high. A small percentage of per-transaction costs would reduce the return significantly. For example, the duration of the transactions indicates that the best performance model for the EUR requires a transaction every eight seconds. This implies a requirement of about 7.5 trades per minute, 427 trades per hour or 4270 trades in the 10-hour period per day.

Interestingly, although we find that applying filters can improve forecasting accuracy, this does not necessarily translate into economic benefit. This is especially so in the case of the EUR and JPY, where the best forecasting models in terms of accuracy are those with modest thresholds while the best return is obtained by the those models without filtering. The contrasts are sharper when the CER measures are considered. Analysis suggests that strategy based on models without filtering can produce higher returns with lower standard deviations. On the

other hand, models with the best forecasting accuracy may perform very badly in economic terms. For example, the best JPY model (threshold of 100 and Est. Win. of 5) in Table 3 achieves a very low return and CER in Table 4. This finding is consistent with Brooks et al. (2001), who study technical trading rules and find that a model that appears to have performed well according to the statistical evaluation criteria does not necessarily generate economic benefit. In the context of this research, the inferior return performance of the best directional forecast model suggests that there is asymmetry in the return for correct and incorrect sign predictions. Although the filtered models can predict direction more accurately, when their prediction is wrong, the size of negative impact is much larger than the size of positive impact when they get it right.

Overall, Table 4 shows that the ACM–ACD model performs well economically. Long–short strategies generate very high accumulated daily returns before transaction costs. Furthermore, the model with the best forecasting accuracy does not produce the best return distribution.

In the next section we consider the impact of transaction costs.

4.2.2 Effect of Transaction costs

Table 5 reports the effect of transaction costs on the long–short strategies based on the model forecasts. We use half a basis point as the measure of per-transaction cost in this illustration. Specifically, half a basis point is taken away from each pre-cost return before accumulating the returns. We can see that this small transaction cost wipes out all profitability obtained for all the EUR and GBP models. The only two models that remain positive are for the JPY with filtering. This demonstration highlights two important points. First, given the high frequency of trading required for these strategies, profitability is difficult to obtain unless very low or zero transaction costs are attainable for these algorithm trades. To our knowledge, the lowest spread in the FX market documented in public literature is one basis point in the direct inter-dealer market (Lyons, 2001). When we apply transaction costs of one basis point, none of the above models produce a positive return. There is thus a significant barrier to entry preventing such models

from being widely used in practice. This could partly explain why such high predictability and high pre-cost profitable opportunities are unarbitraged away. However, this does not negate the potential application of such forecasting models in practice. Since 2001, when Lyons wrote his book, the increased popularity of electronic platforms (BIS, 2010) may have driven transaction costs down to a lower level. More importantly, a model with such high levels of accuracy in predicting quotation changes would be useful to dealers who are obligated to make the market with their high-frequency quotation updates. For example, if a downtick forecast is generated by the model for the USD/EUR pair, a dealer who has a net position in the EUR would want to reduce her position by posting a lower ask (selling) price but not as low as the expected price move. In this way, she would be able to attract order flow from customers buying the EUR and selling at a better price.

Second, the results in Table 5 again show the advantage of data filtering. In the two sections above, we found that filtering raw data with some modest thresholds improves forecasting accuracy. However, it does not necessarily generate high pre-cost profit. When transaction costs are considered, trading frequency plays a much more important role, given its compounding effect. It shows that a second benefit of applying a threshold, in addition to the increase in the information-to-noise ratio in the point process, is to reduce the number of transactions and hence lower the transaction costs of the trading strategy.

Overall, we show that transaction costs have a detrimental effect on the profitability of high-frequency trading strategies. On the one hand, this finding supports market efficiency, in the sense that no abnormal profit can be systematically generated using a prediction model. The existence of predictability is partly driven by the transaction cost, which prevents information being fully reflected in the price[6]. On the other hand, the finding highlights an important condition for the practical application of these prediction models: if they are to be used for trading strategy, profitability is possible only with very low transaction costs.

[6] The previous literature mainly argues that transaction costs cause the slow adjustment to a small mispricing (e.g., Sercu et al., 1995; Panos et al., 1997; Roll et al., 2007 and Oehmke, 2010).

5 CONCLUSION

This paper takes up the challenge of forecasting exchange rate movement tick by tick. Using the ACM–ACD approach (Russell and Engle, 2005), a high-frequency financial econometric model that is capable of measuring a marked point process with irregular arrival times, we examine the accuracy and economic significance of the forecasts in the context of the foreign exchange rate market. Specifically, we study the quotation changes of the three most traded currency pairs: the EUR, JPY and GBP against the USD, in the one-year period of 2010. Our empirical analyses produce the following findings.

First, the ACM–ACD model provides much improved forecasting over the benchmark random walk model and those models reported in the existing literature. The forecast accuracies of the models studied range from 54 to 70%. Correlation analyses and the Wald test on the relationship between predictions and realized quotation movements also confirm the positive and significant prediction power of the model.

Second, we find that a switching strategy based on the directional forecasts could earn substantial returns before transaction costs. For example, the strategies can generate an average daily return, accumulated from intra-day transactions, ranging from 5 to 15% without applying filters. These performances are much better than the buy-and-hold return, which is in the order of about 0.03% daily. However, when transaction costs of half a basis point are applied to each intra-day transaction, the profits quickly turn into loss for most of the model predictions except for the two models for the JPY that use filtered data. If transaction costs of one basis point are considered, none of the models generate positive returns. Therefore, although these models produce high accuracy in predicting the quotation update, these predictions cannot turn into abnormal profit due to the large impact of transaction costs on high-frequency trading. This is consistent with the efficient market hypothesis. The reason why these prediction patterns exist is because it is costly to take these arbitrage opportunities.

However, these findings do not preclude the future application of the ACM–ACD model in practice. The increased use of electronic trading platforms has in general driven transaction

costs down (BIS, 2010). Hence, there is the possibility that some institutions, who can obtain transaction costs as low as half a basis point, will be able to enjoy the economic benefit from this forecast model. Furthermore, the model could certainly provide advantageous information regarding price movement to foreign exchange dealers, when applied to a series of quotation changes that contain multiple dealers' quotes. Dealers are obligated to quote on both sides of the market in order to attract order flow. A good forecast model would assist them in both maximizing the gains from the spread and managing their inventory risk by setting quotations that are of benefit to them, given the forecast price movement in the next tick.

Third, we examine the impact of data filtering on the model's forecasting performance. The ACM–ACD model and especially the ACD model are designed to model an event that arrives in irregular time intervals. The realized sequence such as the quotation changes may contain much noise; for example, there may be quote updates that contain no changes in the quote price. Engle and Russell (1997) point out the importance of filtering the data in empirical applications of the ACD model. We take this point further and examine the impact of higher thresholds on the marked point process. When applying a threshold to the quote change, it effectively redefines the events. It requires that only quote changes larger than the threshold parameters would be recorded as event points, and the marked point process is constructed accordingly.

Our empirical study shows two important impacts of data filtering with a threshold parameter. First, applying a modest threshold would increase the information-to-noise ratio, while too high a threshold would reduce the information content of the point process. There is thus a balance between the reduction of noise and the reduction of information when performing data filtering; seeking the optimal threshold is an empirical challenge in itself. Second, from the point of view of trading strategy, there is an important benefit in using thresholds, in addition to the improved prediction accuracy: it reduces the number of transactions required, since the duration between the signals generated by the model is longer. A smaller number of transactions means lower transaction costs overall; an especially significant

consideration given the compounding effect of per-trade transaction costs on the overall profitability of the strategy.

Overall, we show that the ACM–ACD approach performs good out-of-sample forecasts when applied to the highest-frequency traded financial market: the foreign exchange rate market. Its advanced dynamic modelling structure appears capable of abstracting private information from the publicly observable data. However, it is only possible to benefit economically from these forecasts when very low transaction costs are obtainable. These findings support the efficient market hypothesis in general, while providing empirical evidence that reveals the potential of the practical applications of the model for large institutional FX traders and dealers.

APPENDIX: THE ACM-ACD MODEL SETUP

The detail of the ACM–ACD model settings is discussed below. Following Russell and Engle's (2005) notation, let t_i denote the time that t^{th} market event occurs. In the context of the current research, a market event is a change of the mid-quote price. Let τ_i denote the duration between two market events occurring at times t_{i-1} and t_i. And let y_i denote the price change of this market event.

The choices for construction of the y_i and τ_i series are various. y_i can have two dimensions, comprising only up and down tick price movements; three dimensions, comprising up, down and zero tick movements; or five dimensions, as in Russell and Engle (2005) comprising up two ticks, up one tick, zero ticks, down one tick and down two ticks. τ_i can measure trade duration or quote revision duration. Furthermore, construction of the price movement y_i and duration τ_i series can be conditional on the thresholds of other variables. The most obvious choice of threshold variable is the price change itself. For example, y_i only contains those points with price movements larger than a given number of ticks.

Zhang et al. (2009) point out two advantages of choosing quote change as a market event. First, use of the mid-quote change removes the microstructure effect of bid–ask bounce and, therefore, reduces the noise in the data. Second, forecasting mid-quote change has a practical advantage over forecasting trade price change, in that the forecast midpoint constitutes a useful reference point for investors, regardless of whether they are buyers or sellers.

In addition, our discussion in Section 2 points out that a potential benefit of modelling the dealers' quote change behaviour is that information private to the dealers may be abstracted from duration modelling. Therefore, we choose to model the mid-quote change and duration in our study.

We let y_i denote the direction of change of mid-quote price, giving y_i two dimensions[7]:

y_i =1 if the mid-quote price increases over duration τ_i, and

y_i =-1 if the mid-quote price decreases over duration τ_i.

[7] For simplicity, in this paper we use 'mid-quote price change' and 'price change' interchangeably.

The ACM–ACD model decomposes the joint conditional density of y_i and τ_i into the product of the conditional density of the mark and the marginal density of the arrival times, both conditional on the past information set:

$$f(y_i, \tau_i \mid y^{(i-1)}, \tau^{(i-1)}) = g(y_i \mid y^{(i-1)}, \tau^{(i)}) q(\tau_i \mid y^{(i-1)}, \tau^{(i-1)}), \quad \text{(A1)}$$

Where $g(\cdot)$ denotes the probability density function of the price change y_i conditional on $\tau^{(i)}$ and $y^{(i-1)}$, and $q(\cdot)$ denotes the density function of the ith quote duration conditional on $\tau^{(i-1)}$ and $y^{(i-1)}$.

Russell and Engle (2005) use the autoregressive conditional multinomial (ACM) specification to model $g(\cdot)$. In the two-dimension case, let $\tilde{x}_i$ be a 2*1 vector indicating the possible states of the discrete mid-quote price change y_i. $\tilde{x}_i$ takes the jth column of the 2*2 identity matrix if the jth state occurred. Let $\tilde{\pi}_i$ denote a 2*1 vector of conditional probabilities associated with the states. The jth element of $\tilde{\pi}_i$ corresponds to the probability that the jth element of $\tilde{x}_i$ takes the value 1. $\tilde{\pi}_i$ must satisfy the following two conditions:

1. All elements of $\tilde{\pi}_i$ are non-negative.
2. All columns must sum to unity.

Russell and Engle (2005) propose using an inverse logistic transformation that imposes such conditions directly. Let π_i denote the probability of an up movement of mid-quote price, and define the log of the probability ratios as $h(\pi_i) = \ln(\pi_i/(1-\pi_i))$. The ACM model assumes that $h(\pi_i)$ follows an autoregressive conditional specification:

$$h(\pi_i) = c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1}) \quad \text{(A2)}$$

The conditional probability of an up movement of mid-quote price is easily recovered from the above inverse logistic transformation using the following equation:

$$\pi_i = \frac{\exp[c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1})]}{1 + \exp[c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1})]} \quad \text{(A3)}$$

And the probability of a down movement of mid-quote price is calculated by $1-\pi_i$. Given the initial condition, the entire path of $\tilde{\pi}_i$ can be constructed, and the log likelihood is expressed as:

$$L = \sum_{i=1}^{N} \sum_{j=1}^{K} (\tilde{x}_{ij} \ln(\tilde{\pi}_{ij})) = \sum_{i=1}^{N} \tilde{x}_i' \ln(\tilde{\pi}_i) \quad \text{(A4)}$$

In the above equation, $\tilde{x}_i$ represents the realized observations and $\tilde{\pi}_i$ contains the variables needing to be estimated. In Russell and Engle (2005), durational dynamics are specified by an exponential ACD model. This means that the hazard function is constant. Here, the conditional density function is generalized to the Weibull Log-ACD model. The log likelihood associated with the ith duration is:

$$\ln(q(\tau_i \mid y^{(i-1)}, \tau^{(i-1)})) = \ln(\frac{\gamma}{\tau_i}) + \gamma \ln(\frac{\Gamma(1+1/\gamma)\tau_i)}{\psi_i}) - (\frac{\Gamma(1+1/\gamma)\tau_i)}{\psi_i})^{\gamma} \quad \text{(A5)}$$

Where:

$$\ln(\psi_i) = \omega + \sum_{j=1}^{u} \alpha_j \varepsilon_{i-j} + \sum_{j=1}^{v} \beta_j \ln(\psi_{i-j}) \quad \text{(A6)}$$

The estimation of the ACM–ACD model can be performed by maximizing the sum of the ACM and ACD log-likelihood functions or separately maximizing the two log-likelihood functions.

REFERENCES

Bacchetta, P., van Wincoop, E., 2004. A scapegoat model of exchange-rate fluctuations. American Economics Review 94, 114–118.

Bauwens L, Giot P., 2003. Asymmetric ACD models: introducing price information in ACD models. Empirical Economics 28: 709–731.

Brogaard, J. A., 2010. High frequency trading and its impact on market quality. Working paper, Northwestern University.

BIS, 2010 'The $4 trillion question: what explains FX growth since the 2007 survey?' Bank for International Settlements, http://www.bis.org/publ/qtrpdf/r_qt1012e.htm.

Brooks, C., Rew, A.G., Ritson, S., 2001. A trading strategy based on the lead-lag relationship between the spot index and futures contract for the FTSE100. International Journal of Forecasting 17, 31-44.

Cheung, Y.W., Chinn, M.D., Garcia Pascual, A., 2005. Empirical exchange rate models of the nineties: are any fit to survive? Journal of International Money and Finance 24, 1150–1175.

Curcio, R., Goodhart, C., Guillaume, D., Payne, R., 1997. Do technical trading rules generate profits? Conclusions from the intra-day foreign exchange market. International Journal of Finance and Economics 2, 267–280.

Dacorogna, M.M., Gencay, R., Muller, U., Olsen, R.B., Picktet, O.V. 2001. An introduction to high frequency finance, Academic Press.

Dooley, M.P., Shafer, J.R., 1983. Analysis of short-run exchange rate behavior: March 1973 to November 1981. In: Bigman, T., Taya, T. (Eds.), Exchange Rate and Trade Instability: Causes, Consequences and Remedies. Harper Business, pp. 43–69.

Easley D, O'Hara M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics19, 69–90.

Easley D, O'Hara M., 1992. Time and the Process of Security Price Adjustment, Journal of Finance 47, 577-605.

Engel, C., Mark, N.C.,West, K.D., 2008. Exchange rate models are not as bad as you think. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2007. MIT Press, Cambridge, MA, pp. 381–441.

Engle, R., Russell, J., 1997. Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model. Journal of Empirical Finance 12, 87-212.

Engle R, Russell J., 1998. Autoregressive conditional duration: a new approach for irregularly spaced transaction data. Econometrica 66, 1127–1162.

Evans, M.D.D., 2002. FX trading and exchange rate dynamics. Journal of Finance 57, 2405–2447.

Evans, M.D.D., Lyons, R.K., 2002. Order flow and exchange rate dynamics. Journal of Political Economics 110, 170–180.

Evans, M.D.D., Lyons, R.K., 2005. Meese–Rogoff redux: micro-based exchange-rate forecasting. American Economics Review 95, 405–414.

Evans, M.D.D., Lyons, R.K., 2007. Exchange rate fundamentals and order flow. NBER Working Paper 13151.

Harris, L., 2003. Trading and exchanges: market microstructure for practitioners. Oxford University Press, New York.

Hong, Y., Li, H., Zhao, F., 2007. Can the random walk model be beaten in out-of sample density forecasts? Evidence from intraday foreign exchange rates. Journal of Econometrics 141, 736–776

Levich, R.M., Thomas, L.R., 1993. The significance of technical trading rule profits in the foreign exchange market: a bootstrap approach. Journal of International Money and Finance 12, 451–474.

Lyons, R.K., 1995. Tests of Microstructural Hypotheses in the Foreign Exchange Market. Journal of Financial Economics, October 39, 321-351.

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

Meese, R.A., Rogoff, K.,1983. Empirical exchange rate models of the seventies: do they fit out of sample? Journal of International Economics 14, 3–24.

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

Neely, C.J., Weller, P.A., Dittmar, R., 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. Journal of Financial and Quantitative Analysis 32, 405–426.

Osler, C.L., 2000. Support for resistance: technical analysis and intraday exchange rates. Federal Reserve Bank of New York Policy Review 6, 53–68.

Osler, C.L., 2001. Conditional orders and exchange-rate dynamics: Explaining the success of technicalanalysis. Unpublished manuscript, Federal Reserve Bank of New York.

Payne, R., 2003. Informed trade in spot foreign exchange markets: an empirical investigation. Journal of International Economics 61, 307–329.

Rime, D., Sarno, L., Sojli, E., 2010. Exchange rate forecasting, order flow and Macroeconomic information. Journal of International Economics 80, 72–88.

Russell J, Engle R. 2005. A discrete-state continous-time model of financial transactions prices and times: the autoregressive conditional multinomial-autoregressive conditional duration model. Journal of Business and Economic Statistics 23, 166–180.

Sweeney, R.J., 1986. Beating the foreign exchange market. Journal of Finance 41, 163–182.

Zhang, Q., Cai, C., Keasey, K., 2009. Forecasting using high-frequency data: a comparison of asymmetric financial duration models. Journal of Forecasting 28, 371–386.

TABLES

Table 1 Summary statistics of the price and return of the exchange rate.

This table reports the summary statistics of the price and return of the exchange rate. Panel A reports the range and distribution of the exchange rate during the sample period of 2010. Panel B report the mid-quote percentage changes in basis points. N reports the number of quote changes in the sample.

Currency	Mean	Min	P25	Median	P75	Max	Std	N
Panel A. Price (mid-quote)								
EUR	1.32355	1.19025	1.27540	1.33025	1.36725	1.45800	0.06027	2052284
GBP	1.53786	1.42313	1.50108	1.54100	1.57830	1.64580	0.05145	2890427
JPY	88.0628	80.3350	84.3500	88.7850	91.3350	94.9900	3.89761	1306350
Panel B. Changes in Basis Point								
EUR	-0.00031	-22.24019	-0.74566	-0.34743	0.74685	20.34505	0.77914	2052025
GBP	-0.00005	-26.63635	-0.48042	0.00000	0.48099	29.00512	0.57611	2890168
JPY	-0.00110	-102.33057	-1.07708	0.00000	1.07724	53.32961	1.01469	1306091

Table 2 Summary of realized durations

This table reports a summary of durations for the mid-quote movements in our sample. We report summary of the statistics under different model specifications. The Filter column indicates the threshold parameters used to filter the quotation changes. Only the movement which is greater than the filter parameter is recorded as a quotation change event. The Est Win column reports the number of days are used to estimate the model. All the statistics of duration are measured in second except for the maximum value which is measured in hour. The N column reports the total number of the movement event recorded. The sample period is calendar year 2010.

Currency	Filter	Est. Win.	Mean	Min	Median	Max (hour)	Mode	N
EUR	0	1	5	1	3	0.25	1	2,052,025
	0	5	5	1	3	0.25	1	2,025,474
	1	5	6	1	3	0.27	1	1,681,113
	1.5	5	73	1	26	0.96	1	110,650
GBP	0	1	4	1	2	0.16	1	2,890,168
	0	5	4	1	2	0.16	1	2,840,714
	1	5	7	1	4	0.33	1	1,417,434
	1.5	5	60	1	22	1.72	1	135,122
JPY	0	1	8	1	3	0.33	1	1,306,091
	0	5	8	1	3	0.33	1	1,286,957
	75	5	12	1	5	0.33	1	832,939
	100	5	22	1	7	0.94	1	364,772
	150	5	152	1	37	3.68	1	43,332

Table 3 Direction quality and Signal Correlations

This table reports the percentage of correct directional forecast of the models. A variable, *correct*, is assigned value of 1 when the predicted sign of quote change is the same as that of the realized quotation change and zero otherwise. A prediction signal is generate only when there is a change of the prediction direction. The mean column report the mean of the variable correct for all the signals generate during the sample period. The two confidence interval columns report the approximation of the confidence interval of the random walk at 0.01% and 99.99%, respectively, which is calculated with Equation (???). The correlation column reports the Pearson correlation. The Wald ChiSq column reports the statistics for the Wald test on the predictability of the forecast on the realized quote changes. *** signifies statistical significant at 0.01%. The N column reports the number of signals generated by each model.

Currency	Filter	Est. Win.	Mean	Confidence Interval		Correlation	Wald ChiSq	N
				0.0001	*0.9999*			
EUR	0	1	0.599	0.498	0.502	0.20 ***	43229 ***	1,123,006
	0	5	0.603	0.498	0.502	0.21 ***	49784 ***	1,197,581
	1	5	0.617	0.498	0.502	0.23 ***	53209 ***	987,905
	1.5	5	0.538	0.489	0.511	0.08 ***	165 ***	29,197
GBP	0	1	0.553	0.499	0.501	0.11 ***	17554 ***	1,549,743
	0	5	0.560	0.498	0.502	0.12 ***	15023 ***	1,063,140
	1	5	0.589	0.498	0.502	0.18 ***	21913 ***	693,357
	1.5	5	0.541	0.491	0.509	0.08 ***	295 ***	44,588
JPY	0	1	0.629	0.498	0.502	0.26 ***	48912 ***	750,633
	0	5	0.634	0.498	0.502	0.27 ***	55963 ***	804,943
	75	5	0.676	0.497	0.503	0.35 ***	65123 ***	549,199
	100	5	0.695	0.496	0.504	0.39 ***	34595 ***	240,622
	150	5	0.529	0.484	0.516	0.06 ***	44 ***	13,542

Table 4 Performance of the Long–Short Strategies Based on the Model Forecasts

This table reports the daily accumulated return, certainty equivalent return and durations of transactions. Returns are reports in basis point. CER is calculated using Equation (2) with the risk adverse parameters γ are set to be 1, 3, 5 respectively in the three columns. Durations is measured in second. All model specifications with different filter and estimations are reported. In addition, the returns of buy and hold strategies for each currency of the pairs are reported. The N column reports the number of days.

Currency	Filter	Est. Win.	Return					CER			Duration		N
			Mean	Min	Median	Max	Std	(1)	(3)	(5)	Mean	Median	
EUR	0	1	1107.2	-750	1201	2220	484	865	381	-104	9	5	259
	0	5	1251.6	-133	1251	2220	326	1089	763	437	8	5	255
	1	5	1001.6	-473	1030	1831	345	829	484	139	10	6	255
	1.5	5	9.7	-133	0	321	61	-21	-82	-144	230	72	255
	B&H –EUR		-2.6	-161	-2	178	55	-30	-85	-139	10 Hr		259
	B&H –USD		2.9	-175	2	164	55	-24	-79	-134	10 Hr		259
GBP	0	1	561.9	-1562	516	1983	541	291	-250	-792	7	4	259
	0	5	465.2	-575	248	1637	517	207	-311	-828	8	4	255
	1	5	515.9	-158	476	1514	393	320	-73	-466	14	7	255
	1.5	5	37.8	-100	25	598	83	-4	-87	-171	166	61	255
	B&H –GBP		-0.6	-141	-2	122	52	-27	-79	-132	10 Hr		259
	B&H –USD		0.8	-121	2	143	52	-25	-78	-130	10 Hr		259
JPY	0	1	1314.2	-580	1453	2663	693	968	274	-419	14	6	259
	0	5	1489.3	-78	1600	2693	572	1203	631	59	13	6	255
	75	5	1265.0	-384	1459	2507	702	914	212	-490	18	8	255
	100	5	535.1	-529	574	1407	377	347	-30	-408	33	10	255
	150	5	-1.0	-162	0	530	64	-33	-96	-160	378	80	255
	B&H –USD		-5.7	-216	-6	157	52	-31	-83	-135	10 Hr		259
	B&H –JPY		5.9	-155	6	221	52	-20	-72	-124	10 Hr		259

Table 5. Effect of transaction costs

This table reports the effect of transaction costs on the long-short strategies based on the model forecasts. Half basis point is taken away from each pre-cost return before accumulating them. The daily return, CER and duration of transaction are reported.

Currency	Filter	Est. Win.	Return					CER			Duration		N
			Mean	Min	Median	Max	Std	(1)	(3)	(5)	Mean	Median	
EUR	0	1	-1050	-3168	-975	59	531	-1315	-1846	-2377	9	5	259
	0	5	-1094	-3168	-1029	-165	461	-1325	-1786	-2247	8	5	255
	1	5	-929	-2800	-868	31	438	-1148	-1586	-2023	10	6	255
	1.5	5	-47	-342	-44	132	66	-80	-145	-211	230	72	255
GBP	0	1	-2119	-5282	-1876	-20	955	-2596	-3551	-4506	7	4	259
	0	5	-1440	-4356	-1359	66	1106	-1993	-3099	-4205	8	4	255
	1	5	-817	-1989	-826	49	425	-1029	-1454	-1878	14	7	255
	1.5	5	-49	-285	-38	252	76	-87	-163	-239	166	61	255
JPY	0	1	-209	-3363	-1	713	633	-525	-1159	-1792	14	6	259
	0	5	-177	-3364	41	713	655	-505	-1160	-1815	13	6	255
	75	5	105	-2569	159	941	432	-111	-543	-974	18	8	255
	100	5	46	-1727	39	464	225	-66	-292	-517	33	10	255
	150	5	-27	-490	-13	148	63	-59	-122	-186	378	80	255