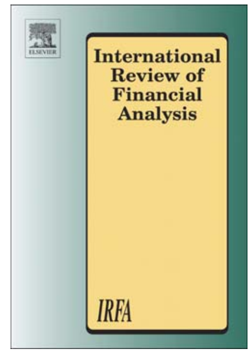

FX technical trading rules can be profitable sometimes!

Nima Zarrabi, Stuart Snaith, Jerry Coakley

PII: S1057-5219(16)30195-8
DOI: doi:10.1016/j.irfa.2016.12.010
Reference: FINANA 1070

To appear in: *International Review of Financial Analysis*

Received date: 14 November 2016
Accepted date: 31 December 2016

FX technical trading rules can be profitable sometimes!

Nima Zarrabi[*] Stuart Snaith[†] Jerry Coakley[‡]

November, 2016

Abstract

This paper investigates the profitability of technical trading rules in the foreign exchange market taking into account data snooping bias and transaction costs. A universe of 7,650 trading rules is applied to six currencies quoted in U.S. dollars over the 1994:3-2014:12 period. The Barras, Scaillet, and Wermers (2010) false discovery rate method is employed to deal with data snooping and it detects almost all outperforming trading rules while keeping the proportion of false discoveries to a pre-specified level. The out-of-sample results reveal a large number of outperforming rules that are profitable over short periods based on the Sharpe ratio. However, they are not consistently profitable and so the overall results are more consistent with the adaptive markets hypothesis.

JEL Classification: C12, C15, F31, G11, G14.

Keywords: Technical trading, False discovery rate, Persistence analysis, Exchange rate.

[*]Essex Business School, University of Essex, Colchester, CO4 3SQ, UK; nzarra@essex.ac.uk.

[†]Gustavson School of Business, University of Victoria, Victoria, V8P 5C2, Canada; ssnaith@uvic.ca.

[‡]Corresponding author: Essex Finance Centre and Essex Business School, University of Essex, Colchester, CO4 3SQ, UK; jcoakley@essex.ac.uk. Coakley gratefully acknowledges support from grant number ES/L011859/1, from the Business and Local Government Data Research Centre, funded by the Economic and Social Research Council.

1 Introduction

Weak-form efficiency implies that technical trading rules (TTRs) based on historical data should not be profitable. There is an ongoing debate about whether TTRs can consistently yield profits or whether any profits realised may reflect just luck rather than the effects of the TTRs themselves. Menkhoff and Taylor (2007) highlight the fact that technical analysis is extremely popular among foreign exchange (FX) market participants and an extensive literature has developed on this topic. One aspect of this literature is based on surveys of technical traders. These surveys, from the first by Taylor and Allen (1992) to the more recent ones by Cheung and Chinn (2001) and Cheung, Chinn, and Pascual (2005) highlight the extensive use of technical analysis in the FX market.

A subset of the TTR literature examines TTR profitability while addressing data snooping bias. However, the focus of many of these papers is on the equity and not the FX market. For example Brock, Brock, Lakonishok, and LeBaron (1992) claim that the deliberate choice of a simple class of rules that has been commonly used for a long period can allay the problem of data snooping. Qi and Wu (2006) and Hsu, Hsu, and Kuan (2010) use White's (2000) reality check (RC) and stepwise-superior predictive ability (SPA) tests, respectively, to address this issue. While the reality check is conservative, Hansen (2005) shows that the SPA is more powerful and less sensitive to the inclusion of poor rules. The stepwise-SPA technique is a multiple test based on the SPA and has been used in recent studies on technical analysis applied to equity markets. This literature includes Hsu and Kuan (2005), Marshall, Cahan, and Cahan (2008) and Shynkevich (2012a,b,c). Recently, Hsu and Taylor (2013) and Coakley, Marzano, and Nankervis (2016) have applied the stepwise-SPA test to TTRs of large samples in the FX markets. Both find little or no evidence for TTR profitability after taking into account transaction costs and data snooping.

This paper contributes to the literature in two respects. Our first contribution is to evaluate whether TTR profitability is robust to data snooping using a new and more powerful data snooping approach. Specifically it employs the false discovery rate (FDR) method proposed by Barras et al. (2010) to check the statistical robustness of our results which determines the balance between wrongly selecting underperforming trading rules (false discoveries) and leaving out truly outperforming ones. To the best of our knowledge, this paper is the first to employ the FDR methodology to the FX market. The motivation is that the previously used data snooping methodologies either can detect only one outperforming rule (the RC and SPA tests) or they are too conservative (the stepwise-RC and stepwise-SPA test). The FDR method's advantage over other methodologies is that it allows for a small (and specific) proportion of false discoveries while detecting almost all outperforming rules. In other words, the FDR controls for the expected proportion of Type I errors among the rejected hypotheses. We use the Sharpe ratio criterion to evaluate the long-run in-sample performance of a set of 7,650 TTRs after implementing the FDR methodology. The results indicate that a large proportion of TTRs (up to 75%) for four currencies is profitable, even after taking into account transaction costs.

Given the evidence on TTR in-sample predictive ability, our second contribution is to perform a comprehensive out-of-sample persistence analysis for the economic significance of TTRs. The motivation is that, as noted by Bajgrowicz and Scaillet (2012), the economic value of trading strategies can be assessed only by considering performance persistence after taking both transaction costs and data snooping bias into account. The outperforming rules are selected at the end of each month based on their FDR and accounting for transaction costs and on that basis a portfolio is constructed. This is adjusted on a monthly basis and used in the subsequent month. Persistence tests are employed for daily U.S. dollar prices for six currencies over the 1995-2014 sample period and the out-of-sample performance of TTRs is evaluated based on their Sharpe ratios.

The results show that, the FDR methodology indicates a large number of outperforming trading rules for all exchange rates over the 1994-2014 sample period. These TTRs generate profits for periods of between one and three months but not over the full sample. The implication is that the Sharpe ratio fluctuates over time, implying that a positive Sharpe ratio is often offset by a negative one. Thus TTRs are sometimes profitable having taken into account both data snooping bias and transaction costs. However, despite the short-run profit opportunities an investor could not consistently make profits by using previously outperforming TTRs over the entire sample period which includes the 2008 banking and financial crisis and its aftermath. These results suggest that rather than FX markets being efficient in the traditional sense, instead they exhibit behavior consistent with adaptive market efficiency (see Lo's (2004)).

This paper proceeds as follows. Section 2 reviews the literature on profitability of technical analysis. Section 3 reviews existing methods to account for data snooping and presents the FDR based approach. Section 4 presents our data and empirical findings. Section 5 concludes.

2 Technical analysis

Technical analysis is particularly popular in FX markets. It uses past price behavior and sometimes volume data to guide trading decisions in asset markets. The philosophy of technical analysis is explained in the seminal textbook of Murphy (1986) in which he points out three principles about the behavior of technical analysts. The first is that an asset's price history incorporates all relevant information which makes the use of asset fundamentals valueless. The second is that asset prices move in trends, implying predictability and also profitability by buying (selling) assets when the price is rising (falling). This is why technicians believe that "the trend is your friend." The third principle is that asset price patterns tend to repeat themselves. In other words, traders tend to react in a similar way when confronted by similar conditions.

Pring (1991) explains that the technical method is basically a reflection of the idea that price movements occur in trends. He argues that such trends are determined by the changing attitudes of traders toward a variety of economic, monetary, political and psychological variables. It is, thus, no surprise that an extensive literature has been developed on the performance of trading rules (see Menkhoff and Taylor (2007) for a review). However, regardless of the popularity of TTRs among practitioners, academics have long been skeptical about the advantages of technical analysis. They claim that the profitability of technical trading rules is contrary to the efficient market hypothesis (EMH) which assumes that no trading rules based on publicly available information should be able to generate profits without bearing any risk.

This paper builds on Bajgrowicz and Scaillet (2012) which focuses on equity markets. They evaluate the performance of 7,846 TTRs over more than hundred years while using the FDR to control for the data snooping bias. They evaluate the out-of-sample performance of the TTRs to find whether an investor could reasonably have predicted which rules would generate superior returns after taking into account transaction costs. The results show a negative out-of-sample performance for TTRs throughout the recent periods, implying that there is no hot hand phenomenon.[1]

2.1 Technical analysis in the FX market

Menkhoff and Taylor (2007) explain that the FX market differs from equity markets in some aspects. First, total turnover in the global FX market is much greater than the turnover of the largest stock exchanges. Second, as mentioned by Sager and Taylor (2006), due to existence of professional traders in FX markets, the impact of individual investors may be neglected without loss of generality. Third, as highlighted by Lyons (2001), the FX market has a much higher share of short term interdealer trading. Finally, there is less confidence among traders in models of fair value in the FX market compared to equity markets (see Taylor, 1994).

Neely, Rapach, Tun, and Zhou (2014) explain that technical analysis dates back at least to 1700 and although modern technical analysis was originally developed in the context of the stock market, foreign currency traders have employed this approach for trading in the modern floating rate era. The studies by Goodman (1979) and Goodhart (1988) are the first to bring the broad use of technical analysis by FX professionals to the attention of academic researchers. An early survey of market participants by Taylor and Allen (1992) shows that almost all chief foreign exchange dealers in London use technical analysis to some extent and also tend to combine it with fundamental analysis. They find that at short horizons (less than a week) traders use technical analysis more frequently than fundamental analysis. In a survey of U.S. FX traders, Cheung and Chinn (2001) find that 30% of participants could be considered as technical traders due to the increasing use of technical analysis.[2]

The extensive literature on the performance of TTRs in the FX market finds support for their profitability. Dooley and Shafer (1984) find evidence for the profitability of filter rules over relatively short horizons. Using the post-Bretton Wood data, Cornell and Dietrich (1978) find that filter and moving average rules are profitable. Sweeney (1986) confirms previous findings on the effectiveness of filter rules on many dollar exchange rates while he considers both transaction costs and risk. These early studies have several drawbacks: First, their results may be subject to the data snooping problem, since the

[1]The hot hand phenomenon refers to the expectation of "streaks" in sequence of hits and misses whose probabilities are independent (Wilke and Barrett (2009)).

[2]See Gehrig and Menkhoff (2004) for similar survey evidence for German FX traders.

early literature is mainly focused on popular rules such as the filter or moving average. Second, the riskiness of TTRs is often ignored. In other words, TTR returns do not necessarily reject market efficiency since returns may reflect risk premiums. Finally, it is hard to interpret the results of early studies since the performance of the rules are often reported as averages across all the rules or all the assets.

The out-of-sample performance of TTRs is assessed by Neely, Weller, and Dittmar (1997) using genetic programming techniques. They find that not only out-of-sample excess returns of the selected rules are economically significant, but also these returns are not compensations for bearing systematic risk. They account for data snooping bias by using the data snooping technique which was first used by Brock et al. (1992), to support their findings on the profitability of technical analysis in FX market. They explain that while central bank intervention generates a source of speculative profit in the market, it is less likely to be the reason for their findings. Instead, they highlight the role of market efficiency as a plausible explanation for their results.

The Qi and Wu (2006) study is the first to use the RC in the FX market to control for data snooping bias. A large number of TTRs are applied to seven daily dollar exchange rates over the 1973-1998 period. They find that the profitability of all seven currencies is significant at the one percent level even after taking into account both data snooping bias and transaction costs. However, their trading rules do not perform as well in their out-of-sample exercise. They suggest that the FX market becomes more efficient over time. The performance of technical analysis as opposed to market efficiency is the main focus of the Neely and Weller (2011) paper. They analyse the intertemporal stability of excess returns to trading rules in the FX market obtained from out-of-sample tests on previously studied rules. Their results show that the excess returns of the 1970s and 1980s were genuine and not just due to data mining (see 3.1). They demonstrate that, although these profit opportunities disappeared by the early 1990s for well-known rules such as the filter and moving average rules, less-studied rules have remained profitable. These findings are consistent with the adaptive market hypothesis by Lo (2004) in which the performance of TTRs is short-lived with prolonged periods of success and failure.

In a recent study, Hsu and Taylor (2013) carry out a large scale investigation of trading rules in the FX market. They use daily data over a maximum of forty years for thirty developed and emerging market currencies and over 21,000 TTRs. They apply the stepwise-SPA test to control for data snooping and find strong evidence for predictability of TTRs in the 1970s and 1980s in developed currencies and in the 2000s in emerging currencies. They show that an equally weighted portfolio of TTRs yields an average compound return of 9% over the last thirty years in the case of emerging currencies. They explain that the predictability and profitability of TTRs are due to short-lived, not-fully-rational behavior and immaturity, instead of autocorrelation, risk permia and central bank intervention.

2.2 Technical trading rules

Technical analysis can be divided into two broad categories: charting and mechanical methods. The former is the older method and graphs the history of prices over a specific period and uses patterns to forecast future movements. Charting is subjective and requires the analyst to be able to find and interpret the patterns. Charting is beyond the scope of this paper and we focus on mechanical methods. Mechanical rules require the analyst to apply rules based on mathematical functions of past and current data.

This paper investigates ten types of trading rules including both well-known and less-studied rules: filter, moving average, linearly weighted moving average, exponentially weighted moving average, moving average convergence-divergence, moving average oscillator, stochastic oscillator, channel breakout, relative strength indicator and bollinger bands. Qi and Wu (2006) point out that choosing too few rules is likely to cause biases in statistical inference due to data mining. On the other hand, Hansen (2003) argues that considering too many irrelevant rules can reduce test power. We study a universe of 7,650 rules which is similar to the number of trading rules used in both the Sullivan, Timmermann, and White (1999) and Bajgrowicz and Scaillet (2012) papers.[3] We use Qi and Wu's (2006) parameters for the filter, moving average and channel breakout rule and, for the others, we select reasonable parameters that lie in the ranges used in the literature. Table 1 provides a summary of the parameters used in the calibration of trading rules. The details of each rule are given in Appendix A.

< Table 1 around here >

[3]Sullivan et al. (1999) study 7,846 trading rules. Bajgrowicz and Scaillet (2012) study the same trading rules and sample but employ a different bootstrap method.

3 Methodology

3.1 Data snooping methods

Neely and Weller (2011) point out that data snooping, data mining and publication bias are three related but distinct problems that could lead to false conclusions about the profitability of TTRs. Data snooping is defined as a situation in which researchers, either consciously or unconsciously, study rules that are already proven to be profitable. Data mining is when researchers consider many rules but only interpret the most successful ones. In other words, negative results are ignored, while positive results are reported. Publication bias is the tendency of journals to accept papers with positive results rather than negative results. Technicians normally backtest rules which implies that they tend to study rules which are already profitable on past data. As Neely and Weller (2011) discuss that it is almost impossible to avoid some data snooping as data sets and rules are limited but there are nevertheless a number of ways to control for it, which include the reality check and its extension, and the FDR which we present in the following subsections.

3.2 White's (2000) reality check and extensions

White's (2000) RC is the first method which exploits the dependence structure of the individual test statistics. The RC approach tests the null hypothesis that the best rule does not outperform the benchmark. One disadvantage of White's method is that once it identifies an outperforming rule, the procedure stops. In other words, when the null hypothesis is rejected, it concludes that there is at least one model which beats the benchmark, but it is not able to identify the complete subset of outperforming rules. This drawback is addressed by Romano and Wolf (2005) who suggest a stepwise multiple testing procedure to asymptotically control the family-wise error rate (FWER) at a chosen level. Using a stepwise-RC method, further outperforming rules can be detected in subsequent steps which makes this method more powerful than the RC.[4]

Another drawback of the RC from which the stepwise-RC also suffers is that when the

[4]When testing multiple hypotheses, each test has type I and type II errors. A measure to determine the overall error rate is FWER which is the probability of making one or more type I errors. Therefore, instead of controlling the probability of a type I error at a chosen level for each test, the FWER is controlled at this level.

hypotheses involve inequality constraints both methods are conservative. This is because they are both based on the least favorable configuration (LFC).[5] Hansen (2005) shows that a test based on the LFC is likely to lose power dramatically when many poor and irrelevant rules are included in the test and therefore develops the SPA test. He modifies White's RC test such that it is not based on the LFC but instead tests whether the best rule beats a benchmark forecast.[6]

Similar to the extension from the RC test to the stepwise-RC test, it is logical to extend the SPA to the stepwise-SPA test to try to detect all outperforming rules. Hsu et al. (2010) prove analytically and by simulations that the stepwise-SPA is more powerful than the stepwise-RC. The main drawback of their proposed method is that the stepwise-SPA test controls for the FWER which guards against any false positives. This implies that the FWER is too conservative and results in many missed findings. However, in practice, investors tend to identify as many outperforming rules as possible since they do not consider the signal of one TTR at a time but a combination of multiple strategies.

3.3 False discovery rate

To address the issue of a too conservative FWER, Benjamini and Hochberg (1995) propose the false discovery rate which is a less conservative method. The idea of the FDR is that by allowing a specific proportion of false discoveries, it significantly improves the power of detecting the outperforming rules. Barras et al. (2010) extend the FDR by introducing the $FDR^{+/-}$, which estimates separately the proportion of false discoveries among trading rules that perform better or worse than the benchmark of no excess returns. For instance, an FDR^{+} of 20% implies that among the rules selected as outperforming, on average 20% do not deliver positive returns while an FDR^{+} of 100% indicates that none of the selected models generates positive performance. Table 2 shows possible outcomes of hypothesis testing. While the FWER is calculated as $FWER=P(V\geq 1)$ where V is the number of type I errors, the FDR can be interpreted as the expected proportion of Type I errors among the rejected hypothesis (R), $FDR=E(\frac{V}{R}|R>0)P(R>0)$. For multiple hypothesis tests, this implies that the FDR is much less conservative compared to the FWER and results in a significant improvement in the power of test.

[5]That is least favourable to the alternative hypothesis. In other words, it is a set of parameter values that minimizes the probability of a correct decision.

[6]As with the RC test once an outperforming rule is identified the test stops.

< Table 2 around here >

The FDR identifies the outperforming rules, even if the performance of the outperforming rules are due to luck, contrary to the stepwise-RC and the stepwise-SPA. Therefore, the FDR detects almost all outperforming rules, while keeping the amount of false discovery at a chosen level which provides investors with a diversified pool of signals.[7]

Benjamini and Hochberg (1995) assume that the multiple hypotheses are independent when they propose the FDR approach. Storey (2003), Storey and Tibshirani (2003), and Storey, Taylor, and Siegmund (2004) show that when the number of tests is large, the FDR still holds under weak dependence of the p-values. Bajgrowicz and Scaillet (2012) define weak dependence as any form of dependence whose effect becomes negligible as the number of tests increases to infinity. In this paper, although TTRs are dependent in small groups, each family of rules acts independently of the others. For example, a five-day MA with a 0.001 band is highly correlated with a five-day MA with a 0.005 band. However, it is less likely to be correlated with a two hundred-day MA. Such type of dependence is called block dependence and meets the weak dependency conditions.

The FDR$^+$ is estimated as follows:

$$\mathrm{FDR}^+_{(\gamma)} = \frac{\hat{F}^+_{(\gamma)}}{\hat{R}^+_{(\gamma)}},$$

where $\hat{R}^+$ is an estimator of R^+ and denotes the number of trading rules selected as significantly outperforming rules:

$$\hat{R}^+_{(\gamma)} = \#(p_k \leq \gamma, SR_k > 0) \quad \text{where} \quad k = 1, \ldots, 7650,$$

where SR_k denotes the Sharpe ratio of each rule k. $\hat{F}^+$ is an estimator of F^+ and denotes the number of trading rules that do not generate genuine performance but have been selected erroneously. The FDR method allows one to estimate π_A^+ and π_A^- which are the proportions of positive and negative TTRs in the population, respectively. π_0 is the proportion of rules without abnormal performance and as explained below, is required

[7]The advantages of the FDR approach over the stepwise-RC are demonstrated by the simulations of Bajgrowicz and Scaillet's (2012) paper where they show that the lack of power of the stepwise-RC comes from the very conservative criteria underlying that method.

for estimating $\hat{F}^+$:

$$\hat{F}^+_{(\gamma)} = \frac{1}{2}\hat{\pi}_0 l\gamma,$$
$$\hat{\pi}_0(\lambda) = \frac{\#(p_k > \lambda; k = 1, \ldots, l)}{l(1-\lambda)},$$

where l is the number of TTRs in our sample, γ is a threshold which is applied to determine the null and alternative p-values, and λ is a tuning parameter.

The estimation procedure for the optimal values of λ (λ^*) and γ (γ^*) is proposed by Storey (2003) and Storey et al. (2004). To determine the value of λ^*, one considers a range of values ($\lambda = 0.05, 0.1, \ldots, 0.95$) and calculates the $\hat{\pi}_0$ for each value. Among them, choose the minimum $\hat{\pi}_0$ and for each possible value of λ and implement 1,000 bootstrap replications of $\hat{\pi}_0(\lambda)$ by drawing with replacement from TTR p-values. We denote them $\hat{\pi}^b_0(\lambda)$ for $b = (1, \ldots, 000)$. The estimated mean squared error of $\hat{\pi}^b_0(\lambda)$ for each value of λ is computed as follows:

$$M\hat{S}E_\lambda = \frac{1}{1000}\sum_{b=1}^{1000}[\hat{\pi}^b_0(\lambda) - \min_\lambda \hat{\pi}_0(\lambda)]^2. \tag{1}$$

The λ^* is determined such that $\lambda^* = \arg\min_\lambda M\hat{S}E_{(\lambda)}$. It is , however, worth noting that Barras et al. (2010) find that $\hat{\pi}_0$ is not overly sensitive to the choice of λ^*.

The value of γ^* is determined in the same way. Consider a range of values ($\gamma = 0.30, 0.35, \ldots, 0.50$) and compute the relative $\hat{\pi}^-_A(\lambda)$ which is the proportion of rules with negative abnormal performance using the following formula:

$$\hat{\pi}^-_A = \frac{T^-_{(\gamma)} - A^-_{(\gamma)}}{l},$$

where $T^-_{(\gamma)}$ denotes the number of alternative rules with negative performance and a p-value smaller than γ, and $A^-_{(\gamma)}$ denotes the number of alternative models with negative performance and a p-value greater than γ. Then find the value of γ which maximises $\hat{\pi}^-_A$ and perform 1,000 bootstrap replications of $\hat{\pi}^-_A$ for each possible value of γ which are denoted by $\hat{\pi}^{b-}_A(\gamma)$ for $b = 1, \ldots, 1000$. Calculate the estimated mean squared error of

$\hat{\pi}_A^-(\gamma)$ for each value of γ as follows:

$$M\hat{S}E_(^-\gamma) = \frac{1}{1000}\sum_{b=1}^{1000}[\hat{\pi}_A^{b-}(\gamma) - \max_\gamma \hat{\pi}_A^-(\gamma)]^2. \quad (2)$$

Therefore, γ^- is calculated such that $\gamma^- = \arg\min_\gamma M\hat{S}E^-_{(\gamma)}$. The same data driven procedure is used to determine γ^+. Note that if $\min_\gamma M\hat{S}E^-_{(\gamma)} < \min_\gamma M\hat{S}E^+_{(\gamma)}$, set $\hat{\pi}_A^-(\gamma^*) = \hat{\pi}_A^-(\gamma^-)$ and to preserve the equality, set $\hat{\pi}_A^+(\gamma^*) = 1 - \hat{\pi}_0 - \hat{\pi}_A^-(\gamma^*)$. Otherwise, set $\hat{\pi}_A^+(\gamma^*) = \hat{\pi}_A^+(\gamma^*)$ and $\hat{\pi}_A^-(\gamma^*) = 1 - \hat{\pi}_0 - \hat{\pi}_A^+(\gamma^*)$. Barras et al. (2010) explain that although this method is entirely data driven, there is some flexibility in the choice of γ^*, as long as it is sufficiently high.

4 Data and empirical analysis

4.1 Data

Daily data are downloaded from Datastream for the following currency pairs relative to the U.S. dollar over the 1994:3-2014:12 period : the British pound (GBP), the Canadian dollar (CAD), the Japanese yen (JPY), the Swedish krona (SEK), the Swiss franc (CHF), and the Norwegian krone (NOK). The sample is chosen to include a mixture of heavily traded currencies (GBP, CAD, JPY, CHF) and less heavily traded currencies (SEK, NOK). This provides a comprehensive test of TTR profitability since the evidence on this is weak (see Hsu and Taylor (2013)). For interest rates, we collect the overnight Eurocurrency interest rates from Datastream.

Table 3 reports descriptive statistics of the daily spot exchange rate returns.[8] The mean return rates show that on average the U.S. dollar depreciates against all exchange rates. The maximum daily depreciation of the U.S. dollar relative to these six currencies lies between 4.47% (GBP) and 6.6% (JPY). The standard deviation of daily changes lies between 0.0053 and 0.0072 which implies substantial daily volatility. The JPY has the smallest Sharpe ratio, while the corresponding number is much larger for the CHF. With the exception of GBP and CHF, all distributions appear to be right skewed, while all daily return distributions have excess kurtosis compared to the normal distribution.

[8]Note that returns are calculated as changes in the natural logarithm of the daily exchange rates.

ACCEPTED MANUSCRIPT

< Table 3 around here >

4.2 FDR portfolio

The FDR level determines the balance between wrongly selecting underperforming trading rules (false positives) and leaving out truly outperforming ones. Barras et al. (2010) explain that if a low $\hat{\text{FDR}}^+$ target is chosen, only a small proportion of "lucky rules" are allowed in the portfolio.[9] On the other hand, a high $\hat{\text{FDR}}^+$ target decreases the expected performance of the portfolio but improves the diversification of the portfolio as more rules are admitted. Bajgrowicz and Scaillet (2012) set $\hat{\text{FDR}}^+ = 10\%$ but they find that results are qualitatively stable for values ranging from 5% to 20%. In this study, $\hat{\text{FDR}}^+$ is set equal to 20%. This implies that eighty percent of the rules included in the portfolio produce genuine performance which results in a pool of multiple outperforming TTRs. The algorithm described in Section 3.3 is used to determine the parameters to get an $\hat{\text{FDR}}^+$ as close as possible to the $\hat{\text{FDR}}^+$ target level.[10] The relative π_A^+ and π_A^-, the proportions of rules with positive and negative performance, respectively, are then estimated.

To evaluate the long-term in-sample performance of the TTRs, the set of 7,650 trading rules is applied over the whole sample. Following Neely and Weller (2011), we assume that the investor either buys or sells in the market and there are no neutral signals. It is worth noting that since neutral signals are not considered, if pooling the signals results in equal number of buy and sell signals, the signal obtained on time $t-1$ is used for time t. Compute excess returns for the trading rules and assume that if a trading rule generates a long (short) signal, the investor converts the borrowed dollar (foreign currency) to foreign currency (dollar) at the closing rate and earns the foreign (U.S.) overnight rate. The excess return is computed as:

$$r_{k,t+1} = [lnS_{t+1} - lnS_t + ln(1 + i_t^*) - ln(1 + i_i)] * signal_k,$$

where i_t and i_t^* denote domestic and foreign interest rates, respectively. The Sharpe ratio is a risk adjustment criterion to check whether technical trading returns compensate

[9]The "lucky rules" term refers to the rules with no genuine performance, in other words, rules with no real predictive power.

[10]Our findings remain consistent when $\hat{\text{FDR}}^+ = 10\%$.

investors for bearing total risk. It is used to measure performance of the TTRs and is calculated as:

$$SR_k = \frac{\bar{r}_k}{\sigma_k},$$

where $\bar{r}_k$ denotes the excess return of the rule k:

$$\bar{r}_k = \frac{1}{N}\sum_{t=L}^{T} r_{k,t+1},$$

and σ_k is the standard deviation of the excess return generated by k-th trading rule:

$$\sigma_k = \sqrt{\frac{1}{N-1}\sum_{t=L}^{T}(r_{k,t+1} - \bar{r}_k)^2}.$$

Figure 1 shows the proportion of outperforming rules (π_A^+), underperforming rules (π_A^-) and those without abnormal performance (π_0). These estimates are long-term in-sample results and obtained in the absence of transaction costs. With the exception of CAD, the FDR finds that an important proportion of the rules exhibit significant predictive power. The proportion of outperforming rules is above 70% for the SEK, CHF and JPY, but the corresponding number is much smaller for the GBP and NOK. To highlight the advantage of the FDR over other methodologies, the number of outperforming rules selected by the FDR is compared with those of the stepwise-SPA test. We find that the stepwise-SPA test does not detect any outperforming rule for any of the currencies. This is consistent with Hsu and Taylor's (2013) results where they find that predictive ability of TTRs for developed currencies disappears since the early 1990s. This could be interpreted as the FWER being too conservative since it controls for not making even one type I error. This would imply that a TTR is selected by the stepwise-SPA test only if its p-value is smaller than $(1-\alpha)^{7,650}$.

< Figure 1 around here >

To check for the robustness of our results, transaction costs are taken into account. Following Chang and Osler (1999), Qi and Wu (2006) and Hsu and Taylor (2013), we use fixed 2.5 basis points one-way transaction costs.[11] Figure 2 shows the results with transaction costs. Strikingly, there is only a negligible decline in the proportion of outperforming

[11]Neely et al. (1997) show that the excess returns earned by trading rules are very sensitive to the level of transaction costs and to the liquidity of the markets. They assume one basis point and two basis

rules which indicates that our previous findings are robust to transaction costs, with the exception of NOK where the outperformance disappears after taking into account these costs. Our results are consistent with the recent literature where transaction costs do not necessarily eliminate TTRs profitability (see Qi and Wu, 2006; Hsu and Taylor, 2013).

< Figure 2 around here >

4.3 Persistence analysis

Persistence analysis is used to evaluate the out-of-sample performance of the outperforming rules. The motivation is that not only is historical outperformance no indication that an investor could have chosen the future outperforming rules ex-ante, but also in practice investors change their strategies in an attempt to adapt to the changes in economic environment. The question is whether an investor could have predicted which trading rule would generate profitable returns. We test whether outperforming rules could have been selected by an investor who has just access to the information that would have been readily available to her. Therefore, the universe of 7,650 TTRs is used for the selected currencies on a monthly basis and the outperforming rules over the month are selected based on the FDR and in the presence of transaction costs. Each currency's outperforming rules are used over the following month and returns from different currencies are pooled with equal weights to measure the total out-of-sample performance of the outperforming rules. To the best of our knowledge, this is the first time this type of persistence analysis is performed on TTRs in the FX.

< Figure 3 around here >

Similar to the in-sample estimation, the $\widehat{\text{FDR}}^+$ target is set to 20%.[12] The results show that the target of 20% allows for a small proportion of false discoveries while resulting in a well-diversified portfolio of trading rules.[13] Figure 3 presents the selected FDR level $\widehat{\text{FDR}}^+$ of each exchange rate at the end of each month. The average $\widehat{\text{FDR}}^+$ is also reported

points as two fixed values of one-way transaction costs. Neely and Weller (2011), on the other hand, use Bloomberg data on one-month forward bid-ask spreads as the basis for estimating transaction costs. Comparing their data with those on actual traders screen, they conclude that actual spreads are roughly one third of the quoted spreads

[12]The main findings remain qualitatively similar for $\widehat{\text{FDR}}^+ = 10\%$.

[13]It is important to note that at each step optimal values for λ and γ are estimated. See Eq. 1 and Eq. 2.

which shows that the average does not always match its target. For example, in the case of the UK the average $\hat{\text{FDR}}^+$ is 56% instead of the targeted 20% which indicates that the proportion of outperforming rules in the population is too low to achieve the FDR target. Table 4 also shows that for the GBP, NOK and CAD more than 20% of the achieved FDR are higher than 70% which implies an increase in the proportion of TTRs included in the portfolio since our selection becomes less restrictive.

Figure 4 demonstrates the number of outperforming trading rules for the currencies at the end of each month. With the exception of GBP, we observe TTRs outperformance for the 1995-2000 period for all currencies where the number of outperforming rules reaches its maximum in the case of JPY. The performance of TTRs is similar across all currencies for the 2000-2005 period where a large number of outperforming rules in one month is often followed by a large number of underperforming rules in the following month, making the corresponding numbers volatile. The performance of technical analysis over the period of the recent financial crisis is different across the currencies. While there seems to be underperformance in the case of SEK and GBP, our findings suggest that the initial outperformance of trading rules for the JPY and NOK breaks down by mid-2008. During this period of high uncertainty, TTRs perform well for the CHF and CAD currency pairs. The outperformance of technical analysis disappears for the post-crisis period across all currencies with the exception of the SEK.

< Figure 4 around here >

Next we seek to find the best performing technical trading family for each currency. The proportion of outperforming rules is computed for ten groups of TTRs. We find that CHB, EWMA, LWMA, MA and MAOS are often selected as the best strategies throughout the sample. The findings show that while MAOS rules appear as the best rule most often for the GBP, SEK and CHF, the best rules are more likely to belong to MA family for the CAD and NOK. In the case of JPY, EWMA rules are chosen more than other groups. Figures showing performance by trading rule family are given in Appendix B.

Building on Figure 4 we now examine which rules have some level of persistence in their outperformance (see Figure 5). The blue bars represent the proportion of TTRs that continues to outperform for three consecutive months. The red line represents the number of outperforming trading rules at each point in time as in Figure 4. Our findings show that while there are periods in which all the best models remain outperforming over three months, the average proportion of the rules that remain in the portfolio lies between 25% to 36%, with the exception of GBP for which the corresponding number is 19%. This highlights the importance of rebalancing our portfolio on a monthly basis. Over the recent crisis the outperformance of the TTRs becomes more short-lived in the case of the GBP and SEK, while there is an improvement in the performance of trading strategies for the CAD, JPY and NOK. In the case of the latter, this shows that over this period of high uncertainty, TTRs are reliable and persistent indicators for these currencies.

< Figure 5 around here >

Having examined the composition of the outperforming monthly portfolios, portfolio performance is evaluated based on the Sharpe ratios. Figure 6 shows the end-of-period annualized Sharpe ratios for the six currencies. Monthly portfolio rebalancing based on the FDR results in positive Sharpe ratios on average with the exception of GBP. The results are in line with previous findings that the performance of TTRs is volatile throughout the sample and periods with a positive Sharpe ratio are followed by those with a negative Sharpe ratio. This explains why the average Sharpe ratio is small across all currencies. It is worth noting that during the recent crisis, TTRs generate a positive Sharpe ratio for the SEK, NOK and CAD, while during the post-crisis period, they generate profit only for the SEK and JPY.

Finally, we aggregate the generated returns from rebalancing the portfolio of TTRs for each currency into one large portfolio. As Figure 7 shows, the annualized Sharpe ratios remain volatile. The findings show that the large portfolio generates a positive Sharpe ratio of 0.0375 during the crisis (2007:01-2009:12) and a positive Sharpe ratio of 0.027 over the whole sample (1996:03-2014:12), however, there is no persistence in profitability. In other words, while there are opportunities for an investor to make profit, they are short-lived. These results are in line with the adaptive market hypothesis by Lo (2004) which modifies the efficient market hypothesis. The adaptive market hypothesis proposes that the forces that drive prices to their efficient levels are weak and do not operate instantaneously. According to Neely and Weller (2011), the adaptive market hypothesis can be regarded as the most plausible explanation of TTR profitability.

< Figures 6 & 7 around here >

5 Conclusion

This paper investigates the profitability of TTRs in the FX market over the 1995-2014 period. A universe of 7,650 trading rules comprising both well-known and more recent rules are applied to six currencies: SEK, CHF, GBP, NOK, JPY and CAD. The FDR methodology is employed to control for data snooping bias. The motivation is that the previously used methodologies to control for data snooping either can only detect one outperforming rule (the RC and SPA tests) or they are too conservatives (the stepwise-RC and stepwise-SPA tests). By contrast, while the FDR allows for a small (and specific) proportion of false findings, it detects almost all possible outperforming rules. Our long-term in-sample results find predictive ability for an important proportion (up to 75%) of the rules. Taking into account transaction costs, we find our results to be robust to the choice of one-way transaction cost of 2.5 basis points with the exception of the NOK.

Persistence analysis is used to address the question whether investors could have predicted outperforming rules ex-ante in the presence of transaction costs where a portfolio of outperforming rules is constructed at the end of each month and out-of-sample performance is evaluated over the following month. The results show that at each step, FDR detects a large number of outperforming rules throughout the sample. This highlights the advantage of FDR since investors tend to combine the signals of multiple strategies rather than considering the signal of one TTR. Our findings suggest that among ten trading rule families, CHB, EWMA, LWAMA, MA and MAOS are the most selected strategies by FDR. However, a study of the portfolio turnover shows that, on average, the proportion of the rules that remain in the portfolio after three rebalancings lies between 0.19 and 0.36. Comparing these findings with the initial in-sample results shows that an investor should update her portfolio frequently to adapt to changes in the economy rather than sticking to a specific set of TTRs.

The performance measure used in this paper is the Sharpe ratio which measures the average excess return per unit total risk. The annualized Sharpe ratios at the end of each month show that the performance of the technical strategies is fairly volatile where a short period with a positive Sharpe ratio is often followed by a period with a negative Sharpe ratio. The persistence analysis shows that if an investor constructed a portfolio of the outperforming rules and updated it on a monthly basis, she would obtain a positive Sharpe ratio for all currencies, with the exception of GBP. Although our findings indicate that there are profitable opportunities for TTRs, their performance fluctuates throughout the sample which implies that their profitability is not persistent over time. These findings support Lo's (2004) Adaptive Market Hypothesis more than the efficient markets hypothesis.

References

Bajgrowicz, P., Scaillet, O., 2012. Technical trading revisited: False discoveries, persistence tests, and transaction costs. Journal of Financial Economics 106, 473–491.

Barras, L., Scaillet, O., Wermers, R., 2010. False discoveries in mutual fund performance: Measuring luck in estimated alphas. The Journal of Finance 65, 179–216.

Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society. Series B (Methodological) 57, 289–300.

Brock, W., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. The Journal of Finance 47, 1731–1764.

Chang, P. K., Osler, C. L., 1999. Methodical madness: Technical analysis and the irrationality of exchange-rate forecasts. The Economic Journal 109, 636–661.

Cheung, Y.-W., Chinn, M. D., 2001. Currency traders and exchange rate dynamics: A survey of the us market. Journal of International Money and Finance 20, 439–471.

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

Coakley, J., Marzano, M., Nankervis, J., 2016. How profitable are fx technical trading rules? International Review of Financial Analysis 45, 273–282.

Cornell, W. B., Dietrich, J. K., 1978. The efficiency of the market for foreign exchange under floating exchange rates. The Review of Economics and Statistics 60, 111–120.

Dooley, M. P., Shafer, J., 1984. Analysis of short-run exchange rate behavior: March 1973 to november 1981. In: Bigman, D., Taya, T. (Eds.), Floating Exchange Rates and the State of World Trade and Payments. BeardBooks, Washington.

Gehrig, T., Menkhoff, L., 2004. The use of flow analysis in foreign exchange: Exploratory evidence. Journal of International Money and Finance 23, 573–594.

Goodhart, C., 1988. The foreign exchange market: a random walk with a dragging anchor. Economica 55, 437–460.

Goodman, S. H., 1979. Foreign exchange rate forecasting techniques: implications for business and policy. The Journal of Finance 34, 415–427.

Hansen, P. R., 2003. Asymptotic tests of composite hypotheses, Working Paper, Brown University Economics.

Hansen, P. R., 2005. A test for superior predictive ability. Journal of Business and Economic Statistics 23, 365–380.

Hsu, P.-H., Hsu, Y.-C., Kuan, C.-M., 2010. Testing the predictive ability of technical analysis using a new stepwise test without data snooping bias. Journal of Empirical Finance 17, 471–484.

Hsu, P.-H., Kuan, C.-M., 2005. Reexamining the profitability of technical analysis with data snooping checks. Journal of Financial Econometrics 3, 606–628.

Hsu, P. H., Taylor, M. P., 2013. Forty years, thirty currencies and 21,000 trading rules: A large-scale, data-snooping robust analysis of technical trading in the foreign exchange market, Working Paper, Thirty Currencies and 21.

Lo, A. W., 2004. The adaptive markets hypothesis. The Journal of Portfolio Management 30, 15–29.

Lyons, R. K., 2001. New perspective on fx markets: Order-flow analysis. International Finance 4, 303–320.

Marshall, B. R., Cahan, R. H., Cahan, J. M., 2008. Does intraday technical analysis in the us equity market have value? Journal of Empirical Finance 15, 199–210.

Menkhoff, L., Taylor, M. P., 2007. The obstinate passion of foreign exchange professionals: Technical analysis. Journal of Economic Literature 45, 936–972.

Murphy, J. A., 1986. Futures fund performance: A test of the effectiveness of technical analysis. Journal of Futures Markets 6, 175–185.

Neely, C., Weller, P., 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.

Neely, C. J., Rapach, D. E., Tun, J., Zhou, G., 2014. Forecasting the equity risk premium: The role of technical indicators. Management Science 60, 1772–1791.

Neely, C. J., Weller, P. A., 2011. Technical analysis in the foreign exchange market, federal Reserve Bank of St. Louis Working Paper.

Pring, M. J., 1991. Technical analysis explained. New-York: McGraw-Hill.

Qi, M., Wu, Y., 2006. Technical trading-rule profitability, data snooping, and reality check: evidence from the foreign exchange market. Journal of Money, Credit and Banking 38, 2135–2158.

Romano, J. P., Wolf, M., 2005. Stepwise multiple testing as formalized data snooping. Econometrica 73, 1237–1282.

Sager, M. J., Taylor, M. P., 2006. Under the microscope: the structure of the foreign exchange market. International Journal of Finance & Economics 11, 81–95.

Shynkevich, A., 2012a. Global industry portfolios and short-term predictability of returns: Is it there? Journal of Financial Markets 15, 438–466.

Shynkevich, A., 2012b. Performance of technical analysis in growth and small cap segments of the us equity market. Journal of Banking & Finance 36, 193–208.

Shynkevich, A., 2012c. Short-term predictability of equity returns along two style dimensions. Journal of Empirical Finance 19, 675–685.

Storey, J. D., 2003. The positive false discovery rate: a bayesian interpretation and the q-value. Annals of Statistics 31, 2013–2035.

Storey, J. D., Taylor, J. E., Siegmund, D., 2004. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 66, 187–205.

Storey, J. D., Tibshirani, R., 2003. Statistical significance for genomewide studies. Proceedings of the National Academy of Sciences 100, 9440–9445.

Sullivan, R., Timmermann, A., White, H., 1999. Data-snooping, technical trading rule performance, and the bootstrap. The Journal of Finance 54, 1647–1691.

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

Taylor, M. P., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11, 304–314.

Taylor, S. J., 1994. Trading futures using a channel rule: A study of the predictive power of technical analysis with currency examples. Journal of Futures Markets 14, 215–235.

White, H., 2000. A reality check for data snooping. Econometrica 68, 1097–1126.

Wilke, A., Barrett, H. C., 2009. The hot hand phenomenon as a cognitive adaptation to clumped resources. Evolution and Human Behavior 30, 161–169.

Table 1. Technical trading rule parameters

This table summarizes technical trading rule parameters used in this paper. Our set of trading rules consists of filter rules (FR), moving average (MA), linearly weighted moving average (LWMA), exponentially weighted moving average (EWMA), moving average convergence-divergence (MACD), moving average oscillator (MAOS), stochastic oscillator (STOC), channel breakout (CHB), relative strength indicator (RSI), and bollinger bands (BOLL).

	Parameters	Description	Value
FR	x	Band for buy signal	0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
	y	Band for sell signal	0.0005, 0.001, 0.005, 0.01, 0.05
	c	Number of days a position is held during which all other signals are ignored	5, 10, 25, 50
	d	Number of days for the time delay filter	2, 3, 4, 5
MA LWMA EWMA	m	Short run moving average	2, 5, 10, 15, 20, 25, 30, 40, 50, 75 100, 125, 150, 200
	n	Long run moving average	5, 10, 15, 20, 25, 30, 40, 50, 75, 100 125, 150, 200, 250
	b	Fixed band multiplication value	0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5
MACD	m	Short run moving average	12, 15, 20
	n	Long run moving average	26, 30, 35
	l	Length of histogram	7, 9, 12, 15
MAOS	m	Short run moving average	2, 5, 10, 15, 20, 25, 30, 40, 50, 75 100, 125, 150, 200
	n	Long run moving average	5, 10, 15, 20, 25, 30, 40, 50, 75, 100 125, 150, 200, 250
	b	Fixed band multiplication value	0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5
STOS	m	Number of days used for minimum calculation	3, 5, 10, 15
	n	Number of days used for maximum calculation	10, 15, 20, 25, 30
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5
CHB	e	Evaluation period	5, 10, 15, 20, 25, 50, 100, 150, 200, 250
	$b1$	Band for buy signals	0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
	$b2$	Band for buy signals	10% to 90% of $b1$
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5
RSI	e	Evaluation period	14, 25, 30, 50
	$Lowerband$	Long run moving average	10, 20, 30, 50
	$Upperband$	Fixed band multiplicative value	60, 70, 80
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5
BOLL	e	Evaluation period	2,5,10,15,20,25,30,40,50,75,100,125 150,200,250
	$nstd$	Number of standard deviations	2, 3, 4
	b	Fixed band multiplicative value	0.0005, 0.001, 0.005, 0.01, 0.05, 0.1
	c	As previous	5, 10, 25, 50
	d	As previous	2, 3, 4, 5

Table 2. Decision making in hypothesis testing

This tables shows possible outcomes from n hypothesis tests. V is the number of type I errors, T is the number of type II errors and R is the number of trading rules selected as significantly outperforming rules. Therefore, $FWER = P(V \geq 1)$ is the probability of at least one type I error, while $\text{FDR}^+ = E(V/R|R > 0)$ is the rate that discoveries are false.

	H_0 accepted	H_0 rejected	Total
H_0 true	U	V	n_0
H_1 false	T	S	$n - n_0$
	$n - R$	R	n

Table 3. Descriptive statistics

This table summarizes statistics for daily changes in the logarithm of exchange rates for the 1994:03-2014:12 period (5,217 observations). Exchange rates are defined as the U.S. dollar price of one unit foreign currency. The Sharpe ratios are annualized.

	GBP	CAD	JPY	NOK	SEK	CHF
Mean*100	0.0022	0.0041	0.0013	0.0041	0.0050	0.0098
Min	-0.0392	-0.0434	-0.0371	-0.0502	-0.0354	-0.0847
Max	0.0447	0.0505	0.0658	0.0646	0.0555	0.0545
Standard deviation	0.0055	0.0053	0.0069	0.0072	0.0072	0.0068
Sharpe ratio	0.0039	0.0078	0.0019	0.0057	0.0070	0.0143
Skewness	-0.0415	0.0966	0.4695	0.0223	0.1685	-0.1371
Kurtosis	7.3438	10.1228	8.0239	8.3122	6.5587	10.2294

Table 4. Achieved false discovery rate

This table classifies the achieved $\hat{\mathrm{FDR}}^{+}$ at the end of each month. Periods with high achieved $\hat{\mathrm{FDR}}^{+}$ show that the proportion of outperforming rules in the population is too low to achieve a 20% $\hat{\mathrm{FDR}}^{+}$ target.

Currency	20%	20-50%	50-70%	>70%
SEK	0.50	0.34	0.04	0.12
CHF	0.34	0.55	0.02	0.09
GBP	0.27	0.27	0.05	0.41
NOK	0.34	0.28	0.05	0.33
JPY	0.41	0.38	0.04	0.17
CAD	0.25	0.44	0.02	0.28

Figure 1. In-sample performance without transaction costs

This figure displays the proportion of in-sample outperforming (π_A^+), neutral (π_0) and underperforming (π_A^-) rules for the 1994:03-2014:12 period. Theses results are obtained in the absence of transaction costs. The optimal λ and γ are computed as explained in Section 3.3.

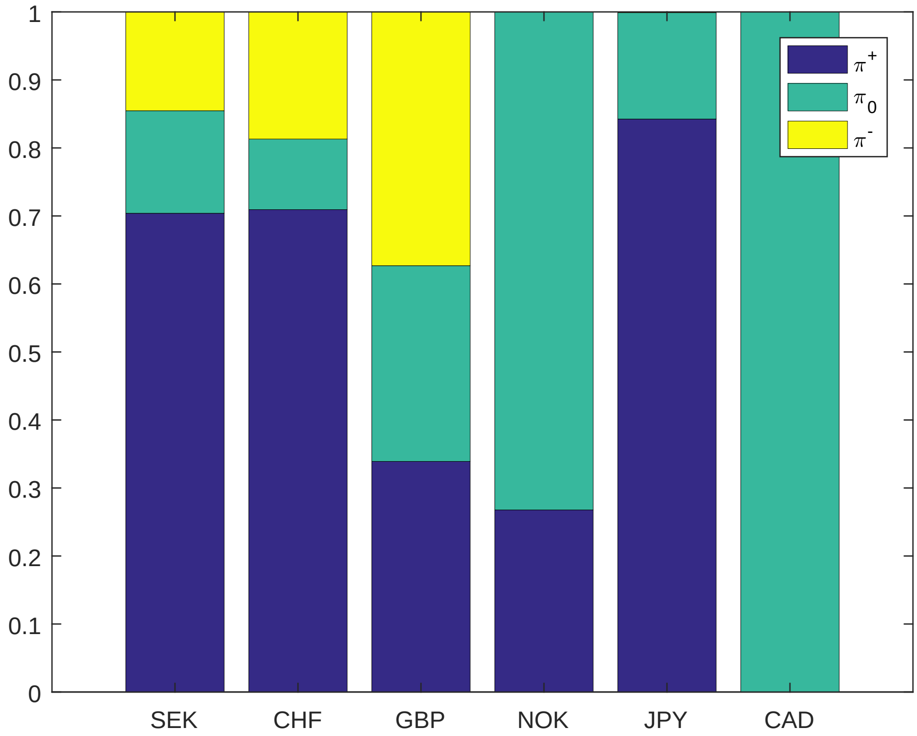

Figure 2. In-sample performance with transaction costs

This figure displays the proportion of in-sample outperforming (π_A^+), neutral (π_0) and underperforming (π_A^-) rules for the 1994:03-2014:12 period. Theses results are obtained in the presence of one-way 2.5 basis point transaction cost. The optimal λ and γ are computed as explained in Section 3.3.

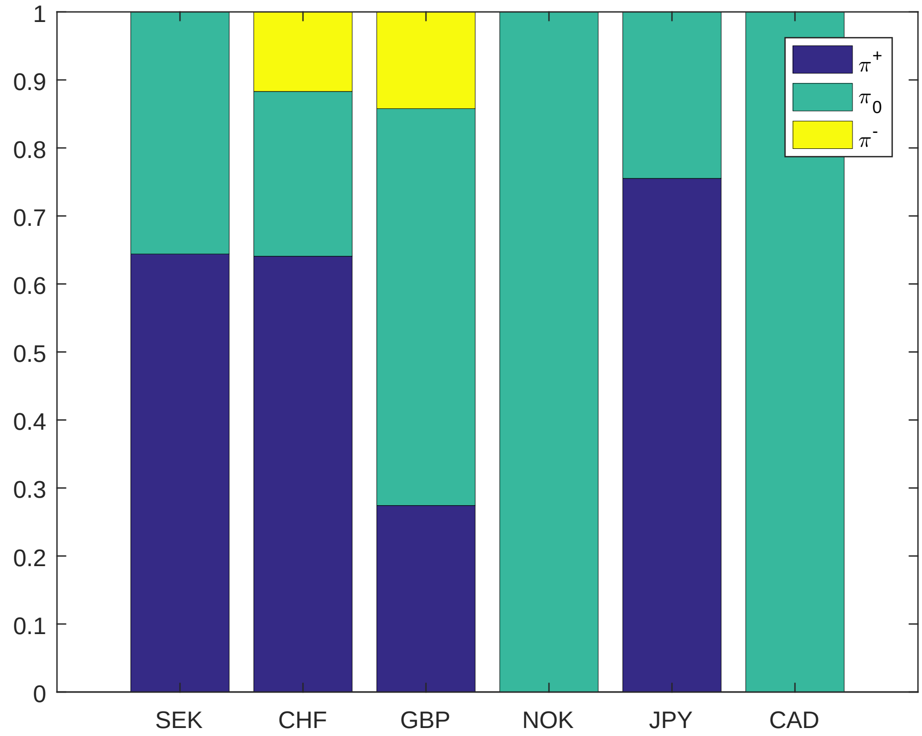

Figure 3. Monthly performance with transaction costs

This figure displays the achieved $\widehat{FDR}^+$ level at the end of each month throughout the sample. The dotted line The average selected $\widehat{FDR}^+$ is also reported for each currency. The target $\widehat{FDR}^+$ is set to 20% (dotted line) which implies that 80% of the rules included in the portfolio generate genuine performance. The average of achieved FDRs is reported in parenthesis. If the selected $\widehat{FDR}^+$ is higher than 20%, it shows that the proportion of outperforming rules in the population is too low to achieve the target. Theses results are obtained in the presence of one-way 2.5 basis points transaction cost.

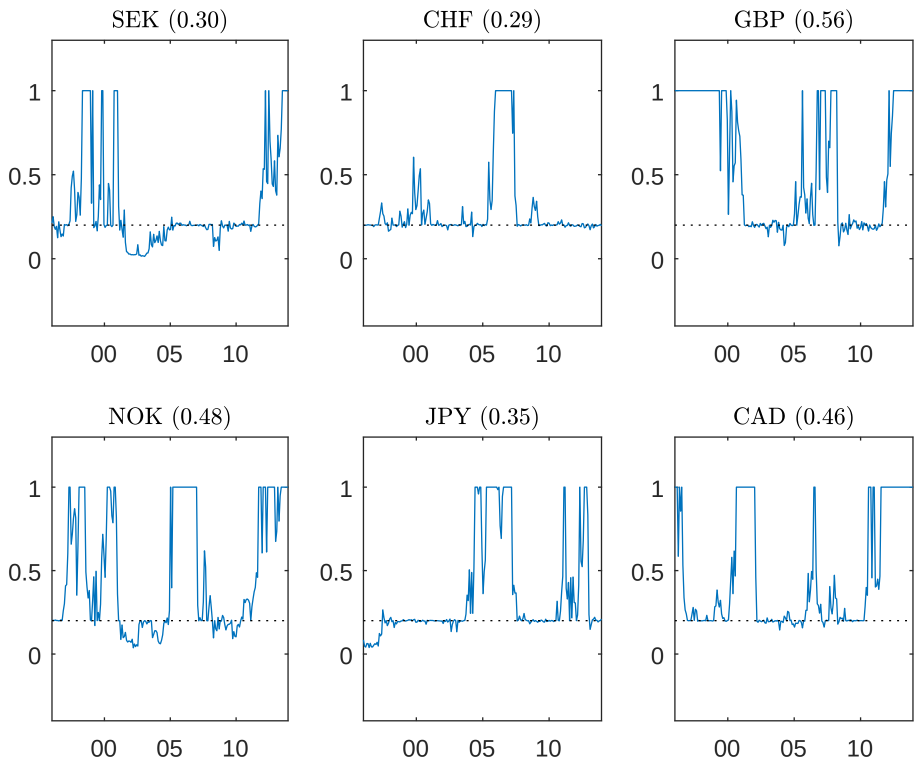

Figure 4. Number of outperforming rules

This figure displays the number of outperforming rules selected by the FDR at the end of each month. The universe of 7,650 trading rules are evaluated on a monthly basis based on the Sharpe ratio criterion. It is important to note that these results are obtained while transaction costs are taken into account and data snooping bias is controlled for.

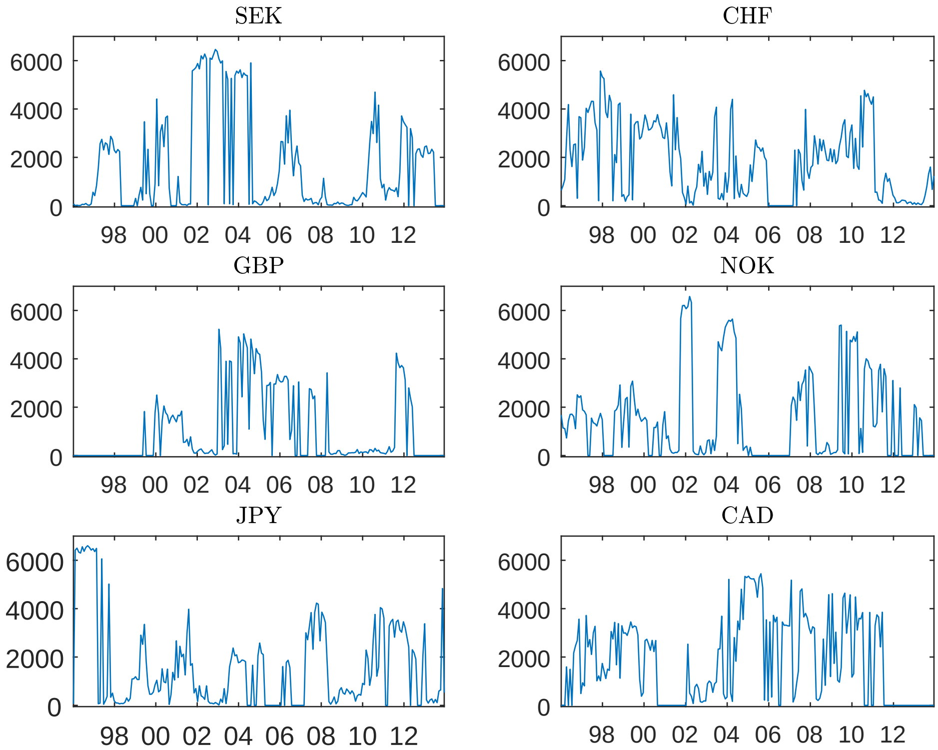

Figure 5. Portfolio turnover

This figure displays results of the portfolio turnover. The blue bars represent the proportion of rules remaining in the portfolio after three rebalancings. The red line represents the number of outperforming rules selected by the FDR. The average proportion of the outperforming rules that remain in the portfolio is reported in parenthesis. These results are obtained when one-way transaction costs of 2.5 basis points are taken into account.

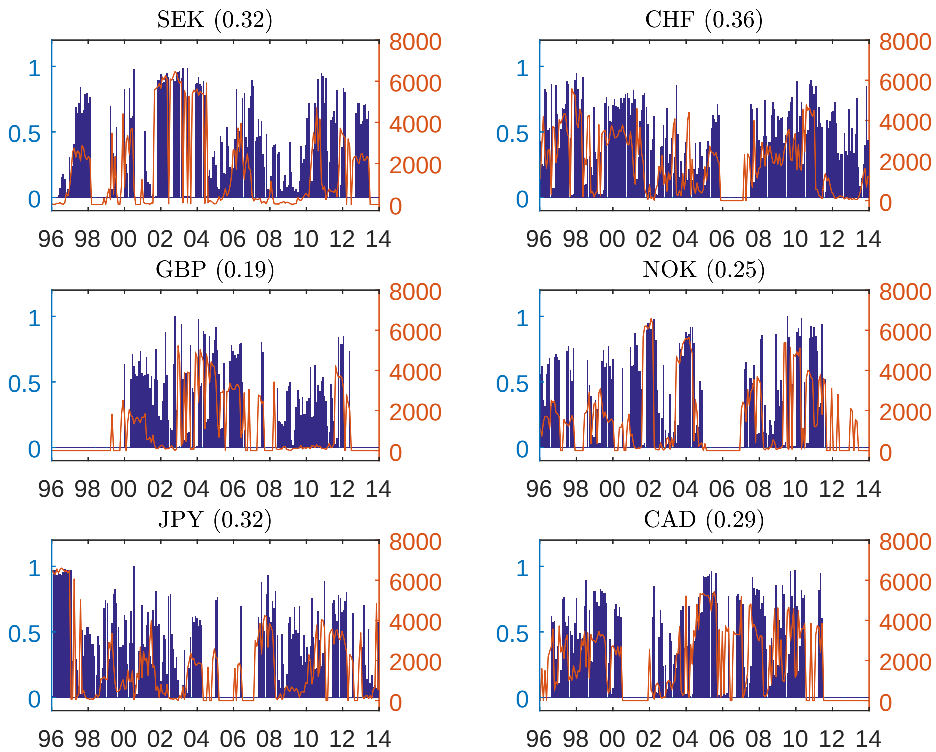

Figure 6. Sharpe ratio of each currency

This figure displays the Sharpe ratio for each currency. The best performing models are selected at the end of each month and their performance is evaluated in the following month based on the Sharpe ratio. The average Sharpe ratio throughout the sample is reported for each currency in parenthesis. All Sharpe ratios are annualized. One-way transaction costs of 2.5 basis points are taken into account.

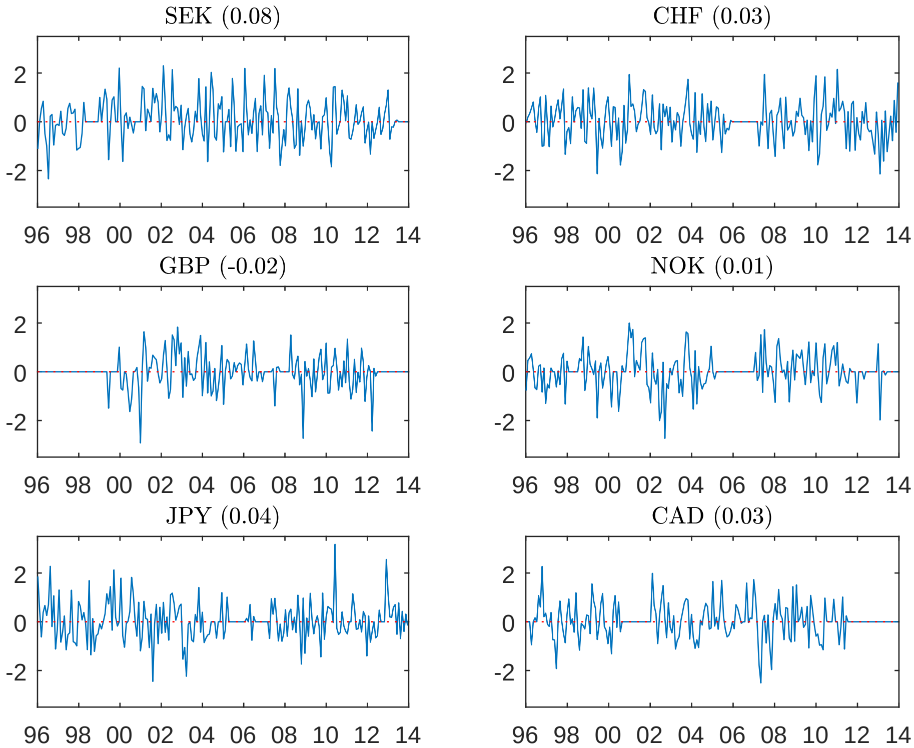

Figure 7. Sharpe ratio of the portfolio

This figure displays the Sharpe ratio for a portfolio consisting of six currencies. The universe of 7,650 trading rule is used every month and the outperforming rules selected by the FDR are used in the following month for each currency. Their out-of-sample performance is evaluated by the Sharpe ratio. The Sharpe ratios are annualized. It is important to note that one-way transaction costs of 2.5 basis points are taken into account.

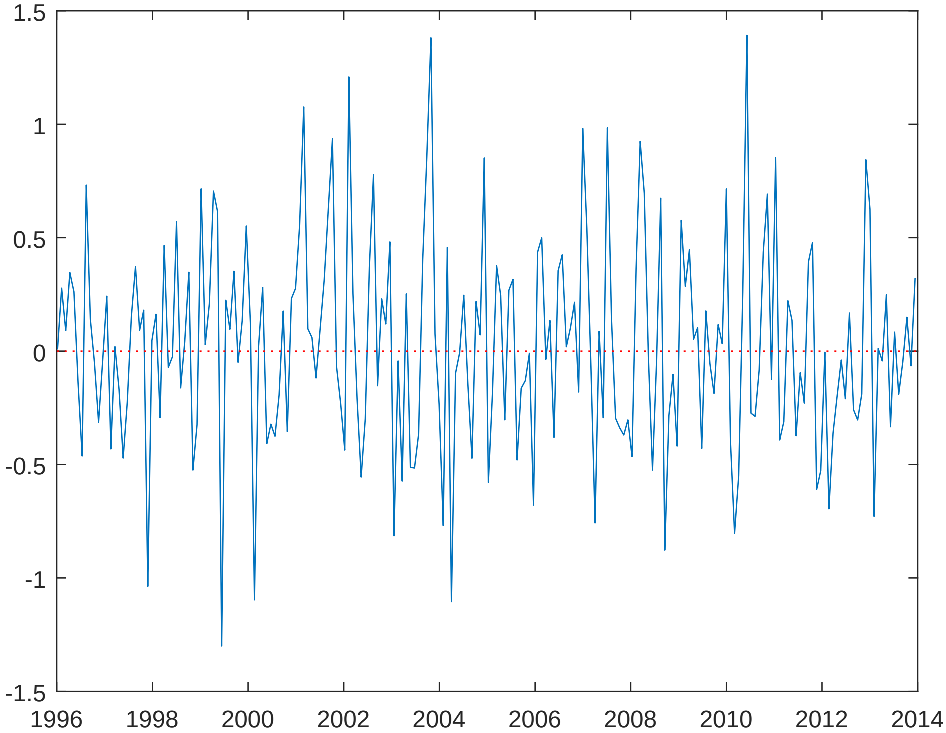

Appendix A Details of technical trading rules

Filter rules (FR)

A filter rule generates a buy signal if the exchange rate rises by $x\%$ or more from its most recent low. The investor borrows the dollar and uses the proceeds to buy the foreign currency. On the other hand, when the exchange rate falls by $y\%$ or more from a subsequent high, the investor shorts the foreign currency and uses the proceeds to purchase the dollar. We define the subsequent high and low as a local maximum and minimum respectively.

Moving average (MA)

A moving average rule utilizes a simple average of past prices. The short (long) MA is defined as a simple average of prices over the previous m (n) days, where $n > m$. When the short MA of a foreign currency is above (below) the long MA by an amount larger than the band with $b\%$, the investor borrows (short sells) the dollar (foreign currency) to purchase the foreign currency (dollar).

Linearly weighted moving average (LWMA)

This is a type of moving average where a higher weighting is assigned to recent price data than in the case of the simple MA. The LWMA is calculated by multiplying each one of the prices within the selected series, by a particular weight which is determined by dividing the position of time periods selected by the sum of the number of time periods. It is worth noting that weights in LWMA decrease in an arithmetic progression. The signal producing process is the same as the simple MA, implying that the rule generates buy (sell) signal when the short MA is above (below) the long MA.

Exponentially weighted moving average (EWMA)

Similar to the LWMA, the EWMA allocates a higher weighting to recent price data than does the simple MA. However, the weights are calculated differently. We define α as the smoothing factor,

$$\alpha = \frac{2}{n+1}$$

where n is the number of observations. To calculate the EWMA at each point in time, we multiply the last price by α and sum it with the product of the last day EWMA and $1-\alpha$. Here, weights decrease exponentially and the signal producing process is the same as the simple MA, implying that the rule generates buy (sell) signal when short MA is above (below) the long MA.

Moving average convergence-divergence (MACD)

The MACD measures the difference between two EWMAs. The representation of the MACD includes another EWMA which acts as a trigger indicator, known as the signal line. In other words, MACD (l, s, m) is an indicator where the MACD series is the difference of a long (l) and short (s) EWMAs, and the signal line is an EWMA of the MACD series with parameter m,

$$\begin{aligned} MACD_t &= (1-\lambda_l)\sum_{i=0}^{\infty}\lambda_l^i P_{t-i} - (1-\lambda_s)\sum_{i=0}^{\infty}\lambda_s^i P_{t-i} \\ Signal\,Line_t &= (1-\lambda_m)\sum_{i=0}^{\infty}\lambda_m^i \delta_t \end{aligned}$$

where $\lambda_l = 1 - \frac{2}{l+1}$, $\lambda_s = 1 - \frac{2}{s+1}$, and $\lambda_m = 1 - \frac{2}{m+1}$. Signals are determined as follows,

$$\begin{aligned} z_t &= +1 \quad \text{if} \quad MACD_t > Signal\,Line_t \\ z_t &= -1 \quad \text{if} \quad MACD_t < Signal\,Line_t \\ z_t &= z_{t-1} \quad\quad \text{otherwise.} \end{aligned}$$

And an investor

$$\begin{aligned} &\text{Buys} \quad \text{if} \quad z_t - z_{t-1} = +2 \\ &\text{Sells} \quad \text{if} \quad z_t - z_{t-1} = -2. \end{aligned}$$

Moving average oscillator (MAOS)

The moving average oscillator is an alternative of the MACD while they are both computed similarly. The MAOS is calculated as the difference between a short MA and a long MA,

$$MA_t = m^{-1}\sum_{i=t-m}^{t-1} P_{i+1} - n^{-1}\sum_{i=t-n}^{t-1} P_{i+1}.$$

Signals are determined as follows

$$
\begin{aligned}
z_t &= +1 \quad \text{if} \quad MA_t > 0 \\
z_t &= -1 \quad \text{if} \quad MA_t < 0 \\
z_t &= z_{t-1} \quad \text{otherwise,}
\end{aligned}
$$

and an investor

$$
\begin{aligned}
&\text{Buys} \quad \text{if} \quad z_t - z_{t-1} = +2 \\
&\text{Sells} \quad \text{if} \quad z_t - z_{t-1} = -2.
\end{aligned}
$$

Stochastic oscillator (STOS)

The stochastic oscillator is a momentum indicator which shows the location of the price relative to the high-low range over a specific period. The stochastic oscillator is calculated as

$$
\begin{aligned}
\%K_t &= 100 * \frac{P_t - min(P_{it-m}^{t-1})}{max(P_{it=t-m}^{t-1}) - min(P_{ii=t-m}^{t-1})} \\
\%D &= \frac{1}{n} \sum_{i=t-m}^{n} \%K_{t-i+1}.
\end{aligned}
$$

Buy and sell signals are determined as

$$
\begin{aligned}
&\text{Buy} \quad \text{if} \quad \%K_{t-1} < \%D_{t-1} \quad \text{and} \quad \%K_t > \%D_t \\
&\text{Sell} \quad \text{if} \quad \%K_{t-1} > \%D_{t-1} \quad \text{and} \quad \%K_t < \%D_t
\end{aligned}
$$

Channel breakout (CHB)

A channel occurs when the high price over a specific period is within $x\%$ of the low price over the same period. Therefore, an investor buys (sells) when the price goes above (below) the channel by $b\%$.

Relative strength indicator (RSI)

The RSI compares the magnitude of recent gains to recent losses to determine overbought and oversold conditions of a currency. The RSI is calculated as

$$
RSI = 100 - \frac{100}{1 + RS^*},
$$

where

$$
RS = \frac{average\, of\, x\, days'\, up\, prices}{average\, of\, x\, days'\, down\, prices}.
$$

The RSI ranges from 0 to 100%. If it is above 70%, the currency is considered as overbought and the RSI generates a sell signal. On the other hand, if the RSI is below 30%, the currency is considered as oversold and buy signal is determined.

Bollinger band (BOLL)

Bollinger band is a band which is two standard deviations away from a simple moving average. It is developed by John Bollinger in the 1980s. Since standard deviation is a measure of volatility, the more volatile the market is, the wider the band gets. An investor buys (sells) the FX when the price is above (below) the upper (lower) band.

Appendix B Figures

Figure B1. Outperforming TTRs for SEK

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of Swedish Krona (SEK) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).

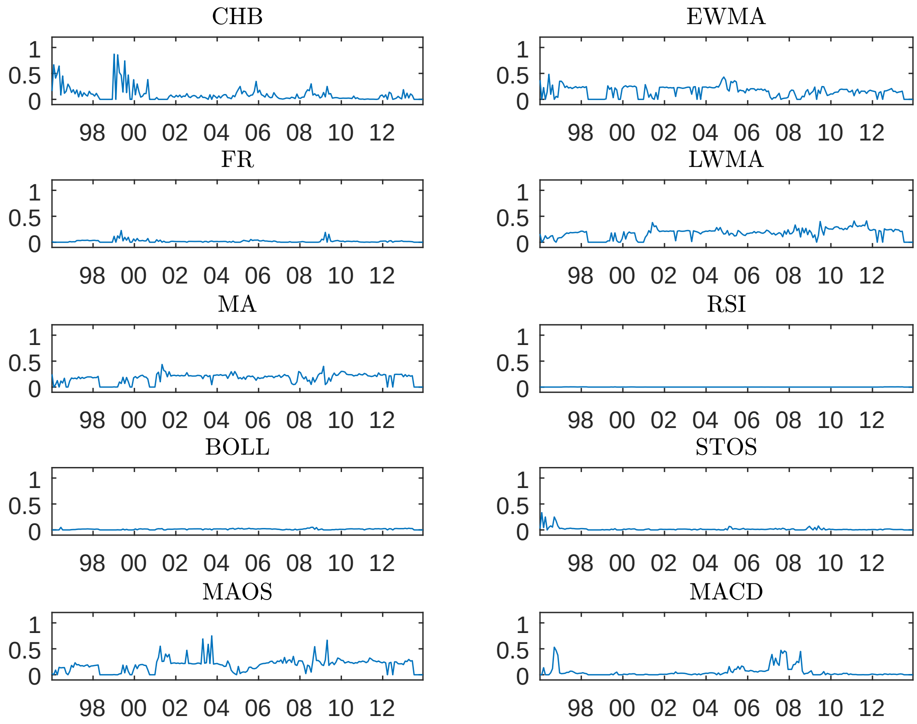

Figure B2. Outperforming TTRs for CHF

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of Swiss Franc (CHF) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).

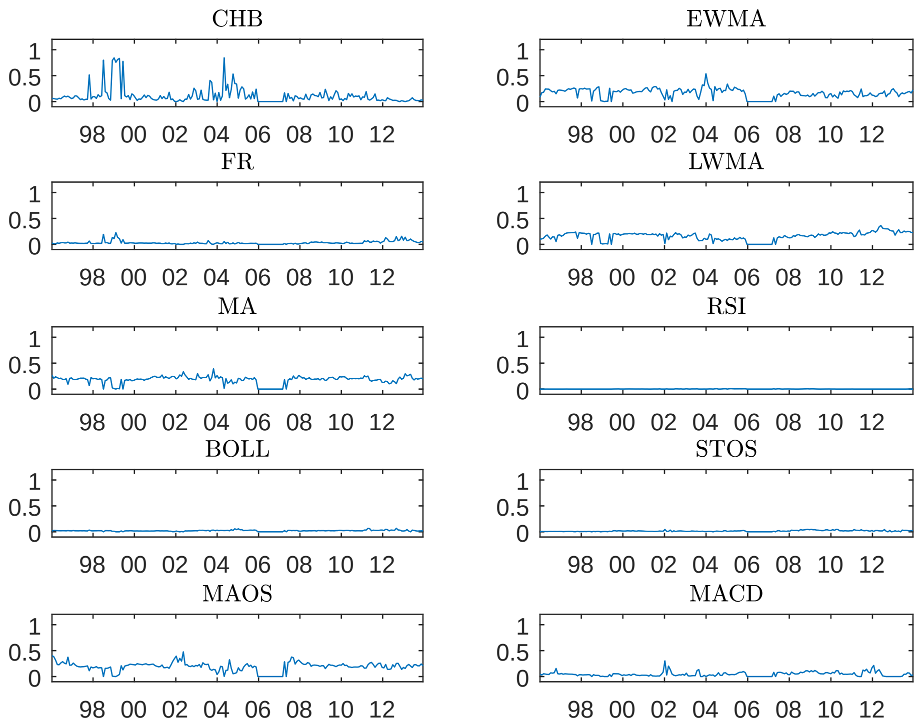

Figure B3. Outperforming TTRs for GBP

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of British Pound (GBP) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).

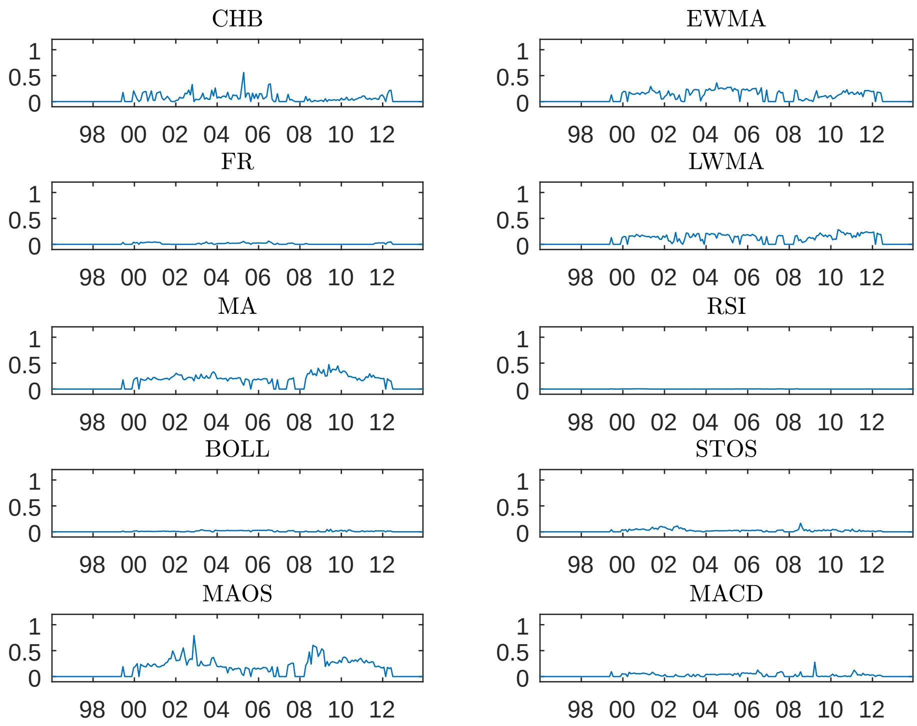

Figure B4. Outperforming TTRs for NOK

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of Norwegian Krone (NOK) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).

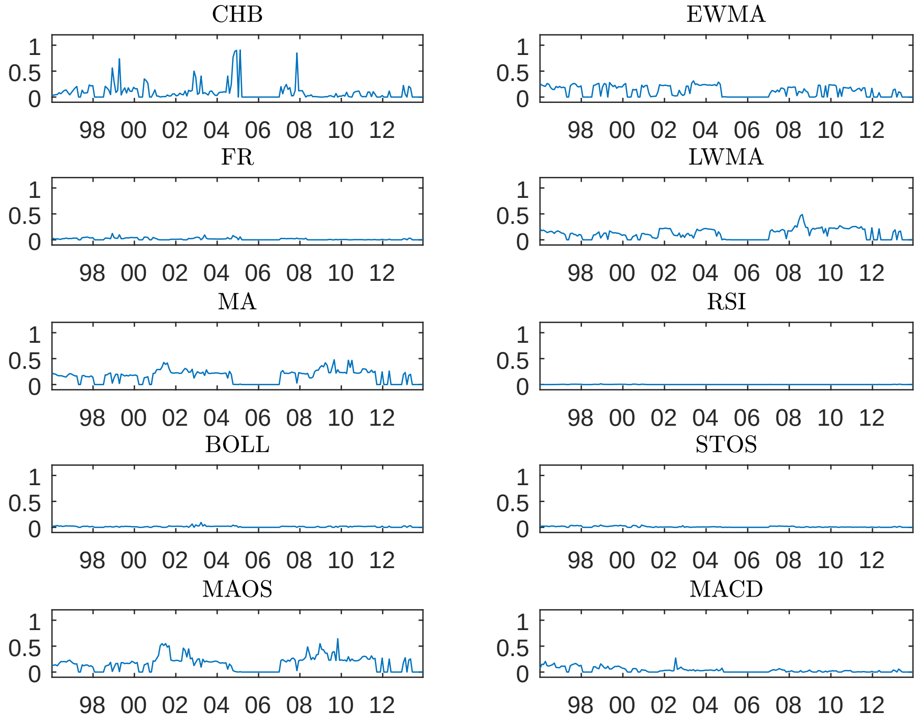

Figure B5. Outperforming TTRs for JPY

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of Japanese Yen (JPY) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).

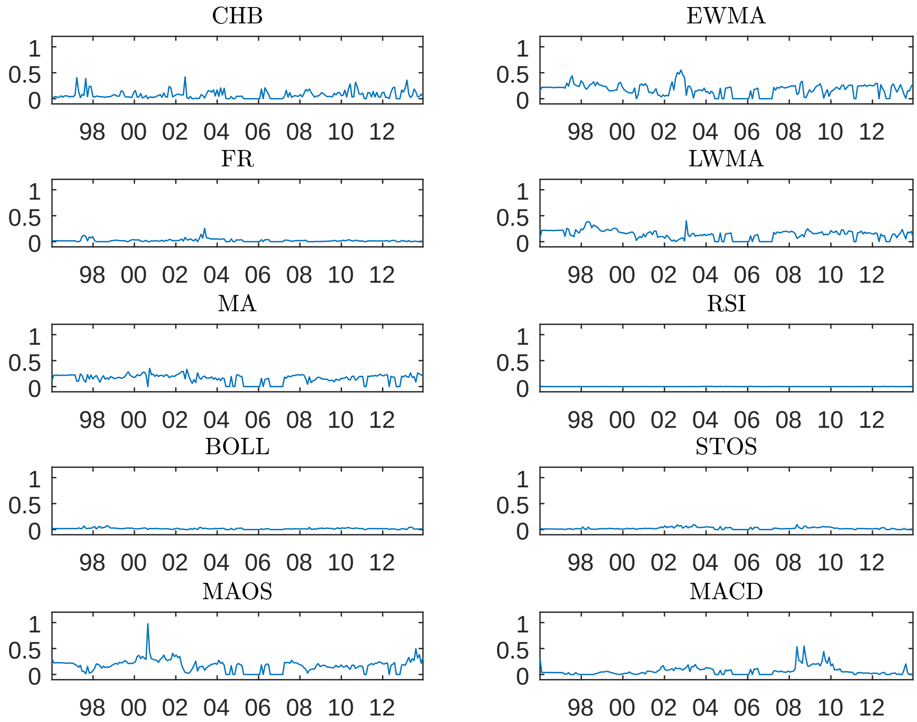

Figure B6. Outperforming TTRs for CAD

This figure displays the proportion of outperforming rules selected by the FDR throughout the sample in the case of Canadian Dollar (CAD) for ten trading strategy families: channel breakout (CHB), exponentially weighted moving average (EWMA), filter rule (FR), linearly weighted moving average (LWMA), moving average (MA), relative strength index (RSI), bollinger band (BOLL), stochastic oscillator (STOS), moving average oscillator (MAOS), moving average convergence-divergence (MACD).