

Technical Trading: Is it Still Beating the Foreign Exchange Market?

Po-Hsuan Hsu

School of Economics and Finance, University of Hong Kong

Mark P. Taylor

Warwick Business School, University of Warwick

Zigan Wang

School of Economics and Finance, University of Hong Kong

Abstract: We carry out a large-scale investigation of technical trading rules in the foreign exchange market, using daily data over 45 years for 30 developed and emerging market currencies. Employing a stepwise test to counter data-snooping bias and examining over 21,000 technical rules, we find evidence of substantial predictability and excess profitability in both developed and emerging currencies, measured against a variety of performance metrics. We cross-validate our results using out-of-sample analysis. We find time-series and cross-sectional variation in sub-periods and cultural and/or geographic groups, respectively, suggesting that temporarily not-fully-rational behavior and market immaturity generate technical predictability and potential excess profitability.

JEL Classifications: F31, C53, G15

Keywords: Foreign exchange; technical analysis; trading rules; data-snooping bias.

Acknowledgements: We thank the London branch of BlackRock for providing us with data on daily midday quotations of foreign exchange rates and short-term interest rates. We are grateful to Peter Reinhard Hansen, Campbell Harvey, Yan Liu, Christopher Neely, Thomas Sargent, Giorgio Valente, Paul Weller and conference participants at the 2014 Asian Meeting of the Econometric Society (Taipei), as well as our editor (Alan Taylor) and two anonymous reviewers, for insightful and constructive comments on an earlier version of this paper. Any errors which remain are the sole responsibility of the authors. Corresponding author: Mark P. Taylor, Warwick Business School, University of Warwick, Coventry, CV4 7AL, United Kingdom. Tel: +44 2476 524 523. E-mail: mark.taylor@wbs.ac.uk.

1. Introduction

Technical analysis (sometimes alternatively referred to as chartist analysis) is a set of techniques for deriving trading recommendations for financial assets by analyzing the time-series history of the particular asset price either graphically or mathematically. Although technical analysis is not rooted in underlying economic or financial theory (the ‘fundamentals’), the widespread use of technical analysis among financial practitioners in financial markets in general and in the foreign exchange market in particular is well documented (e.g., Frankel and Froot, 1990; Allen and Taylor, 1990; Taylor and Allen, 1992; Cheung and Chinn, 2001). Following pioneering early work by Cornell and Dietrich (1978) and Sweeney (1986), which appeared to show that technical trading could ‘beat the foreign exchange market’, the predictive ability and profitability of technical analysis in the foreign exchange market have been the subject of extensive analysis, most recently as a branch of behavioural finance and economics (e.g., Azzopardi, 2010).¹ Indeed, a recent literature survey on the topic (Menkhoff and Taylor, 2007) concludes that the ‘obstinate passion’ of foreign exchange professionals for technical analysis is an intrinsic part of the behaviour of practitioners in this market. Given this, and in the wake of the global financial crisis, an understanding of the drivers of international financial markets, from the perspective of both economic fundamentals and behavioural considerations, is clearly of high interest.

Nevertheless, a comprehensive and up-to-date analysis of technical analysis in the foreign exchange market still seems to be lacking, since most previous studies of this issue tend to consider only a small number of currencies, short sample periods, limited sets of technical trading rules, simple performance metrics, and basic testing methods which may be subject to data-snooping bias.² As a result, the intriguing question of whether technical analysis can beat the foreign exchange market calls for a large-scale investigation with an appropriate empirical design.

Moreover, even if the predictability and excess profitability of technical analysis exist with statistical significance for certain currencies at certain times, as some studies appear to show (Menkhoff and Taylor, 2007), a further question arises, namely *why* should technical analysis work in the foreign exchange market? Menkhoff and Taylor (2007) categorize the various explanations proposed in the literature into four positions: (i) technical analysis indicates not-fully-rational behavior or investor psychology and market sentiment (e.g.

¹ An incomplete list of studies in this area includes Allen and Taylor (1990); Taylor and Allen (1992); Levich and Thomas (1993); Kho (1996); Neely, Weller, and Dittmar (1997); LeBaron (1999); Gencay (1999); Chang and Osler (1999); Neely (2002); Okunev and White (2003); Qi and Wu (2006); Neely, Weller, and Ulrich (2009). See Menkhoff and Taylor (2007) and Neely and Weller (2012) for literature surveys. The academic literature on technical analysis in the equity market and financial markets in general is also very large; see, for example, Lo, Mamaysky and Wang (2000) and the references cited therein.

² Data-snooping bias arises whenever researchers continue searching for predictive models or rules but conduct only individual tests for each trial using the same data set without considering the fact that all models or rules should be tested together for their significance. An early criticism of such an approach is given by Jensen and Benington’s (1970) comment on Levy (1967): ‘Likewise, given enough computer time, we are sure that we can find a mechanical trading rule which ‘works’ on a table of random numbers... (p.470).’ Although there exist in the literature different labels for such bias (see Section 5, below for further discussion), we use the name ‘data snooping’, following the usage of Lo and MacKinlay (1990), Sullivan, Timmermann, and White (1999), White (2000), and Schwert (2003).

Frankel and Froot, 1990; Taylor and Allen, 1992; Oberlechner and Osler, 2012); (ii) technical analysis exploits or reinforces movements in the market caused by official intervention (e.g. Sweeney, 1986; LeBaron, 1999); (iii) technical analysis serves as a tool for processing information about fundamental influences on exchange rates (e.g. Treynor and Ferguson, 1985; Brown and Jennings, 1989; Blume, Easley, and O'Hara, 1994; Osler, 2003; Kavajecz and Odders-White, 2004; Zhu and Zhou, 2009); and, lastly, (iv) the excess profitability of technical analysis may be simply attributed to risk premia (e.g. Cornell and Dietrich, 1978; Kho, 1996). Previous empirical studies have not to date reached a conclusive verdict on these issues and a more complete examination is also, therefore, warranted from this perspective.

In this paper, we perform the most comprehensive study of technical trading rules in the foreign exchange market to date in order to assess the predictability of such rules and to provide further insights on what it is that may make them at times profitable. In addition, we also check a set of 'stylized facts' that may be gleaned from the literature on technical analysis in the foreign exchange market (Menkhoff and Taylor, 2007), such as that it may have diminished in profitability over time, that it may be more profitable with more volatile currencies, and that transaction costs do not necessarily eliminate its excess profitability.

Our study analyses daily data over a maximum of 45 years (1971-2015) for 30 U.S. dollar exchange rates, covering both emerging and developed markets, which we use to examine the predictive ability and excess profitability of over 21,000 technical trading rules. In constructing our tests, we examine the investment performance of foreign currency traders who use the U.S. dollar as home currency, and calculate their profits based on spot foreign exchange rate and the differentials of interest rates in the U.S. and foreign countries.^{3,4} We also consider a range of performance metrics which summarize the overall performance of trading rules as well as splitting this into a dynamic, 'market-timing' component and a static, 'tilt' or 'buy-and-hold' component. In order to eliminate data-snooping bias from our analysis, we employ a stepwise test developed by a series of methodological studies including White (2000), Romano and Wolf (2005), Hansen (2005), and Hsu, Hsu, and Kuan (2010). This testing method is powerful in identifying predictive or profitable technical trading rules from among a very large set of trading rules without data-snooping bias, and

³ Although we mainly report results based on U.S. dollar-based exchange rates and assume investors are U.S. based, our results are robust to using alternative home currencies and alternative exchange rate base currencies (namely euro, U.K. pound, Japanese yen, and Swiss franc); these results are discussed in Sections 7.3 and 7.4, and full details are given in Tables 2 to 17 in the Online Appendix.

⁴ Note that we could alternatively have used forward exchange rates in the trading strategy. However, for most trading rules in our strategy classes, the investor does not know the exact date to close the position when he or she opens it, and therefore it is difficult for them to choose a forward rate with specific maturity to trade ex-ante. Of course, the trader could trade the overnight forward rate and simply roll that position forward on days when the position is held by effecting an overnight swap (i.e. simultaneous spot and forward transaction). However, as long as covered interest rate parity holds (Taylor, 1987, 1989), this would in general be equivalent to holding the position in cash and paying the net interest differential, while the latter is perhaps slightly simpler. Where it could make a difference, however, is where there may be significant deviations from covered interest parity due to market imperfections such as capital controls, default risk, or liquidity interruptions, which may be the case for some currencies, especially emerging market currencies, during some periods.

thus allows us to make appropriate statistical inferences.⁵ Further, we provide an out-of-sample analysis of technical trading rules that is extremely stringent in the sense that it uses nearly four years of daily data that did not exist when the previous version of this paper was completed.

The rest of the paper is organized as follows. In the next section we discuss the coverage of our data set and provide some basic descriptive statistics, while in Section 3 we briefly discuss the various families of technical trading rules under consideration.⁶ In Section 4 we describe the various metrics we consider in order to assess the performance of technical trading rules. A main contribution of the paper is the analysis of a large number of technical trading rules with a large number of exchange rates over long periods of time while controlling statistically for the fact that we deliberately search for the best-performing rules over our data set in order to avoid data-snooping bias. Because the literature on controlling for data-snooping is highly technical and has largely appeared in the econometric and finance literature, in Section 5 we provide an intuitive and largely non-technical overview of the key developments in this area as a key to understanding our empirical methodology. In Section 6 we report our main empirical results on the performance of technical trading rules. In Section 7 we report the results of a number of robustness checks, including changes in base and home currency, alternative sub-sample periods, break-even transaction costs, and out-of-sample performance analysis. Finally, in Section 8 we provide some concluding comments.

2. Data

We consider daily data on foreign exchange rates between the U.S. dollar and 30 foreign currencies, including nine developed market currencies (Australian dollar, Canadian dollar, German mark/euro, Japanese yen, New Zealand dollar, Norwegian krone, Swedish krona, Swiss franc, and U.K. pound) and 21 emerging market currencies (Argentine peso, Brazilian real, Chilean peso, Colombian peso, Czech koruna, Hungarian forint, Indian rupee, Indonesian rupiah, Israeli shekel, Korean won, Mexican peso, Philippine peso, Polish zloty, Romanian new leu, Russian ruble, Singaporean dollar, Slovak koruna, South African rand, Taiwanese dollar, Thai baht, and Turkish lira). The sample periods for developed market currencies start from January 4 1971 and end on September 11 2015, while the sample

⁵ The reality check test proposed by White (2000) is the first formal testing method that corrects data-snooping bias for large-scale joint test problems. His method was then improved by Hansen (2005), Romano and Wolf (2005), and Hsu, Hsu, and Kuan (2010) to identify predictive models in large-scale, multiple testing problems. These tests have been used to examine the technical trading rule predictiveness and profitability in stock market indexes (Sullivan, Timmermann, and White, 1999; Hsu and Kuan, 2005), foreign exchange rates (Qi and Wu, 2006), futures markets (Park and Irwin, 2010), and exchange traded funds (Hsu, Hsu, and Kuan, 2010). We recognize the existence of other testing methods in handling data-snooping bias, including the false discovery rate methodology (Barras, Scaillet, and Wermers, 2010) and the wild bootstrap reality check of Clark and McCracken (2012). The former is used by Bajgrowicz and Scaillet (2012) to test technical predictability in the Dow Jones Industrial Average index, while the latter is used by Neely, Rapach, Tu, and Zhou (2014) to examine whether technical indicators forecast equity risk premium. Harvey and Liu (2014) and Harvey, Liu, and Zhu (2016) review the recent development in handling multiple testing problems in evaluating trading strategies. We provide an intuitive overview of methods for correcting for data-snooping bias in Section 5.

⁶ A detailed description of all technical trading rules considered is given in Appendix A.

periods for emerging market countries start from various dates due to data availability.⁷ Israel has the earliest starting date among emerging market currencies (January 1978) and is followed by South Africa (January 1981), Singapore (January 1982) and Taiwan (October 1983); all emerging market data end in September 2015. Our data on exchange rates and short-term interest rates were kindly supplied by the London branch of the asset manager BlackRock and are based on midday quotations in the London market.

To measure the investment performance in currency trading of an investor based on U.S. dollars, we first calculate the daily gross return (without interest rates) from buying one unit of a foreign currency and holding it for one day as $r_t = \ln(s_t / s_{t-1})$, where s_t denotes the spot foreign exchange rate (U.S. dollars per unit of foreign currency) on day t . $s_t / s_{t-1} > 1$ indicates that the foreign currency appreciates against the U.S. dollar. We start by considering daily gross returns without adjustment for transaction costs.

Table 1 reports summary statistics of the daily returns on all foreign currencies and short-term interest rates. Among the nine developed currencies, the Swiss franc appreciates the most on average (1.3 basis points per day or 3.25% per year) and the New Zealand dollar depreciates the most (0.5 basis points per day or 1.25% per year). Among emerging market currencies, the Slovak koruna appreciates the most on average (0.54 basis points per day or 1.35% per year), while the Turkish lira depreciates the most (10.7 basis points or 26.75% per year).

Interest rates are, of course, a major concern for technical currency traders since they affect the overall return from currency trading, even if technical analysts will typically only analyze exchange rate data in determining an exchange rate trading rule. We convert the annualized short-term interest rate, i_t^a , into a daily interest rate i_t using the formula

$i_t = \ln(1 + i_t^a) / 360$. Table 1 shows that daily short-term interest rates available for daily trading in developed countries range from 0.79 basis points per day (or 2.8% per year) to 3.03 basis points (or 10.9% per year). It is also found that short-term interest rates vary greatly across emerging countries. The highest average short rate is as high as 13.1 basis points per day (or 47.2% per year) in Turkey since 1990, while the lowest average short rate is as low as 0.23 basis points per day (0.8% per year) in Chile since 1994.

We also find that emerging market currencies are in general more volatile than developed currencies. The most and least volatile currencies among developed countries, in terms of standard deviation of daily gross returns, are the Swiss franc (0.75%) and Canadian dollar (0.40%). Four emerging currencies are associated with 1% or higher standard deviations. The most volatile currencies are the Indonesian rupiah and Russian ruble (1.40%), and the most stable currency is the Taiwanese dollar (0.29%). The observation that emerging market currencies are more volatile than developed country currencies may be attributed to

⁷ Since we require both exchange rates and short-term interest rates to calculate currency investment returns, the sample periods for emerging currencies start from the date when both exchange rates and interest rates are available.

many reasons such as lower liquidity and greater variability in underlying fundamentals such as growth rates, terms of trade shocks and monetary policy.

We also considered the relationship between mean excess returns (gross daily returns in excess of short rate differentials, equivalent to holding period return) and the standard deviation of daily returns and found no significant correlation (see Figure 1 in the Online Appendix for a scatter plot).

Another important dimension of exchange rate fluctuations is the existence of trends that are reflected in persistent return series. In our sample, the first-order autocorrelation coefficients of daily returns from developed country exchange rates are low, ranging from -0.019 to 0.062. The returns from the emerging market exchange rates present higher diversity in persistence: five emerging currencies carry first-order autocorrelation coefficients in excess of 0.1.⁸ However, we could detect no significant correlation between mean excess returns and the persistence of returns (see Figure 2 in the Online Appendix for a scatter plot of mean excess returns against the first-order autocorrelation coefficients of returns).

3. Technical Trading Rules

Technical analysis can be performed in qualitative form, relying mainly on the analysis of charts of past price behaviour and loose inductive reasoning that attempts to identify particular patterns in the data,⁹ or it can be strictly quantitative, by constructing trading signals or forecasts through a quantitative analysis of time series data (Allen and Taylor, 1992).¹⁰ In this paper, we are concerned with analysing the excess profitability of quantitative technical trading rules as they are objective and readily computable.

We construct the following five classes of technical trading rules that are commonly used by traders (Taylor and Allen, 1992; Menkhoff and Taylor, 2007): *oscillator trading rules*, or ‘overbought-oversold’ indicators, which attempt to identify imminent market corrections after rapid exchange rate movements; *filter trading rules*, which attempt to follow trends by buying (selling) a currency whenever it has risen (fallen) by a given percentage; *moving average trading rules*, which attempt to ride trends and identify imminent breaks in trend by examining the behavior of the exchange rate relative to a moving average of a given length, or by analyzing the interaction of two or three moving averages of different

⁸ The highest autocorrelation coefficient occurs in the Russian ruble (0.231) and the lowest autocorrelation coefficient occurs in the Mexican peso (-0.135); however, both currencies are subject to substantial management.

⁹ Some attempts have been made by researchers to identify technical patterns from market charts in a systematic manner (Levy, 1971; Chang and Osler, 1999; Lo, Mamaysky, and Wang, 2000). Nevertheless, technical charting remains a very subjective tool as the same figure may in practice give two analysts entirely different inspiration.

¹⁰ In many financial markets, technical analysts will supplement the price data with transactions volume data. This is generally not possible in the foreign exchange market, however, due to its decentralized nature. Nevertheless, there is some anecdotal evidence that some analysts may combine technical trading rules with proprietary data on foreign exchange order flow, although the evidence on the usefulness of the latter for foreign exchange prediction is mixed (Sager and Taylor, 2008).

lengths; ¹¹*support-resistance trading rules*, which are based on the premise that a breach of a support or resistance level (lower and upper bounds through which the exchange rate appears to have difficulty in penetrating) will trigger further rapid exchange rate movement in the same direction; and *channel breakout trading rules*, which seek to identify time-varying support and resistance levels, or a ‘channel of fluctuation’ on the presumption that, once breached, further rapid exchange rate movement in the same direction will ensue.

By considering a number of variants of each trading rule and a range of different plausible parameterizations of each variant (see e.g., Sullivan, Timmermann, and White, 1999; White, 2000), we obtain a very large number of possible trading rules. In Appendix A, we provide precise details of each trading rule, of its variants and of the various parameterizations considered. This leads us to consider a total of 21,195 distinct technical trading rules, including 2,835 filter rules, 12,870 moving average rules, 1,890 support-resistance signals, 3,000 channel breakout rules and 600 oscillator trading rules.

4. Returns and Performance Metrics

4.1 Excess returns

The daily excess return from buying one unit of foreign currency (against the U.S. dollar) and holding it for one day is defined as

$$r_t \equiv \ln(s_t / s_{t-1}) - \ln[(1 + i_{t-1}) / (1 + i_{t-1}^*)], \quad (1)$$

where i_{t-1} and i_{t-1}^* denote daily interest rates on U.S. dollar deposits and foreign currency deposits on day $t - 1$, respectively, and s_t and s_{t-1} denote the spot foreign exchange rate (U.S. dollars per unit of a foreign currency) on days t and $t - 1$, respectively. The excess return is thus made up of the appreciation of the foreign currency relative to the domestic currency (U.S. dollar) over the holding period, $\ln(s_t / s_{t-1})$, minus the interest carry associated with borrowing one unit of domestic currency and lending one unit of foreign currency overnight, $\ln[(1 + i_{t-1}) / (1 + i_{t-1}^*)]$. For an investor committing their own funds,

$\ln(s_t / s_{t-1})$ represents the gross return and $\ln[(1 + i_{t-1}) / (1 + i_{t-1}^*)]$ represents the benchmark return, while for an investor who starts with zero funds and borrows domestic currency in order to invest, the gross return is $\ln(s_t / s_{t-1}) - \ln[(1 + i_{t-1}) / (1 + i_{t-1}^*)]$ and the benchmark return is zero since there is zero commitment of funds (in which case the gross return and the excess return coincide). When r_t is negative, a positive return could have been made by shorting the foreign currency, i.e. selling one unit of foreign currency against domestic

¹¹ Moving average trading rules are closely related to ‘momentum’ trading strategies, which are essentially trend-following strategies.

currency overnight.

More generally, the daily excess return of the j -th technical trading rule is defined as

$$R_{j,t} \equiv S_{j,t-1} r_t, \quad (2)$$

for $j=1, \dots, J$, where $S_{j,t-1}$ denotes the daily position guided by j -th technical trading rule (details of which are given in Appendix A), which is determined by all historical prices tracking back from the closing spot rate of day $t - 1$. We shall mostly think of a position $S_{j,t-1}$ in a currency as taking a value of either +1 (long the foreign currency, short the U.S. dollar), -1 (short the foreign currency, long the U.S. dollar), or else 0 (neutral) based on the information set at time $t - 1$, although it is also possible that a technical trading rule may generate buy or sell signals with different intensity, indicating that a long or short position should be initiated but at a level which is less than the total risk budget.

The discussion so far has assumed zero transaction costs; in practice these may be significant, especially when trading emerging market currencies (Burnside, Eichenbaum, and Rebelo, 2007; Ramadorai, 2008). Indeed, a technical trading rule may predict exchange rate movements in the sense of generating significantly positive excess returns but still not be profitable once the excess returns are adjusted for transaction costs (Timmermann and Granger, 2004). If transaction costs are attributed to the existence of a bid-ask spread in spot exchange rates and interest rates then, following Neely and Weller (2013), we can estimate them from the bid-ask spread in forward exchange rates, since forward rates are in practice calculated by traders as the spot rate plus the ‘forward points’, which in turn are calculated from the interest rate differential.¹²

Even so, however, there is an issue that posted (and therefore recorded) bid-ask spreads are indicative only and will tend to be larger than effective spreads at rates that are actually traded (Lyons, 2001; Neely and Weller, 2013). Following Neely and Weller (2013), who use informal survey evidence from foreign exchange market traders in an attempt to resolve this issue, we therefore use one third of the quoted one-month forward rate bid-ask spread in each currency published on Bloomberg as an estimate of one-way transaction costs on any particular day. For periods before the Bloomberg data is available, we also follow Neely and Weller (2013): for developed country currencies we set the transaction cost at a flat 5 basis points in the 1970s, 4 basis points in the 1980s and 3 basis points in the 1990s, and for emerging market currencies we set the daily cost at one third of the average of the first 500 bid-ask observations available on Bloomberg. Over the full sample periods, this resulted in average one-way transaction costs for developed country currencies of just under 4.5 basis points, while the corresponding figure for emerging market currencies is just under 21 basis points (full details are provided in Table 1 of the Online Appendix).

¹² Equivalently, this follows from the covered interest parity condition when bid-ask spreads are introduced; see, e.g. Taylor (1987, 1989).

4.2 Performance Metrics

Our first performance metric is the mean excess return (after allowing for transaction costs) of the j -th technical trading rule, which is defined simply as

$$\bar{R}_j \equiv 1/T \sum_{t=1}^T R_{j,t}, \quad (3)$$

and which is the simplest performance metric. Its major shortcoming is that it does not take into account the riskiness of the trading rule in terms of the volatility of returns.¹³ Our second measure is the ex post Sharpe ratio (SR), which is a standard performance metric in the finance industry and measures units of average excess return per unit of risk with the latter measured as the standard deviation of excess returns (Sharpe, 1966).¹⁴ The Sharpe ratio of the j -th technical trading rule is defined as

$$SR_j \equiv \bar{R}_j / \sigma_j, \quad (4)$$

where σ_j is the standard deviation of excess returns generated by the j -th trading rule and is based on the heteroskedasticity and autocorrelation consistent (HAC) estimator of Politis and Romano (1994).

Now, a standard method in the asset management industry for assessing the skill inherent in a strategy is to decompose a performance metric into a component due to ‘tilting’ (i.e. the component due to being on average long or short an asset, which could be replicated by a simple buy-and-hold strategy) and a component due to ‘market timing’ (i.e. the component due to timing trades to increase profits rather than just tilting). Thus, we can split daily excess returns into a tilt component (average foreign currency position over the whole sample period times foreign currency return for period t) and a market timing component (the remainder) as follows:

$$R_{j,t} \equiv R(\text{Tilt})_{j,t} + R(\text{Tim})_{j,t}, \quad (5)$$

¹³ Volatility is only a very crude measure of risk, of course, since modern asset pricing specifies risks as arising from the covariances of returns with the sources of risks in the economy, and these risk factors are in turn the stochastic processes that drive the stochastic discount factor that prices all assets.

¹⁴ The Sharpe ratio is a more informative metric than the mean excess return as it adjusts mean excess returns by the associated volatility. Suppose, for example that we found that two trading rules $TR1$ and $TR2$ have the same mean return but the Sharpe ratio of $TR1$ is twice that of $TR2$. By doubling the size of the positions taken by $TR1$, $TR1$ could have earned twice the average return of $TR2$ for the same level of risk (as measured by the volatility of returns) as $TR2$, since doubling the positions taken will double both the mean return and the standard deviation of returns. Scaling up the size of the positions taken is often referred to as leveraging the strategy, or scaling up the risk budget, since it scales up the standard deviation of excess returns and the volatility of standard deviations—i.e. risk—by the same factor. Although the Sharpe ratio does adjust for standard deviation, it is possible that one currency provides consistently high returns with low volatility against the U.S. dollar due to country-specific risk premia. Profits from investing in foreign currencies, including interest differentials, may simply reflect risk compensation because these currencies are associated with fundamental uncertainty such as unexpected government intervention or restricted repatriation of funds (e.g., Cornell and Dietrich, 1978; Hodrick and Srivastava, 1986; Froot and Thaler, 1990). Perhaps the simplest approach to measuring country-specific risk premia is through calculating the returns from a simple buy-and-hold position in the foreign currency (Sweeney, 1986), on the supposition that this must represent the compensation to a currency investor for holding risky foreign currency, which is one definition of a country risk premium to compensate various uncertainty including the peso problem. This assumes, of course, that risk premia remain constant over time.

where the tilt and timing components, $R(Tilt)_{j,t}$ and $R(Tim)_{j,t}$, respectively, are defined as:¹⁵

$$R(Tilt)_{j,t} \equiv \left[\frac{1}{T} \sum_{t=1}^T S_{j,t-1} \right] r_t \quad (6)$$

and

$$R(Tim)_{j,t} \equiv R_{j,t} - \left[\frac{1}{T} \sum_{t=1}^T S_{j,t-1} \right] r_t. \quad (7)$$

A simple performance metric based on the timing component of excess returns is simply the time-series mean of this component:

$$\overline{R(Tim)}_j \equiv \left[\overline{R}_j - \left[\frac{1}{T} \sum_{t=1}^T S_{j,t-1} \right] \frac{\sum_{t=1}^T r_t}{T} \right]. \quad (8)$$

The relative mean excess return, $\overline{R(Tim)}_j$, subtracts the average foreign currency position times the average foreign currency excess return from holding the foreign currency over the period, and is our third performance metric. In particular, it penalizes trading rules that have a high tilt component and which may simply collect risk premia by riding a trend appreciation or depreciation without timing well the trades into or out of the currency to exploit changes in direction.

Note that $\overline{R(Tim)}_j$ is very similar to the X -statistic introduced by Sweeney (1986) in his assessment of technical foreign exchange trading rules. With foreign currency holdings normalized to zero or plus or minus one, $\overline{R(Tim)}_j$ and X_j are in fact equivalent. If, however, daily positions are allowed to take another value the two will differ. We therefore propose the mean excess return to market timing, $\overline{R(Tim)}_j$, as a generalization of the X -statistic and our third suggested performance metric.¹⁶

Both the X -statistic nor our $\overline{R(Tim)}_j$ statistic are adjusted for risk in the sense of subtracting out a constant risk premium associated with holding foreign currency. Neither of them is adjusted for risk in the sense of the volatility of returns, however. Nevertheless, as

¹⁵ As discussed in the previous footnote, the tilt (i.e. buy-and-hold) component of the excess return of a foreign exchange trading rule may be interpreted as capturing the country risk premium corresponding to the average position in the foreign currency, so that the timing component serves as indicator of performance adjusted for country risk premium. Thus, in basing performance metrics on $R(Tim)_{j,t}$, the assumption is that a high market timing component of the excess return indicates that a trading rule provides returns in excess of risk premia associated with country-specific risk factors including the peso problem (measured with the tilt component). Further, one might argue that a profitable trading rule with a high timing component is in some sense more skillful than one with a high tilt component, since timing involves actively buying and selling the foreign currency while tilting is by definition more passive.

¹⁶ Sweeney (1986) effectively interprets the tilt component of the excess return as a (constant) risk premium associated with holding the foreign (i.e. non-U.S.) currency – see the previous two footnotes.

with the mean excess return, we can easily adjust our measure of market timing performance by calculating the Sharpe ratio of the excess return relative to market timing:

$$SR(Tim)_j \equiv \bar{R}(Tim)_j / \sigma(Tim)_j, \quad (9)$$

where σ_j is the HAC estimator of the standard deviation of excess returns relative to market timing for the j -th trading rule.

Overall, therefore, we suggest four different performance criteria for the j -th trading rule: the mean excess return (\bar{R}_j), the Sharpe ratio (SR_j), the mean excess return relative to timing ($\bar{R}(Tim)_j$), and the Sharpe ratio for market timing ($SR(Tim)_j$). Henceforth, we shall also make a distinction between the *predictability* and *excess profitability* of a trading rule. Typically, researchers use the term ‘predictable’ in the sense of statistical predictability, often using metrics such as mean square forecast error. In the present context, this is most closely related to generating a significantly positive mean excess return, and so we shall refer to trading rules that produce a significantly positive mean excess return as providing predictability or being predicitive. On the other hand, ‘profitable’ is more often used in connection with risk-adjusted returns, and so we use this term to relate to cases where a significantly positive Sharpe ratio is produced.

5. Empirical Methods: Avoiding Data-Snooping Bias

In this section we provide a largely non-technical and intuitive overview of the literature on methods for avoiding data-snooping bias, as a means of providing an intuitive exposition of our empirical methods, a full technical exposition of which is given in Appendix B.

To answer the intriguing question whether technical analysis is significantly profitable (i.e. can ‘beat the market’), our empirical strategy is to test if there exist technical trading rules that generate significantly positive performance metrics, as defined above. In addition, we are also interested in understanding the characteristics of any outperforming rules (i.e., rules that are statistically significantly profitable). Classical statistical inference is based on rejecting the null hypothesis if the likelihood of the observed data under the null hypothesis is low. Searching among competing model specifications or trading rules implicitly involves increasing the number of hypotheses tested as poorly performing models or rules are discarded. The problem of multiplicity arises from the fact that as we increase the number of hypotheses being tested (even implicitly), we also increase the likelihood of a rare event and, therefore, the likelihood of incorrectly rejecting the null hypothesis of interest in each competing model or trading rule (i.e., making a Type I error). Put another way, any good performance detected by rejecting the individual null hypothesis may not really be statistically significant but just based on luck, which has been maximized because of an

extensive specification search. In our case, given that we are searching among over 21,000 trading rules, a skeptic might say that they would have been surprised if we had *not* found any that performed extremely well, while perhaps quietly entertaining the notion that “if you torture the data long enough, it’ll confess to anything.”¹⁷

Applied researchers will recognize this problem as data mining, or over-fitting the data. Concern with the problem of data mining or, as it is now more commonly called, data snooping (because of the increased use of the former term to describe analysis based on so-called ‘big data’), has a long history in applied economics and finance (see e.g. Leamer, 1978 and the references therein) and there have been a number of important developments in this area over the past 15 years or so, which we draw on in our empirical work in order to mitigate the problem of data-snooping bias.

Let $\Theta = (\theta_1, \theta_2, \dots, \theta_J)$ denote the $I \times J$ vector in which the j -th element θ_j denotes the performance metric (e.g. Sharpe ratio or mean excess return) of the j -th trading rule for $j=1, \dots, J$ (in our case, $J = 21,195$). Traditionally, a researcher might choose the maximal element of Θ , $\max_{j=1, \dots, J} \theta_j = \theta_i$ say, and test the null hypothesis that this element is zero:

$$H_0 : \theta_i = 0. \tag{10}$$

A test of the null hypothesis (10) is generally regarded as an “individual test”. Because the performance metric will be constructed relative to an implicit or explicit benchmark, a test of the null hypothesis (10) amounts to a test of equal performance with the benchmark.

White (2000), however, points out that classical statistical inference based on individual testing applied to (10) will not take into account that θ_i is the maximal element of Θ , which will affect its statistical distribution: since θ_i has been chosen to be as large as possible after a search among J alternatives, where J may be very large, the assumed nominal significance of the test will tend to understate the true probability of a Type I error. In other words, the test will be biased towards rejection of the null hypothesis because of data snooping.

Further, White also notes that while a test of the null hypothesis (10) is a test for equality of performance relative to a benchmark, we may more properly wish to test for *superior* performance relative to the benchmark, which implies testing the null hypothesis in a joint testing framework:

$$H_0 : \theta_i \leq 0, \tag{11}$$

a rejection of which implies accepting the alternative hypothesis:

¹⁷ This saying, or something similar, is often attributed to Economics Nobel Laureate Ronald Coase; see, e.g., Leamer (1983).

$$H_1 : \theta_i > 0. \tag{12}$$

- i.e., in our case, that the trading rule corresponding to the maximal performance metric is significantly superior to the benchmark.

To account for these issues, White (2000) proposes a ‘reality check’ test, which tests the composite null hypothesis (11) based on the joint distribution of all elements of Θ . Apart from the complexity of modelling a high-dimensional joint distribution, the introduction of a composite null hypothesis (i.e. one involving an inequality as opposed to an equality) is highly problematic because distributions in composite hypothesis testing typically depend on so-called nuisance parameters, such that the distribution of the test statistic under the null hypothesis is not unique. White therefore suggests estimating the empirical distribution of the reality check test statistic through bootstrapping, which is a method in which blocks of the data set (in order to preserve any serial correlation) are randomly sampled and then joined together to form a pseudo time series of the same length and with similar properties to the true time series, from which the vector of performance metrics is estimated and the i -th element of that vector (where i is fixed to correspond to the original maximal element of Θ , θ_i) is stored. After this has been done a large number of times, a set of the maximal differences in performance metrics between the original and the pseudo time series, ordered by magnitude, is taken as the empirical distribution of the maximal performance metric and is used to construct the marginal significance level or p -value of the original statistic. White terms this test the bootstrap reality check (BRC).

While the development of the reality check was a landmark step forward in this literature, Hansen (2005) notes that the BRC statistic will tend to have low power to reject the null hypothesis (i.e. detect a superior trading rule) in cases where J is very large and many poorly performing trading rules are involved.¹⁸ Hansen’s (2005) test for superior predictive ability (SPA test) improves the BRC essentially by weighting the performance metrics in constructing a test statistic such that poor performers are given lower weight.¹⁹

Both White’s (2000) BRC test and Hansen’s (2005) SPA test for a *single* (maximal) significantly outperforming trading rule. In practice, one may wish to identify *all* statistically significantly outperforming trading rules (i.e., all rules that reject the null hypothesis). To address this, Romano and Wolf (2005) propose tests based on all elements of the performance metric vector $\Theta=(\theta_1, \theta_2, \dots \theta_J)$ in a multiple test framework:

¹⁸ One way to see the intuition behind this is as follows. One very simple (albeit conservative) way to correct for choosing the maximal performance metric – i.e. the one with the smallest p -value – from among a set of J alternatives is to reduce the chosen nominal significance level to a fraction of $1/J$ of what it would be had only one trading rule been considered; this is the so-called Bonferroni bound test. Thus, if we let p_i denote the smallest of the J p -values (corresponding to the maximal element of Θ , θ_i) then the Bonferroni bound test, at nominal significance level α , rejects the null hypothesis if $p_i < \alpha/J$. If we include a sufficiently high number of poorly performing trading rules that have high p -values and so do not affect p_i but only increase J , then p_i will never be small enough to reject the null hypothesis.

¹⁹ With respect to the remarks made in the previous footnote, this is analogous to using a modified denominator when defining the appropriate critical value, i.e. α/J^* for some $J^* \leq J$.

$$H_0^{(j)} : \theta_j \leq 0. \tag{13}$$

Now, one can in fact imagine doing this using White's (2000) reality check framework by using the bootstrap to construct the empirical distribution of each element of Θ rather than just the maximal element, and then using these to test the family of null hypotheses (13). Romano and Wolf show, however, that greater test power can be obtained by following a stepwise multiple (StepM) testing procedure as follows.²⁰ In the first step, the joint empirical distribution of all J trading rule performance metrics is calculated, in a framework similar to that of White (2000), and those judged statistically significant at a given nominal significance level (i.e. for which $H_0^{(j)} : \theta_j \leq 0$ is rejected) are recorded. In the second step, the statistically significant trading rules in the first step are excluded and the procedure is repeated; because the trading rules may be correlated, this may result in rules that in the first step appeared insignificant now becoming significant. The steps are then repeated until no significant trading rules remain.²¹

Finally, Hsu, Hsu, and Kuan (2010) propose a Stepwise SPA testing procedure that effectively combines the best features of Hansen's SPA test procedure and Romano and Wolf's StepM test procedure by minimizing the influence of poor performers on the power of the tests while identifying as many statistically significant trading rules as possible. This method, which we apply in the present analysis, is designed for large-scale multiple testing problems with potential data-snooping bias and is a powerful method of identifying as many significant rules as possible given an exact significance or Type I error level.²² In particular, this test allows us to jointly test each individual null hypothesis, $H_0^{(j)}$, such that the rejection of the j -th individual null hypothesis indicates that the j -th technical rule is significantly profitable, free of data-snooping bias.

We give the precise technical details of the implementation of our Stepwise SPA test in Appendix B.

In our empirical analysis for each foreign currency in a sample period, we report the number and the lowest p -values of the technical rules that are rejected by the Stepwise SPA test based on a significance level of 10% or lower.²³

²⁰ Recall that the power of a test is the probability of rejecting a false null hypothesis; in the present context it is the probability of detecting a set of profitable trading rules whose profitability is not just due to chance.

²¹ Hansen (2005) and Romano and Wolf (2005) also introduce other technical refinements to improve the power characteristics of their test procedures which need not detain us in this non-technical overview.

²² Technically, the error for which we control in such a multiple testing framework is the family-wise error, defined as the probability of rejecting at least one correct null hypothesis. That is, when we impose a 10% significance level in our testing, we expect a 10% chance of wrongly identifying any ineffective technical rules as profitable ones.

²³ Although we shall generally highlight and distinguish between statistical significance at the 5% and 10% levels, unless explicitly stated otherwise, we shall state that a test statistic is significant if it is significant at the 10% level or lower.

6. The Empirical Performance of Technical Trading Rules

6.1 Mean excess return and Sharpe ratio

To examine the predictability of exchange rates using technical analysis, i.e. whether technical trading rules can generate significantly positive mean excess returns, we focus on two indicators generated from the stepwise test: the first is the number of predictive rules that produce significantly positive performance metrics,²⁴ while the second is the performance metric and the associated p -value of the best rule that provides the highest performance metric among all rules.

Panel A of Table 2 reports the test results based on mean excess returns and suggests that technical rules do indeed forecast foreign exchange rate movements in a general sense. Based on performance in mean excess returns, fifteen out of thirty currencies are predictable at the 10% significance level (i.e., the number of currencies with at least one asterisk) and eleven currencies are predictable at the 5% significance level (i.e., the number of currencies with two asterisks).

Of these, five out of the nine developed currencies (i.e. 56%) are found to be predictable: the German mark/euro, Japanese yen, New Zealand dollar, Swedish krona, and Swiss franc. The New Zealand dollar appears to be the most predictable developed currency on this metric in that there exist 199 significantly predictive rules for this currency. The economic magnitude of the predictability is also substantial: the annualized excess returns on the best performing rules that are statistically significant at the 10% level or lower for developed currencies are clustered in a tight range from 6.5% (German mark/euro) to 7.7% (Japanese yen), with an average of 6.9% per annum.

The evidence for technical rule predictiveness, based on mean excess returns, is also strong in emerging currencies. Ten out of 21 emerging market currencies (i.e. 48%) are predictable at the 10% level or better, and seven of these are predictable at the 5% level. Specifically, there are 2,086, 427, and 88 significantly predictive rules in the Taiwanese dollar, Colombian peso, and Russian ruble, respectively. The annual returns generated by the best technical rule in those ten currencies range from 5.2% (Indian rupee) to as spectacularly high as 16.2% (Russian ruble), 12.7% (Korean won) and 11.6% (Colombian peso), with an average of 9.5% per annum.

We then examine the excess profitability of exchange rates based on technical analysis with simple allowance for risk by analyzing whether technical trading rules can generate a significantly positive Sharpe ratio that adjusts for risk related to volatility of returns. Panel B of Table 2 reports the test results based on the Sharpe ratio,²⁵ and suggests that technical excess profitability in foreign exchange trading remains significant when adjustment is made for risk. Three out of nine developed currencies remain significantly profitable on this

²⁴ According to Timmermann and Granger (2004), the existence of a ‘thick’ set of outperforming models can be regarded as strong evidence for predictability.

²⁵ The Sharpe ratio reported in Panel B has been annualized following LeBaron (2002). That is, we multiply the daily Sharpe ratio (i.e. the mean excess return divided by the standard deviation) by $\sqrt{250}$.

criterion at the 10% significance level or better, with Sharpe ratios tightly clustered: 0.652 (New Zealand dollar), 0.681 (German mark/euro) and 0.752 (Japanese yen). The Japanese yen not only has the highest and most significant Sharpe ratio, there exist 31 outperforming technical rules for the Japanese yen, compared to four for the German mark/euro and two for the New Zealand dollar.²⁶

Among our 21 emerging market exchange rates, nine are profitable using technical trading rules at the 10% level when we use the Sharpe ratio criterion, of which four are profitable at the 5% level. While the Taiwanese dollar is still strongly profitable on this metric, with 1,170 significantly outperforming technical rules and a top Sharpe ratio of 1.149 (p -value = 0.01), the Colombian peso is overall most profitable, with 237 outperforming rules and a top Sharpe ratio of 1.279. The significant Sharpe ratios are also impressively high, ranging from 0.677 for Israel, with four of them exceeding 1.0 and with an average of 0.972.

It is also interesting to examine which technical trading rules are the best performing for developed and emerging market currencies using these two criteria. Among developed country currencies, there is a penchant for moving average rules to be the best performing. Of the five cases that produce statistically significant mean excess returns, the highest performing rules in three of the cases are moving average rules while one is a filter rule and one is a channel-breakout rule. When the Sharpe ratio criterion is used, however, the highest performing rules for all of the significant (and indeed all of the insignificant) cases are moving average rules. Amongst these, the triple moving average rule identified by Lequeux and Acar (1998) (coded MA5 in Table 2) as popular among foreign exchange traders is most often the best performing rule, in five out of nine cases, and is statistically significant for one of three developed country currencies (New Zealand dollar) that achieve significance.

For emerging market currencies, of the ten cases where significant mean excess returns are generated, in five cases the highest performing strategies are filter rules and five are support-resistance level rules. There is again some indication of moving average rules becoming the highest performing strategies when the Sharpe ratio criterion is considered, with four of nine significant cases based on the Sharpe ratio relating to moving averages, although in three cases a support-resistance rule is the highest performing, in one case it is a channel-breakout rule and in one case it is an oscillator (overbought/oversold) indicator trading rule. There is also again a tendency for the triple moving average to be the best performing technical trading rule, in nine out of 21 cases, and supplies four of the nine cases of statistical significance (Brazilian real, Chilean peso, Israeli shekel and Turkish lira).

Table 3 repeats some of the information in Table 2 for the statistically significant cases, but also adds important information for these cases. In particular, Panel A of Table 3 reports only the cases where the mean excess return is significant (corresponding to the starred elements of Panel A of Table 2) and also reports the Sharpe ratio associated with each

²⁶ Our results for developed currencies are therefore largely consistent with the empirical findings of Qi and Wu (2006), whose study is based on a smaller set of technical trading rules and the one-step joint test of White (2000).

strategy. Conversely, Panel B of Table 3 reports only the cases selected by a significant Sharpe ratio and also reports the associated mean excess returns.

In terms of overall performance we can see from Table 3 that, for both developed country and emerging market currencies, there is a slight difference in choosing the trading rule based on significant mean excess returns as opposed to Sharpe ratio: while the results for the German mark/euro and Japanese yen are almost the same regardless of which selection criterion is used, for the New Zealand dollar choosing the trading rule based on mean excess return would have yielded average returns of 7.1% per annum with a Sharpe ratio of 0.596, compared to returns of only 3.3% per annum and Sharpe ratio only slightly higher, at 0.652, when the Sharpe ratio criterion is used to select the trading rule.²⁷ In addition, Sweden and Switzerland would be excluded completely using the significant Sharpe ratio criterion, whereas the mean excess return criterion identifies trading strategies that yield significant mean excess returns for each currency of over 6% per annum with associated Sharpe ratios of around 0.6.

For emerging market currencies, Table 3 shows that in only one of the eight cases which yield both significant mean excess returns and a significant Sharpe ratio—namely the Philippine peso—, do the two selection criteria yield similar results. In six of the other cases the mean excess returns are much higher without a major reduction in Sharpe ratio when significant mean excess return is used as the selection criterion (Brazilian real, Chilean peso, Colombian peso, Indian rupee, Israeli shekel, and Korean won), and in one (Taiwanese dollar) the Sharpe ratio is appreciably higher for little reduction in mean excess return when the Sharpe ratio criterion is used to select the best significant trading rule. Russia and South Africa are not selected at all by the significant Sharpe ratio criterion, but yield significant mean excess returns of 16.2% and 9.3% per annum respectively with respectable Sharpe ratios of 0.723 and 0.609 respectively. The significant Sharpe ratio criterion selects the Turkish lira (while the significant highest mean excess return does not) and yields average excess returns of 4.7% per annum with a Sharpe ratio of 0.852.

Overall, therefore, we find significant evidence of the existence of predictive and profitable technical trading rules among both developed and emerging market currencies after allowing for data-snooping bias, using either the significant highest mean excess return or the significant highest Sharpe ratio as the performance criteria.

In a close and wide reading of the empirical literature on technical trading rules, Menkhoff and Taylor (2007) highlight a set of stylized facts concerning the importance and profitability of technical currency analysis, one of which is that technical analysis tends to work better with volatile currencies. Figure 1 is a scatter plot of the highest significant mean excess return to technical trading for each currency (as reported in the first column of Panel A of Table 3) against the annualized standard deviation of the volatility of daily returns for

²⁷ Of course, if the volatility of excess returns is a sufficient measure of risk, then a trading rule should always be selected on the basis of the Sharpe ratio, since the strategy with the highest Sharpe ratio can always be leveraged to yield a higher mean return for the same level of risk than any strategy with a lower Sharpe ratio; see footnote 14.

those countries (as reported in Table 1). The scatter plot does indeed suggest a significant positive relationship and the correlation coefficient is 0.81. The estimated linear regression line, also plotted in Figure 1, yields a highly significant and positive slope coefficient and a very good fit (especially considering these are exchange rate returns in cross section):

$$\begin{aligned} \tilde{R}_i = 1.245 + 4.174 \sigma_i, \quad r^2 = 0.66, \\ (0.796) \quad (5.005) \end{aligned} \tag{14}$$

where \tilde{R}_i denotes the fitted value of the highest significant annualized mean excess return among all technical trading rules for country i , (first column Panel A of Table 3), σ_i denotes the annualized standard deviation of daily currency returns (Table 1), r^2 denotes the coefficient of determination, and figures in parentheses are t -ratios. The strength of this relationship is striking in terms of both statistical and economic significance: regression equation (14) suggests that a 100 basis points increase in the annualized standard deviation of daily exchange rate movements is associated with a 417 basis points increase in the annualized mean excess return to technical trading.

However, if we take the same countries and plot for each country the Sharpe ratio of the trading rule which yields the highest significant mean excess return (the second column of Panel A of Table 3) against the same measure of exchange rate volatility then the scatterplot appears to be random (see Figure 3 in the Online Appendix) and any positive correlation disappears. Indeed the correlation coefficient is -0.13 and the regression line is:

$$\begin{aligned} \tilde{S}\tilde{R}_i = 0.816 - 0.033 \sigma_i, \quad r^2 = 0.01, \\ (6.220) \quad (-0.473) \end{aligned} \tag{15}$$

where $\tilde{S}\tilde{R}_i$ denotes the Sharpe ratio associated with the technical trading rule yielding the highest significant mean excess return for country i .

These results suggest that while higher significant returns are indeed associated with more volatile currencies, as seen in regression equation (14), selecting a trading rule based on significant highest excess mean return does not lead to systematically higher risk being taken relative to the mean excess returns -- otherwise there would be a statistically significant negative correlation between the Sharpe ratio of those trading rules and volatility, instead of an insignificant one, as seen in regression equation (15). The key point here is that we are considering trading rules chosen by considering whether they have *significant* highest mean excess returns: one reason why highest mean excess returns may be insignificant is because the excess returns series is too volatile relative to the mean of the series.²⁸

²⁸ In other words, the standard error of the estimated mean excess return will be high and so a t -ratio statistic

If we consider trading rules selected by the highest significant Sharpe ratio, as in Panel B of Table 3, then we find that there is no significant correlation either between Sharpe ratios (second column of Panel B) and currency volatility, or between the associated mean excess returns (first column of Panel B) and currency volatility (see Figures 4 and 5 in the Online Appendix for scatter plots).

These results therefore support but also qualify the stylized fact highlighted by Menkhoff and Taylor (2007) that technical analysis tends to work better with volatile currencies: trading rules may indeed be more predictive (in the sense of producing significant mean excess returns) for more volatile currencies, but there is no significant relationship between volatility and technical rule excess profitability (i.e. after adjusting for risk). Thus, while high returns may have been generated by technical trading rules applied to some volatile currencies, similarly high returns with similar levels of risk (in the sense of the volatility of returns) could also have been generated by technical strategies applied to other, relatively less volatile currencies by appropriately leveraging the strategy (i.e. taking larger positions).

6.2 The performance of technical trading rules adjusted for tilt

We next examine, in Table 4, the technical predictability and excess profitability using two timing criteria, namely the mean excess return and the Sharpe ratio each adjusted for tilt. As discussed above, these criteria strip trading performance of the buy-and-hold element. This may be viewed as taking out the ‘low hanging fruit’ of riding an average trend over the sample period as opposed to timing trades more skillfully, or (as in Sweeney, 1986) it may be viewed as stripping out a constant risk premium associated with holding foreign currency.

Comparing the results relative to the timing criteria reported in Table 4 with those based on standard criteria, in Table 2, we can see that there is in fact little difference in most cases. For developed currencies there is no qualitative and little quantitative difference at all, while for emerging market currencies Brazil fails to yield significant results once the tilt component of returns is removed and the Singaporean and South African Sharpe ratios become significant once tilt is removed. There is also no clear tendency for the timing criteria to select different trading rules for either developed or emerging market currencies.

These results may be interpreted in three ways. First, as a simple robustness check on the empirical results reported above, they demonstrate that the results are indeed fairly robust using alternative performance metrics. Second, interpreting the tilt component as a constant risk premium for holding foreign currency, they suggest that such risk premia are not the main driving force explaining the excess profitability of technical analysis. Third, they demonstrate that successful technical trading rules contain inherent skill in timing market trades and do not for the most part rely on simply riding trends without skillfully timing market entry and exit.

for its significance will be small. This is also why the significant highest mean excess returns also have high associated Sharpe ratios, given that the Sharpe ratio is proportional to the t -ratio of the mean – see footnote 33.

7. Robustness Checks

In this section we report the results of a number of robustness checks on our analysis.

7.1 Break-even transaction costs

All of the technical trading performance results reported so far have been generated after adjusting the returns to the various trading strategies according to transaction costs based on the forward rate spread, as discussed above. In this section, following Bessembinder and Chan (1995) and Neely and Weller (2003), we consider break-even transaction costs—i.e. the level of one-way transaction costs that would reduce excess profitability to exactly zero. Examining break-even transactions costs can be viewed as a robustness check on our results since the greater they exceed what might credibly be conceived of as a reasonable estimate of actual transaction costs, the more robust the results may be considered.

Table 5 reports break-even costs (in basis points) and the number of trades for the best-performing significant technical rules from Table 2, as well as information on estimated transaction costs. The first three columns relate to trading rules selected using the significant mean excess return criterion, the next three columns relate to trading rules selected using the significant Sharpe ratio criterion, and the final column lists the mean estimated transaction costs as estimated from the forward rate bid-ask spread, as discussed above, following Neely and Weller (2013). Within each block of three columns relating to mean excess return-selected rules or Sharpe ratio-selected rules, the first column gives the constant mean one-way transaction costs that would drive overall mean excess returns to zero; the second column (which we have headed ‘break-even ratio’) gives the ratio of the break-even transaction cost to actual estimated mean transaction cost; and the third column lists the number of trades triggered by the best-performing significant trading rule over the full sample period.

Consider first the case of significant mean excess return-selected rules for developed countries; the number of trades triggered ranges from 159 (Swedish krona) to 917 (Swiss franc), while for Sharpe ratio-selected rules it ranges from 24 (New Zealand dollar) to 731 (German mark/euro). Given that this analysis is based on daily data for developed country currencies over a period spanning some 45 years, from January 1971 to September 2015, with a total of 11,660 trading days, this suggests that the best-performing rules are quite judicious.

Although the emerging market currencies have much smaller sample periods, the same is by and large also true among emerging market currencies for rules selected by significant mean excess return. However, the trading rules selected by the Sharpe ratio criterion tend to split into three groups: high frequency (Korean won) with some 2,500 trades, medium frequency (Colombian peso, Philippine peso and Taiwan dollar) with around 150 or so trades, and low frequency (Brazilian real, Chilean peso, Indian rupee, Israeli shekel and Turkish lira) with 15 or fewer trades.

It is noteworthy that the break-even transaction costs in every case far exceed the actual estimated transaction costs. Of course, since we tested for superior predictive ability based on

returns adjusted for estimated transaction costs, this is hardly surprising since the gap between break-even average transaction costs and average estimated transaction costs is simply an alternative way of viewing the excess profitability of the best-performing rules. Nevertheless, the *size* of the gap between break-even and estimated transaction costs may be viewed as a robustness check on our assumptions.

Consider first rules selected by significant mean excess returns. For developed country currencies, the break-even ratio ranges from about 6.6 (Swiss franc), to nearly 43 (Japanese yen). For emerging market currencies the range is even greater, from in the region of about three (Taiwanese dollar, Israeli shekel and Russian ruble) to about 170 (Chilean peso).

While these ratios may be viewed as high, the break-even ratios for rules selected by significant Sharpe ratios are in nearly every case even higher: among developed country currencies they range between just above ten (German mark/euro) to about 100 (New Zealand dollar). For emerging market currencies in only one case is the break-even ratio less than ten (Korean won, the high-frequency traded currency, at just over three), while for other emerging market currencies it ranges from about 13 (Taiwanese dollar) to well over 500 (Chilean peso).

Overall, the results reported in Table 5 deliver the following implications. First, the best-performing technical trading rules do not rely on frequent trading for developed country currencies, although among emerging market currencies there is a greater range of trading frequency of the best-performing rules. Second, it is clear that technical predictability can be converted to excess profitability given a reasonable level of transaction costs in both developed and emerging currencies, even among currencies that are traded at relatively high frequency by the best-performing rules. Our results therefore strongly support the stylized fact highlighted by Menkhoff and Taylor (2007) that transaction costs do not necessarily eliminate the possibility of making significant profits using technical trading rules.

7.2 *Sub-period analysis*

In a discussion of the empirical literature on technical analysis in foreign exchange markets, Menkhoff and Taylor (2007) suggest that technical trading rules may have become less profitable over time, perhaps reflecting an increase in informational efficiency. Our long and broad sample of data series allows us to investigate the issue of the time-varying predictability and excess profitability of technical trading rules in the foreign exchange market.

We split the whole sample period into eight subsamples, 1972-1976, 1977-1981, 1982-1986, 1987-1991, 1992-1996, 1997-2001, 2002-2007, and 2008-2015.²⁹ Given that the first year of data (1971) is required to initialise some of the trading rules, our eight subsamples largely divide the full sample into five-year periods, with the exception of the last two. However, they coincide with some natural break points related to important events

²⁹ For emerging market currencies, we only conduct tests for available subsamples. Our sample splitting follows Levich and Thomas (1993), which splits the entire sample period into three five-year subsample periods for time-series analysis.

in the foreign exchange market, including the Maastricht Treaty in 1992, the East Asian currency crisis in 1997, the appearance of physical euro in 2002 and the global financial crisis of 2008.

For each currency in every subsample period, we conduct the stepwise test based on the mean excess return criterion, and report the numbers of significantly predictive rules in Table 6.³⁰ In the upper part of the table, we observe that developed currencies are more predictable with technical analysts in earlier subsample periods. In fact, five out of nine developed currencies are predictable in the first sample period (1972-1976). There are, however, only a few predictive rules found in the 1980s, and none since the 1992-1996 period. The disappearance of the significant predictability of developed currencies using technical trading rules since the early 1990s, noted by Menkhoff and Taylor (2007), is consistent with the findings of LeBaron (2002), Olson (2004), and Neely, Weller, and Ulrich (2009), among others. Nevertheless, we have to interpret the results of no significance since the 1990s with caution: it is possible that some profitable rules are not identified due to the low power of our tests with short sub-periods. Thus, a safer statement based on the sub-period analysis is that technical predictability certainly *used to* be strong in developed country foreign exchange rates but has decayed in more recent decades.

The lower part of Table 6 shows the time-series variation of technical predictability for emerging market currencies. Emerging market currencies also seem to be more significantly predictable in earlier periods, such as the Israeli shekel in the 1982-1986 period, the Taiwanese dollar and Philippine peso in the 1987-1991 period, and the Chilean peso in the 1997-2001 period. Consistent with our finding in developed currencies, the predictability of technical trading rules in emerging markets also decreases over time, however. Nevertheless, and notwithstanding the caveat concerning the reduced power of the stepwise test in the sub-sample, emerging market currencies still appear to be more predictable using technical trading rules than developed markets in the most recent periods.

Our finding of a downward trend in the performance of technical rules among both developed country and emerging market currencies supports the unavoidable ‘self-destruction’ process discussed by Timmermann and Granger (2004) and Timmermann (2008): specific technical rules generate profits from uncovering and exploiting important information or market sentiment and such excess profitability will sooner or later be identified and effectively arbitrated away by technical analysts.³¹ It also supports the closely related Adaptive Markets Hypothesis (Lo, 2004; Neely, Weller and Ulrich, 2009), according to which traders exploit and diminish the returns to trading strategies over time at a rate determined by the speed with which more and more traders assimilate and deploy the strategies. Under the Adaptive Markets Hypothesis, however, the market does not necessarily converge on the rational expectations equilibrium of the Efficient Markets Hypothesis not only because with zero profitable trading opportunities there will be no incentive to gather

³⁰ Tables based on other performance metrics present similar patterns and, thus, are unreported.

³¹ In stock markets, Schwert (2003) and McLean and Pontiff (2016) find that seemingly profitable patterns documented by academic papers greatly lose their profitability out of sample.

information, which would lead to a collapse of financial markets (Grossman and Stiglitz, 1980), but also because new profit opportunities are constantly being created in an evolutionary environment where institutions, business conditions and behavioural influences on the markets change over time. In this connection, it is also worth noting that the downward trend in the performance of technical trading rules does not appear to be uniformly monotonic over time: in Table 6 both India and Philippines show an increase in the number of technical rules with significantly positive mean excess returns in some periods, which is consistent with the prediction of the Adaptive Markets Hypothesis that trading rule performance may occasionally wax as well as wane as the trading environment shifts (Lo, 2004: 24):

Note also that, in a general sense, the sub-period analysis does not appear to support the proposition that central bank intervention leads to technical predictability,³² since we do not observe particularly strong predictability in currencies during periods with intensive central bank intervention such as during the recent global financial crisis (Melvin and Taylor, 2009).

7.3 Alternative home currencies

As a further robustness check on our results, we considered the performance of technical trading rules when investors have home currencies different from the U.S. dollar. In particular, we considered trading rules defined on U.S. dollar-based exchange rates, but with the Japanese yen, German mark/euro, U.K. pound and Swiss franc defined as the home currency in the following sense. The first trade requires the home currency (e.g. Swiss franc) to be sold to initiate the first position in either the U.S. dollar or a third currency. Every time a position is closed out, the home currency is bought and is then sold again to take the position implied by the next trade. Thus, the process is extremely conservative, and effectively marks every trade to market in the home currency, with all trades in this process incurring transaction costs, which in turn are conservatively estimated as the transaction cost of each currency to the U.S. dollar plus the average transaction cost of the home currency to the U.S. dollar.

In order to conserve space, the full set of empirical results relating to these alternative home currencies are given in Tables 2 to 9 in the Online Appendix, but we can summarise them here by saying that the results are largely qualitatively unchanged in all dimensions from those using the U.S. dollar as the home currency.

7.4 Alternative base currencies

In the next robustness check we considered the Japanese yen, German mark/euro, U.K. pound, and Swiss franc as alternative base currencies; that is to say, we constructed the cross-rates for all currencies against these base currencies, rather than against the U.S. dollar, and looked for profitable technical trading rules among those cross-rates. We again used a

³² This implication is in accordance with the findings of other authors that technical trading profits cannot be attributed purely to central bank intervention (e.g., Cornell and Dietrich, 1978; Neely and Weller, 2001; Neely, 2002).

conservative estimate of transaction costs by assuming investors need to move in and out of the U.S. dollar to effect the cross rate.

Again, to conserve space, the full detailed results are reported in Tables 10 to 17 in the Online Appendix, but may be summarised here as follows.

First, there is a general diminution in the number of statistically significant technical trading rules using alternative base currencies. Using mean excess returns as the selection criterion, only two developed country currencies and four emerging market currencies yield a significant trading rule with the Japanese yen as the base currency; two developed country currencies and three emerging market currencies with the German mark/euro as the base currency are found to be significant; two developed country currencies and two emerging market currencies with the U.K. pound as the base currency; and there are no significant trading rules with the Swiss franc as the base currency. Similar results are reported using the significant Sharpe ratio as the selection criterion.

These results, together with our main results reported earlier, appear to correlate broadly with the amount of foreign exchange market turnover in these base currencies. According to the latest Bank for International Settlements (2013) triennial central bank survey, for example, in April 2013 currency trades with the U.S. dollar as base currency accounted for 86.9% of average daily turnover in the global foreign exchange market, while trades with the euro as base currency accounted for 9.4% and those with the yen as base currency accounted for 2.0%, with trades against any other currency as base currency accounting for the remaining 1.7%.

The fact that technical analysis appears to be less profitable among the less traded base currency exchange rates presents, in some ways, something of a puzzle. However, the Adaptive Markets Hypothesis might perhaps be adduced in order to explain this phenomenon through an argument that, in the framework and language of Lo (2004), a sparsely populated market will be less subject to evolutionary shifts and therefore less profit opportunities may arise. In any event, it is clear that this issue is worthy of further research.

7.5. Out-of-sample performance

Another common solution to address the data-snooping or over-fitting issue is to employ out-of-sample analysis to examine if the best model from the in-sample estimation indeed performs well out of sample (Lo and MacKinlay, 1990; Sullivan, Timmermann, and White, 1999; Harvey, Liu, and Zhu, 2016). Out-of-sample forecast evaluation is not, however, without its pitfalls. In particular, while statistical methods may be used to allow for sampling variability over the forecast period, it is more difficult to control for structural shifts. From their reading of the empirical literature, Menkhoff and Taylor (2007) concluded that the predictability and excess profitability of technical trading rules has been diminishing over time, and this appears to be borne out by the empirical analysis reported above in Section 7.2. Thus, we might indeed expect technical trading rules to perform worse out of sample if that trend has continued. Nevertheless, notwithstanding these caveats, a simple test as to whether technical trading rules selected using the test for superior predictive ability continue to

perform well out of sample would seem to be a valid robustness check on our results, albeit a very stringent one.

We therefore provide a glimpse of the kind of investment performance that could have been obtained over an out-of-sample period from January 1 2012 to September 11 2015. Although our full-sample analysis to this point has been based on data up to September 2015, we can resample to choose the statistically significant best-performing rules using data up to the end of 2011 and then examine how these rules perform over the ensuing until the end of the sample in September 2015, a total of 1,817 trading days.³³ The results of this exercise are given in Table 7.

Consider first the results for developed country currencies. Using data up to the end of 2011, the significant mean excess return criterion selects trading rules for six currencies (Australian dollar, German mark/euro, Japanese yen, New Zealand dollar, Swedish krona and Swiss franc). In the post-sample period, all of these rules except one yield positive mean excess annualised returns, three of which are greater than 2%, two of which are less than 1% and one of which (for the Swiss franc) is negative; for the five positive mean excess return cases, the out-of-sample Sharpe ratios range from about 0.1 to about 0.4. Using the significant Sharpe ratio criterion, three developed country currencies are selected (German mark, Japanese yen and New Zealand dollar), and all of them achieve positive out-of-sample mean excess returns (with two of them in excess of 2% per annum) and two of them achieve Sharpe ratios as high as 0.3 and 0.5.

For emerging market currencies, ten currencies are selected in sample by the significant mean excess returns criterion, and eight of them retain positive mean excess returns out of sample, with out-of-sample mean excess returns as high as around 19%, 8%, 7%, and 5%, coupled with equally impressive Sharpe ratios of around 2.0, 0.9, 0.9 and 1.1, although two currencies record out-of-sample mean excess returns of around -9% and -3% per annum. Using the significant Sharpe ratio criterion, nine emerging market currencies are selected in sample, seven of which yield positive performance post sample, much of which is again impressive—in three cases exceeding 18%, 9% and 7%, with associated Sharpe ratios around 1.9, 1.3 and 0.6, respectively.

While these results tend to suggest a strong post-sample performance of technical trading rules, especially for emerging market currencies, a clearer overall picture can be obtained by examining the performance of a portfolio of technical trading rules over in-sample and out-of-sample periods. Accordingly, we constructed three such portfolios (one for developed country currencies, one for emerging market currencies and one combining the first two) using the technical trading rules and currencies identified as significant on pre-2012 data in terms of Sharpe ratios (i.e. as shown in Table 7: German mark/euro, Japanese yen, New Zealand dollar, Brazilian real, Chilean peso, Colombian peso, Indian rupee, Korean won, Philippine peso, Taiwanese dollar, Romanian new leu and Turkish lira).

³³ In fact, the first version of this paper submitted to this journal only used data up to 2011, so that selecting the significant best-performing trading rules trading rules using data up to 2011 and evaluating their performance using data from 2012 onwards may in fact be construed as an extremely stringent out-of-sample evaluation.

No attempt at portfolio optimization was made. In the first period we invest one dollar and simply allocate this evenly across all currencies in the particular portfolio. The allocation to each currency then grows daily at the rate of the daily mean excess return of the best technical trading rule for that currency and is compounded over time, so that gains and losses are rolled over, and the portfolio return is the sum of the returns to each of its constituent currencies. Effectively, therefore, the excess return for each portfolio is simply the unweighted average of the excess return to each of its constituent currencies and the value of the portfolio is simply the value of \$1 compounded at that rate of return.

The out-of-sample performance of these three portfolios, over the period January 1 2012 to September 11 2015, is reported in Table 8. Although all three portfolios generate positive mean excess returns out of sample, the developed country currency portfolio generates a Sharpe ratio of only 0.24 and annualized mean excess returns of 1.35%, which is not statistically significantly different from zero at standard significance levels, although it should be noted that the power of this test will be low with less than four years of out-of-sample data.³⁴ On the other hand, the emerging market currency portfolio yields an outstanding out-of-sample Sharpe ratio of 1.81 and mean excess returns of 7.3% per annum, even after allowing for transaction costs; notwithstanding the low test power, the *t*-ratio, at 3.5, strongly rejects the null hypothesis of zero mean excess returns out of sample.

These results are therefore consistent with the general tenor of the research reported in this paper and elsewhere (Menkhoff and Taylor, 2007), such that the profitability of technical analysis has diminished over time, particularly among developed country exchange rates, but that it may remain significantly profitable among emerging market currencies. Moreover, Table 8 shows that even a combined portfolio including both developed country and emerging market currencies would have yielded a strongly significant mean excess return of 4.6% per annum over the out-of-sample period, together with a very healthy Sharpe ratio of 1.2.

Figures 2-4 give a visual impression of the in-sample and out-of-sample performance of the three portfolios by tracking the cumulative value of \$1 invested over time. Consider Figure 2, relating to the developed country currency portfolio. Because there is in fact little trading in the portfolio prior to 1974, in this figure we have graphed the cumulative value over time of a \$1 investment on January 1 1974. Thus, a \$1 investment in the developed country currency portfolio at the beginning of 1974 would have been worth \$11.58 at the end

³⁴ The approximate *t*-ratios for the null hypothesis of zero mean excess returns in the post-sample period are computed as the square root of the sample period length (in years) times the annualized Sharpe ratio. This result is independent of the distribution of returns and follows from the Central Limit Theorem which states that whenever a random sample of size T is taken from a distribution with mean μ and variance σ^2 , then the sample mean will be approximately normally distributed with mean μ and variance σ^2/T . Given 3.7 years of out-of-sample data, however, $\sqrt{T}=1.92$, so this would require a Sharpe ratio of 0.85 or more for mean excess returns to be significantly different from zero in a one-sided test at even the 10% level (*t*-ratio of 1.64), which is a level not achieved in-sample or over the full sample for any trading rule for the developed country currencies in our data set. The low power of our out-of sample tests may be further illustrated by noting that, given that a number of the best trading rules for developed country currencies have a Sharpe ratio in the region of 0.6 (Table 3), even if this continued in the out-of-sample period, we would need $(1.64/0.6)^2=7.5$ years of data (double the length of our out-of-sample period) in order to reject the null hypothesis of zero mean excess returns at the 10% level. On the other hand, any significant results that are found will hold *a fortiori*.

of 2011, and would have increased to \$12.19 by September 11 2015, an increase of some 1,100%. Even allowing for the fact that the US consumer price index (CPI) has increased by a factor of about 4.8 over this period, this represents a very respectable average annual real return of about 2%. Figure 2 shows that although the portfolio performed well during the 2007-2008 global financial crisis, there was a significant falling off in performance after about 2010, and while there is some pickup in this performance in the out-of-sample period, from 2012 onwards, returns to the portfolio remain volatile, reducing both the mean excess return and the Sharpe ratio in the out-of-sample period, as shown in Table 8.

Figure 3 shows the cumulative growth in value of \$1 invested in the emerging market currency portfolio on July 4 1994 (the date when data on all emerging market currencies in the portfolio become available). The \$1 investment would have grown to \$4.06 by the end of 2011 and increased to \$5.36 by September 11 2015, representing an average real return, after adjusting for US CPI inflation, of nearly 6% per annum. The strong performance of the emerging market portfolio, especially out of sample, is quite striking. The portfolio shows a steady performance which, if anything, appears to improve during the period of the 2001-2008 global financial crisis and which even accelerates during the out-of-sample period.

When we consider the combined portfolio, Figure 4 shows that an initial \$1 investment on July 4 1994 reaches \$3.07 at the end of the in-sample period, and further rises to \$3.65 at the end of the out-of-sample period, representing an average real return to a U.S.-based investor of 4% per annum.

Overall, the out-of-sample analysis supports the view that technical currency analysis has diminished in profitability for developed country currencies over time, but may still be significantly highly profitable for emerging market currencies and may still provide an important source of return with acceptable risk within an investment portfolio.

8. Conclusion

A number of researchers have reported the widespread use and significant excess profitability of technical analysis in the foreign exchange market since the 1970s. However, much of the previous evidence is based on small sets of technical trading rules applied to a handful of currencies, assuming low transaction costs and with simple performance metrics and testing methods that may often have been subject to data-snooping bias. As a result, the long-debated issue of whether and why technical analysis beats the foreign exchange market has not been satisfactorily answered and several intriguing issues call for further exploration.

In the research reported in the present study, we accepted this challenge and carried out a large-scale investigation of the profitability of technical trading rules across a large set of developed country and emerging market currencies in long sample periods, in an analysis of a very large number of technical trading rules. Further, we employed recently developed testing methods which, while computationally extremely burdensome for the large-scale investigation at hand because they involve very large numbers of bootstrap replications, are extremely attractive in the present application as they allow us to make formal statistical

inferences for several important research questions while safeguarding against data-snooping bias.

The results of this research project are extremely interesting and shed new light on a number of important issues, including a number of stylized facts concerning technical currency analysis previously highlighted by Menkhoff and Taylor (2007).

In particular, we conclude that technical analysis indeed has predictive power for both developed and emerging currencies, in terms of being able to generate significant mean excess returns and impressive Sharpe ratios, but that emerging market currencies are in general more predictable with technical analysis than are developed country currencies. Moreover, this excess profitability is not in general wiped out when realistic allowance is made for transaction costs, although the profitability of technical analysis among both emerging market and developed country currencies does appear to have been declining over time. These general findings were also supported by an out-of-sample analysis, which in particular revealed the very impressive performance of technical analysis for some emerging market currencies during the out-of-sample period 2012-2015.

We also found, again supporting the stylized facts highlighted by Menkhoff and Taylor (2007), that the predictiveness of technical currency trading rules, in terms of the magnitude of significant mean excess returns, is significantly positively correlated with the volatility of a currency. This, however, is only half of the story, since volatile currencies will typically be associated with more volatile returns: once a simple allowance for risk is made by adjusting mean excess returns for the volatility of returns (i.e. using the Sharpe ratio as the performance metric), then the relationship between technical analysis performance and volatility disappears. This implies that although technical trading may generate higher significant mean excess returns with more volatile currencies, similarly high returns with similar levels of risk (in the sense of the volatility of returns) can be generated by technical strategies applied to relatively less volatile currencies by appropriately leveraging the trading strategy.

Our empirical investigation also suggests that technical excess profitability in the foreign exchange market cannot be simply attributed to constant risk premia, market crises or central bank intervention. Instead, our analysis favors the explanation of temporarily not-fully-rational behavior as the basis for technical predictability and excess profitability: this is consistent with our finding that the profitability of technical analysis has diminished over time in developed country currencies but remains to some extent, especially in certain emerging currencies. Our empirical evidence is also consistent with the Adaptive Markets Hypothesis (Lo, 2004; Neely, Weller and Ulrich, 2009), which holds that market participants increasingly exploit and therefore diminish the returns to technical trading rules over time, with the speed of diminution depending on the speed with which the market learns about and exploits trading strategies.

We conclude, therefore, that market maturity, and the associated degree of informational efficiency, is an important determinant of technical predictability and profitability in the foreign exchange market.

Appendix A: Details of Technical Trading Rules

A.1 Oscillator trading rules

One technical device that is widely used in the foreign exchange market (Taylor and Allen, 1992) is the ‘overbought/oversold’ indicator, or oscillator, although it has rarely been discussed in the academic literature. Oscillators are measures designed to indicate that price movements in a particular direction have recently been too rapid and that a correction in the opposite direction is imminent; they may take a number of precise forms. One popular form is the relative strength indicator (RSI) (Levy, 1967; Wilder, 1978), which is defined as:

$$RSI_t(h) = 100 \left[\frac{U_t(h)}{U_t(h) + D_t(h)} \right], \quad (\text{A1})$$

where $U_t(h)$ denotes the cumulated ‘up movement’ (i.e. the close-to-close increase on a day when the exchange rate has closed higher than the previous day’s closing rate) over the previous h days, and $D_t(h)$ denotes the cumulated absolute ‘down movement’ (the absolute close-to-close decrease on a day when the exchange rate has closed lower than the previous day’s closing rate) over the same period:³⁵

$$U_t(h) = \sum_{j=1}^h \mathbf{1}(s_{t-j} - s_{t-1-j} > 0)(s_{t-j} - s_{t-1-j}) \quad (\text{A2})$$

$$D_t(h) = \sum_{j=1}^h \mathbf{1}(s_{t-j} - s_{t-1-j} < 0) |s_{t-j} - s_{t-1-j}|, \quad (\text{A3})$$

where $\mathbf{1}(\cdot)$ is an indicator variable that takes the value one when the statement in parentheses is true and zero otherwise. The RSI thus attempts to measure the strength of ‘up movements’ relative to the strength of ‘down movements’ and is normalised to lie between 0 and 100. Common values at which a particular currency is deemed to have been overbought (signalling an imminent downward correction) or oversold (signalling an imminent upward correction) are 70 and 30, respectively (see, e.g., Henderson, 2002). Note that the RSI is a kind of ‘reversal’ indicator, since it is designed to anticipate a reversal in trend.

A standard oscillator trading rule based on the RSI may be expressed as follows:

O1: If $RSI_t(h)$ moves above $50+v$ for at least d days and then subsequently moves below $50+v$, go short the currency. If $RSI_t(h)$ moves below $50-v$ for at least d days and then subsequently moves above $50-v$, go long the currency.

A variation on the standard RSI trading rule imposes a pre-specified holding period for a position:

³⁵ Some expositions of the RSI define U_t and D_t in terms of average rather than cumulated up and down movements. This is equivalent to our definition, however, since it just involves dividing by the total number of days and this factor cancels out when the RSI is calculated.

O2: If $RSI_t(h)$ moves above $50+v$ for at least d days and then subsequently moves below $50+v$, go short the currency for k days and then neutralize the position. If $RSI_t(h)$ moves below $50-v$ for at least d days and then subsequently moves above $50-v$, go long the currency for k days and then neutralize the position.

Note that a trading signal is *not* generated when the RSI *enters* the overbought or oversold region (i.e. goes above $50+v$ or below $50-v$) but as it *exits* the region (i.e. crosses $50+v$ from above or crosses $50-v$ from below). This is because the currency can remain overbought or oversold for long periods and may become even more overbought or oversold for a while. The oscillator trading rule is designed to allow the spot rate to continue moving in the desired direction until a trend change becomes evident.

We consider a total of 600 oscillator trading rules. The specifications of oscillator trading rules we consider are described as below:

$h \in \{5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$, $\#h = 10$.

$v \in \{10, 15, 20, 25\}$, $\#v = 4$.

$d \in \{1, 2, 5\}$, $\#d = 3$.

$k \in \{1, 5, 10, 25\}$, $\#k = 4$.

O1: This trading rule has three parameters, h , v and d . The number of versions of *O1* considered is therefore $\#O1 = \#h \times \#v \times \#d = 10 \times 4 \times 3 = 120$.

O2: This trading rule has four parameters, h , v , d and k . The number of versions of *O2* considered is therefore $\#O2 = \#h \times \#v \times \#d \times \#k = 10 \times 4 \times 3 \times 4 = 480$.

A.2 Filter rules

The filter rule is one of the simplest techniques of technical analysis (Alexander, 1961) and in the present context basically involves buying a currency against another currency whenever the exchange rate has risen by more than a given percentage above its most recent low and selling it when the rate drops by more than a given percentage below its most recent high. A simple filter rule for trading a currency may be stated as follows:

F1: If the daily closing spot exchange rate of the currency moves up at least x percent above its most recent low and remains so for d days, go long the currency. If the daily closing spot rate moves down below its most recent high at least x percent and remains so for d days, go short the currency.

In this context, we define the most recent high (low) as the most recent closing price that is greater (less) than the j previous daily closing prices, for a given value of j .

A variation on this basic filter rule allows for neutral positions rather than requiring that positions always be either long or short after the first position is initiated:

F2: If the daily closing spot exchange rate of the currency moves up at least x percent above its most recent low and remains so for $d(x)$ days, go long the currency until its daily closing spot rate moves down at least y percent below the subsequent high and remains so for $d(y)$

days, at which time sell the currency and neutralize the long position. If the daily closing spot exchange rate of the currency moves down at least x percent below its most recent high and remains so for $d(x)$ days, go short the currency until its daily closing spot rate moves up at least y percent above the subsequent low and remains so for $d(y)$ days, at which time buy the currency and neutralize the short position. y is less than x . $d(y)$ is less than or equal $d(x)$.

Subsequent lows and highs in this context are again defined with respect to the previous j days.

A third, more naïve variation on the filter rule imposes a pre-specified holding period that is adhered to ignoring all other signals:

F3: If the daily closing spot exchange rate of the currency moves up at least x percent above its most recent low and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot exchange rate of the currency moves down at least x percent below its most recent high and remains so for d days, go short the currency for k days and then neutralize the position.

In our empirical work, we examine a total of 2,835 filter rules, the parameterization of which is described as below:

$x \in \{0.05, 0.1, 0.5, 1.0, 5.0, 10.0, 20.0\}$ in %, # $x = 7$.

$y \in \{0.05, 0.1, 0.5, 1.0, 5.0, 10.0, 20.0\}$ in %. Given $y < x$, the number of x - y combinations:

$$\#(x-y) = \sum_{j=1}^6 j = 21.$$

$d \in \{0, 1, 2, 3, 4, 5\}$, # $d = 6$.

$d(x) \in \{0, 1, 2, 3, 4, 5\}$, $d(y) \in \{0, 1, 2, 3, 4\}$. Given $d(y) \leq d(x)$, the number of $d(x)$ - $d(y)$

combinations: $\#[d(x)-d(y)] = \sum_{j=1}^5 j = 15$.

$j \in \{1, 2, 5, 10, 20\}$, # $j = 5$.

$k \in \{5, 10, 15, 20, 25\}$, # $k = 5$.

F1: This trading rule has three parameters, x , d and j . The number of versions of *F1* considered is therefore $\#F1 = \#x \times \#d \times \#j = 7 \times 6 \times 5 = 210$.

F2: This trading rule has five parameters, x , y , $d(x)$, $d(y)$ and j . The number of versions of *F2* considered is therefore $\#F2 = \#(x-y) \times \#[d(x)-d(y)] \times \#j = 21 \times 15 \times 5 = 1,575$.

F3: This trading rule has four parameters, x , d , j and k . The number of versions of *F3* considered is therefore $\#F3 = \#x \times \#d \times \#k \times \#j = 7 \times 6 \times 5 \times 5 = 1,050$.

A.3 Moving average trading rules

Moving average trading rules can range in construction from very simple to complex, and have been among the most widely used of technical trading rules in the foreign exchange market (Taylor and Allen, 1992). This family of rules attempts to ride trends and identify imminent breaks in trend or the emergence of new trends. In a simple (single) moving average trading rule, for example, the local trend is approximated by the moving average and a break in trend, which is a trading signal to initiate a new position or neutralise the current

position, is indicated when the moving average is crossed by the spot rate. Traders often use a short-term moving average in place of the spot rate in rules of this kind, in which case an imminent upward break in trend might be signalled by a short moving average intersecting from below a longer moving average. Conversely, a downward break in trend would be signalled by the short moving average crossing the long moving average from above.

If we define the simple moving average of the exchange rate over j days as

$$MA_t(j) = \frac{1}{j} \sum_{i=0}^{j-1} s_{t-i} , \quad (\text{A4})$$

then a single moving average trading rule may be expressed as follows:

MA1: If the daily closing spot exchange rate of the currency moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency until its daily closing spot rate moves down at least x percent below $MA_t(q)$ and remains so for d days, at which time go short the currency. If the daily closing spot exchange rate of the currency moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency until the daily closing spot rate moves up at least x percent above $MA_t(q)$ and remains so for d days, at which time go long the currency.

As before, a simple variation that pre-specifies the holding period, ignoring all other signals during the holding period, would be as follows:

MA2: If the daily closing spot exchange rate of the currency moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot exchange rate of the currency moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency for k days and then neutralize the position.

A double moving average trading rule may be formulated in the following fashion:

MA3: If $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency until $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, at which time go short the currency. If $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency until $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, at which time go long the currency. p is less than q .

And the pre-specified holding period version is of the form:

MA4: If $MA_t(p)$ moves up at least x percent above $MA_t(q)$ and remains so for d days, go long the currency for k days and then neutralize the position. If $MA_t(p)$ moves down at least x percent below $MA_t(q)$ and remains so for d days, go short the currency for k days and then neutralize the position. p is less than q .

Finally, consider a triple moving average rule (Lequeux and Acar, 1998):³⁶

MA5: If the daily closing spot exchange rate of the currency moves up at least x percent above any two of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go long the currency with one third of the risk budget (currency position= $+1/3$). If the daily closing spot rate of the currency moves up at least x percent above all three of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go long the currency with the whole of the risk budget (currency position= $+1$). If the daily closing spot exchange rate of the currency moves down at least x percent below any two of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go short the currency with one third of the risk budget (currency position= $-1/3$). If the daily closing spot rate moves down at least x percent below all three of $MA_t(n)$, $MA_t(p)$ and $MA_t(q)$ and remains so for d days, go short the currency with the whole of the risk budget (currency position= -1). n is less than p , which is less than q .

Clearly, the trading rule never indicates a neutral position—for example if the spot rate is above only one moving average, it must be below two of them and a one third short position is indicated.

We consider a total of 12,870 moving average trading rules, the details of the specifications are described as below:

$q \in \{2, 5, 10, 15, 20, 25, 50, 100, 150, 200, 250\}$, $\#q = 11$.

$p \in \{2, 5, 10, 15, 20, 25, 50, 100, 150, 200\}$. Given $p < q$, the number of p - q combinations:

$$\#(p-q) = \sum_{j=1}^{10} j = 55.$$

$n \in \{2, 5, 10, 15, 20, 25, 50, 100, 150\}$. Given $n < p < q$, the number of n - p - q combinations:

$$\#(n-p-q) = \sum_{J=1}^9 \sum_{j=1}^J j = 165.$$

$x \in \{0, 0.05, 0.1, 0.5, 1.0, 5.0\}$ in %, $\#x = 6$.

$d \in \{0, 2, 3, 4, 5\}$, $\#d = 5$.

$k \in \{5, 10, 25\}$, $\#k = 3$.

MA1: This trading rule has three parameters, q , x and d . The number of versions of *MA1* considered is therefore $\#MA1 = \#q \times \#x \times \#d = 11 \times 6 \times 5 = 330$.

MA2: This trading rule has four parameters, q , x , d and k . The number of versions of *MA2* considered is therefore $\#MA2 = \#q \times \#x \times \#d \times \#k = 11 \times 6 \times 5 \times 3 = 990$.

MA3: This trading rule has four parameters, p , q , x and d . The number of versions of *MA3* considered is therefore $\#MA3 = \#(p-q) \times \#x \times \#d = 55 \times 6 \times 5 = 1,650$.

MA4: This trading rule has five parameters, p , q , x , d and k . The number of versions of *MA4* considered is therefore $\#MA4 = \#(p-q) \times \#x \times \#d \times \#k = 55 \times 6 \times 5 \times 3 = 4,950$.

³⁶ Lequeux and Acar (1998) report that the correlation of excess returns generated by a particular parameterization of this strategy with those of the median Commodity Trading Advisor (CTA) was approximately 0.85 over their simulation period. This trading rule therefore appears to replicate well a key constituent of the investment strategy of an informed and active group trading in the foreign exchange market that is known predominantly to use technical trading rules (Sager and Taylor, 2006).

MA5: This trading rule has five parameters, n , p , q , x and d . The number of versions of MA5 considered is therefore $\#MA5 = \#(n-p-q) \times \#x \times \#d = 165 \times 6 \times 5 = 4,950$.

A.4 Support-resistance trading rules

Support-resistance trading rules attempt to identify levels of the exchange rate above which the rate appears to have difficulty rising (a resistance level) and levels below which the rate appears to have difficulty penetrating (a support level). These trading rules are based on the premise that a breach of a support or resistance level will trigger further rapid exchange rate movement in the same direction. In construction, support-resistance trading rules are similar to filter rules except that a trading signal is generated when the rate moves beyond a support or resistance level by a certain percentage, rather than beyond a recent high or a recent low.

The support and resistance levels have to be pre-specified. We define a resistance level as the highest closing rate of the j previous closing rates and we define a support level as the lowest closing rate of j previous closing rates:

SR1: If the daily closing spot exchange rate of the currency moves up at least x percent above the highest closing of the j previous closing rates and remains so for d days, go long the currency. If the daily closing spot rate moves down at least x percent below the lowest closing of the j previous closing rates and remains so for d days, go short the currency.

The pre-specified holding period version of the support-resistance rule is also analogous to the corresponding filter rule:

SR2: If the daily closing spot exchange rate of the currency moves up at least x percent above the highest closing of the j previous closing rates and remains so for d days, go long the currency for k days and then neutralize the position. If the daily closing spot rate moves down at least x percent below the lowest closing of the j previous closing rates and remains so for d days, go short the currency for k days and then neutralize the position.

The details of our parameterization of a total of 1,890 support-resistance trading rules are provided as below:

$x \in \{0.05, 0.1, 0.5, 1.0, 2.5, 5.0, 10.0\}$ in %, $\#x = 7$.

$d \in \{0, 1, 2, 3, 4, 5\}$, $\#d = 6$.

$j \in \{2, 5, 10, 15, 20, 25, 50, 100, 250\}$, $\#j = 9$.

$k \in \{1, 5, 10, 25\}$, $\#k = 4$.

SR1: This trading rule has three parameters, x , d and j . The number of versions of *SR1* considered is therefore $\#SR1 = \#x \times \#d \times \#j = 7 \times 6 \times 9 = 378$.

SR2: This trading rule has four parameters, x , d , j and k . The number of versions of *SR2* considered is therefore $\#SR2 = \#x \times \#d \times \#j \times \#k = 7 \times 6 \times 9 \times 4 = 1,512$.

A.5 Channel breakout trading rules

A trading channel for an exchange rate is perceived as occurring when the highest level of the daily closing exchange rate over a previous given period is within a given range of the

lowest level over the previous given period so that, in a sense, there are time-varying support and resistance levels that appear to be drifting together within a certain range. These time-varying support and resistance levels are the lower and upper bounds of the trading channel. Under a channel breakout rule, a trading signal is generated when a trading channel occurs and the exchange rate penetrates beyond the upper or lower bound of the channel. The assumption is that once the channel is breached, there will be sustained movement of the exchange rate in the same direction.

A $c\%$ trading channel for an exchange rate may be defined as occurring when the high level of the daily closing exchange rate over the previous j days is within $c\%$ of the low over the previous j days so that, in a sense, there are time-varying support and resistance levels that appear to be drifting together with about $c\%$ or less separation. These time-varying support and resistance levels are the lower and upper bounds of the trading channel. The upper bound of the trading channel on a particular day will be $c\%$ above the low of the previous j days and the lower bound will be $c\%$ below the high of the previous j days.

CB1: If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves up at least x percent above the upper bound of the channel and remains so for d days, go long the currency. If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves down at least x percent below the lower bound of the channel and remains so for d days, go short the currency.

And the pre-specified holding period version is:

CB2: If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves up at least x percent above the upper bound of the channel and remains so for d days, go long the currency for k days and then neutralize the position. If a $c\%$ trading channel exists and if the daily closing spot exchange rate of the currency moves down at least x percent below the lower bound of the channel and remains so for d days, go short the currency for k days and then neutralize the position..

We analyse a total of 3,000 channel breakout rules, the parameterization of which is described as below:

$x \in \{0.05, 0.1, 0.5, 1.0, 5.0\}$ in %, $\#x = 5$.

$d \in \{0, 1, 2\}$, $\#d = 3$.

$j \in \{5, 10, 15, 20, 25, 50, 100, 200\}$, $\#j = 8$.

$c \in \{0.1, 0.5, 1.0, 5.0, 10.0\}$ in %, $\#c = 5$.

$k \in \{1, 5, 10, 25\}$, $\#k = 4$.

CB1: This trading rule has four parameters, x , d , j and c . The number of versions of *CB1* considered is therefore $\#CB1 = \#x \times \#d \times \#j \times \#c = 5 \times 3 \times 8 \times 5 = 600$.

CB2: This trading rule has five parameters, x , d , j , c and k . The number of versions of *CB2* considered is therefore $\#CB2 = \#x \times \#d \times \#j \times \#c \times \#k = 5 \times 3 \times 8 \times 5 \times 4 = 2,400$.

A.6 Total technical trading rules

The total number of different calibrations of trading rules considered is therefore:

$$(\#O1+\#O2) + (\#F1+ \#F2+ \#F3) + (\#MA1+ \#MA2+ \#MA3+ \#MA4+ \#MA5) + (\#SR1+ \#SR2) + (\#CB1+ \#CB2) = 600 + 2,835 + 12,870 + 1,890 + 3,000 = 21,195.$$

Appendix B: Implementation of the Stepwise SPA test

The implementation of the stepwise SPA test with a pre-specified Type I error level α for foreign exchange rate returns in a specific sample period is as follows:

1. Compute the return matrix \mathbf{R} , in which each element R_{jt} denotes the excess return provided by each of J technical rules on each of T days ($j=1, \dots, J; t=1, \dots, T$).
2. Compute the vector of performance metrics $\boldsymbol{\Theta}=(\theta_1, \theta_2, \dots, \theta_J)$, based on \mathbf{R} .
3. Resample \mathbf{R} using the stationary bootstrap method of Politis and Romano (1994), with pre-specified parameter Q , for B times, and label each resample as \mathbf{R}_b , $b = 1, \dots, B$.³⁷
4. For each b , compute the vector of performance metrics $\boldsymbol{\Theta}_b=(\theta_{1b}, \theta_{2b}, \dots, \theta_{Jb})$, based on resample \mathbf{R}_b and set the loop indicator $n = 1$.
5. Construct an empirical null distribution for the test statistics as follows:
 - 5.1. For each b , for each of the J trading rules, compute $s_{bj} = T^{1/2} \max_{j=1, \dots, J} [\theta_{jb} - \theta_j + \boldsymbol{\iota}(T^{1/2}\theta_j \leq -\sigma_j [2\log\log(T)]^{1/2})]$, $j=1, \dots, J$, where $\boldsymbol{\iota}(E)$ denotes the indicator function of the event E and σ_j denotes the standard deviation of the original return series of the j -th technical rule.
 - 5.2. Collect all $\{s_{bj}\}_{b=1, \dots, B}$, rank them in descending order and then collect its $(1-\alpha)$ -th quantile as $q_j(\alpha)$.
6. Compare each technical rule's $T^{1/2}\theta_j$ to $q_j(\alpha_0)$, and treat the j -th null hypothesis as rejected at the n -th step if $T^{1/2}\theta_j > q_j(\alpha_0)$. Record all information of these rejected trading rules and label them as being rejected at the n -th step. Then, restart from Step 5, let $\theta_j = 0$ and $\theta_{jb} = 0$ for all rejected hypotheses j , and change the loop indicator from n to $n + 1$. However, if no technical rule is rejected given $q_j(\alpha_0)$, i.e. $T^{1/2}\theta_j \leq q_j(\alpha_0)$ for remaining j , then stop and go to Step 7.
7. Finally, restore the original θ_j from \mathbf{R} and estimate each technical rule's marginal p -value, p_j , as the percentile of $T^{1/2}\theta_j$ in the last $\{s_{bj}\}_{b=1, \dots, B}$ as an empirical null distribution.
8. Compare each technical rule's p_j to α . If $p_j < \alpha$, we claim that technical rule j is

³⁷ The stationary bootstrap of Politis and Romano (1994) involves choosing the lengths of the blocks of data according to a random number chosen from the geometric distribution. Consider a series of independent trials that each can be classified as success or failure (e.g. drawing a white ball, with replacement, from an urn containing a fixed number of white balls and a fixed number of red balls). The geometric distribution gives the probability of a random variable X taking a positive integer value k and is defined as the probability that the first occurrence of success requires exactly k trials, each with individual success probability Q (the proportion of white balls in the urn). If the probability of success in each trial is Q , then the probability that the k -th trial (out of k trials) is the first success is $Pr(X = k) = (1-Q)^{k-1}Q$, $k = 1, 2, 3, \dots$. The expected value of a geometrically distributed random variable X is $1/Q$ and the variance is $(1-Q)/Q^2$. Politis and Romano (1994) exploit these properties to ensure that the pseudo time series generated by their bootstrap method is stationary.

significantly profitable for the foreign exchange rates in the sample period at a significance level of α .

In our empirical tests, we set $\alpha = 0.10$, $Q = 0.9$, and $B = 500$ following the literature.³⁸ Moreover, we skip the observations of all currencies in the first year of their available sample periods in order to initialize some of the longer-term technical trading rules because they require up to 250 trading days to generate the first signals.

³⁸ We have also performed a range of tests based on different parameter settings and have found similar results to those reported in the text.

References

- Alexander, S. S. 1961. Price Movements in Speculative Markets: Trends or Random Walks. *Industrial Management Review* 2: 7-26.
- Allen, H. L., and M. P. Taylor. 1990. Charts, Noise and Fundamentals in the London Foreign Exchange Market. *Economic Journal* 100: 49-59.
- Allen, H. L., and M. P. Taylor. 1992. Chartist Analysis. In: Milgate, M., P. Newman, and J. Eatwell (Eds). *The New Palgrave Dictionary of Money and Finance*. MacMillan, London: 339-342.
- Azzopardi, P. V. 2010. *Behavioural Technical Analysis: An Introduction to Behavioural Finance and its Role in Technical Analysis*. Harriman House, London.
- Bajgrowicz, P., and O. Scaillet. 2012. Technical Trading Revisited: False Discoveries, Persistence Tests, and Transaction Costs. *Journal of Financial Economics* 106: 473-491.
- Bank for International Settlements. 2013. *Triennial Central Bank Survey: Foreign exchange turnover in April 2013*. Bank for International Settlements, Basel.
- Barras, L., O. Scaillet, and R. Wermers. 2010. False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas. *Journal of Finance* 65: 179-216.
- Bessembinder, H., and K. Chan. 1995. The Profitability of Technical Trading Rules in the Asian Stock Markets. *Pacific-Basin Finance Journal* 3: 257-284.
- Blume, L., D. Easley, and M. O'Hara. 1994. Market Statistics and Technical Analysis: The Role of Volume. *Journal of Finance* 49: 153-181.
- Brown, D. P., and R. H. Jennings. 1989. On Technical Analysis, *Review of Financial Studies* 2: 527-551.
- Burnside, C., M. Eichenbaum, and S. Rebelo. 2007. The Returns to Currency Speculation in Emerging Markets. *American Economic Review Papers and Proceedings* 97: 333-338.
- Chang, P. H. K., and C. L. Osler. 1999. Methodical Madness: Technical Analysis and the Irrationality of Exchange-rate Forecasts. *Economic Journal* 109: 636-661.
- Cheung, Y.-W., and M. D. Chinn. 2001. Currency Traders and Exchange Rate Dynamics: A Survey of the US market. *Journal of International Money and Finance* 20: 439-471.
- Clark, T. E., and M. W. McCracken. 2012. Reality Checks and Nested Forecast Model Comparisons. *Journal of Business and Economic Statistics* 30: 53-66.
- Cornell, W. B., and J. K. Dietrich. 1978. The Efficiency of the Market for Foreign Exchange under Floating Exchange Rates. *Review of Economics and Statistics* 60: 111-120.
- Frankel, J. A., and K. A. Froot. 1990. Chartists, Fundamentalists, and Trading in the Foreign Exchange Market. *American Economic Review, Papers and Proceedings* 80: 181-185.
- Froot, K. A., and R. H. Thaler. 1990. Anomalies: Foreign Exchange. *Journal of Economic*

Perspectives 4: 179-192.

Gencay, R. 1999. Linear, Non-linear and Essential Foreign Exchange Rate Prediction with Simple Technical Trading Rules. *Journal of International Economics* 47: 91-107.

Grossman, S.J., and J.E Stiglitz. 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review* 70: 393-408.

Hansen, P. R. 2005. A Test for Superior Predictive Ability. *Journal of Business and Economic Statistics* 23: 365-380.

Harvey, C. R., and Y. Liu. 2014. Evaluating Trading Strategies. *Journal of Portfolio Management* 40: 108-118.

Harvey, C. R., Y. Liu, and H. Zhu. 2016 and the Cross-Section of Expected Returns. *Review of Financial Studies* 29: 5-68.

Henderson, C. 2002. *Currency Strategy: The Practitioner's Guide to Currency Investing, Hedging and Forecasting*. John Wiley and Sons, Chichester.

Hodrick, R. J., and S. Srivastava. 1986. The Covariation of Risk Premiums and Expected Future Spot Exchange Rates. *Journal of International Money and Finance* 5: S5-S21.

Hsu, P.-H., Y.-C. Hsu, and C.-M. Kuan. 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., and C.-M. Kuan. 2005. Reexamining the Profitability of Technical Analysis with Data-snooping Checks. *Journal of Financial Econometrics* 3: 606-628.

Jensen, M. C., and G. A. Benington. 1970. Random Walks and Technical Theories: Some Additional Evidence. *Journal of Finance* 25: 469-482.

Kavajecz, K. A., and E. R. Odders-White. 2004. Technical Analysis and Liquidity Provision. *Review of Financial Studies* 17: 1043-1071.

Kho, B.-C. 1996. Time-varying Risk-premia, Volatility, and Technical Trading Rule Profits: Evidence from Foreign Currency Futures Markets. *Journal of Financial Economics* 41: 249-290.

Leamer, E. E. 1978. *Specification Searches: Ad Hoc Inference with Nonexperimental Data*. John Wiley and Sons, New York, NY.

Leamer, E. E. 1983. Let's Take the Con Out of Econometrics. *American Economic Review* 73: 31-43.

LeBaron, B. 1999. Technical Trading Rule Profitability and Foreign Exchange Intervention. *Journal of International Economics* 49: 125-143.

LeBaron, B. 2002. Technical Trading Profitability in Foreign Exchange Markets in the 1990s, Working Paper, Brandeis University.

- Lequeux, P., and E. Acar. 1998. A Dynamic Index for Managed Currencies Funds Using CME Currency Contracts. *European Journal of Finance* 4: 311-330.
- Levich, R. M., and L. R. Thomas. 1993. The Significance of Technical Trading-rule Profits in the Foreign Exchange Market: A Bootstrap Approach. *Journal of International Money and Finance* 12: 451-474.
- Levy, R. A. 1967. Relative Strength as a Criterion for Investment Selection. *Journal of Finance* 22: 595-610.
- Levy, R.A. 1971. The Predictive Significance of Five-point Chart Patterns. *Journal of Business* 44: 316-323.
- Lo, A. W. 2004. The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. *Journal of Portfolio Management* 30: 15-29.
- Lo, A. W., and A. C. MacKinlay. 1990. Data-snooping Biases in Tests of Financial Asset Pricing Models. *Review of Financial Studies* 3: 431-467.
- Lo, A. W., H. Mamaysky, and J. Wang. 2000. Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation. *Journal of Finance* 55: 1705-1765.
- Lyons, R. K. 2001. *The Microstructure Approach to Exchange Rates*. MIT Press, Cambridge, MA.
- McLean, R. D., and J. Pontiff. 2016. Does Academic Research Destroy Stock Return Predictability? *Journal of Finance* 71: 5-32.
- Melvin, M., and M. P. Taylor. 2009. The Crisis in the Foreign Exchange Market. *Journal of International Money and Finance* 28: 1317-1330.
- Menkhoff, L., and M. P. Taylor. 2007. The Obstinate Passion of Foreign Exchange Professionals: Technical Analysis. *Journal of Economic Literature* 45: 936-972.
- Neely, C. J. 2002. The Temporal Pattern of Trading Rule Returns and Exchange Rate Intervention: Intervention Does Not Generate Technical Trading Profits. *Journal of International Economics* 58: 211-232.
- Neely, C. J., D. E. Rapach, J. Tu, and G. Zhou. 2014. Forecasting the Equity Risk Premium: The Role of Technical Indicators, *Management Science* 60: 1772-1791.
- Neely, C. J., P. A. Weller, and R. Dittmar. 1997. Is Technical Analysis Profitable in the Foreign Exchange Market? A Genetic Programming Approach. *Journal of Financial and Quantitative Analysis* 32: 405-436.
- Neely, C. J., and P. A. Weller. 2001. Technical Analysis and Central Bank Intervention. *Journal of International Money and Finance* 20: 949-970.
- Neely, C. J., and P. A. Weller. 2003. Intraday Technical Trading in the Foreign Exchange Market. *Journal of International Money and Finance* 22: 223-237.

- Neely, C.J., and P.A. Weller, P.A. 2012. Technical Analysis in the Foreign Exchange Market. In: J. James, I.W. Marsh and L. Sarno, (Eds). *Handbook of Exchange Rates*. John Wiley, Hoboken, NJ, John Wiley: 343-374.
- Neely, C. J., and P. A. Weller. 2013. Lessons From the Evolution of Foreign Exchange Trading Strategies. *Journal of Banking & Finance*, 37: 3783-3798.
- Neely, C. J., P. A. Weller, and M. Ulrich. 2009. The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market. *Journal of Financial and Quantitative Analysis* 44: 467-488.
- Oberlechner, T., and C. L. Osler. 2012. Survival of Overconfidence in Currency Markets. *Journal of Financial and Quantitative Analysis* 47: 91-113.
- Okunev, J., and D. White. 2003. Do Momentum-based Strategies Still Work in Foreign Currency Markets? *Journal of Financial and Quantitative Analysis* 38: 425-447.
- Olson, D., 2004. Have Trading Rule Profits in the Currency Markets Declined Over Time? *Journal of Banking and Finance* 28: 85-105.
- Osler, C. L. 2003. Currency Orders and Exchange Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis. *Journal of Finance* 58: 1791-1819.
- Park, C.-H., and Irwin, S. H. 2010. A Reality Check on Technical Trading Rule Profits in the U.S. Futures Markets. *Journal of Futures Markets* 30: 633-659.
- Politis, D. N., and J. P. Romano. 1994. The Stationary Bootstrap. *Journal of the American Statistical Association* 89: 1303-1313.
- Qi, M., and Y. Wu. 2006. Technical Trading-rule Profitability, Data Snooping, and Reality Check: Evidence from the Foreign Exchange Market. *Journal of Money, Credit and Banking* 30: 2135-2158.
- Ramadorai, T. 2008. What Determines Transactions Costs in Foreign Exchange Markets? *International Journal of Finance and Economics* 13: 14-25.
- Romano, J. P., and M. Wolf. 2005. Stepwise Multiple Testing as Formalized Data Snooping, *Econometrica* 73: 1237-1282.
- Sager, M. J., and M. P. Taylor. 2006. Under the Microscope: The Structure of the Foreign Exchange Market. *International Journal of Finance and Economics* 11: 81-95.
- Sager, M. J., and M. P. Taylor. 2008. Commercially Available Order Flow Data and Exchange Rate Movements: Caveat Emptor. *Journal of Money, Credit and Banking* 40: 583-625.
- Schwert, G. W. 2003. Anomalies and Market Efficiency, In: Constantinides, G. M., M. Harris, and R. M. Stulz (Eds.) *Handbook of the Economics of Finance*. Elsevier, Amsterdam: 937-972.
- Sharpe, W. F. 1966. Mutual Fund Performance. *Journal of Business* 39: 119-138.

- Sullivan, R., A. Timmermann, and H. White. 1999. Data Snooping, Technical Trading Rule Performance, and the Bootstrap. *Journal of Finance* 54: 1647-1691.
- Sweeney, R. J. 1986. Beating the Foreign Exchange Market. *Journal of Finance* 41: 163-182.
- Taylor, M. P., and H. L. Allen. 1992. The Use of Technical Analysis in the Foreign Exchange Market. *Journal of International Money and Finance* 11: 304-314.
- Taylor, M. P. 1987. Covered Interest Parity: A High Frequency, High Quality Data Study. *Economica* 54: 429-438.
- Taylor, M. P. 1989. Covered Interest Arbitrage and Market Turbulence. *Economic Journal* 99: 376-391.
- Timmermann, A. 2008. Elusive Return Predictability. *International Journal of Forecasting* 24: 1-18.
- Timmermann, A., and C. W. J. Granger. 2004. Efficient Market Theory and Forecasting. *International Journal of Forecasting* 20: 15-27.
- Treynor, J. L., and R. Ferguson. 1985. In Defense of Technical Analysis. *Journal of Finance* 40: 757-775.
- White, H. 2000. A Reality Check for Data Snooping. *Econometrica* 68: 1097-1126.
- Wilder, J. W. Jr. 1978. *New Concepts in Technical Trading Systems*. Hunter Publishing Company, Greensboro, NC.
- Zhu, Y., and G. Zhou. 2009. Technical Analysis: An Asset Allocation Perspective on the Use of Moving Averages. *Journal of Financial Economics* 92: 519-544.

Table 1
Summary statistics of daily returns on foreign currencies and daily short rates

Countries	Gross returns on foreign currencies					Short-term interest rates					Sample period
	Mean (%)	Max	Min	Std. dev.	1 st auto.	Mean (%)	Max (%)	Min (%)	Std. dev. (%)	1 st auto.	
Developed											
Australia	-0.0041	0.1073	-0.1925	0.0069	0.001	0.0300	0.2441	0.0025	0.0160	0.968	1/4/1971-9/11/2015
Canada	-0.0023	0.0505	-0.0434	0.0040	0.010	0.0269	0.0852	0.0008	0.0174	1.000	1/4/1971-9/11/2015
Germany/E.U.	-0.0026	0.0462	-0.0421	0.0059	0.025	0.0150	0.0491	-0.0011	0.0104	1.000	1/4/1971-9/11/2015
Japan	0.0094	0.0950	-0.0626	0.0064	0.014	0.0174	0.0580	0.0000	0.0178	1.000	1/4/1971-9/11/2015
New Zealand	-0.0049	0.0995	-0.2050	0.0074	0.002	0.0303	0.1837	0.0079	0.0148	0.996	1/4/1971-9/11/2015
Norway	-0.0013	0.0646	-0.0682	0.0065	-0.011	0.0253	0.1960	0.0045	0.0149	0.982	1/4/1971-9/11/2015
Sweden	-0.0043	0.0555	-0.1507	0.0066	-0.004	0.0274	0.2986	-0.0006	0.0189	0.994	1/4/1971-9/11/2015
Switzerland	0.0129	0.1532	-0.0762	0.0075	-0.019	0.0079	0.3251	-0.0832	0.0142	0.410	1/4/1971-9/11/2015
U.K.	-0.0039	0.0467	-0.0392	0.0058	0.062	0.0252	0.0657	0.0020	0.0153	0.978	1/4/1971-9/11/2015
U.S.	-	-	-	-	-	0.0202	0.0665	0.0000	0.0135	0.999	1/4/1971-9/11/2015
Emerging											
Argentina	-0.0360	0.1712	-0.3418	0.0088	0.029	0.0420	0.5920	0.0051	0.0542	0.979	4/1/1991-9/11/2015
Brazil	-0.0245	0.1178	-0.1080	0.0096	0.060	0.0689	0.3989	0.0264	0.0411	0.972	7/4/1994-9/11/2015
Chile	-0.0082	0.1114	-0.1160	0.0064	-0.052	0.0023	0.0305	0.0000	0.0025	0.574	1/3/1994-9/11/2015
Colombia	-0.0256	0.0562	-0.0508	0.0058	0.124	0.0508	0.1297	0.0132	0.0355	1.000	1/3/1992-9/11/2015
Mexico	-0.0390	0.2231	-0.2231	0.0118	-0.135	0.0642	0.4839	0.0109	0.0624	0.995	1/1/1987-9/11/2015
India	-0.0202	0.0376	-0.1281	0.0048	-0.035	0.0311	0.2778	0.0002	0.0208	0.848	1/1/1991-9/11/2015
Indonesia	-0.0328	0.2361	-0.3576	0.0140	0.036	0.0469	0.2934	0.0000	0.0347	0.984	1/2/1986-9/11/2015
Korea	-0.0073	0.2012	-0.1809	0.0083	0.228	0.0257	0.0966	0.0057	0.0199	0.997	1/3/1992-9/11/2015
Philippines	-0.0110	0.1015	-0.0860	0.0047	0.121	0.0324	0.2802	0.0021	0.0234	0.911	1/2/1987-9/11/2015
Singapore	0.0041	0.0414	-0.0276	0.0033	-0.054	0.0086	0.0259	0.0001	0.0071	0.991	1/4/1982-9/11/2015
Taiwan	0.0023	0.0430	-0.0420	0.0029	-0.058	0.0195	0.0690	0.0012	0.0135	0.998	10/3/1983-9/11/2015
Thailand	-0.0055	0.0635	-0.2077	0.0059	-0.029	0.0167	0.0948	0.0005	0.0162	0.993	1/2/1991-9/11/2015
Czech	0.0031	0.0522	-0.0707	0.0077	-0.040	0.0181	0.3755	0.0004	0.0196	0.957	4/22/1992-9/11/2015
Hungary	-0.0205	0.0520	-0.0842	0.0084	-0.026	0.0516	0.1191	0.0078	0.0297	1.000	6/3/1991-9/11/2015
Israel	-0.0795	0.0645	-0.1725	0.0061	0.083	0.0912	0.9012	0.0001	0.1372	0.999	1/3/1978-9/11/2015
Poland	-0.0145	0.0670	-0.0715	0.0083	-0.017	0.0365	0.1132	0.0030	0.0282	0.961	6/4/1993-9/11/2015
Romania	-0.0340	0.0953	-0.1139	0.0074	0.022	0.0784	0.5238	0.0011	0.0779	0.992	1/1/1997-9/11/2015
Russia	-0.0628	0.2779	-0.3863	0.0140	0.231	0.0447	0.4583	0.0031	0.0619	0.974	9/1/1994-9/11/2015
Slovak	0.0054	0.0462	-0.1097	0.0062	0.010	0.0229	0.2954	0.0002	0.0241	0.969	4/27/1993-9/11/2015
South Africa	-0.0324	0.1440	-0.1030	0.0096	-0.001	0.0437	0.0840	0.0000	0.0162	0.988	1/2/1981-9/11/2015
Turkey	-0.1072	0.1256	-0.3348	0.0111	0.133	0.1309	1.4754	0.0186	0.0960	0.965	1/2/1990-9/11/2015

We report descriptive statistics of daily returns on holding foreign currencies against U.S. dollar and daily log short-term interest rates in all countries. Developed countries are listed in alphabetical order, and emerging countries are in alphabetical order within three cultural/geographic groups (Latin America, Asia and Europe).

Table 2**The predictive ability and excess profitability of technical trading rules**

Countries	A. Mean excess return			B. Sharpe ratio		
	# predictive rules	Highest return (%) (<i>p</i> -values)	Best rule	# profitable rules	Highest ratio (<i>p</i> -values)	Best rule
Developed						
Australia	0	6.1 (0.14)	F3	0	0.589 (0.21)	MA5
Canada	0	2.9 (0.51)	F3	0	0.529 (0.37)	MA5
Germany/E.U.	38	6.5 (0.00)**	MA1	4	0.681 (0.05)**	MA1
Japan	65	7.7 (0.01)**	CB1	31	0.752 (0.01)**	MA4
New Zealand	199	7.1 (0.04)**	F3	2	0.652 (0.07)*	MA5
Norway	0	4.8 (0.49)	MA5	0	0.516 (0.48)	MA5
Sweden	24	6.4 (0.05)**	MA4	0	0.600 (0.15)	MA4
Switzerland	1	6.6 (0.10)*	MA2	0	0.561 (0.24)	MA2
U.K.	0	5.3 (0.12)	F3	0	0.582 (0.32)	MA5
%Sig. dvlpd	56%			33%		
Emerging						
<i>Latin America</i>						
Argentina	0	5.9 (0.83)	MA5	0	0.547 (0.41)	MA5
Brazil	1	13.5 (0.09)*	SR1	7	1.097 (0.07)*	MA5
Chile	9	8.9 (0.02)**	SR1	4	0.980 (0.07)*	MA5
Colombia	427	11.6 (0.00)**	SR1	237	1.279 (0.01)**	SR1
Mexico	0	5.0 (1.00)	F1	0	0.526 (0.30)	MA5
<i>Asia</i>						
India	9	5.7 (0.06)*	F3	1	0.833 (0.10)*	O2
Indonesia	0	13.4 (0.22)	MA1	0	0.619 (0.49)	CB2
Korea	8	12.7 (0.04)**	F1	14	1.059 (0.04)**	CB2
Philippines	37	6.1 (0.03)**	SR1	6	0.825 (0.10)*	SR1
Singapore	0	2.8 (0.41)	MA2	0	0.522 (0.53)	MA2
Taiwan	2086	5.2 (0.00)**	F1	1170	1.149 (0.01)**	SR1
Thailand	0	9.6 (0.28)	MA2	0	0.934 (0.99)	F1
<i>Europe</i>						
Czech	0	7.0 (0.67)	MA4	0	0.615 (0.70)	MA5
Hungary	0	5.9 (0.94)	MA5	0	0.878 (0.13)	F2
Israel	4	5.8 (0.03)**	F3	1	0.677 (0.09)*	MA5
Poland	0	6.1 (0.96)	F3	0	0.546 (0.81)	MA5
Romania	0	7.2 (0.92)	O1	0	0.787 (0.16)	MA5
Russia	88	16.2 (0.04)**	F1	0	0.714 (0.39)	F1
Slovak	0	6.8 (0.31)	SR2	0	0.715 (0.40)	SR2
South Africa	1	9.3 (0.10)*	SR1	0	0.610 (0.29)	SR1
Turkey	0	10.0 (0.33)	F1	8	0.852 (0.05)**	MA5
%Sig. dvlpd	48%			43%		

We impose historical transaction costs in returns and examine the performance of total 21,195 technical rules over available sample periods. We implement the stepwise test to inspect if there exist technical rules that are able to provide significantly positive performance. We consider mean excess return and Sharpe ratio as two performance metrics. “#predictive rules” denotes the number of technical rules that provide significantly positive mean excess returns at the 10% significance level and “#profitable rules” denotes the number of technical rules that provide significantly positive Sharpe ratios at the 10% significance level. “Highest return/ratio” denotes the best rules’ mean excess returns and Sharpe ratios with *p*-values in parentheses. The best rules are defined as technical rules providing the highest performance metric among all trading rules in the sample period; see Appendix A for details of the various trading rules and a key to the codes (F3, MA1, etc). All mean excess returns and Sharpe ratios have been annualized. “***” and “**” denote statistical significance at the 5% and 10% levels, respectively. The sample periods for each currency are reported in Table 1.

Table 3**The annualized excess return and annualized Sharpe ratio of the best rules**

Criteria	A. Selected by		B. Selected by	
	significant highest		significant highest	
	mean excess return		Sharpe ratio	
	Annualized	Sharpe	Annualized	Sharpe
	Return (%)	Ratio	Return (%)	Ratio
Developed				
Australia	-	-	-	-
Canada	-	-	-	-
Germany/E.U.	6.5	0.652	6.2	0.681
Japan	7.7	0.734	7.4	0.752
New Zealand	7.1	0.596	3.3	0.652
Norway	-	-	-	-
Sweden	6.4	0.600	-	-
Switzerland	6.6	0.561	-	-
U.K.	-	-	-	-
Emerging				
<i>Latin America</i>				
Argentina	-	-	-	-
Brazil	13.5	0.887	1.2	1.097
Chile	8.9	0.878	5.0	0.980
Colombia	11.6	1.171	1.2	1.279
Mexico	-	-	-	-
<i>Asia</i>				
India	5.7	0.764	0.9	0.833
Indonesia	-	-	-	-
Korea	12.7	0.947	6.4	1.059
Philippines	6.1	0.824	6.0	0.825
Singapore	-	-	-	-
Taiwan	5.2	0.798	4.8	1.149
Thailand	-	-	-	-
<i>Europe</i>				
Czech	-	-	-	-
Hungary	-	-	-	-
Israel	5.8	0.621	2.1	0.677
Poland	-	-	-	-
Romania	-	-	-	-
Russia	16.2	0.723	-	-
Slovak	-	-	-	-
South Africa	9.3	0.609	-	-
Turkey	-	-	4.7	0.852

We report the annualized return and the annualized daily Sharpe ratio over available sample periods. The best rules are defined as technical rules providing the highest mean excess return or Sharpe ratio among all trading rules in the whole sample period. We impose historical transaction costs in computation. The sample periods are reported in Table 1.

Table 4
The predictive ability and profitability of technical trading rules: timing criteria

Countries	A. Mean excess return (Tim)			B. Sharpe ratio (Tim)		
	# predictive rules	Highest return (%) (<i>p</i> -values)	Best rule	# predictive rules	Highest ratio (<i>p</i> -values)	Best rule
Developed						
Australia	0	6.1 (0.16)	F3	0	0.587 (0.23)	MA5
Canada	0	2.9 (0.43)	MA4	0	0.551 (0.35)	MA5
Germany/E.U.	44	6.5 (0.00)**	F3	3	0.676 (0.05)**	MA1
Japan	98	7.7 (0.01)**	SR2	16	0.754 (0.02)**	SR2
New Zealand	260	7.1 (0.02)**	MA4	2	0.659 (0.05)**	MA5
Norway	0	4.8 (0.57)	MA2	0	0.574 (0.30)	MA5
Sweden	26	6.4 (0.05)**	MA2	0	0.605 (0.18)	MA5
Switzerland	1	6.6 (0.07)*	MA2	0	0.564 (0.29)	MA2
U.K.	0	4.8 (0.32)	MA2	0	0.584 (0.30)	MA5
%Sig. dvlpd.	56%			33%		
Emerging						
<i>Latin America</i>						
Argentina	0	6.3 (0.78)	MA5	0	0.595 (0.36)	MA5
Brazil	0	13.5 (0.17)	SR1	0	0.889 (0.34)	MA5
Chile	3	8.3 (0.04)**	F1	0	0.857 (0.17)	MA5
Colombia	660	12.1 (0.00)**	SR1	422	1.347 (0.01)**	SR1
Mexico	0	5.0 (1.00)	F1	0	0.428 (0.59)	MA5
<i>Asia</i>						
India	13	5.7 (0.06)*	F3	1	0.842 (0.10)*	O2
Indonesia	0	13.4 (0.22)	MA1	0	0.702 (0.26)	MA1
Korea	8	12.7 (0.04)**	F1	7	1.060 (0.07)*	CB2
Philippines	108	6.4 (0.01)**	SR1	27	0.909 (0.04)**	SR1
Singapore	0	2.9 (0.27)	MA2	1	0.689 (0.04)**	MA5
Taiwan	2205	5.2 (0.00)**	F1	1063	1.150 (0.01)**	SR1
Thailand	0	9.7 (0.30)	MA2	0	0.937 (1.00)	F1
<i>Europe</i>						
Czech	0	7.0 (0.65)	MA4	0	0.588 (0.79)	MA5
Hungary	0	6.1 (0.91)	MA3	0	0.876 (0.13)	F2
Israel	17	6.1 (0.01)**	F3	1	0.684 (0.05)**	MA5
Poland	0	6.4 (0.87)	MA4	0	0.543 (0.82)	MA5
Romania	0	7.2 (0.99)	O1	0	0.626 (0.50)	MA5
Russia	1	16.2 (0.10)*	F1	0	0.769 (0.34)	SR1
Slovak	0	6.8 (0.47)	SR2	0	0.646 (0.56)	SR2
South Africa	1	9.3 (0.09)*	SR1	7	0.708 (0.04)**	O1
Turkey	0	10.0 (0.36)	F1	5	0.824 (0.08)*	MA5
%Sig. emrgng.	43%			43%		

We impose historical transaction costs in returns and examine the performance of total 21,195 technical rules over available sample periods. We implement the stepwise test to inspect if there exist technical rules that are able to provide significantly positive performance. We consider mean excess return and Sharpe ratio for daily returns with the tilt component removed, corresponding to equations (8) and (9) in the text, as two performance metrics. “#predictive rules” denotes the number of technical rules that provide significantly positive mean excess returns and Sharpe ratios for tilt-adjusted returns at the 10% significance level. “Highest return/ratio” denotes the best rules’ adjusted mean excess returns and adjusted Sharpe ratios with *p*-values in parentheses. The best rules are defined as technical rules providing the highest performance metric among all trading rules in the sample period; see Appendix A for details of the various trading rules and a key to the codes (F3, MA1, etc). All mean excess returns and Sharpe ratios have been annualized. “**” and “*” denote statistical significance at the 5% and 10% levels, respectively. The sample periods for each currency are reported in Table 1.

Table 5
Break-even transaction costs for predictive technical rules

Criteria	Mean excess return			Sharpe ratio			Mean estimated transaction cost
	Cost (bps)	Break-even ratio	#trades	Cost (bps)	Break-even ratio	#trades	Cost (bps)
Developed							
Australia	-	-	-	-	-	-	-
Canada	-	-	-	-	-	-	-
Germany/E.U.	19.0	9.10	825	21.9	10.49	731	2.088
Japan	57.9	42.89	295	39.2	29.04	524	1.350
New Zealand	73.5	9.51	236	770.0	99.64	24	7.728
Norway	-	-	-	-	-	-	-
Sweden	138.7	24.61	159	-	-	-	5.635
Switzerland	26.9	6.58	917	-	-	-	4.090
U.K.	-	-	-	-	-	-	-
Emerging							
Argentina	-	-	-	-	-	-	-
Brazil	1590.3	61.32	10	1630.0	62.85	15	25.935
Chile	577.0	170.16	17	1810.0	533.45	7	3.393
Colombia	86.6	18.54	159	98.8	21.15	149	4.671
Mexico	-	-	-	-	-	-	-
India	55.9	23.30	177	243.1	101.33	10	2.399
Indonesia	-	-	-	-	-	-	-
Korea	22.4	8.38	751	8.7	3.26	2500	2.672
Philippines	75.2	13.26	128	75.3	13.28	128	5.672
Singapore	-	-	-	-	-	-	-
Taiwan	10.5	2.74	734	50.0	13.05	165	3.831
Thailand	-	-	-	-	-	-	-
Czech	-	-	-	-	-	-	-
Hungary	-	-	-	-	-	-	-
Israel	75.5	3.54	300	1330.0	62.42	9	21.307
Poland	-	-	-	-	-	-	-
Romania	-	-	-	-	-	-	-
Russia	44.7	2.63	537	-	-	-	16.974
Slovak	-	-	-	-	-	-	-
South Africa	1174.7	67.45	14	-	-	-	17.394
Turkey	-	-	-	3540.0	55.04	4	64.317

We report the highest one-way break-even transaction costs (in basis points) that will reduce the performance metrics of the most predictive rules (from Tables 2 and 3) to zero. Four performance metrics are considered in this table: Mean excess return, Sharpe ratio, mean excess return (Tim), and Sharpe ratio (Tim). “#trades” denotes the number of trades triggered by each trading rule over the sample period. “-” denotes that, for given the foreign exchange rate and performance metric, there does not exist any significantly profitable technical rule. The sample periods for each currency are reported in Table 1.

Table 6

The number of technical rules with significantly positive mean excess returns in eight subsample periods

Subsample	With real historical transaction costs							
	1972 -1976	1977 -1981	1982 -1986	1987 -1991	1992 -1996	1997 -2001	2002 -2007	2008 -2015
Developed								
Australia	0	0	35	0	0	0	0	0
Canada	2	0	0	553	0	0	0	0
Germany/E.U.	0	0	112	0	0	0	0	0
Japan	29	29	0	1	0	0	0	0
New Zealand	0	31	0	0	0	0	0	0
Norway	1	0	0	5	0	0	0	0
Sweden	92	0	2	0	0	0	0	0
Switzerland	0	0	0	0	0	0	0	0
U.K.	1103	0	0	0	0	0	0	0
Avg.(developed)	136.3	6.7	16.6	62.1	0.0	0.0	0.0	0.0
%Sig. dvlpd.	56%	22%	33%	33%	0%	0%	0%	0%
Emerging								
Argentina						0	0	0
Brazil						0	0	0
Chile						676	0	0
Colombia						0	60	30
Mexico					0	0	0	0
India					0	0	171	0
Indonesia				0	225	0	0	64
Korea					0	0	1	0
Philippines				137	0	0	709	0
Singapore			0	0	0	0	0	0
Taiwan				2200	218	0	0	0
Thailand						0	4	0
Czech						0	0	0
Hungary						0	0	0
Israel			1989	0	0	0	0	0
Poland						0	0	0
Romania							0	0
Russia						1	0	0
Slovak						0	0	0
South Africa			0	0	0	0	0	0
Turkey						0	167	0
Ave. (emerging)			663.0	389.5	44.3	33.9	53.0	4.5
%Sig.(emerging)			33%	33%	20%	10%	29%	10%

This table reports the numbers of technical rules (out of a total of 21,195) that provide significantly positive mean excess returns (based on the stepwise test) over eight subsample periods: 1972-1976, 1977-1981, 1982-1986, 1987-1991, 1992-1996, 1997-2001, 2002-2007, and 2008-2015. We design the subsample periods based on historical events, including the Maastricht Treaty in 1992, Asian currency crisis in 1997, the appearance of the physical euro in 2002, and the global banking crisis of 2008. Historical transaction costs are imposed in returns for the test of profitability.

Table 7**Out-of-sample annualized mean excess return and annualized Sharpe ratio of the best significant in-sample rules**

Criteria	Selected by significant in-sample mean excess return			Selected by significant in-sample Sharpe ratio		
	Annualized Return (%)	Sharpe Ratio	Rule	Annualized Return (%)	Sharpe Ratio	Rule
Developed						
Australia	3.629	0.414	F3	-	-	
Canada	-	-		-	-	
Germany/E.U.	0.843	0.141	MA1	0.141	0.027	MA1
Japan	2.275	0.294	MA4	2.275	0.294	MA4
New Zealand	3.894	0.408	MA2	2.307	0.536	MA5
Norway	-	-		-	-	
Sweden	0.519	0.102	MA4	-	-	
Switzerland	-3.933	-0.292	MA2	-	-	
U.K.	-	-		-	-	
Emerging						
<i>Latin America</i>						
Argentina	-	-		-	-	
Brazil	-	-		7.369	0.633	MA5
Chile	7.047	0.919	SR1	9.003	1.337	MA5
Colombia	19.078	1.966	SR1	18.488	1.910	SR1
Mexico	-	-		-	-	
<i>Asia</i>						
India	8.247	0.898	F3	0.000	0.000	O2
Indonesia	-	-		-	-	
Korea	0.842	0.166	F1	0.947	0.254	CB2
Philippines	2.154	0.530	F3	2.154	0.530	F3
Singapore	-	-		-	-	
Taiwan	0.990	0.264	F1	1.698	0.471	MA2
Thailand	5.432	1.139	MA2	-	-	
<i>Europe</i>						
Czech	-	-		-	-	
Hungary	-	-		-	-	
Israel	2.465	0.398	F3	-	-	
Poland	-	-		-	-	
Romania	-	-		-1.614	-0.489	MA5
Russia	-8.598	-0.290	F1	-	-	
Slovak	-	-		-	-	
South Africa	-	-		-	-	
Turkey	-2.985	-0.268	F1	0.808	0.268	MA5


We report the out-of-sample annualized return and annualized Sharpe ratio for the period January 1, 2012 to September 11 2015, based on the best significant rules (highest mean excess return or Sharpe ratio among all trading rules) computed for the sample period between the earliest available data for each currency and December 31, 2011. See Appendix A for details of the various trading rules and a key to the codes (F3, MA1, etc). We impose historical transaction costs in computation. The earliest dates available for each currency are reported in Table 1.

Table 8**Out-of-sample performance of technical trading rule portfolios**

Developed country currency portfolio		Emerging market currency portfolio		Combined portfolio	
Annualized Return (%)	Sharpe Ratio	Annualized Return (%)	Sharpe Ratio	Annualized Return (%)	Sharpe Ratio
1.350	0.244	7.321	1.808	4.604	1.220
[0.469]		[3.4817]		[2.349]	


We report out-of-sample test annualized mean excess return and annualized Sharpe ratio for the period January 1 2012 to September 11 2015, for portfolios composed of the best rules computed from the sample period between the earliest available data for each currency and December 31, 2011. The best rules are defined as technical rules providing the highest Sharpe ratio among all trading rules in the in-sample (pre-2012) period. The rules and currencies for each portfolio are as given in Table 7, namely, Developed country currency portfolio: German mark/euro (MA1), Japanese yen (MA4), New Zealand dollar (MA5); Emerging market currency portfolio: Brazilian real (MA5), Chilean peso (MA5), Colombian peso (SR1), Indian rupee (O2), Korean won (CB2), Philippine peso (F3), Taiwanese dollar (MA2), Romanian new leu (MA5), Turkish lira (MA5); Combined portfolio: Developed country currency portfolio plus Emerging market currency portfolio. We impose historical transaction costs in computation. The earliest dates available are reported in Table 1. Figures in brackets are approximate t -ratios for the null hypothesis of zero mean excess returns in the post-sample period, computed as the square root of the sample period length (in years) times the annualized Sharpe ratio (this result is independent of the distribution of returns and follows from the Central Limit Theorem which states that whenever a random sample of size T is taken from a distribution with mean μ and variance σ^2 , then the sample mean will be approximately normally distributed with mean μ and variance σ^2/T).

Figure 1: The correlation between mean excess return and standard deviation (currencies with significant highest mean excess return)




We report scatterplot between mean excess return and annualized standard deviation for currencies with significant highest mean excess return. The regression line is equation (14). Standard deviation data is from Table 1. Data on significant highest mean excess returns is taken from Table 3.

Figure 2: The compounded cumulative excess return of a portfolio of best Sharpe ratio rules in the in-sample and out-of-sample periods: developed currencies




The compounded cumulative return to \$1 invested on January 4 1974 in an equally-weighted portfolio using the in-sample best Sharpe ratio rules of profitable developed country currencies. The best rules are defined as technical rules providing the highest Sharpe ratios among all trading rules in the whole in-sample period. We impose estimated historical transaction costs. The returns on the right side of the vertical line illustrate out-of-sample performance.

Figure 3: The compounded cumulative excess return of a portfolio of best Sharpe ratio rules in the in-sample and out-of-sample periods: emerging market currencies



The compounded cumulative return to \$1 invested on September 1 1997 in an equally-weighted portfolio using the in-sample significant best Sharpe ratio rules of emerging market currencies. The best rules are defined as technical rules providing the highest Sharpe ratios among all trading rules in the whole in-sample period. We impose estimated historical transaction costs. The returns on the right side of the vertical line illustrate out-of-sample performance.

Figure 4: The compounded cumulative excess return of a portfolio of best Sharpe ratio rules in the in-sample and out-of-sample periods: combined portfolio



The compounded cumulative return to \$1 invested in an equally-weighted portfolio using the in-sample significant best Sharpe ratio rules of emerging markets currencies and developed country currencies. The best rules are defined as technical rules providing the highest Sharpe ratios among all trading rules in the whole in-sample period. We impose estimated historical transaction costs. The returns on the right side of the vertical line illustrate out-of-sample performance.