

Can abnormal returns be earned on bandwidth-bounded currencies? Evidence from a genetic algorithm

Pedro Godinho¹

ABSTRACT

Most of the studies about the Foreign Exchange market (Forex) analyse the behaviour of currencies that are allowed to float freely (or almost freely), but some currencies are still bounded by bandwidths (either disclosed or undisclosed). In this paper, I try to find out whether two bandwidth-bounded currencies, the Hong Kong dollar (HKD) and the Singapore dollar (SGD), present opportunities for abnormal returns. I consider a set of trading rules, and I use a genetic algorithm to optimise both the subset of rules to be used and their parameters, using real market data. I use four pairs of currencies, two of them involving currencies bounded by bandwidths and two others involving only free-floating currencies. I compare the results obtained for the different pairs, both in terms of profitability and in terms of the types of the rules that are used. Evidence of profitability is more consistent for the only pair including a bandwidth bounded currency without a narrow price band, the USD/SGD. Trend reversing rules are preferred for this currency pair, while the preferred type of rule seems to depend on the pair of currencies when free-floating currencies are considered. In the case of the USD/HKD, the small number of price changes, as well as the price stability (possibly consequences of a narrow price band) do not allow me to obtain conclusive results.

1. INTRODUCTION

THE FOREIGN EXCHANGE MARKET (Forex) is very liquid, open to a very large number of participants and the information that is relevant for determining the exchange rates is accessible to most participants. Therefore, according to the Efficient Market Hypothesis (EMH; Fama, 1970), strategies based on technical analysis should not lead to abnormal returns. However, technical analysis is widely used in the Forex (Menkhoff and Taylor, 2007,

Neely *et al.*, 2009), and some studies show that, at least until the early 1990s, it generated significant excess returns (Park and Irwin, 2007, Neely *et al.*, 2009, Sweeney, 1986, Levich and Thomas, 1993, LeBaron, 1999, among others). Such abnormal returns present a challenge to the efficiency of the Forex market, and may be explained by some particular characteristics of this market, such as the existence of several major players whose objectives may differ greatly from those of profit maximising economic agents (LeBaron, 1999).

Evidence shows that, by the mid 1990s, profit opportunities presented by technical trading rules seemed to have declined significantly (Park and Irwin, 2007, Neely *et al.*, 2009, LeBaron, 2002, Olson, 2004, Schulmeister, 2008, among others). Thus, the market inefficiencies found in previous years may have been temporary.

As noted by Neely *et al.* (2009), these findings may be consistent with the Adaptive Market Hypothesis (AMH; Lo, 2004). According to the AMH, individuals make choices based on their past experience, and they learn by receiving negative or positive reinforcement from the outcomes of their decisions. Therefore, profitable investment strategies may cease to generate excess returns because they become more widely used, and other strategies may become profitable.

Data snooping has been an important concern regarding studies of the profitability of technical trading rules. It may be the case that the trading rules that are most used nowadays, as well as their respective parameter ranges, are the 'survivors' of a broader set of trading rules and parameters, that is, the ones that performed better in the past. It may also be the case that this outperformance is due to luck, instead of an intrinsic ability to forecast future price movements. In fact, if you consider a large enough set of randomly generated trading rules and a given period, you will get some rules that beat the benchmark in that period. If you then ignore the rules that failed to beat the benchmark, the ones that remain will seem good investment rules. However, if you consider that you started with a much larger set, you realise that the fact some rules beat the benchmark may not be statistically significant, and there may be no reason to believe that such rules will continue to outperform. Sullivan *et al.* (1999) address this issue by using a test based on 'White's reality check' (White, 2000). They consider a set of 7846 trading rules, each of them based on a single technical indicator. In spite of many of such rules outperforming the benchmark, Sullivan *et al.* are unable to reject the null hypothesis that none of them performs better than the benchmark. This means that, taking into account a broad set of trading rules, the outperformance of some of these rules may be due to a selection bias: choosing the ones that perform better and ignoring the ones with a worse performance.

Technical analysts usually resort to trading strategies based on multiple indicators instead of a single one (like the ones used by Sullivan *et al.*, 1999). If we consider the possible combinations of multiple indicators, the space of possible strategies becomes too large to allow us to test each combi-

nation individually, and we must resort to some kind of optimisation procedure. Mendes *et al.* (2010) argue that, in such cases, it is also important to develop tools that allow the adaptation of investment strategies to evolving conditions. Evolutionary algorithms in general, and genetic algorithms in particular, may be suitable optimisation tools for defining such strategies. In order to avoid data snooping issues, such algorithms divide the data set into a training set and a test set. The training set is used in order to optimise a trading strategy (that is, for choosing the combination of indicators to be used and the respective parameters). The trading strategy is then applied to the test set in order to find out whether the rules can be exploited to achieve out-of-sample trading profits.

Several authors have applied evolutionary algorithms to the definition of the most profitable technical trading strategies. Dempster and Jones (2001) use genetic programming to develop a trading system for the USD/GBP² currency pair, based on combinations of different technical indicators. They are able to find evidence of abnormal returns only when they perform an adaptive reoptimisation of the strategy. Neely *et al.* (1997) use a genetic programme to find trading rules for the foreign exchange market. The authors achieve consistent excess returns, even after considering transaction costs. Neely and Weller (2003) use a similar approach to derive intraday trading strategies, but they are unable to find evidence of excess returns, after transaction costs and trading hours are taken into account. Dunis *et al.* (1999) use genetic algorithms to optimise intraday trading models for the DEM/JPY and USD/DEM currency pairs. The authors note that several strategies show in-sample profits but fail to make positive out-of-sample returns. Brabazon and O'Neill (2004) use grammatical evolution to generate trading rules based on daily frequency data for three pairs of currencies. The authors consider two test periods for each training series, and conclude that the evolved strategies outperform the buy-and-hold strategy in five of the six test sets, after allowing for trading, slippage and net interest costs. Mendes *et al.* (2010) use a genetic algorithm to optimise a trending system for the EUR/USD pair. The authors use four different intraday data frequencies, ranging from one minute to one hour, and consider test series of different lengths. The system achieves moderate average profits in the test series for most data frequencies.

Studies about the results achieved by technical analysis, and applications of evolutionary algorithms to the optimisation of Forex trading systems, usually resort to pairs of currencies that are allowed to float freely, that is, without explicitly defined floating bands. However, a few currencies are still linked to other currencies or baskets of currencies. This is the case of the HKD, which is pegged to the USD, being allowed to float between 7.75 and 7.85 HKD to each USD. In the case of the SGD, the exchange rate is monitored by the Monetary Authority of Singapore, being allowed to float within an undisclosed bandwidth of the central parity of a concealed basket of currencies representing the country's most important trading partners and competitors.

In this paper I test whether these floating bands create exploitable opportunities for abnormal profits. In order to do so, I will analyse the profitability of trading strategies based on a set of technical indicators, optimised by a genetic algorithm that defines the rules' weights and parameters. I will consider two pairs of currencies that involve a bandwidth-bounded currency - the USD/HKD and the USD/SGD — and two pairs of currencies for which no floating bands exist — the EUR/USD and the GBP/USD — and I will compare the profits achieved for the different pairs. I perform the analysis using the tick-by-tick bid and ask quotes for contract-for-difference (CFD) style contracts defined for these pairs of currencies. In order to avoid data snooping issues, I use different data sets to generate the trading strategies and to test their performance.

The analysis of the characteristics of the most profitable strategies is an issue of significant practical interest. Schulmeister (2008) concludes that the profitability of technical currency trading is due to the exploitation of persistent trends. However, such analysis is seldom performed through applications of evolutionary algorithms to the Forex market. In this paper, I also analyse whether the rules chosen for different pairs have significantly different characteristics. In order to perform such an assessment, I consider a set of four rules that seek to identify trends that are present and are expected to continue (termed trend-following rules); and another set of four rules that aim at identifying trends that should be ending, and are expected to be reversing (termed trend-reversing rules). For each pair of currencies, I try to assess whether a given set of rules is particularly effective, or whether it is hampering the results. *A priori*, I expect trend-reversing rules to be preferred in the case of bandwidth bounded currencies, since a large change in the quotes of a pair containing such a currency will usually imply that one of the bands is being approached, and therefore lead to a reversion of the trend. I have no reason to expect that any one type of rule will be preferred in the case of free-floating currencies.

The paper is structured as follows. After this introduction, Section 2 describes the structure and evaluation of the trading strategies. Section 3 outlines the genetic algorithm. Section 4 describes the data series used in the analysis, as well as the empirical tests. Section 5 presents and analyses the results obtained. The conclusions are presented in Section 6.

2. THE TRADING STRATEGIES

The analysis of trading strategies looks at the average exchange rate of the different periods, and four different technical indicators: the exponential Moving Average (MA), the Relative Strength Index (RSI), the stochastic oscillator (%k) and a three period moving average of the stochastic oscillator (%d). These are fairly standard indicators, and a detailed presentation can be found in any book about technical analysis (e.g, Meyers, 1989).

Both the average exchange rate and these indicators are used to define eight rules: four are aimed at finding trends that are in place (trend-following rules), and the other four are aimed at finding extreme price movements that are likely to be followed by movements in the opposite direction (trend-reversing rules). Each rule may provide a neutral result, support a buy signal or support a sell signal. These rules are given weights, and are then used for generating buy and sell signals. The signals are then used for opening positions, and for setting the prices for closing positions. The next sub-sections describe the process in greater detail.

a. Trend-following rules

Rule 1 — Short-term moving average: The first trend-following rule is based on a Short-term exponential Moving Average (SMA). The average exchange rate of the considered period is compared with the SMA and, if it is higher, the rule supports a buy signal, if it is lower the rule supports a sell signal (if the SMA is exactly equal to the average price, the rule supports neither a buy nor a sell signal). The number of periods used to calculate the SMA is a parameter to be optimised, and may be any number between 3 and 30.

Rule 2 — Short-term moving average: The second trend-following rule is very similar to the first one, but it is based on a Long-term exponential Moving Average (LMA). An average exchange rate higher than the LMA supports a buy signal, and a lower one supports a sell signal. The number of periods used to calculate the LMA is a parameter to be optimised, and may be any number between 31 and 200.

Rule 3 — Relative Strength Index (RSI): The RSI is an indicator that may take values between 0 and 100. A high RSI indicates that an increase in the exchange rate is taking place, whereas a low RSI points to a decrease in the exchange rate. Rule 3 defines a threshold parameter to be optimised (*rsi_threshold*), which may take values between 51 and 70. If the RSI is larger than this parameter, the rule supports a buy signal; if $RSI < 100 - rsi_threshold$ the rule supports a sell signal. The number of periods to be used for calculating the RSI is also a parameter to be optimised, and may take values between 9 and 60.

Rule 4 — Stochastic oscillator: In order to detect trends, both the stochastic oscillator (%k) and its three-period moving average (%d) are calculated. Having $\%k > \%d$ provides support for a buy signal, whereas having $\%k < \%d$ supports a sell signal. The number of periods used to calculate the stochastic oscillator is a parameter to be optimised, which may take values between 5 and 60.

b. Trend-reversing rules

Rule 5 — Price moving too fast: The first trend reversing rule tries to find situations in which the price is moving too fast, and the trend is therefore expect-

ed to reverse. The rule defines two parameters to be optimised: a price change (*pips_change*), which may take values between 20 and 200 pips³, and a number of periods (*mov_fast_period*) which may take values from 3 to 300 periods. If the exchange rate increased by *pips_change* or more in the last *mov_fast_period* periods, then the rule supports a sell signal; if the exchange rate decreased by at least such an amount in that number of periods, then the rule supports a buy signal.

Rule 6 — Price far from moving average: Rule 6 compares the exchange rate with its long term exponential moving average. A large difference between them is an indication of a recent significant change from previous levels, and is taken as a signal of a short term trend that may be reversed. The rule defines two parameters to be optimised: a price difference (*pips_diff*), which may take values between 20 and 60 pips, and a number of periods for calculating the exponential moving average, which may take values from 31 to 200 periods. If the exchange rate is above the moving average by at least *pips_diff*, then the rule supports a sell signal; if the exchange rate is below the moving average by at least such an amount, then the rule supports a buy signal.

Rule 7 — RSI too extreme: Rule 7 tries to find extreme values of the RSI, signalling trends that may be reversed. The rule defines two parameters to be optimised: the number of periods for calculating the RSI, which may take values between 9 and 60, and a threshold (*rsi_threshold2*), which may take values between 71 and 95. If $RSI > rsi_threshold2$, then the rule supports a sell signal; if $RSI < 100 - rsi_threshold2$, then the rule supports a buy signal.

Rule 8 — Stochastic oscillator too extreme: Rule 8 tries to find extreme values of the stochastic oscillator. The rule defines the number of periods for calculating the stochastic oscillator %k, between 5 and 60, and a threshold (*sto_threshold*), between 70 and 95. If $\%k > sto_threshold$, then the rule supports a sell signal; if $\%k < 100 - sto_threshold$, then the rule supports a buy signal.

c. Generating buy and sell signals

Buy and sell signals are generated according to the support provided by the rules. In the optimisation process, each rule is given a weight between 0 and 11. The trend-following and trend-reversing rules are considered separately and, for each of these types of rules, a total support for opening a long (short) position is calculated as the sum of the weights of the rules that provide support for opening a long (short) position. So, at each period we will have four total support scores:

- Total support provided by trend following rules for opening a long position;
- Total support provided by trend following rules for opening a short position;
- Total support provided by trend reversing rules for opening a long position;

- Total support provided by trend reversing rules for opening a short position.

The generation of buy and sell signals depends on these total support scores. It was defined that a total support of at least 10, provided by either type of rule (trend-following or trend-reversing) is necessary for generating a signal. For each type of rule, only the position with largest total support is considered and, in case of contradictory indications given by different types of rules, trend-reversing rules are given priority over trend-following ones (the reason being that trend-reversing rules can be considered an extreme case of finding a trend that is being followed).

As an example, consider that the weights defined by the optimisation process were 1, 1, 8, 3 for trend-following rules 1-4 and 0, 11, 5, 5 for trend-reversing rules 5-8, respectively. Assume that rules 1, 2 and 4 and 7 provide support for a buy signal, and the other rules provide support for neither a buy nor a sell signal. The total support provided by trend-following rules for a buy signal is $1+1+3=5$, and the total support provided by trend-reversing rules is 5 (both provide null support for a sell signal). Since no type of rule provides a support of at least 10 for either buy or sell signals, no signal would be generated. Now assume that rules 2, 3 and 4 provide support for a buy signal, rules 7 and 8 provide support for a sell signal, and the other rules provide support for neither a buy nor a sell signal. The total support of trend-following rules for a buy signal is 12, and the total support of trend-reversing rules for a sell signal is 10. So, trend following rules would generate a buy signal whereas trend-reversing ones would generate a sell signal. Since trend-reversing rules are given priority, a sell signal would be generated.

Notice that such a system of non-normalised weights provides a large amount of flexibility since, at one extreme, the support of a single rule may be enough to generate a signal and, at the other extreme, it may allow all rules of a given type to be disregarded completely, if they do not contribute to maximising the risk-adjusted profits. In the middle of these two extremes, such a system allows signals to be generated from a very large number of combinations of support from different rules.

The reader may notice that, in spite of a score of 10 being enough to generate a signal, each single rule may have a score as large as 11. This was defined in this way in order to make it easier for the optimisation procedure (the genetic algorithm) to detect cases in which a rule should be enough to generate a signal (the cases in which a rule is very profitable by itself).

Finally, I also stress that the space of potential trading strategies that I take into account is extremely large - even if I ignore the difference between scores of 10 and 11 in each rule, I obtain about 3×10^{36} possible parameter combinations. This number is much larger than the number of strategies that may be extensively tested — for example, Sullivan *et al.* (1999) considered 7846 trading rules, and that may be already considered a very large set. The potential space of trading rules that I consider is much larger, and it would be impracticable to test all of them in order to choose the best one. So, I consid-

er a heuristic optimisation procedure — a genetic algorithm — to choose a ‘good’ trading rule instead of aiming to find the optimal one.

d. Using signals to simulate trades

I assume that, at each moment, there is at most one lot of 10,000 units of the base currency bought (or sold) for the pair of currencies being considered. This means that the system will be at one of three possible states: long (having bought 10,000 units of the base currency and sold the quote currency), short (having sold 10,000 units of the base currency and bought the quote currency), or neutral (having neither bought nor sold any currency).

The way in which the system reacts to a buy or sell signal depends on the state of the system. Let me start by considering the system in a neutral state. For the sake of clarity and economy of space, I will only explain how the system reacts to a buy signal (the reaction to a sell signal will be, essentially, symmetrical). When a buy signal is generated in a period in which the system is neutral, the system will try to open a long position in the next period, resorting to a parameter, *entry_pips*, which has a value between 0 and 20, defined by the optimisation process. The system simulates a buying order at a limit price, in which the limit is the result of adding *entry_pips* to the maximum ask price that occurred in the previous period (that is, in the period the signal was generated).⁴ If it would be possible to enter a long position with such a limit order, the system will consider the opening of that position at a price that mimics the price that would be achieved under real trading conditions (that is, the opening price depends on the way the quotes evolved in the period the position was opened). If it is not possible to open a long position with such a limit order, the system considers the order expired at the final of the period.

Whenever a position is opened, the system assumes that two associated orders are instantly issued: a take profit order and a stop loss order. The take-profit price of a long position is defined by adding a *take_profit_margin* to the entry price, whereas the stop-loss is defined by subtracting a *stop_loss_margin* from the entry price. Both the *take_profit_margin* and the *stop_loss_margin* are defined in the optimisation process, and both parameters may take values between 15 and 150 pips. An open position will thus be closed when the take-profit or the stop-loss prices are reached.

When the system is in a long position, a subsequent buy signal will be ignored. The same thing will happen to a sell signal occurring when the system is in a short position. A sell signal when the system is long, and a buy signal when the system is short, will result in a tightening of the stop-loss price (thus making it more probable that the position will be closed in a short time). In the case of a long position, if the system is long and a sell signal is issued in a given period, the stop loss price immediately becomes equal to the minimum bid price that occurred in that period, minus one pip (the higher stop

loss price that would not have triggered the closing of the position in that period).

Finally, open positions are also closed at the end of the trading day (assumed to be 17.00 ET). In fact, in a real setting, a position that is kept through the end of the trading day will be subject to the rollover rates. In order to perform a realistic evaluation of the trading strategies and not to have to take such rollover rates into account, open positions are automatically closed at the end of the trading day, at the prevailing price.

3. THE GENETIC ALGORITHM

In this work, the optimisation of the parameters of the trading strategy resorts to a genetic algorithm. A genetic algorithm is population-based heuristic used to find solutions for optimisation problems. The working of a genetic algorithm is inspired by the process of natural evolution: it starts with an initial set of solutions to the problem at hand (the initial 'population', often generated randomly) and makes it evolve in a way that aims to find better solutions, using techniques like selection, crossover, mutation, elitism and introduction of migrants. The algorithm stops when a given criterion is satisfied (usually a criterion based on the number of generations). I will now provide a brief outline of the genetic algorithm.

a. The population

The goal of the genetic algorithm is to find a set of parameters of the trading strategy that maximise an evaluation function (taking into account the profit and risk) in a training series. Each set of parameters is a solution, also referred to as an 'individual', and at each time the algorithm maintains a 'population' of solutions with a pre-defined size of 100 individuals.

Each parameter that composes a solution is termed a 'gene'. In this application, the set of genes defines a 'chromosome', and an individual is therefore defined by the respective chromosome.

A chromosome will have 24 genes, which were already mentioned in the previous section. They are:

- The eight weights of the rules;
- The 13 parameters that define the indicators used in the rules;
- The margins used to define the limit price of the orders, the stop loss price and the take profit price (that is, *entry_pips*, *stop_loss_margin* and *take_profit_margin*).

b. Evaluation function

Each individual in the population is evaluated according to an evaluation function. The result of the application of the evaluation function to an individual is usually termed the individual's fitness, and it is used to define the way the individual intervenes in the operations of the genetic algorithm.

The evaluation function that was defined takes into account the profit and the risk of the strategy in the training series. In this application, I used the Stirling ratio as the evaluation function. This ratio is defined as the quotient between the profit and the maximum drawdown (the maximum peak-to-trough decline in the value of the investment):

$$\text{Stirling ratio} = \frac{\text{profit}}{\text{maximum drawdown}} \quad (1)$$

The reason for using the maximum drawdown as a risk measure is related to the type of contract being considered. I am considering highly leveraged contracts and, in such cases, large losses cannot be tolerated even for a short time, since they may lead to the depletion of the available capital. In these cases, it is usual to use drawdown-based measures for the risk, since they explicitly account for the maximum decline in investment value. Hryshko and Downs (2004) also use the Stirling ratio in their analysis, and Brazabon and O'Neill (2004) and Dempster and Jones (2001) are examples of other authors that resort to drawdown-based measures in their analysis.

c. Operations

The initial population is generated randomly, and it evolves generation by generation. A set of operations, usual in genetic algorithms, were used in this application: selection, crossover, mutation, elitism and introduction of migrants. I will now briefly describe these operations.

Selection: All individuals that have not been involved in a minimum of one deal for 1000 periods are discarded. The remaining individuals are selected for breeding a new generation.

Crossover: Among those individuals that were chosen by the selection operator, two are randomly chosen as parents to generate an individual for the new generation. The probability of an individual being chosen is proportional to its fitness, and it is possible that one parent is equal to the other, in a given pair. In order to define the new individual, a cutting point n ($1 \leq n < 24$) is randomly chosen. The new individual inherits genes 1 to n from one parent and the remaining genes from the other parent.

Mutation: For each individual of the new generation, the value of each gene has a one per cent probability of being replaced by a new random value.

Introduction of migrants: An individual that is discarded in one generation — that is, an individual that is not involved in a minimum of one deal for 1000 periods — is replaced, in the next generation, by a migrant — a new individual generated randomly.

Elitism: The best individual of a given generation is always carried over to the next generation.

The evolution takes place until a pre-defined number of 50 generations is reached. The best individual of the last generation represents the best strategy that the genetic algorithm was able to identify.

4. DATA AND EMPIRICAL TESTS

I used data series from the Gain Capital market operator, available at <http://ratedata.gaincapital.com>. I considered four pairs of currencies: USD/HKD, USD/SGD, EUR/USD and GBP/USD. The data refer to the period from the beginning of September 2008 to the end of February 2011, since no data concerning previous periods is available for the USD/HKD and USD/SGD pairs. Original data contain tick-by-tick bid and ask quotes. The number of data points (in millions) was 0.485, 1.91, 10.5 and 15.0 for the USD/HKD, USD/SGD, EUR/USD and GBP/USD. This means that there are many more quote changes for the EUR/USD and GBP/USD than for the USD/HKD and USD/SGD pairs, reflecting the higher liquidity of the former currency pairs.

These original data were aggregated into time series with five and 15 minutes frequencies, which were used for defining the trading strategies. The original tick-by-tick data were then used to define exactly which trades take place. Figures 1-4 depict the behaviour of the data series used in this paper.

Fig. 1 - Quotes for the USD/HKD pair, for the considered period. Fig. 2 - Quotes for the USD/SGD pair, for the considered period.





Fig. 3 - Quotes for the EUR/USD pair, for the considered period. Fig. 4 - Quotes for the GBP/USD pair, for the considered period.



Time series were separated into training sets and test sets, and two different analyses were performed. The first one (denoted by A1) used a training set comprising data from September 15, 2008⁵ to the end of 2009, and a test set from the beginning of 2010 to the end of February 2011. The other one

(denoted by A2) splits the test period into two series, and uses different training series for both periods. The original training set, comprising data from September 15, 2008 to the end of 2009, is used generate trading rules that are applied to a test set that goes from the beginning of 2010 to the end of June 2010. A training set for the period from March 15, 2009 to the end of June 2010 is used to generate trading rules that are applied to a test set that goes from the beginning of July 2010 to the end of February 2011. The comparison of the results achieved in the analyses A1 and A2 will provide some indication of whether or not the use of training periods closer to the periods of application of the rules (in this case, the periods of the test sets) will lead to better performance of the trading rules.

The application used to perform the tests was built from scratch in Visual C++. In all cases, the genetic algorithm was run 50 times, each time considering a population of 100 individuals and stopping after 50 generations. The best individual of the last generation was then selected, and the corresponding strategy was applied to the test series. The results that will be presented in the next section are the averages for the 50 runs.

Since the genetic algorithm will generate a different trading strategy in each run, it may be questioned whether or not the average profits thus obtained can be considered statistically significant. In order to address this issue, I perform a one-tailed t-test to the average profit of the test series, with the null hypothesis that the average profit is smaller or equal to zero. When I am able to reject the null hypothesis, the conclusion is that a trading strategy generated in the training period and applied to the test period will lead to a positive expected profit in the latter period.

The goal of the tests was to assess whether the pairs of currencies including a bandwidth-bounded currency performed differently from the ones composed by free-floating currencies, both in terms of profitability and in terms of the types of rules that lead to better results. It may seem that the analysis of the preferred types of rules may be based on the weights of the different rules, but such analysis may be biased. In fact, it may be the case that trend-reversing rules have systematically larger weights than trend-following ones, not because they have more predictive ability but because of over-fitting — rules may be over-fitted to identify a small number of very profitable trends occurring in the training series, and such over-fitting is more probable with trend-reversing rules than with trend-following ones, since the former rule-type may identify the trends before the start, leading to larger profits. Therefore, a simple analysis of the weights may not identify correctly the most profitable types of rules.

The analysis of the preferred types of rules was based on the performance of the different types of rules in the test series. The algorithm was run in the training series considering the trend-following and trend-reversing rules separately, and the results obtained in the test series were compared for both types of rules. This way, it is possible to find out which type of rule leads to better results.

Finally, in the cases in which I am able to attain out-of-sample profits from the trading strategies, I try to find out whether the profits occur in a given sub-period or if they are divided all over the test period. I divide the test series into five sub-periods P₁-P₅: sub-periods P₁ to P₄ correspond to the four trimesters of 2010 and sub-period P₅ comprises January and February of 2011. Then I apply the trading strategies generated by analysis A2 to each of these sub-periods, in order to analyse the profits obtained in each of them.

5. RESULTS AND DISCUSSION

Initially, the algorithm was run with both types of rules, in order to find out whether the different pairs had significantly different profitability. Since the type of contracts being considered is highly leveraged, I do not take into account the opportunity cost of the required capital.

Table 1 shows the results obtained in the training series, for analysis A1, and the average for the two training series, in the case of A2. As expected, the results in the training series are always positive. The profits are higher for the EUR/USD and GBP/USD pairs, showing that the trading strategies take advantage of the frequent price changes and significant price swings. The profits are quite low for the USD/HKD pair, reflecting the lower frequency of price changes, and the stability of the prices during the training period (as can be confirmed in Figure 1).

Table 1 - Results obtained in the training series (average for 50 runs).

Currency pair	Freq.	Analysis A1			Analysis A2			
		Eval. Function	Pips	Pips/Cont. size	Eval. Function	Pips	Pips/Cont. size	
With transaction costs	USD/HKD	5 m	4.40	270	0.35%	4.46	243	0.31%
		15 m	6.28	450	0.58%	8.19	511	0.66%
	USD/SGD	5 m	2.80	390	2.77%	5.33	571	4.06%
		15 m	3.97	535	3.80%	6.67	694	4.92%
	EUR/USD	5 m	13.7	2058	15.2%	14.1	1941	14.3%
		15 m	16.6	2654	19.5%	19.6	2833	20.9%
GBP/USD	5 m	11.3	2412	15.4%	11.5	2281	14.5%	
	15 m	10.1	2187	13.9%	10.3	2109	13.4%	
Without transaction costs	USD/HKD	5 m	11.6	718	0.93%	10.6	586	0.75%
		15 m	10.7	769	0.99%	14.0	874	1.13%
	USD/SGD	5 m	11.4	1655	11.7%	14.2	1619	11.5%
		15 m	8.91	1220	8.66%	12.7	1371	9.73%
	EUR/USD	5 m	32.6	4791	35.3%	28.9	3910	28.8%
		15 m	35.5	5603	41.3%	38.6	5537	40.8%
GBP/USD	5 m	22.3	4752	30.3%	20.6	4054	25.8%	
	15 m	24.4	5457	34.8%	23.6	4903	31.3%	

Freq.: Frequency used in the application of the rules; Eval. Function: Value of the evaluation function (1); Pips: Profit/loss, in pips; Pips/Cont. Size: Number of pips gained or lost, in relation to the average contract size. In Analysis A2, the values that are presented for the evaluation function and the profit are the averages over the two runs (in the training series, both runs share a portion of the data).

Table 2 shows the average rule weights that were obtained for the different pairs of currencies in analysis A1 (the results corresponding to analysis A2 show a similar pattern). In the table, it can be seen trend-reversing rules have a larger average weight than trend-following ones, for all currencies. As argued before, such larger weight may be either due to a greater usefulness of trend-reversing rules or to a larger degree of over-fitting by this kind of rules. A conclusion may only be drawn when the two types of rules are considered separately.

Table 2 - Rule weights in analysis A1 (average for 50 runs).

<i>Rule</i>	1	2	3	4	5	6	7	8
USD/HKD	3.3	5.2	4.8	3.7	6.0	5.6	5.2	6.0
	Avg. Trend-following: 4.3				Avg. Trend-reversing: 5.1			
USD/SGD	1.8	2.2	2.6	2.0	6.7	5.1	4.0	4.6
	Avg. Trend-following: 2.2				Avg. Trend-reversing: 5.1			
EUR/USD	2.5	2.7	4.4	2.4	7.3	5.4	4.1	5.4
	Avg. Trend-following: 3.0				Avg. Trend-reversing: 5.1			
GBP/USD	1.8	2.8	2.6	1.8	5.6	3.9	2.8	5.8
	Avg. Trend-following: 2.2				Avg. Trend-reversing: 5.1			

Avg. Trend-following: Average weight of trend-following rules; Avg. Trend-reversing: Average weight of trend-reversing rules.

In order to assess whether the generated strategies really have some ability to generate profits, we must look at the results obtained in the test series, which are shown in Table 3. From this table we conclude that, when transaction costs are taken into account, the results are generally negative. The only exception seems to be the USD/SGD pair in the 15 minutes frequency, for which the system reaches statistically significant profits. When transaction costs are ignored, the results change significantly. The results are always positive for the USD/HKD and USD/SGD pairs, and are mostly positive for the GBP/USD. This means that, for these pairs, the derived strategies do have some power in predicting the direction of price changes, although not enough to pay the transaction costs. However, for the EUR/USD the results remain negative, meaning that the derived rules fail to identify the direction of future price changes.

When I compare the results of analyses A1 and A2 in Table 3, I conclude that in most cases analysis A2 leads to better results, although the difference is not large. This seems to indicate that the use of training periods closer to the periods of application of the rules leads to a slightly better performance of the derived strategies.

Table 3 - Results obtained in the test series, when using all rules (average for 50 runs).

Currency pair	Freq.	Analysis A1			Analysis A2				
		Eval. Function	Pips	Pips/Cont. size	Eval. Function	Pips	Pips/Cont. size		
With transaction costs	USD/HKD	5 m	-3.18	-313 (-10.59)	-0.40%	-4.07	-397 (-11.80)	-0.51%	
		15 m	-3.66	-319 (-7.83)	-0.41%	-4.90	-466 (-10.27)	-0.60%	
	USD/SGD	5 m	-0.099	-35.9 (-0.86)	-0.25%	-0.043	-14.6 (-0.35)	-0.10%	
		15 m	1.59	188 (6.48***)	1.34%	1.46	151 (3.96***)	1.07%	
	EUR/USD	5 m	-16.1	-2335 (-10.94)	-17.2%	-12.6	-2015 (-13.69)	-14.8%	
		15 m	-18.9	-2935 (-22.52)	-10.1%	-19.2	-3027 (-26.65)	-22.3%	
	GBP/USD	5 m	-0.741	-183 (-0.95)	-1.16%	0.306	61.7 (0.59)	0.39%	
		15 m	-10.7	-2396 (-6.48)	-15.3%	-6.13	-1276 (-6.87)	-8.13%	
	Without transaction costs	USD/HKD	5 m	0.612	58.1 (3.24***)	0.07%	0.834	75.5 (2.42***)	0.10%
			15 m	0.834	79.5 (2.68***)	0.10%	2.56	268 (6.76***)	0.35%
		USD/SGD	5 m	2.43	296 (6.50***)	2.10%	3.55	424 (9.98***)	3.01%
			15 m	2.86	347 (9.28***)	2.46%	4.15	474 (11.81***)	3.37%
EUR/USD		5 m	-4.39	-633 (-3.59)	-4.66%	-3.77	-621 (-4.76)	-4.58%	
		15 m	-6.39	-1005 (-10.19)	-7.40%	-4.80	-768 (-6.28)	-5.65%	
GBP/USD		5 m	3.87	749 (7.50***)	4.78%	4.16	833 (10.38***)	5.31%	
		15 m	-1.13	-348 (-1.92)	-2.22%	3.47	701 (4.79***)	4.47%	

Freq.: Frequency used in the application of the rules; Eval. Function: Value of the evaluation function (1); Pips: Profit/loss, in pips; Pips/Cont. Size: Number of pips gained or lost, in relation to the average contract size. In parenthesis: t-statistic for the null hypothesis that the average profit is smaller or equal to zero (one-tailed test). *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level.

Table 4 - Results obtained in the test series (avge for 50 runs), using only trend-following rules (left side of the table), and using only trend-reversing rules (right side of the table).

Currency pair	Freq.	Rules 1-4 (Trend-following)				Rules 5-8 (Trend-reversing)					
		Analysis A1		Analysis A2		Analysis A1		Analysis A2			
		Pips	Pips/ Cont. size	Pips	Pips/ Cont. size	Pips	Pips/ Cont. size	Pips	Pips/ Cont. size		
With trans. costs	USD/HKD	5 m	-146 (-7.09)	-0.19%	-374 (-14.55)	-0.48%	-346 (-7.93)	-0.45%	-485 (-12.87)	-0.62%	
		15 m	-48 (-2.53)	-0.12%	-158 (-5.41)	-0.20%	-391 (-6.71)	-0.50%	-200 (-6.61)	-0.26%	
	USD/SGD	5 m	-385 (-6.50)	-2.73%	-399 (-9.11)	-2.83%	52 (1.07)	0.37%	56 (1.32*)	0.39%	
		15 m	-406 (-12.00)	-2.88%	-381 (-14.12)	-2.71%	206 (5.37***)	1.46%	223 (6.95***)	1.58%	
	EUR/USD	5 m	-804 (-8.32)	-5.92%	-1179 (-11.57)	-8.68%	-2768 (-16.45)	-20.4%	-1983 (-16.87)	-14.6%	
		15 m	-709 (-6.87)	-5.22%	-722 (-5.26)	-5.32%	-2687 (-20.88)	-19.8%	-2656 (-23.27)	-19.6%	
	GBP/USD	5 m	-2145 (-11.02)	-13.7%	-2270 (-13.32)	-14.5%	-175 (-1.25)	-1.11%	27 (0.32)	0.17%	
		15 m	-4838 (-22.40)	-30.8%	-3676 (-24.59)	-23.4%	523 (6.13***)	3.33%	800 (6.93***)	5.10%	
	Without trans. costs	USD/HKD	5 m	94 (4.45***)	0.12%	141 (4.28***)	0.18%	282 (5.41***)	0.36%	61 (1.63*)	0.08%
			15 m	154 (6.64***)	0.42%	345 (13.85***)	0.44%	60 (1.59*)	0.08%	102 (4.07***)	0.13%
		USD/SGD	5 m	-177 (-4.16)	-1.26%	-136 (-3.64)	-0.97%	478 (8.34***)	3.39%	482 (12.86***)	3.42%
			15 m	-68 (-2.39)	-0.48%	-67 (-2.76)	-0.48%	392 (13.66***)	2.78%	467 (13.06***)	3.32%
EUR/USD		5 m	58 (0.55)	0.42%	-181 (-1.46)	-1.34%	-1092 (-9.51)	-8.04%	-770 (-8.35)	-5.67%	
		15 m	399 (4.60***)	2.94%	497 (4.33***)	3.66%	-1039 (-9.87)	-7.65%	-937 (-8.68)	-6.90%	
GBP/USD		5 m	-354 (-3.08)	-2.25%	-301 (-2.79)	-1.92%	686 (8.32***)	4.37%	595 (12.09***)	3.79%	
		15 m	-2554 (-21.82)	-16.3%	-1186 (-12.48)	-7.56%	1010 (11.12***)	6.44%	1214 (11.90***)	7.73%	

Freq.: Frequency used in the application of the rules; Eval. Function: Value of the evaluation function (1); Pips: Profit/loss, in pips; Pips/Cont. Size: Number of pips gained or lost, in relation to the average contract size. In parenthesis: t-statistic for the null hypothesis that the average profit is smaller or equal to zero (one-tailed test). *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level.

Table 5 - Profits (in pips) using analysis A2 when the test series is split into five sub-periods (average for 50 runs), for cases in which the profit was positive for the whole period.

		Sub-period							
Rules	Currency pair	Freq.	P_1	P_2	P_3	P_4	P_5		
With trans. costs	All	USD/SGD	15 m	-5 (-1.15)	140 (6.51***)	-26 (-1.87)	22 (0.93)	21 (2.41***)	
		GBP/USD	5 m	-116 (-1.88)	301 (4.80***)	-143 (-6.16)	-11 (-0.36)	31 (1.75**)	
	Trend reversing	USD/SGD	5 m	-26 (-3.37)	83 (3.22***)	26 (1.60*)	-62 (-2.97)	35 (3.66***)	
			15 m	-9 (-1.25)	158 (7.42***)	0 (-0.03)	47 (4.12***)	27 (3.86***)	
		GBP/USD	5 m	-40 (-1.06)	262 (4.39***)	-165 (-5.57)	-59 (-2.00)	29 (2.16**)	
			15 m	51 (1.30)	400 (7.85***)	-182 (-5.15)	206 (4.80***)	120 (6.37***)	
	Without trans. costs	All	USD/HKD	5 m	35 (7.32***)	-16 (-0.93)	-10 (-0.72)	24 (3.92***)	43 (3.24***)
				15 m	26 (3.48***)	27 (1.71**)	24 (1.32*)	105 (5.67***)	86 (4.40***)
			USD/SGD	5 m	5 (0.63)	205 (7.36***)	85 (6.28***)	72 (3.03***)	56 (6.73***)
				15 m	5 (1.32*)	228 (9.50***)	38 (3.20***)	155 (6.95***)	48 (4.25***)
GBP/USD		5 m	22 (0.49)	553 (8.52***)	60 (2.47***)	114 (3.56***)	85 (5.13***)		
		15 m	186 (3.42***)	160 (1.91**)	-19 (-0.27)	223 (2.70***)	151 (3.47***)		
Trend following		USD/HKD	5 m	29 (10.86***)	-67 (-5.48)	121 (7.14***)	37 (3.99***)	20 (1.50*)	
			15 m	23 (6.19***)	-5 (-0.56)	101 (6.72***)	103 (7.39***)	123 (8.26***)	
		EUR/USD	15 m	240 (4.93***)	-46 (-0.83)	-3 (-0.05)	332 (5.08***)	-26 (-0.65)	
		USD/HKD	5m	-10 (-1.26)	77 (4.37***)	-19 (-0.65)	21 (2.92***)	-9 (-0.75)	
15 m	-29 (-4.92)		73 (3.80***)	41 (3.78***)	9 (2.26**)	7 (0.84)			
Trend reversing	USD/SGD	5m	4 (0.56)	271 (9.30***)	79 (5.18***)	78 (4.42***)	51 (5.24***)		
		15 m	0 (-0.02)	272 (14.87***)	28 (3.35***)	132 (7.83***)	35 (4.07***)		
	GBP/USD	5m	91 (2.78***)	504 (8.58***)	53 (1.93**)	66 (2.48***)	85 (5.28***)		
		15 m	127 (3.53***)	542 (9.38***)	29 (0.71)	334 (6.95***)	181 (9.09***)		

Freq.: Frequency used in the application of the rules; P1: 1st trimester of 2010; P2: 2nd trimester of 2010; P3: 3rd trimester of 2010; P4: 4th trimester of 2010; P5: January and February 2011. In parenthesis: t-statistic for the null hypothesis that the average profit is smaller or equal to zero (one-tailed test). *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level.

In order to find out which types of rules are most useful for each pair of currencies, the algorithm was run again, first considering only trend-following rules and then considering only trend-reversing rules. The profits/losses achieved in the test series are shown in Table 4.

For each currency pair, there is a type of rule that leads to clearly better results. In the case of the USD/HKD and the EUR/USD trend-following rules lead to better results. In the case of the USD/SGD and the GBP/USD, trend-reversing rules are more profitable. So, I am not able to find significant differences between bandwidth-bounded and free-floating currencies in the types of rules that lead to better results.

When I was able to attain out-of-sample profits from the trading strategies, I divided the test series into five sub-periods P1-P5, as described in the previous section. I then applied the trading rules generated by analysis A2 to each of these sub-periods. The results obtained are presented in Table 5.

Let me now analyse each pair of currencies. The USD/HKD has the smaller difference of profitability between trend-following and trend-reversing rules. This may be due to the fact of the quotes being relatively stable for this pair, and the number of total quote changes being relatively small for the considered period. Such price stability and a small number of price changes may themselves be the result of a narrow price band, as is the case of this currency pair. Also, for this currency pair, I am not able to identify a sub-period in which the trading rules perform particularly well or particularly badly.

For the remaining pairs of currencies, the results seem conclusive enough. In the case of the USD/SGD, the use of trend-reversing rules always leads to profits, mostly of high statistical significance, whether or not transaction costs are taken into account. As for trend-following rules, they always lead to losses, whether or not transaction costs are taken into account. So, trend-reversing rules are more adequate for this pair of currencies, and seem able to lead to consistently positive results. This kind of behaviour seems consistent with a bandwidth-bounded currency, since a large change in the quotes of a pair containing such a currency will usually imply that one of the bands is being approached, and therefore lead to a reversion of the trend. When analysing the performance of trading strategies in the different sub-periods, I notice that profits are particularly high in sub-period P2 (second trimester of 2010), and particularly low in sub-period P1 (first trimester of 2010). By looking at Figure 2, we may see that in the second trimester of 2010 the currency quotes were more stable than in second, with larger price swings in the second trimester. These price swings seem to have created profit opportunities for trend-reversing rules, while the price stability in P1 made it more difficult to achieve profits.

When both rules were used, the EUR/USD was the only pair to have consistently negative results in the test series even without taking into account transaction costs. This result contrasted with the fact that this pair of currencies has the best results in the training series, when both types of

rules were considered. The results presented in Table 4 shed some light on this apparent contradiction. In fact, the use of trend-reversing rules leads to significant losses, while the use of trend-following rules generally leads to positive results (if transaction costs are ignored). So, trend-reversing rules are probably finding unique profitable patterns in the training data, which do not repeat in the test data. Therefore, such rules cause significant losses in the test series. We may conclude that trend-following rules perform better for this currency pair.

For the GBP/USD pair, trend-following rules always lead to losses. Trend-reversing rules, on the other hand, always achieve positive results of high statistical significance when transaction costs are ignored, and lead to mostly positive results when transaction costs are taken into account. Therefore, trend-reversing rules are more adequate for this pair. Similarly to the USD/SGD pair, profits are larger in sub-period P2, which corresponds to the second trimester of 2010 (although profits are mostly split throughout all sub-periods). The second trimester of 2010 marks the reversion from a downward trend started near the end of 2009 to an upward trend. This reversion seems to have created profit opportunities for the trend-reversing rules, leading to larger profits.

6. CONCLUSIONS

In this paper, I test whether the existence of currency floating bands create exploitable opportunities for abnormal profits, and whether the most profitable trading strategies have different characteristics for bandwidth-bounded currencies and free-floating ones.

The genetic algorithm that was defined for these tests seems to be able to generate positive returns in the case of the USD/SGD, even when transaction costs are taken into account. The results are more significant if only trend-reversing rules are considered (trend-reversing rules could be expected to lead to larger profits when currencies are bounded by bandwidths). In the case of the USD/HKD, there is no evidence of the genetic algorithm being able to achieve profits, if transaction costs are taken into account. However, the results achieved for this pair of currencies may be negatively influenced by the relative stability of the price of this currency pair, and by the small number of quote changes in the considered period. Also notice that the considered rules intended to be generic, not exploiting any currency-specific information. If the specific (known) bands of the HKD were explicitly considered in the rule definition, the results might have been different.

In the pairs of free-floating currencies, there is no indication of the genetic algorithm being able to achieve positive returns for the EUR/USD (when transaction costs are considered), and there is some limited indication that it might be able to achieve profits in the GBP/USD pair, but only if it just uses trend-reversing rules. So, the evidence of being able to achieve profits is somewhat more significant in the case of the USD/SGD, a bandwidth-bound-

ed currency. However, the existence of a narrow price band in a bandwidth-bounded currency, as is the case of the USD/HKD, may lead to price stability, and erase the opportunities for achieving profits.

In the analysis of the preferred types of rules, I find that trend-reversing rules are preferred in the case of the USD/SGD pair, as expected initially, but I have no such indication for the USD/HKD pair. However, once again I must point out that the results achieved for this pair may be influenced by the price band being very narrow. For the other pairs, I find that different types of rules are preferred for the EUR/USD and for the GBP/USD. So, there seems to be no preferred type of rule for the free-floating currencies, while the only pair involving a bandwidth-bounded currency without a very narrow band favours trend-reversing rules, as expected.

Accepted for publication: 24th January 2012

ENDNOTES

1. Grupo de Estudos Monetários e Financeiros (GEMF) and Faculty of Economics - University of Coimbra, Av. Dias da Silva, 165, 3004-512 Coimbra - Portugal. E-mail: pgodinho@fe.uc.pt. The author is indebted to an anonymous referee, whose comments contributed significantly to improve the paper.

2. In a currency pair, the first currency is the base currency and the second is the quote currency. This paper follows the currency codes defined by ISO 4217. DEM is the German mark, EUR is the euro, GBP is the pound sterling, HKD is the Hong Kong dollar, JPY is the Japanese yen, SGD is the Singapore dollar and USD is the United States dollar.

3. A pip (percentage in point) is a measure of price changes that is commonly used in the Forex market. In the case of the currency pairs considered in this paper, it represents a unit change in the fourth decimal place of the exchange rate. If an exchange rate increases from 1.0200 to 1.0203 it is said to have increased by three pips.

4. In the case of a short position, the limit price is the minimum bid price of the previous period minus entry_pips.

5. The reason for starting at this date, instead of using the data series from the beginning, concerns the need of a warming up period for the technical indicators.

REFERENCES

Brabazon A and O'Neill M (2004) 'Evolving Technical Trading Rules for Spot Foreign-Exchange Markets Using Grammatical Evolution', *Computational Management Science*, 1(3), 311-327.

Dempster M A and Jones C M (2001) 'A Real Time Adaptive Trading System Using Genetic Programming', *Quantitative Finance*, 1(4), 397-413.

- Dunis C, Harris A and Leong S (1999) 'Optimising Intraday Trading Models with Genetic Algorithms', *Neural Network World*, 9(3), 193-223.
- Fama E F (1970) 'Efficient Capital Markets: A Review of Theory and Empirical Work', *The Journal of Finance*, 25(2), 383-417.
- Hryshko A and Downs T (2004) 'System for foreign exchange trading using genetic algorithms and reinforcement learning', *International Journal of Systems Science*, 35(13), 763-774.
- LeBaron B (1999) 'Technical trading rule profitability and foreign exchange intervention', *Journal of International Economics*, 49(1), 125-143.
- LeBaron B (2002) 'Technical Trading Profitability in Foreign Exchange Markets in the 1990s', Brandeis University, working paper.
- Levich R M and 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(5), 451-474.
- Lo A W (2004) 'The Adaptive Markets Hypothesis', *The Journal of Portfolio Management*, 30(5), 15-29.
- Mendes L, Dias J and Godinho P (2010) 'The Profitability of Technical Analysis in the EUR/USD Currency Pair: Evidence from a Trading System', Proceedings of the 6th Portuguese Finance Network Conference, 1-3 de July 2010, Ponta Delgada, Portugal.
- Menkhoff L and Taylor M P (2007) 'The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis', *Journal of Economic Literature*, 45(4), 936-972.
- Meyers T A (1989) *The Technical Analysis Course*, New York NY: McGraw-Hill.
- Neely C and Weller P (2003) 'Intraday Technical Trading in the Foreign Exchange Market', *Journal of International Money and Finance*, 22(2), 223-237.
- Neely C, Weller P and Dittmar R (1997) 'Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach', *Journal of Financial and Quantitative Analysis*, 32(4), 405-426.
- Neely C J, Weller P A, and Ulrich J M (2009) 'The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market', *Journal of Financial and Quantitative Analysis*, 44(2), 467-488.
- Olson D (2004) 'Have Trading Rule Profits in the Currency Markets Declined Over Time?', *Journal of Banking & Finance*, 28(1), 85-105.
- Park C-H and Irwin S H (2007) 'What Do We Know About the Profitability of Technical Analysis?', *Journal of Economic Surveys*, 21(4), 786-826.
- Schulmeister S (2008) 'Components of the Profitability of Technical Currency Trading', *Applied Financial Economics*, 18(11), 917-930.
- Sullivan R, Timmermann A and White H (1999) 'Data-Snooping, Technical Trading Rule Performance, and the Bootstrap', *The Journal of Finance*, 54(5), 1647-1691.

P Godinho

Sweeney R J (1986) 'Beating the Foreign Exchange Market', *The Journal of Finance*, 41(1), 163-182.

White H (2000) 'A Reality Check for Data Snooping', *Econometrica*, 68(5), 1097-1126.