Developing High Frequency Foreign Exchange Trading Systems

Vanstone, Bruce J

Developing High Frequency Foreign Exchange Trading Systems

Dr. Bruce Vanstone

Bond University Computational Finance Research Group

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.

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. 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. 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.

Analysis and Discussion

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.

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.

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. Standard trading system evaluation metrics are reported.

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.

Acknowledgements

Professor Gavin Finnie and Tobias Hahn both contributed significantly to this research work.

The data used in this research was supplied by the Securities Industry Research Centre of Asia-Pacific (SIRCA) on behalf or Reuters and the ASX.

I also gratefully acknowledge the financial assistance provided by Bond University through the Vice Chancellors Research Grant (Early Career Researcher).

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
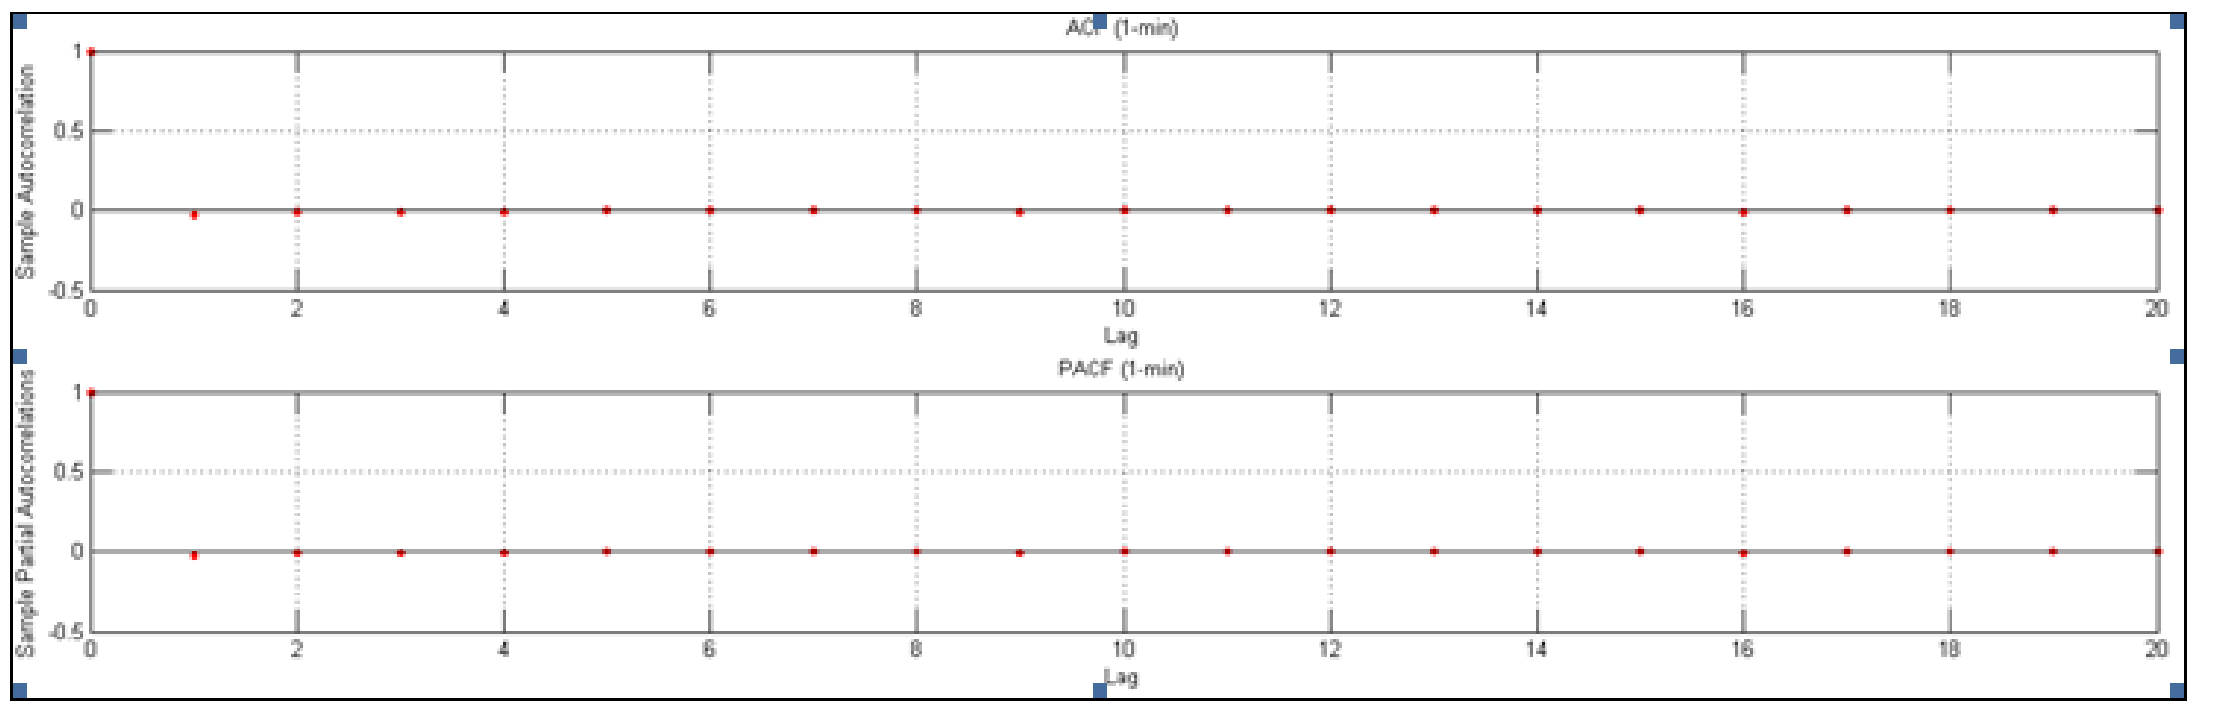

Figure 1 Autocorrelation of 1-min EURUSD returns 2008
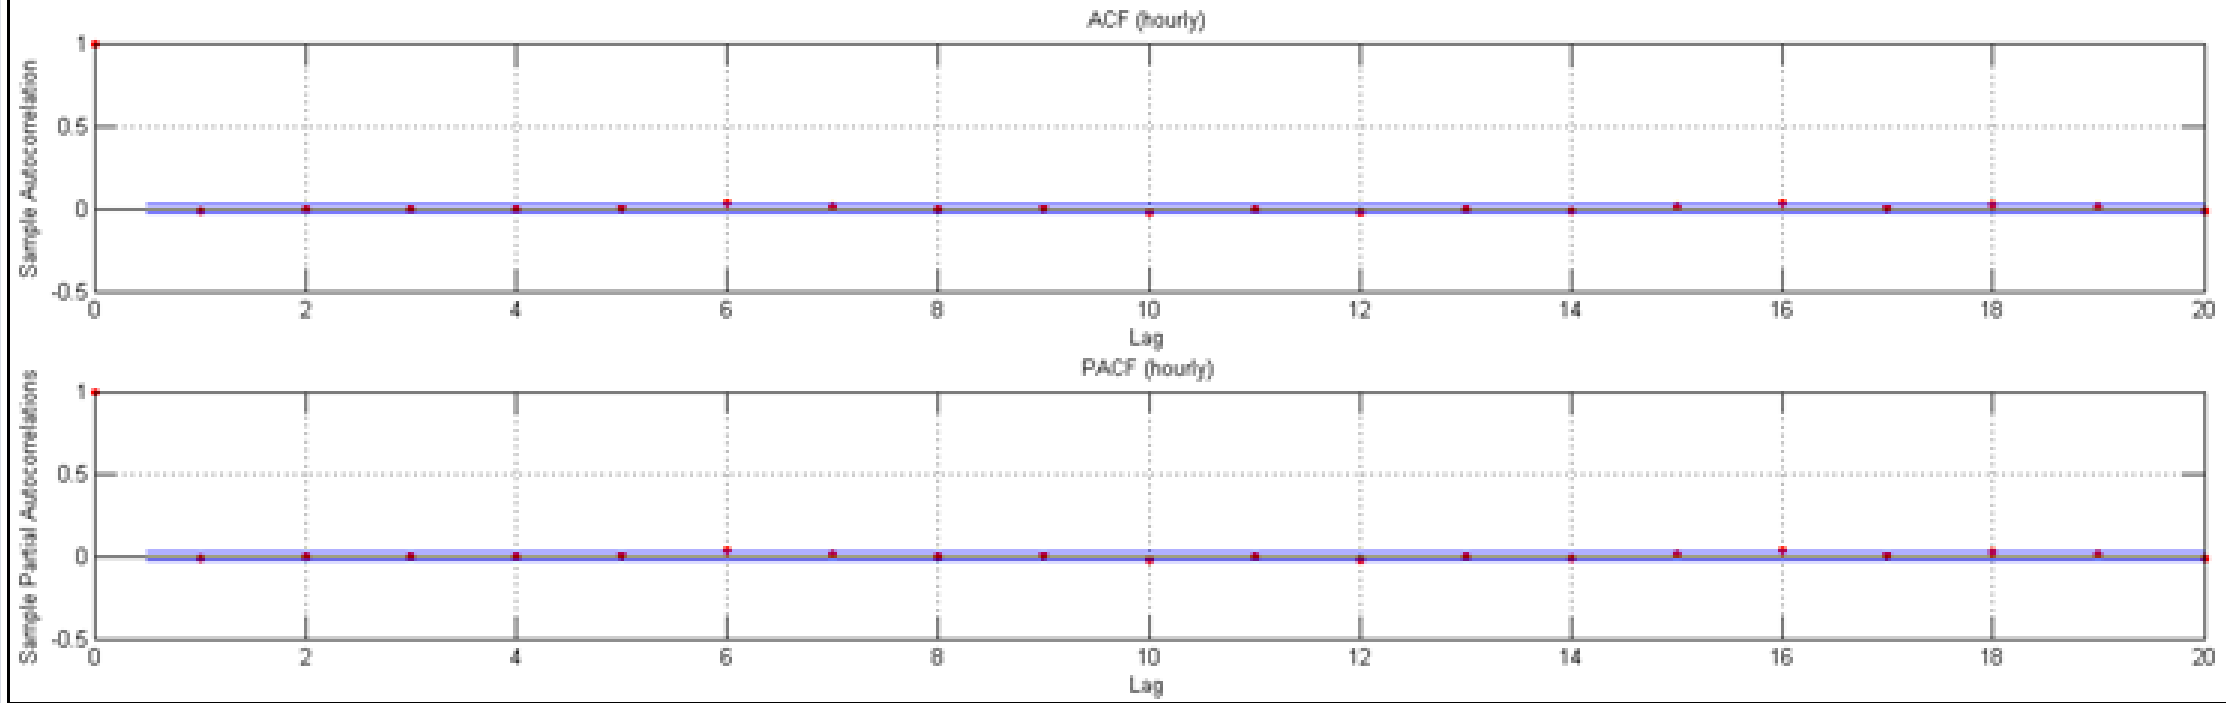

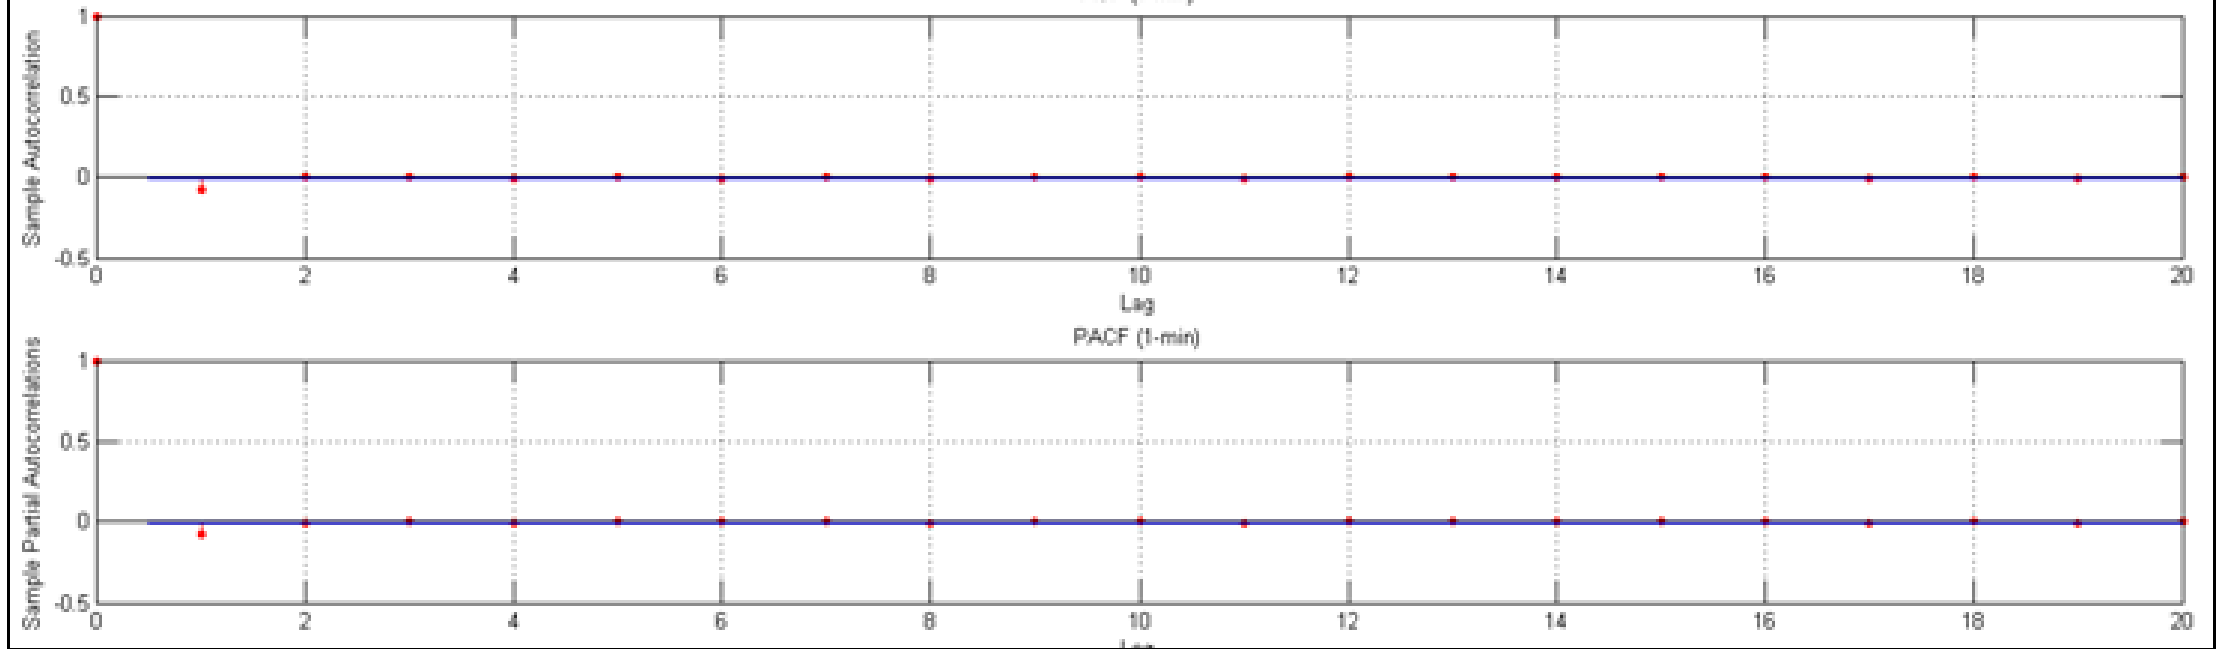Figure 2 Autocorrelation of hourly EURUSD returns 2008

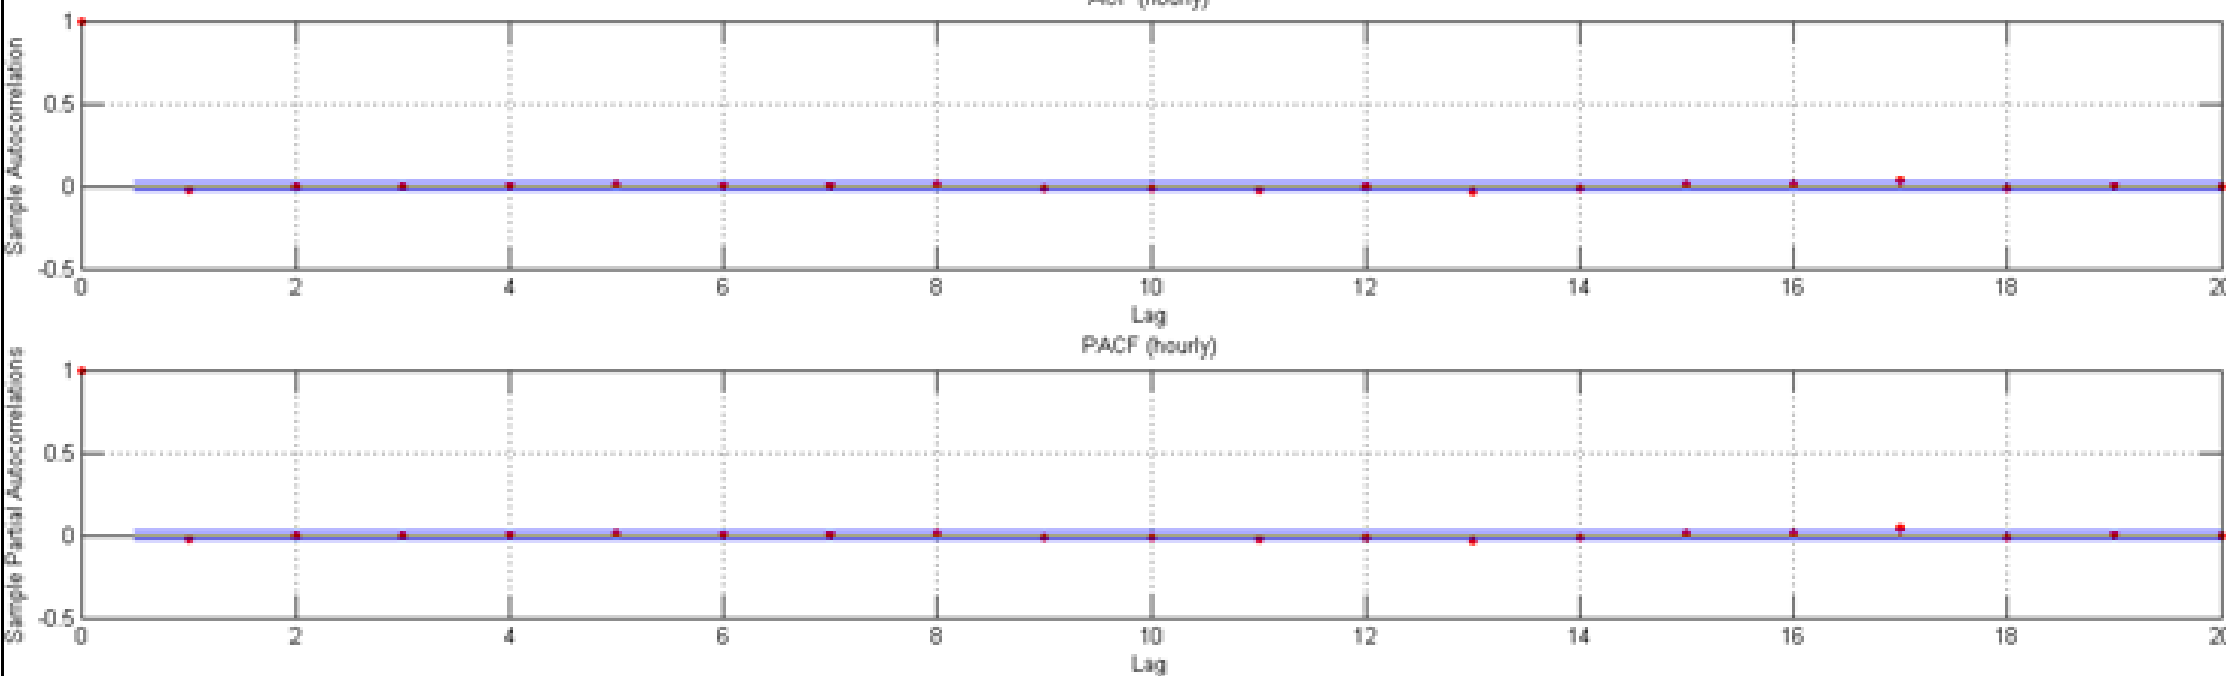Figure 3 Autocorrelation of 1-min EURUSD returns 2009

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

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

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