

Developing High Frequency Foreign Exchange Trading Systems

Bruce Vanstone
Bond University
Gold Coast,
Queensland,
Australia
bvanston@bond.edu.au

Tobias Hahn
Bond University
Gold Coast,
Queensland,
Australia
thahn@bond.edu.au

Gavin Finnie
Bond University
Gold Coast,
Queensland,
Australia
gfinnie@bond.edu.au

ABSTRACT

The foreign exchange (FX) spot markets are well suited to high frequency trading. They are highly liquid, allow leverage, and trade 24 hours a day, 5 days a week. This paper documents and tests the stylized facts known about high-frequency FX markets. It then postulates a high frequency trading system on the basis of these stylized facts. Benchmarking confirms the robustness of the approach, demonstrating the role algorithmic trading has to play in higher frequency trading environments.

1. INTRODUCTION

The FX markets provide an outstanding opportunity for currency trading and speculation. According to the Bank for International Settlements, the daily turnover of the FX markets is in excess of \$4 trillion [1]. FX markets are of particular interest to speculators not just for their size and correspondingly high liquidity, but also for their low transaction costs and common use of leverage. Further, there is a proliferation of software providers and brokers who provide access to these markets to allow high-frequency, intraday trading. The opportunity to trade a relatively stable, highly leveraged currency contract allows intraday traders to exploit small price movements using high frequency trading algorithms.

In finance, the term 'stylized facts' is used to denote those 'facts' discovered through empirical research. Given the shortage of formal academic models to describe the behaviour of higher frequency financial time series, there is a heavy reliance on stylized facts within the trading discipline.

This paper reviews and tests commonly accepted stylized facts known about higher frequency FX time series data, using data from the 1-minute timeframe for the EURUSD currency pair. The EURUSD is the most heavily traded currency in the FX markets, and accounts for some 28% of the spot market [1]. Considering the extraordinary level of trading in this instrument, it should be one of the most 'efficient' instruments available.

After confirming the stylized facts within the EURUSD high-frequency price time series, the paper proposes simple algorithmic trading rules to profit from the observed behaviour. These rules are then benchmarked in-sample, and tested out-of-sample, to demonstrate the robustness of the approach.

2. LITERATURE REVIEW

Stylized facts concerning financial time series data are normally characterized according to dependence, distribution, presence of non-linearity and scaling effects. Research into stylized facts for high frequency FX data shows that the well accepted empirical regularities of daily data and weekly data often do not hold for intraday data [2].

With regard to dependence, high frequency time series data from the FX markets normally displays extremely high negative first-order autocorrelation [3]. Highly negative autocorrelation is documented in high frequency timeframes including tick-by-tick FX data [4], and 1 minute FX returns [2]. Negative serial correlation is reported for other high frequency asset time series, for example, significant negative serial correlation is reported in Italian stock index futures for time periods smaller than 20 minutes [5].

Within financial time series in general, distributions of returns are approximately symmetric, and have high kurtosis. Distributions become increasingly fat tailed as the data frequency increases [6].

Non-linearity is an important characteristic of financial time series. Tests of five major FX rates finds no evidence of linear correlation, however, substantial evidence of non-linearity is found [7]. Tests of 10 different GBP exchange rate pairs also documents non-linearity in many of the pairs [8].

Scaling effects refers to the absolute size of returns as a function of the time interval in which they are measured. Scaling laws provide a relationship between time interval and average volatility as a power of the absolute return of that

interval. There is much evidence of the existence of power laws within high frequency FX data. Scaling is reported in the mean absolute changes of logarithm prices [9], self-similarity is reported in USD/DEM [10] and multi-fractal scaling in DEM/USD [11]. Scaling is reported in exchange rates [12] and exchange rate volatilities [13]. Further support for FX scaling laws is provided by Andersen et al. [14], Dacorogna [2] and Xu & Gencay [15].

To summarize the stylized facts concerning high frequency FX data, there is strong evidence of negative autocorrelation, there is support for the existence of non-linearity and scaling laws, and evidence of fatter tails in higher frequency timeframes.

3. METHODOLOGY

This aim of this paper is to demonstrate the process of creating high frequency algorithmic FX trading systems. When creating algorithmic trading systems, the two primary concerns are credibility and robustness.

The methodology addresses these concerns by splitting the available data into two partitions, namely, in-sample and out-of-sample.

Initially, the in-sample data is tested for market efficiency and randomness. A standard working assumption for financial market data is that the market is efficient and that price changes are random. Under these conditions, there is no expectation that a successful trading system can be developed. Therefore, it is necessary to test the in-sample data for market efficiency and randomness. A pre-condition to developing a trading system is that either of the assumptions of market efficiency or randomness are not demonstrable in the underlying data.

Market efficiency is commonly tested using Lo & MacKinlay's Variance Ratio [16], and randomness can be tested using the runs test [17]. The runs test is conducted on the returns data for each time interval, by splitting the returns firstly on their sign, and secondly on their direction. This allows an assessment of the degree to which the next period returns may be dependent on the previous period return or the previous period direction.

To ensure credibility, it is desirable to create a trading system whose rules are tightly aligned with a stylized fact. The in-sample data can be used to confirm the stylized facts, and this leads to the postulation and testing of trading rules. Considering the strong evidence of negative autocorrelation in high frequency time series data this is the stylized fact that will be tested. As the data is in the 1-minute timeframe, it can easily be converted into other, lower frequency timeframes. Typical timeframes for data used in intraday trading algorithms are 1-minute, and 1-hourly, and the in-sample data will be tested for negative autocorrelation in each of these timeframes.

Autocorrelation describes the relationship between values of a time series that are lagged k periods apart. In this paper, all autocorrelations are performed on the time series of returns. Autocorrelation results are normally presented by ACF and PACF diagrams, the difference being that, at a given lag, the PACF result removes the effects of prior lags.

Once the stylized fact is confirmed, rules can be stated and benchmarked in-sample. They can then be tested in the out-of-sample data, and the results from the in-sample and out-of-sample testing can be compared to ensure minimal deviation. It is consistency between the in-sample and out-of-sample results that supports robustness.

The data used in this paper is 1-minute price data from the EURUSD currency pair over the timeframe 1st January 2008 to 31st December 2011. The in-sample data covers the two year period 1st January 2008 to 31st December 2009. The out-of-sample data covers the two year period 1st January 2010 to 31st December 2011. The data was sourced from Thomson-Reuters TRTH database provided by SIRCA [18].

Testing and comparison of the in-sample and out-of-sample data can be done algorithmically by coding rules into computer programs, and testing these rules against the underlying datasets. Using computers to detect and implement signals is typical of the implementation of high frequency trading algorithms.

It is difficult to find expected measures of returns for high-frequency trading. As there are no pre-established benchmarks, it seems sensible to compare the overall performance metrics of the rules to the naïve buy-and-hold benchmark, as there are no academically accepted high-frequency competitor models. This allows for an assessment of the performance of the rules over the out-of-sample data period, and gives some guidance toward interpreting overall profitability. The same metrics can be used to compare the performance of the rules in the in-sample and out-of-sample datasets, which allows inference of robustness and likely future reliability. This technique is typical in the evaluation of trading systems in lower frequency timeframes [19, 20], and similar metrics are used here.

In this research, MATLAB is used to implement models and conduct the statistical testing.

4. ANALYSIS OF DATA

To allow greater granularity for testing market efficiency and autocorrelation, the in-sample data is further subdivided into two in-sample partitions. Partition 1 covers the year 1st January 2008 to 31st December 2008, and partition 2 covers the period 1st January 2009 to 31st December 2009. As well as allowing for greater granularity, splitting the in-

sample data into two partitions allows us to ensure that the market efficiency and randomness tests do not fail simply due to the fact that the GFC occurred during the first part of the in-sample period. Certainly the majority of the GFC effect occurred in partition 1, and consistent results between partition 1 and partition 2 would strengthen any in-sample findings.

Table 1 summarizes the results of the market efficiency and randomness tests carried out on the two in-sample partitions. For each test, the table shows where the null hypothesis was rejected and the relevant p-value. The null hypothesis for the variance ratio test is that the series tested is a random walk, and the null hypothesis for the runs tests is that the values come in a random order. In both partitions, the 1-minute returns fail the test for market efficiency and both tests for randomness. In 1 hourly returns partition 1 fails one test for randomness, and partition 2 fails both tests for randomness.

Period	Timeframe	Variance Ratio Test, p-value	Runs Test (+ve/-ve), p-value	Runs Test (ud), p-value
2008 – In-sample partition 1	1 minute returns	Rejected, 0.00 **	Rejected, 0.00 **	Rejected, 0.00 **
2008 – In-sample partition 1	1 hourly returns	0.24	Rejected, 0.00 **	0.89
2009 – In-sample partition 2	1 minute returns	Rejected, 0.00 **	Rejected, 0.00 **	Rejected, 0.00 **
2009 – In-sample partition 2	1 hourly returns	0.08	Rejected, 0.00 **	Rejected, 0.02 **

Table 1 Summary of market efficiency and randomness test results

Figure 1 and Figure 2 show the ACF and PACF autocorrelation graphs of the in-sample data for in-sample partition 1 (2008), in each of the chosen timeframes. Figure 3 to Figure 4 show the ACF and PACF autocorrelation graphs of the in-sample data for in-sample partition 2 (2009), in each of the chosen timeframes. These findings are summarized in Table 2, which shows the autocorrelation result and statistical significance of the result for each partition and timeframe tested. Statistical significance is reported using the Ljung-Box Q-test, a portmanteau test which assesses the null hypothesis that a series of residuals exhibits no autocorrelation for a fixed number of lags.




Figure 1 Autocorrelation of 1-min EURUSD returns 2008




Figure 2 Autocorrelation of hourly EURUSD returns 2008




Figure 3 Autocorrelation of 1-min EURUSD returns 2009




Figure 4 Autocorrelation of hourly EURUSD returns 2009

Period	Timeframe	Autocorrelation Direction (Lag 1)	p-value
2008 – In-sample partition 1	1 minute returns	Negative	0.00 **
2008 – In-sample partition 1	1 hourly returns	Negative	0.02 **
2009 – In-sample partition 2	1 minute returns	Negative	0.00 **
2009 – In-sample partition 2	1 hourly returns	Negative	0.01 **

Table 2 Summary of Autocorrelation results

Table 2 confirms the existence of the previously documented stylized facts concerning negative autocorrelation in high frequency FX data. Considering the prior results of Table 1, there is sufficient evidence that this data is suitable for the development of a high-frequency algorithmic trading system in the 1-minute timeframe.

The next step is the postulation of trading system rules. Speculative traders do not trade every possible opportunity that presents itself, typically, they wait for localized extrema events to occur. For this reason, and based on the evidence in this paper, the following simple buy/sell rule is proposed.

Buy: Take a long position (1 contract) when the 1-minute closing price closes above the lowest price over the last week

Sell: Close the position after no new high has been reached within the last 5 minutes

By waiting until the price of the last minute is lower than the price over the last week, we are effectively waiting for a localized extrema event to occur. The periods used in both the buying rule and the selling rule were chosen arbitrarily. However further testing shows the approach is robust to the selection of many different periods. For the buying period, positive results exist for multi-day timeframes greater than 1 day. For the selling period, positive results exist for small holding timeframes, ranging from 5 minutes to 1 hour. The results for these further tests are not included in this paper, but are available from the authors on request.

A trading simulation can be conducted on these rules over the in-sample and out-of-sample period. Transaction costs are extremely low in the FX markets, and typical costs from a large scale broker are used in these simulation results [21]. Standard trading system evaluation metrics are reported [19].

Metric	Simulation	Buy-and-Hold
Raw Profit	\$1,595.00	\$577.50
Number of Trades	135	1
Average minutes per trade	25.98	740,652
Winning percentage	64.44%	100.00%
Profit Factor	1.19	N/A
Recovery Factor	0.63	0.02
Payoff Ratio	0.65	0.00
Maximum Drawdown	-\$2,535.00	-\$37,060.00

Table 3 In-sample trading simulation metrics vs naive Buy-and-Hold

Metric	Simulation	Buy-and-Hold
Raw Profit	\$1,710.00	-\$13,682.50
Number of Trades	164	1
Average minutes per trade	29.16	748,702
Winning percentage	57.93%	0.00%
Profit Factor	1.22	0.00
Recovery Factor	0.73	0.00
Payoff Ratio	0.89	0.00
Maximum Drawdown	-\$2,330.00	-\$27,010.00

Table 4 Out-of-sample trading simulation metrics vs naive Buy-and-Hold

5. DISCUSSION OF RESULTS

Metrics for the in-sample and out-of-sample simulations are remarkably similar, despite very large differences in the buy-and-hold results for each period. During the in-sample period the buy-and-hold result was slightly positive, while during the out-of-sample period, the buy-and-hold result is large and negative. The stability of the simulation results suggests the strategy is quite robust.

The periods used in both the buying rule and the selling rule were chosen arbitrarily. However, further testing shows this approach is robust to the selection of many different buying and selling periods. For the buying period, positive results exist for multi-day timeframes greater than 1 day. For the selling period, positive results exist for small holding timeframes, ranging from 5 minutes to 1 hour. The results for these further tests are not included in this paper, but are available from the authors on request.

6. CONCLUSION

This paper has demonstrated the process of creating high frequency algorithmic FX trading systems. When creating algorithmic trading systems, the two primary concerns are credibility and robustness. The issue of credibility was addressed by postulating trading rules which were aligned with the stylized facts known about higher-frequency price timeseries data.

The issue of robustness is addressed through the methodology used to create algorithmic trading systems. Data used for testing is split into in-sample and out-of-sample partitions, and the in-sample data is used to test stylized facts, and postulate trading rules. The rules are then benchmarked firstly on the in-sample data, and then on the out-of-sample data. The results of the two simulations are compared for consistency, and hence, robustness.

7. FUTURE WORK

The rules postulated and tested in this paper are best seen as a simple trading test of a stylized fact within high frequency FX data. Further work is required to confirm these results across a wider range of currency pairs, and time periods. From a trading perspective, rules such as these could eventually be used to trade a basket of currencies in the high-frequency timeframe.

From a practical perspective, further work is required to identify more and better trading opportunities, by testing different timeframes, and different trading filters. It is expected that machine learning techniques may be able to assist in this regard.

8. ACKNOWLEDGEMENTS

The data used in this research was supplied by the Securities Industry Research Centre of Asia-Pacific (SIRCA) on behalf of Reuters and the ASX. The authors also gratefully acknowledge the financial assistance provided by Bond University through the Vice Chancellors Research Grant.

REFERENCES

1. BIS Triennial Central Bank Survey: *Foreign exchange and derivatives market activity in April 2010: Preliminary Results*. 2010. **September 2010**, 1-21.
2. Dacoronga, M.M., et al., *An introduction to High-Frequency Finance* 2001, San Diego, CA: Academic Press.
3. Zhou, B., *High-frequency data and volatility in foreign-exchange rates*. Journal of Business and Economic Statistics, 1996. **14**(1): p. 45-52.
4. Cont, R., *Empirical properties of asset returns: Stylized facts and statistical issues*. Quantitative Finance, 2001. **1**(2): p. 223-236.
5. Bianco, S. and R. Reno, *Dynamics of intraday serial correlation in the Italian futures market*. The Journal of Futures Markets, 2006. **26**(1): p. 61-84.
6. Sewell, M. *Characterization of Financial Time Series*. UCL Research Note RN/11/01, 2011.
7. Hsieh, D.A., *Testing for non-linear dependence in daily foreign exchange rates*. Journal of Business, 1989. **62**(3): p. 339-368.
8. Brooks, C., *Testing for non-linearity in daily sterling exchange rates*. Applied Financial Economics, 1996. **6**(4): p. 307-317.
9. Muller, U.A., et al., *Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis*. Journal of Banking and Finance, 1990. **14**(6): p. 1189-1208.
10. Evertsz, C.J., *Fractal geometry of financial time series*. Fractals, 1995. **3**(3): p. 609-616.
11. Fisher, A., L. Calvet, and B. Mandelbrot, *Multifractality of deutschemark / US dollar exchange rates*. Discussion paper 1166, Cowles Foundation, 1997.
12. Galluccio, S., et al., *Scaling in currency exchange*. Physica A, 1997. **245**(3-4): p. 423-426.
13. Gencay, R., F. Selcuk, and B. Whitcher, *Scaling properties of foreign exchange volatility*. Physica A, 2001. **289**(1-2): p. 249-266.
14. Andersen, T.G., et al., *The distribution of realized exchange rate volatility*. Journal of the American Statistical Association, 2001. **96**(453): p. 42-55.
15. Xu, Z. and R. Gencay, *Scaling, self-similarity and multifractality in FX markets*. Physica A, 2003. **323**: p. 578-590.
16. Lo, A.W. and C. MacKinlay, *Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test*. Review of Financial Studies, 1988. **1**: p. 41-66.
17. Aldridge, I., *High-Frequency Trading* 2010: Wiley.
18. SIRCA. 2012; Available from: www.sirca.org.au.
19. Vanstone, B. and G. Finnie, *An Empirical Methodology for developing Stockmarket Trading Systems using Artificial Neural Networks*. Expert Systems with Applications, 2009. **36**: p. 6668-6680.
20. Vanstone, B. and T. Hahn, *Designing Stockmarket Trading Systems (with and without Soft Computing)* 2010, London: Harriman House
21. 2011; Available from: (Interactive Brokers Commissions) <http://www.interactivebrokers.com/en/p.php?f=commission>.