



Research Division
Federal Reserve Bank of St. Louis
Working Paper Series



**Lessons from the Evolution of Foreign
Exchange Trading Strategies**

Christopher J. Neely
and
Paul A. Weller

Working Paper 2011-021D
<http://research.stlouisfed.org/wp/2011/2011-021.pdf>

September 2011
Revised April 2013

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

Lessons from the Evolution of Foreign Exchange Trading Strategies

Christopher J. Neely ^{a,*}, Paul A. Weller ^b

^a Federal Reserve Bank of St. Louis, St. Louis, MO, USA

^b University of Iowa, Iowa City, IA, USA

This version: April 9, 2013

Abstract

The adaptive markets hypothesis posits that trading strategies evolve as traders adapt their behavior to changing circumstances. This paper studies the evolution of trading strategies for a hypothetical trader who chooses portfolios from foreign exchange (forex) technical rules in major and emerging markets, the carry trade, and U.S. equities. The results show that a backtesting procedure to choose optimal portfolios improves upon the performance of nonadaptive rules. We also find that forex trading alone dramatically outperforms the S&P 500, with much larger Sharpe ratios over the whole sample, but there is little gain to coordinating forex and equity strategies, which explains why practitioners consider these tools separately. Forex trading returns dip significantly in the 1990s but recover by the end of the decade and have been markedly superior to an equity position since 1998. Overall, trading rule returns still exist in forex markets—with substantial stability in the types of rules—though they have migrated to emerging markets to a considerable degree.

JEL classification: F31; G14; G11; G15

Keywords: Exchange rate; Technical analysis; Technical trading; Carry trade; Efficient markets hypothesis; Adaptive markets hypothesis

*Corresponding author. Send correspondence to Chris Neely, Box 442, Federal Reserve Bank of St. Louis, St. Louis, MO 63166-0442; e-mail: neely@stls.frb.org; phone: +1-314-444-8568; fax: +1-314-444-8731. Paul Weller's email: Paul-Weller@uiowa.edu; phone: +1-319-335-0948. Christopher J. Neely is an assistant vice president and economist at the Federal Reserve Bank of St. Louis. Paul A. Weller is the John F. Murray Professor of Finance Emeritus at the University of Iowa. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks.

1. Introduction

The literature on technical analysis has established that simple technical trading rules on dollar exchange rates provided 15 years of positive, risk-adjusted returns during the 1970s and 1980s before those returns were extinguished (Levich and Thomas, 1993; LeBaron, 2002; Olson, 2004).¹ More recently, more complex and less studied rules have produced more modest returns for a similar length of time (Neely et al., 2009). Researchers have extensively investigated explanations that rely on risk adjustment and/or central bank intervention but found that these do not plausibly justify the observed excess returns produced by simple technical trading rules, nor can data mining explain the apparent profitability of technical analysis (Neely et al., 2009).

Andrew Lo's (2004) adaptive markets hypothesis (AMH) offers a plausible explanation for this technical trading puzzle, however. The AMH posits that profit opportunities will generally exist in financial markets but that learning and competition will gradually erode these opportunities as they become known. A core principle of the AMH is that traders learn over time, adapting their behavior to changing circumstances. This suggests that one should expect to see an evolution of strategies and desired investment currencies. In the context of technical trading in the foreign exchange market, a number of studies have confirmed the prediction that profits associated with particular rules will gradually decline as more traders learn about them.

But another important prediction of the AMH, that adaptive trading strategies will show superior performance to simple fixed rules, has been largely ignored. The present paper focuses on examining this prediction. Ideally, one might like to examine the evolution of technical trading strategies by directly looking at the trading records of technicians. As these data are not

¹ Menkhoff and Taylor (2007) and Neely and Weller (2012) review the literature on technical analysis in the foreign exchange market from different perspectives.

available, an alternative approach is to consider how a hypothetical trader would have adapted to changing market conditions using simple rules of thumb. Traders face a number of practical problems as they choose strategies to maximize their welfare. How to choose rules, individually or as part of a portfolio? How to combine technical rules in foreign exchange (forex) with carry trade or equity strategies? In practice, traders must make these choices by backtesting rules on existing data. In this paper we model adaptive behavior in terms of a simple backtesting procedure applied to a group of commonly studied technical and carry-trade rules in tradable currencies.² Although these rules are not necessarily the most effective or popular rules today, we prefer to consider families of commonly studied rules to maintain continuity with the previous literature and ensure that all rules would be known to traders throughout the sample.

Specifically, we investigate whether a hypothetical trader could use past performance of trading rule-currency pairs—i.e., combinations of a specific trading rule applied to a particular exchange rate—to predict future performance and construct a dynamic trading strategy superior to individual trading rules. To mimic the decision process of a forex trader, we construct a dynamic strategy as follows: We start with a pool of rule-currency pairs (including carry trades) and rank them at month t according to the Sharpe ratio over some past time window.³ We then

² Researchers have independently examined both technical trading rules and the carry trade (Brunnermeier et al., 2009; Jordà and Taylor, 2009; Farhi et al., 2009; Burnside et al., 2011a,b; Menkhoff et al., 2012a,b) and practitioners widely use both sorts of trading strategies, but researchers have done little comparison between them (Menkhoff et al., 2012b).

³ Given that none of the returns appear to have systematic risk, the Sharpe ratios allow one to easily compare performance from strategies with differing volatility. Ingersoll et al. (2007) demonstrate how a clever fund manager can dynamically manipulate his portfolio to maximize his Sharpe ratio. The manager essentially reduces (increases) the size of his investments after a successful (unsuccessful) investment run to increase the relative weight of more

form portfolios of the highest-ranked N rules and measure the return to the portfolio over month $t + 1$. Each month individual rule-currency pairs are re-ranked and the results of the ex ante ranking are allowed to determine the composition of the portfolio for the next month.

In addition, we investigate whether such a trader would benefit from an adaptive approach to diversification. Given the well-documented fact that currency trading rule returns typically display very low correlation with stock market returns, one would expect that combining equity with a dynamic currency trading strategy would substantially improve over the latter.

What does our trader learn? Backtesting works well. Past performance clearly *does* predict the future: Rule-currency pairs that are more highly ranked in backtesting have higher ex post Sharpe ratios. Indeed, the Sharpe ratio of the dynamic trading strategy is much superior to that of the S&P 500. The success of backtesting supports the prediction that an adaptive trading strategy fares better than using fixed rules. It also suggests that the positive results in the literature are not due to data mining. The backtesting methodology is fairly robust to the selection window. Both ex ante optimal and $1/N$ portfolios produce very good Sharpe ratios in every subsample, well exceeding the average of their constituent strategies. The ex ante optimal and $1/N$ forex combinations are similarly profitable over the entire sample, with no statistically significant differences in profitability when other portfolio characteristics —i.e., number of strategies, weighting—are held constant.

The research does, however, confirm a dip in the profitability of major investment currencies in the 1990s and a switch to emerging market currencies in the 1990s. In contrast, the types of

positive outcomes. The dynamic strategies studied in this paper do not change leverage over time and so the Sharpe ratios calculated here are not subject to this problem. Therefore, we focus on Sharpe ratios as our metric for rule/strategy performance.

rules chosen show few noticeable time trends, with the following exceptions: the channel rules become somewhat less frequently used over the sample and the carry trade becomes much more frequently used after the mid-1990s.

There is almost no payoff to diversifying across equities and currencies. We show that this finding is consistent with the observed levels of excess return and volatility in currency and equity markets. Given the substantially higher Sharpe ratio of the dynamic currency strategy, the equity allocation in the optimally diversified portfolio is rather small and so equity's impact on performance is also very small, even ignoring parameter uncertainty and sampling error. This lack of benefit to active diversification is consistent with the prevalence of the previously puzzling "compartmentalization" of forex and equity trading activities by practitioners.

We also find that the selection strategies do not select the bilateral carry trades in the top-ranked rules until the mid-1990s. The fact that carry trade strategies did not measure up well to the best-ranked technical rules might in part explain the almost complete lack of academic interest in the carry trade before 2006. For example, Google Scholar reports only 5 articles with the phrase "carry trade" in the title from 1990 through 2005 but 98 since 2005. We surmise that a combination of time to accumulate data, time to write articles and time to publish them explains the delay between the initial success of the carry trade and publication of articles on the topic.

In studying how a trader would have learned about the properties of adaptive rules, our paper differs from the vast majority of research on technical trading. Early papers considered the profitability of simple nonadaptive (static) technical rules (e.g., Sweeney, 1986), or the statistical significance of this profitability (e.g., Levich and Thomas, 1993). Later papers evaluated more complex nonadaptive rules (Osler, 2003, 2005) or considered explanations for the profitability of nonadaptive rules, such as central bank intervention (LeBaron, 1999; Neely, 2002) or data

mining (Neely et al., 2009). Neely et al. (2009), for example, ruled out data mining as an explanation for technical rule success by examining the true, ex post out-of-sample profitability of several sets of fixed, nonadaptive rules from previous papers. Several papers have looked at time variation in the profitability of nonadaptive rules (Levich and Thomas, 1993; Neely et al., 2009).

We wish to emphasize, however, that this paper does not test the AMH. We believe that existing evidence suggests that the AMH is the most plausible explanation for the changing patterns of profitability in forex markets but we recognize that this remains a hypothesis. Rather, we examine the actions of a hypothetical trader to discover what such a trader would have learned and whether those lessons are consistent with observed patterns in the forex market.

Two studies examine trading strategies with adaptive features, although they differ from our approach in important respects. Olson (2004) dynamically selects the best moving average rule for each of 18 developed market currencies in successive 5-year periods from 1971–2000 and then tests these in successive 5-year out-of-sample periods. He finds that returns declined from the 1970s to about zero in the 1990s. Okunev and White (2003) construct momentum strategies by using moving averages to identify the strongest and weakest momentum currencies. The strategies thus switch between different currencies over time. The authors find that the returns generated by these momentum strategies appear to have been more persistent, at least until the end of their sample in 2000.

2. Methodology

We examine the performance of portfolios of technical trading rules that are rebalanced monthly by applying a performance criterion. We use a standard pool of rules that we consider representative of those that the academic literature has investigated: 7 filter rules, 3 moving

average rules, 3 momentum rules, 3 channel rules, and 1 type of carry trade rule.⁴ Although these rules are not necessarily the most sophisticated and popular rules in current use, we believe that they are appropriate for several reasons: 1) traders had knowledge of these rules over the whole sample; 2) their use allows comparisons with the previous literature; and 3) using commonly known and tested rules insulates us from the danger of rule snooping.

A filter rule generates a buy signal for a foreign currency when the exchange rate (domestic price of foreign currency) has risen by more than y percent above its most recent low. It generates a sell signal when the exchange rate has fallen by more than the same percentage from its most recent high. Thus,

$$z_t = \begin{cases} 1 & \text{if } S_t \geq n_t(1 + y) \\ -1 & \text{if } S_t \leq x_t(1 - y), \\ z_{t-1} & \text{otherwise,} \end{cases}$$

where z_t is an indicator variable that takes the value +1 for a long position and -1 for a short position. We denote the exchange rate at date t (domestic currency per unit of foreign currency) by S_t ; n_t is the most recent local minimum and x_t the most recent local maximum. The seven filter rules have filter sizes (y) of 0.005, 0.01, 0.02, 0.03, 0.04, 0.05, and 0.1.

A moving average rule generates a buy signal when a short-horizon moving average of past exchange rates crosses a long-horizon moving average from below. It generates a sell signal when the short moving average crosses the long moving average from above. We denote these rules by $vma(S, L)$, where S and L are the number of days in the short and long moving averages, respectively. The moving average rules are $vma(1, 5)$, $vma(5, 20)$, and $vma(1, 200)$. Thus,

⁴ Dooley and Shafer (1984) and Sweeney (1986) look at filter rules; Levich and Thomas (1993) look at both filter and moving average rules; Jegadeesh and Titman (1993) consider momentum rules in equities, citing Bernard (1984) on the topic; and Taylor (1994) tests channel rules, for example.

vma(1, 5) compares the current exchange rate with its 5-day moving average and records a buy (sell) signal if the exchange rate is currently above (below) its 5-day moving average.

Our momentum rules imply a long (short) position in an exchange rate when the n -day cumulative return is positive (negative). We consider windows of 5, 20 and 60 days for the momentum rules.

A channel rule counsels to buy (sell) if the price exceeds (is less than) the maximum (minimum) over the previous n days plus (minus) the band of inaction (x).⁵ Thus,

$$z_t = \begin{cases} 1 & \text{if } S_t \geq \max(S_{t-1}, S_{t-2}, \dots, S_{t-n})(1+x) \\ -1 & \text{if } S_t \leq \min(S_{t-1}, S_{t-2}, \dots, S_{t-n})(1-x), \\ z_{t-1} & \text{otherwise.} \end{cases}$$

We set n to be 5, 10, and 20, and x to be 0.001 for all rules.

Finally, we consider a bilateral carry trade rule applied to each exchange rate, as in Burnside et al. (2011a). For each currency pair, these rules take a long position in the currency with the higher overnight interest rate and a short position in the other currency.

We thus generate a pool of 17 bilateral rules applied to 21 dollar exchange rates and 19 cross-rates, which Table 1 lists. The series for the DEM was spliced with that for the EUR after January 1, 1999. For simplicity we refer to this series throughout as the EUR. The exchange rate series are added to the sample as data become available and the respective series can be realistically traded. The next Section of the paper discusses the data more fully.

⁵ We define the channel rule following Taylor (1994). Sullivan et al. (1999) instead call this rule a “support-and-resistance” rule. Sullivan et al.’s (1999) definition of the channel rule is similar to Taylor’s (1994), but the rule is conditioned on a formed channel—that is, the minimum and maximum over the last n days must be within a certain distance of each other.

We sort all currency-rule pairs with at least 250 days of data (since the beginning of the respective samples) by Sharpe ratio. There is a maximum of $(17 \times 40 =)$ 680 rules on any given day, but missing data for some exchange rates often leave fewer than half that number of currency-rule pairs. The ranking and rebalancing procedures are performed every 20 business days. Thus, the top-ranked portfolio's returns will be generated by a given trading rule applied to a particular currency for a minimum of 20 days, at which point it may (or may not) be replaced by another rule applied to the same or a different currency.

In any study of trading performance—especially when using exotic currencies—it is important to pay close attention to transaction costs. Rules and strategies that may appear to be profitable when such costs are ignored turn out not to be attractive once the appropriate adjustments have been made. The impact of transaction costs depends both on their magnitude and on the frequency with which positions are changed. For example, in research on intraday technical trading strategies Neely and Weller (2003) found that realistic transaction costs eliminated very high gross excess returns in the case of four highly liquid currencies, the German mark, the Japanese yen, the British pound and the Swiss franc. This result was driven by the high trading frequencies for the rules considered. The size of the spread plays a particularly important role for emerging market currencies. Burnside et al. (2007) found that bid-ask spreads for emerging market currencies over the period 1997 to 2006 were between two and four times as large as those for developed market currencies. Thus using the same transaction cost for all currencies will exaggerate the relative profitability of trading in emerging market currencies.

In order to account for variation in transaction costs both over currencies and over time we used Bloomberg data on one-month forward bid-ask spreads as the basis for estimating transaction costs. Correspondence with several foreign exchange traders and with the head of the

foreign exchange department of a commercial bank led us to believe, however, that the quoted spreads appear to substantially overestimate the spreads actually available to traders. After comparing spreads from Bloomberg with those on actual trader's screens and then discussing the size of spreads with traders, we concluded that actual spreads were roughly one third of the quoted spreads. Therefore, we calculated transaction costs as follows. Before the spread data from Bloomberg were available (December 1995) the cost of a one-way trade for advanced countries (UK, Germany, Switzerland, Australia, Canada, Sweden, Norway, New Zealand and Japan) was set at 5 basis points in the 1970s, 4 basis points in the 1980s and 3 basis points in the 1990s. For all other countries it was set at one third of the average of the first 500 bid-ask observations.⁶ Once Bloomberg data become available, we use the figure of one third of the quoted one-month forward spread. Deliverable forwards are available for all countries but Russia, Brazil, Peru, Chile and Taiwan, for which we have only non-deliverable forwards. For cross-rate transaction costs, we use the maximum of the two transaction costs against the dollar. We use a minimum of one basis point transaction cost for all currencies. Figure 1 shows the estimated transaction costs for each currency over time. The greater magnitude and volatility of these costs for emerging market currencies is readily apparent.

3. Data

Table 1 shows the complete set of exchange rates that were used, as well as the starting and ending dates for which they were available to trade in our sample. All exchange rates are from the Haver daily or intdaily databases. The original source for most of the exchange rates is the Board of Governors of the Federal Reserve System statistical release H.10 (Foreign Exchange

⁶ The costs during the 1970s and 1980s are consistent with triangular arbitrage estimates originally done by Frenkel and Levich (1975, 1977) and McCormick (1979), and used by Sweeney (1986) and Levich and Thomas (1993).

Rates), but some emerging market exchange rates are from the Wall Street Journal.⁷ The HUF/CHF and ILS/EUR rates are originally sourced from the National Bank of Hungary and Financial Times, respectively. The DEM/USD return series was spliced with the EUR/USD return series at the date of the introduction of the euro, January 1, 1999.

We take a conservative view of the periods in which emerging markets currencies can be traded. To avoid periods in which capital controls or market disruption would have prevented actual trading, we restricted the start of simulated trading for a number of currencies: the South African rand (April 1, 1995), Brazilian real (May 1, 1999), Mexican peso (January 1, 1996), New Zealand dollar (August 1, 1987), Turkish lira (January 1, 2002), Peruvian nuevo sol (April 1, 1996), Israeli shekel (January 1, 1995) and Taiwanese dollar (January 1, 1998).⁸ The Bank for International Settlements (BIS) provided most of the interest rate data, which were mostly overnight money market rates. For several countries, overnight interbank or money market interest rate series were obtained from their central banks: Australia, Europe, Russia, the United States, and the United Kingdom. Japan's interest rate was constructed by splicing three series: one from the Bank of Japan and two from the BIS. Swiss and Japanese interest rate data exhibited a few (small) negative values early in the data and in the most recent period. We set these interest rate observations to zero for return calculations.

4. The performance criterion

We now turn to the measure of excess return, which is the performance criterion we use in

⁷ Exchange rates in the H.10 release are quoted at noon ET, while the Wall Street Journal reports prices at New York close.

⁸ A dual exchange rate system was in operation for the rand until March 1995 (Farrell and Todani, 2004). De Zwart et al. (2009) provide information on the tradability of these currencies.

conjunction with the Sharpe ratio for both technical trading rules and the carry trade. We first distinguish between technical trading “rules” and technical trading “strategies.” Examples of a technical trading *rule* are a 1% filter applied to the Japanese yen or a moving average rule vma(5, 20) applied to the Swiss franc. A technical trading *strategy* uses some selection criterion to switch between individual rule-currency pairs.

The rules/strategies we consider switch between long and short positions in the domestic and foreign currencies. If a trading rule signals a long position in the foreign currency at date t , the trader borrows the domestic currency at the domestic interest rate, converts it to foreign currency at the exchange rate for date t and earns the foreign overnight rate. We denote the domestic (foreign) overnight interest rate by i_t (i_t^*). Then the excess return, R_{t+1} , to a long position in foreign currency is given by

$$R_{t+1} = \frac{S_{t+1} (1+i_t^*)}{S_t (1+i_t)}. \quad (1)$$

We denote the continuously compounded (log) excess return by $z_t r_{t+1}$, where z_t is an indicator variable taking the value +1 for a long position and -1 for a short position, and r_{t+1} is defined as

$$r_{t+1} = \ln S_{t+1} - \ln S_t + \ln(1+i_t^*) - \ln(1+i_t). \quad (2)$$

The cumulative excess return from a single round-trip trade (go long at date t , go short at date $t+k$), with one-way proportional transaction cost c_t , is

$$r_{t,t+k} = \sum_{i=1}^k r_{t+i} + \ln(1-c_{t+k}) - \ln(1+c_t) \quad (3)$$

Note that a trading strategy may incur transaction costs even when individual trading rules do not, and conversely. This will happen if a strategy requires a switch between two rules holding different positions but the rules themselves signal no change of position. In this case, the strategy incurs a transaction cost but the individual rules do not. If, on the other hand, a strategy dictates a

switch from a rule requiring—let us say, a long position at time t to a different rule requiring a long position in the same currency at time $t + 1$ —then no transaction cost is incurred, even though one or both individual rules may have signaled a change of position from time t to $t + 1$.

5. Results

5.1. Average rule performance

As a benchmark for comparison, Table 2 presents the average performance of all rules by individual currency. That is, for each exchange rate, we construct an equally weighted portfolio consisting of the 17 bilateral rules over the available data. For most currencies the net annual returns are modestly positive—in the range of 0 to 5%—but eight are negative. About a quarter of the exchange rates produce statistically significant positive net returns. The mean Sharpe ratio over all rules and exchange rates is 0.17.⁹ Average trading frequency is modest, ranging from about 11 to 21 trades per year.

5.2. Ex ante strategy performance

Of course, choosing an almost-random group of trading rules and currencies would not be a sensible trading strategy. Some rules may consistently outperform others or the level of performance may vary, with certain rules doing well for a while and then declining. In practice, traders seek to exploit such patterns by choosing rules that “backtest” well. In other words, traders choose rules on the basis of past performance. To emulate this behavior, we construct ex ante portfolios with the simple procedure described in Section 1. After an initial period of 500 business days, we commence the following selection procedure each month (20-day period). We rank all rules according to Sharpe ratio over a selection window at the current date. We then

⁹ The statistical significance of this mean Sharpe ratio cannot directly be tested with the mean standard error of the Sharpe ratios in the next column. Table 4 displays standard errors for Sharpe ratios for various portfolios.

measure the performance of N ranked strategies over the next month in an out-of-sample test. To investigate the impact of time variation on rule profitability, we investigated three lengths of selection windows: the full available sample and the 1000- and 500-observation periods prior to the portfolio construction date. In the interests of brevity, we will only present results from the selection window using all ex ante data but will discuss differences with the shorter selection windows when appropriate. We emphasize that all portfolios are constructed with only ex ante information, thus ensuring that traders could have implemented the strategies. Having measured and ranked the N rules by their past performance each month, we then label portfolio strategies according to the rank, n , of the rule. Thus the strategy corresponding to $n = 1$ selects the top-ranked rule every 20 days. The strategy corresponding to $n = 2$ selects the second-ranked rule every 20 days, and so on. Thus, strategies with small values of n will switch between rules that have had relatively high Sharpe ratios over previous data. The composition of these ex ante strategies will vary with the profitability of rule-exchange rate pairs over time, as markets gradually adapt and agents arbitrage away previously profitable trading opportunities.

Table 3 details the performance of the top 10 ex ante strategies. Thus, portfolio 1 describes the performance of the strategy for which trades are determined each period by the signals of the top-ranked rule. Portfolio 2 describes the performance of the strategy using the signals of the second-ranked rule, and so on. Over the full out-of-sample sample period (April 1975–December 2012), the best ex ante strategy earns a gross annual excess return of 10.08%. Since the strategy trades 13.76 times a year, transaction costs lower the gross return to a net return of 9.40%. The associated Sharpe ratio is a very healthy 0.78. Figure 2, which plots the Sharpe ratios for the top 662 ranked strategies, reveals that higher-ranked strategies tend to have better net excess returns and Sharpe ratios. A graph of net excess returns by portfolio rank is almost identical. As rank

declines, return also declines and becomes more volatile across ranks; this supports the hypothesis that the ranking and selection procedures do indeed improve performance.

Figure 3 illustrates a striking pattern of trade frequency across rank. The top-ranked strategies have the lowest trade frequency, with portfolio 1 trading only 13.76 times a year. Trade frequency rises to reach a local plateau of almost 35 trades for strategies 100 to 230 and then declines to a lower plateau until about strategy 600, when it rises sharply again. Note that the trading strategies almost always trade more than individual rules (see Table 2) because of changes made at rebalancing periods when the strategy often switches rules/positions. The use of filter rules probably influences the peaked pattern in trade frequency. As in Neely et al. (2009), filter rules of intermediate size generate the highest excess returns. These rules trade less frequently than small filters and more frequently than large ones. In addition small filter rules outperform large ones. These facts partially explain the pattern in Figure 3.

We next consider the performance of the strategies over time. Figure 4 shows the net annual excess return over time for the top 5 strategies and for the corresponding 1/N portfolio. The consistent profitability until the early 1990s emerges clearly, as does the more recent improvement in performance. The first conclusion we can draw from these results is that although a strategy of switching between rules and currencies may mitigate the 1980-1995 decline in profitability of individual rules, it does not eliminate it. The second conclusion is that profitability returns in the late 1990s. The portfolio of the top 5 ranked strategies has positive Sharpe ratios in 12 of the 16 years from 1997 through 2012 and its average during that period is a very respectable 0.78. Third, the portfolio provides clear diversification benefits. The annual standard deviations of the top 5 individual strategies over the whole sample ranged from 9.79 to 11.42 percent but the annual standard deviation of the portfolio was only 6.91 percent.

5.3. Currency portfolios and diversification

A stylized fact in the literature on technical trading in currency markets is that returns to individual rules and portfolios of rules are uncorrelated with stock returns (e.g., Neely et al., 1997; Neely and Weller, 1999). Therefore, one would expect significant diversification benefits from combining the returns from a technical trading strategy and a stock market index. One possible approach is to consider the performance of an ex ante optimally weighted portfolio for a mean-variance investor. However, DeMiguel et al. (2009) argue that the naïve $1/N$ allocation rule is more robust and outperforms the optimally weighted portfolio in the context of stock portfolios because means and covariances of returns are imprecisely estimated. This issue has not been investigated in the context of forex rates, however. It is therefore of interest to be able to compare the performance of naïve and optimal portfolios of rules.

We form ex ante optimal portfolios as follows. At each date t , we choose the ex ante best N ($N = 10$ and $N = 50$) individual rules according to their Sharpe ratios. We calculate the mean annual excess return and the covariance matrix of the returns to these forex rules and the S&P 500 long position over the previous 500 observations. (Note that this is not the same as the covariance matrix of the trading strategy returns because the identities of the rules making up the strategy change over time.) So, for example, if $N = 2$ and the best 2 rules according to the selection criterion at time t are “GBP filter 0.005” and “CHF vma(1,5),” then we calculate the mean and covariance matrix for those 2 rule-currency pairs and the buy-and-hold equity position over the previous 500 observations. Denoting the covariance matrix by V_t and the mean return by μ_t , we obtain portfolio weights

$$w_t = V_t^{-1} \mu_t. \quad (4)$$

We set negative weights to zero and scale the weight vector to sum to 1. If the non-negativity constraint is not binding, then these weights maximize the Sharpe ratio of a portfolio consisting of the N rules and the buy-and-hold equity position. Next we compute the return to a portfolio consisting of the N forex rules and the equity position with optimal weights over period $t + 1$. We also construct a naïve portfolio consisting of equal weights attached to each of the N rules and the equity position. We also consider treating the equity position separately from the forex strategies in two ways: We consider a 50-50 split between the $1/N$ portfolio of forex strategies and the equity position; and we consider a 50-50 split between ex ante optimal portfolio of forex strategies and the equity position.

We construct 12 different portfolios that vary according to (1) whether they use the top 10 or 50 trading strategies; (2) how the forex strategies are combined with each other and with equity. Table 4 displays the results for these 12 portfolios. For ease of reference, we label the various portfolios as follows:

- $1/(N+1)$ (naïve) weights on each of N currency strategies and S&P 500 NE
- $1/(2N)$ weights on each of N currency strategies and $\frac{1}{2}$ weight on the S&P 500 NH
- $1/N$ weights on each of N currency strategies and zero weight on the S&P 500 NZ
- Optimal weights on each of N currency strategies and the S&P 500 OE
- $\frac{1}{2}$ optimal weights on each of N currency strategies and $\frac{1}{2}$ weight on the S&P 500 OH
- Optimal weights on each of N currency strategies and zero weight on the S&P 500 OZ

To distinguish between portfolios with 10 and 50 foreign exchange strategies, we write, for example, NE-10 or NE-50.

Both the portfolios with $N = 10$ and $N = 50$ perform very well in almost all subsamples. Table 4 shows that the portfolios NE and OE have Sharpe ratios ranging from 0.72 to 0.94 over

the whole sample period (1975–2012). The OE-10 and OE-50 portfolios typically have fairly similar performances for both the full sample and for the three subsamples. Over the full sample, OE-10 has a Sharpe ratio of 0.72 compared with a value of 0.89 for OE-50. However, there is no evidence of significant diversification benefit from combining the currency portfolio strategies with equity. The overall performance of OZ and NZ portfolios is not markedly different from that of the OE and NE portfolios, though the differences between like combinations, e.g., OZ-10 vs. OE-10, are sometimes significant because the series are highly positively correlated. These facts strongly indicate that the high Sharpe ratios are attributable entirely to the currency portfolio strategies. The largest difference in OE/OZ or NE/NZ Sharpe ratios occurs during the bull market subsample of 1988–99, when the OE-10 portfolio, with a Sharpe ratio of 0.28, significantly outperformed the OZ-10 portfolio, which had a Sharpe ratio of 0.10.

The absence of any marked diversification benefit from combining the currency portfolios with equity might appear surprising in light of the fact that they show slightly negative correlations. The top 10 forex strategies have daily correlations between -0.04 and 0.01 with the S&P 500 total return series over the full sample and 7 of these 10 correlations are negative. Nonetheless the lack of diversification benefit is perfectly consistent with the measured levels of return and volatility. Over the full sample, the net returns to equity and the dynamic trading strategy OZ-10 are 6.04 and 4.20 percent respectively, but the Sharpe ratios are 0.37 and 0.66 because the forex returns are much less volatile.¹⁰ The standard deviation of annualized net returns of the S&P 500 is 17.43 percent, whereas for OZ-10 it is only 6.05 percent.

To illustrate how such numbers translate into portfolio weights, consider an example in

¹⁰ Serban (2010) notes the superiority of Sharpe ratios from a forex strategy that combines momentum and mean-reversion elements with an equity position.

which equity and the dynamic strategy earn the same annual return of 5%, and the annual standard deviations of the equity portfolio and dynamic strategy are 15% and 5%, respectively. If the two return series are uncorrelated, then the optimal equity portfolio weight is 0.1. However, the Sharpe ratio of the optimally diversified portfolio is only 5.4% higher than that of the low-volatility dynamic strategy return. If we were to adopt a Bayesian perspective to account for parameter uncertainty, the improvement from diversification would be even smaller. The intuition for the very marginal benefit from diversification is as follows: Excess returns for the two investment strategies are fairly similar, whereas Sharpe ratios are dramatically different because equity returns are much more volatile than currency returns. This means that there are only very modest benefits to diversification even when the two return series are uncorrelated.

Whether or not one finds benefits to diversification depends on the choice of baseline portfolio. Levich and Pojarliev (2011) report that investors with a global equity exposure gain significant diversification by adding returns generated by currency managers. This is certainly what we find for a baseline S&P 500 portfolio. Our result is stronger in that it says that there is no advantage to adding equity exposure to our adaptive forex trading strategies.

Another result of interest is that the OZ and NZ portfolios substantially outperform equity alone. The last panel of Table 4 shows that the Sharpe ratio of the S&P 500 over the full sample is 0.37, whereas OZ-10, OZ-50, NZ-10 and NZ-50 have Sharpe ratios of 0.66, 0.89, 0.80 and 0.92. The latter three Sharpe ratios are significantly higher than that of the S&P 500. Only over the strong (mostly) bull market of 1988–99 does equity outperform the OZ and NZ portfolios. In the other two samples, the currency portfolios clearly outperform equity. For example, over the last 13 years (2000–2012) the OZ-10 (50) Sharpe ratio is 0.54 (0.81) but the ratio for the S&P 500 is only 0.09. The differences in the first subsample are statistically significant but those in

the latter two subsamples are generally not. Consistent with the fact that the currency strategies outperform equity, optimal and 1/N combinations of the forex strategies with equity (OE and NE) produce significantly higher Sharpe ratios, at the 10 percent level, than the 50-50 portfolios (OH and NH).

In contrast to results in equity markets, there is little evidence to suggest that naïve (1/N) portfolios of forex trading strategies outperform optimal portfolios in terms of Sharpe ratios. That is, the average Sharpe ratio produced by the NE, NH, and NZ portfolios is only modestly higher than the average Sharpe ratio produced by the OE, OH, and OZ portfolios. The average full sample improvements in Sharpe ratios for naïve portfolios over the optimal portfolios are only about 0.1 and 0.03 for the 10- and 50-strategy portfolios, respectively, and none of these differences are statistically significant.

Figure 5 shows the time series of rolling Sharpe ratios for several of the top 10 strategy portfolios, both with and without equity, as well as the rolling Sharpe ratio to a buy-and-hold position in the S&P 500. The top (center) panel displays 1-year rolling Sharpe ratios from the OE-10 and OZ-10 (NE-10 and NZ-10) strategy portfolios from 1976 to 2012. Contrary to the general perception in the literature, forex technical trading rules tend to perform at least as well from 2000–12 as from 1990–1999, although the differences are not statistically significant. The bottom panel of Figure 5 displays the 1-year rolling Sharpe ratios to the S&P 500. The ratios are quite variable and show no obvious trend.

To investigate the effect of emerging market currencies on the recovery in profitability after 1997, we redid the strategy selection exercise with only currencies from developed countries. Table 5 shows that when only non-emerging market currencies are used, about half of the 10-strategy and 50-strategy portfolios earn negative excess returns in the final sample (2000-2012)

and none of the portfolios earn statistically significant positive Sharpe ratios. This result is consistent with the literature and the results of Pukthuanthong-Le et al. (2007) and Pukthuanthong-Le and Thomas (2008), who find that emerging market currencies appear to provide profit opportunities to technical rules.

5.4. Currency portfolio composition

Our findings support the view that traders could have used backtesting to improve on the performance of individual trading rules by switching rule/currency compositions. In other words, rules can be reliably ranked according to expected future performance, and these rankings change over time (see Figure 2). How does the composition of the portfolio strategy vary over time? Table 6 presents the frequency with which different rules appeared in the top 5 ranked portfolios. The carry trade applied to the TRY was the overall “winner” in that it was used 14.3% of the time in the top-ranked portfolio. This is a striking illustration of the importance of the carry trade, since the TRY did not start trading in our sample until 2002. The CLP carry trade was the next most frequently used rule in the top portfolio, with a frequency of 11.2%.

Moving average, filter, momentum, channel rules, and the carry trade all appear among the most-used rules in the top portfolio, and both developed and emerging market currencies are represented. However, the analysis for the full sample masks substantial variation across subsamples. Some of this variation is driven mechanically by the fact that data for some emerging markets are either not available or cannot be used for certain (earlier) periods because of the presence of capital controls or other restrictions on market activity. Table 7 reproduces the information for the top-ranked portfolio divided into four distinct subperiods. The GBP ch(10) rule during the first subperiod (1973-82) was dominant; it was used 45.9% of the time. The next most frequently used rules were the CAD/GBP mom(20) and EUR ch(10), which were used 18.4

and 15.3% of the time, respectively. The EUR ch(10) rule continued to be frequently used in the second subperiod (1983-92), where it was used 24.8% of the time. It is not until the third subperiod (1993–2002) that emerging market currencies start to acquire a more prominent role. The CLP carry trade and CLP 0.02 filter are used 41.2% and 10.7% of the time respectively. Over the most recent subperiod (2003 – 2012) rules applied to emerging market currencies become completely dominant.

Figure 6 shows the frequency with which various rule classes appeared in the top-ranked portfolios. The top panel depicts the proportion of appearances in the top 10 ex ante trading strategies of rules from each group k :

$$s_k = \frac{\sum_{j=1}^{10} \sum_{i \in k} n_{ij}}{10 * N} \quad (5)$$

where n_{ij} is the number of days rule i was the top j th strategy and N is the total number of days in the given sample. The s_k should sum to one. The bottom panel adjusts for the fact that larger rule groups would have a better chance of being represented in the top 10 trading strategies. The adjusted rule score (\tilde{s}_k) controls for the size of the groups by dividing s_k by the number of rules in each group (n_k) and then normalizing the results to sum to one.

$$\tilde{s}_k = \left(\frac{s_k}{n_k} \right) / \left(\sum_k \frac{s_k}{n_k} \right) \quad (6)$$

Over the whole sample, channel rules dominate, whereas from the mid-1990s the carry trade takes the top spot until near the end of the sample when there is a striking decline in its use. This suggests that the effects of learning and competition might have already come into play.

The rule group prevalence seems to be reasonably stable over time with a few caveats. First, the channel rules, momentum rules and MA rules tend to decline in importance toward the end of the 1990s, recovering more recently. Second, the small filter rules and the large filter rules each

have a brief upsurge of frequency—peaking in 1990 and 2002, respectively—before declining again. Third, and most remarkably, the carry trade is unimportant until the mid-1990s.

How frequently were different types of exchange rates used in the best 10 strategies over time? To summarize the prevalence of exchange rates in the best strategies over time, we divide the currencies into 5 currency groups, shown in Table 1. The advanced market exchange rates consist of the AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD, and SEK; developing Europe consists of the CZK, HUF, PLN, RUB, TRY and HUF/CHF; the Latin American group consists of the BRL, CLP, MXN, PEN and the JPY/MXN; the Other group consists of ILS, TWD, ZAR and ILS/EUR; and the Advanced cross rates consist of non-emerging cross rates. Figure 7 shows the prevalence of exchange rate groups in the top 10 trading strategies. The lower panel of Figure 7 adjusts the frequency of each group’s representation by dividing by the number of exchange rates in the group and then normalizing the frequencies to sum to 1. It is constructed similarly to the lower pane of Figure 6, as described in (6).

Exchange rates from advanced economies dominate the top 10 ex ante trading strategies in the early part of the sample because there were no developing currencies in our data sample before the early to mid-1990s. Consistent with Lee and Mathur (1996), cross rates tend to be used with lower frequency in the top trading strategies throughout the whole sample. In the late 1990s, currencies from Latin America began to dominate the top 10 ex ante strategies and have maintained that position until very recently.¹¹

¹¹ Lee, Gleason and Mathur (2001) discern mixed results for MA and channel rules for 13 Latin American exchange rates using an earlier sample.

6. Discussion and conclusion

The “efficient markets hypothesis” holds that no trading strategy should be able to generate unusual profits based on publicly available information—such as past prices—except by bearing unusual risk. Previous research has established that the standard approach to risk adjustment using the CAPM cannot explain the observed positive excess returns to technical trading in currency markets. This is a consequence of the very low and sometimes negative correlation between returns to technical trading rules and stock market returns. The long-term profitability of technical strategies in the forex market suggests that the adaptive markets hypothesis would better describe market functioning. Adaptive behavior allows for the possibility that profit opportunities persist for considerable periods of time. Eventually, however, traders learn about these opportunities and compete them away. A number of studies of the forex market have confirmed this prediction. However, researchers have paid little attention to the distinct question of whether an adaptive trading strategy can outperform a nonadaptive strategy. Previous research has very largely focused on nonadaptive strategies, namely fixed trading rules or fixed portfolios of these rules. The contribution of this paper is to examine the performance of explicitly adaptive trading strategies and to compare them to nonadaptive strategies.

We draw several conclusions from our analysis. First, a portfolio trading strategy that switches between different rule-currency pairs according to past Sharpe ratios improves substantially on the average performance of the rule-currency pairs (Figure 2). That is, backtesting is an effective adaptive strategy because rule-currency performance is persistent. Second, there are benefits to diversifying among forex trading strategies: The optimal and 1/N currency portfolio strategies (OZ-10, OZ-50, NZ-10 and NZ-50) clearly outperform almost all strategies based on using a single currency rule at a time. They also turn out to be very

significantly superior to a pure equity portfolio (S&P 500) in terms of Sharpe ratios (Table 4). But the portfolio strategy optimally combined with equity generally does not markedly improve on the portfolio strategy on its own. The naïve strategies that combine portfolios split evenly between equity and a currency strategy (OH and NH) are generally inferior to the currency-only portfolio strategies (OZ and NZ) and the results are statistically significant for the 50-rule portfolios. The lack of a diversification benefit may help to explain why firms typically treat their forex and equity positions separately. There is little or no advantage to be gained from coordinating them.

Although the performance of the currency portfolio strategies has fluctuated, with a noticeable dip in the 1990s, Sharpe ratios have rebounded over the most recent decade (Figures 4 and 5). This observation sharply contrasts with the evidence from other studies that the profitability of individual technical trading rules had disappeared by the early 1990s. It lends support to the prediction of the Adaptive Markets Hypothesis that adaptive strategies will outperform nonadaptive strategies. The rebound in optimal rule profitability since 1998 coincides with a strong shift in the optimal strategies away from major currencies to emerging markets, first in Latin America in the late 1990s and then recently to developing Europe (Figure 7).

The types of rules used by the optimal rule portfolios are fairly stable over time (Figure 6). Channel rules, momentum rules and MA rules decline somewhat in importance after the mid-1980s and small and large filter rules each become more important for a time before declining again. The most interesting change, however, is that the carry trade becomes prominent only after 1995. This shortly predates a surge in academic and practitioner interest in carry-trade rules. The relatively poor performance of the carry trade compared with the best technical strategies prior to 1999 might explain the dearth of interest in the carry trade until recently.

Acknowledgements

The authors thank an anonymous referee and participants at presentations at the Federal Reserve Bank of St. Louis, Colorado State University Department of Finance, the Midwest Finance Association, Rutgers University Department of Economics, the SNDE meetings and the Society for Quantitative Finance for helpful comments and Brett Fawley for excellent research assistance. The authors are responsible for any errors.

References

- Bernard, A., 1984. How to Use the Value Line Investment Survey: A Subscriber's Guide. Value Line, New York.
- Brunnermeier, M.K., Nagel, S., Pedersen, L.H., 2009. Carry trades and currency crashes. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2008. University of Chicago Press, Chicago.
- Burnside, A.C., Eichenbaum, M.S., Rebelo, S., 2007. The returns to currency speculation in emerging markets. *American Economic Review* 97, 333-338.
- Burnside, A.C., Eichenbaum, M.S., Kleshchelski, I., Rebelo, S., 2011a. Do peso problems explain the returns to the carry trade? *Review of Financial Studies* 24, 853–891.
- Burnside, A.C., Eichenbaum, M.S., Rebelo, S., 2011b. Carry trade and momentum in currency markets. NBER working paper 16942.
- DeMiguel, V., Garlappi, L., Raman U., 2009. Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies* 22, 1915-1953.

- De Zwart, G., Markwat, T., Swinkels, L., van Dijk, D., 2009. The economic value of fundamental and technical information in emerging currency markets. *Journal of International Money and Finance* 28, 581-604.
- Dooley, M.P., Shafer, J., 1984. Analysis of short-run exchange rate behavior: March 1973 to November 1981. In: Bigman, D., Taya, T. (Eds.), *Floating Exchange Rates and the State of World Trade Payments*. Ballinger Publishing, Cambridge, MA.
- Farhi, E., Fraiberger, S.P., Gabaix, X., Ranciere, R., Verdelhan, A., 2009. Crash risk in currency markets. NBER working paper 15062.
- Farrell, G. N., Todani K.R., 2004. Capital flows, exchange control regulations and exchange rate policy: The South African experience. South African Reserve Bank working paper.
- Frenkel, J. A., Levich, R. M., 1975. Covered interest arbitrage: Unexploited profits? *Journal of Political Economy* 83, 325-338.
- Frenkel, J. A., Levich, R. M., 1977. Transaction costs and interest arbitrage: Tranquil versus turbulent periods. *Journal of Political Economy* 85, 1209-1226.
- Ingersoll, J., Spiegel, M., Goetzmann, W., Welch, I., 2007. Portfolio performance manipulation and manipulation-proof performance. *Review of Financial Studies* 20, 1503-1546.
- Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48, 65-91.
- Jordà, Ò., Taylor, A.M., 2009. The carry trade and fundamentals: Nothing to fear but FEER itself. NBER working paper 15518.
- 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.
- Lee, Chun I., Gleason, K.C., Mathur, I., 2001. Trading rule profits in Latin American currency spot rates. *International Review of Financial Analysis* 10, 135–156.
- Lee, Chun I., Mathur, I., 1996. Trading rule profits in European currency spot cross-rates. *Journal of Banking and Finance* 20, 949–962.
- Levich, R.M., Pojarliev, M., 2011. Are all currency managers equal? *Journal of Portfolio Management* 37, 42-53.
- Levich, R.M., Thomas, L.R. III, 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.
- Lo, A.W., 2004. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management* 30, 15-29.
- McCormick, F., 1979. Covered interest arbitrage: Unexploited profits: Comment. *Journal of Political Economy* 87, 411-417.
- Menkhoff L., Sarno L., Schmeling M., Schrimpf A., 2012a. Carry trades and global foreign exchange volatility. *Journal of Finance* 67, 681-718.
- Menkhoff, L., Sarno, L., Schmeling, M., Schrimpf, A., 2012b. Currency momentum strategies. *Journal of Financial Economics* 106, 660-684.
- Menkhoff, L., Taylor, M.P., 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 rule profits. *Journal of International Economics* 58, 211–232.
- Neely, C.J., Weller, P.A., 1999. Technical trading rules in the European Monetary System. *Journal of International Money and Finance* 18, 429–458.
- Neely, C.J., Weller, P.A., 2003. Intraday technical trading in the foreign exchange market. *Journal of International Money and Finance* 22, 223-237.
- Neely, C.J., Weller, P.A., 2012. Technical analysis in the foreign exchange market. In: James, J., Marsh, I.W., Lucio Sarno, L. (Eds.), *Handbook of Exchange Rates*. John Wiley, Hoboken, NJ.
- Neely, C.J., Weller, P.A., Dittmar, R., 1997. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. *Journal of Financial and Quantitative Analysis* 32, 405-426.
- Neely, C.J., Weller, P.A., Ulrich, J.M., 2009. The adaptive markets hypothesis: Evidence from the foreign exchange market. *Journal of Financial and Quantitative Analysis* 44, 467–488.
- Okunev, J., White, D.R., 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, 1791-1820.
- Osler, C.L., 2005. Stop-loss orders and price cascades in currency markets. *Journal of International Money and Finance* 24, 219-241.

- Pukthuanthong-Le, K., Levich, R.M., Thomas, L.R. III, 2007. Do foreign exchange markets still trend? *Journal of Portfolio Management* 34, 114–118.
- Pukthuanthong-Le, K., Thomas, L.R. III, 2008. Weak-form efficiency in currency markets. *Financial Analysts Journal* 64, 31–52.
- Serban, A., 2010. Combining mean reversion and momentum trading strategies in foreign exchange markets. *Journal of Banking and Finance* 34, 2720–2727.
- Sullivan, R., Timmermann, A., White, H., 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, S.J., 1994. Trading futures using a channel rule: A study of the predictive power of technical analysis with currency examples. *Journal of Futures Markets* 14, 215–235.

Table 1

Data description

Currency Group	Country	Currency abbreviation versus the USD	# of trading obs	Trading start date	Trading end date	Mean TC	STD of Annualized FX Return
Advanced	Australia	AUD	9008	4/7/1976	12/31/2012	3.1	11.6
Advanced	Canada	CAD	9344	1/2/1975	12/31/2012	2.9	6.7
Advanced	Euro Area	EUR	9717	4/3/1973	12/31/2012	3.0	10.6
Advanced	Japan	JPY	9599	4/3/1973	12/28/2012	3.0	10.5
Advanced	New Zealand	NZD	6027	8/3/1987	12/31/2012	3.9	12.4
Advanced	Norway	NOK	6515	1/2/1986	12/31/2012	3.4	11.6
Advanced	Sweden	SEK	7278	1/3/1983	12/28/2012	3.3	11.4
Advanced	Switzerland	CHF	9697	4/3/1973	12/31/2012	3.1	12.0
Advanced	UK	GBP	9338	1/2/1975	12/31/2012	2.9	9.9
Dev. Europe	Czech Republic	CZK	5049	1/5/1993	12/31/2012	5.2	12.4
Dev. Europe	Hungary	HUF	4466	1/2/1995	12/28/2012	10.3	14.3
Dev. Europe	Hungary/Switzerland	HUF_CHF	4165	1/3/1996	12/28/2012	10.5	12.0
Dev. Europe	Poland	PLN	3918	2/24/1997	12/31/2012	7.1	14.6
Dev. Europe	Russia	RUB	3055	8/1/2000	12/28/2012	3.6	7.4
Dev. Europe	Turkey	TRY	2769	1/2/2002	12/31/2012	12.9	15.4
Latin America	Brazil	BRL	3330	5/3/1999	12/31/2012	6.0	16.8
Latin America	Chile	CLP	4359	6/1/1995	12/28/2012	5.9	9.5
Latin America	Japan/Mexico	JPY_MXN	3887	1/4/1996	12/28/2012	4.6	16.9
Latin America	Mexico	MXN	4220	1/4/1996	12/31/2012	4.6	10.5
Latin America	Peru	PEN	4252	4/1/1996	12/31/2012	5.3	5.0
Other	Israel	ILS	3750	7/20/1998	12/31/2012	8.1	7.8
Other	Israel/Euro Area	ILS_EUR	2552	1/2/2003	12/31/2012	8.5	10.2
Other	South Africa	ZAR	4394	4/3/1995	12/31/2012	8.7	16.4
Other	Taiwan	TWD	3605	1/5/1998	12/28/2012	5.0	5.3
Adv. Cross Rates	Switzerland/UK	CHF_GBP	9169	1/3/1975	12/31/2012	3.0	9.8
Adv. Cross Rates	Australia/UK	AUD_GBP	8920	4/7/1976	12/31/2012	3.2	12.4
Adv. Cross Rates	Canada/UK	CAD_GBP	9217	1/2/1975	12/31/2012	3.0	10.3
Adv. Cross Rates	Japan/UK	JPY_GBP	8982	1/2/1975	12/28/2012	3.0	12.2
Adv. Cross Rates	Euro Area/UK	EUR_GBP	9187	1/2/1975	12/31/2012	3.0	8.1
Adv. Cross Rates	Australia/Switzerland	AUD_CHF	8848	4/7/1976	12/31/2012	3.2	14.4
Adv. Cross Rates	Canada/Switzerland	CAD_CHF	9150	1/3/1975	12/31/2012	3.1	12.4
Adv. Cross Rates	Japan/Switzerland	JPY_CHF	9338	4/3/1973	12/28/2012	3.2	11.7
Adv. Cross Rates	Euro Area/Switzerland	EUR_CHF	9602	4/3/1973	12/31/2012	3.2	5.9
Adv. Cross Rates	Canada/Australia	CAD_AUD	8894	4/7/1976	12/31/2012	3.2	10.3
Adv. Cross Rates	Japan/Australia	JPY_AUD	8633	4/7/1976	12/28/2012	3.2	15.3
Adv. Cross Rates	Euro Area/Australia	EUR_AUD	8861	4/7/1976	12/31/2012	3.2	12.8
Adv. Cross Rates	Japan/Canada	JPY_CAD	8968	1/2/1975	12/28/2012	3.1	12.7
Adv. Cross Rates	Euro Area/Canada	EUR_CAD	9158	1/2/1975	12/31/2012	3.1	10.7
Adv. Cross Rates	Japan/Euro Area	JPY_EUR	9347	4/3/1973	12/28/2012	3.1	11.3
Adv. Cross Rates	New Zealand/Australia	NZD_AUD	5943	8/3/1987	12/31/2012	3.9	8.7

Notes: The table depicts the 21 exchange rates versus the USD and 19 non-USD cross rates used in our sample along with the starting and ending dates of the samples, number of trading dates, average transaction cost, and standard deviation of annualized log returns.

Table 2

Average trading rule statistics by foreign exchange rate

Currency	Gross AR	Net AR	Net AR t-statistic	Sharpe	Sharpe (s.e.)	Trades per year per rule	Observations
GBP	2.36	1.91	2.18	0.34	0.16	16.79	9337
CHF	3.50	2.95	2.81	0.46	0.16	18.09	9696
AUD	2.01	1.50	1.42	0.25	0.16	16.73	9007
CAD	0.59	0.18	0.32	0.05	0.16	14.68	9343
SEK	1.71	1.13	0.99	0.19	0.18	17.88	7277
JPY	3.53	3.07	3.46	0.53	0.15	15.98	9598
ZAR	2.27	0.42	0.19	0.05	0.25	19.89	4393
CZK	2.69	1.71	1.15	0.27	0.23	18.02	5048
RUB	3.78	3.26	2.86	0.64	0.22	12.96	3054
EUR	4.21	3.73	4.12	0.64	0.15	16.53	9716
BRL	5.25	4.03	1.44	0.42	0.26	19.48	3329
HUF	-0.28	-1.48	-1.21	-0.29	0.25	12	4465
MXN	-0.52	-1.37	-0.98	-0.25	0.27	17.72	4219
NZD	1.28	0.57	0.41	0.09	0.21	18.42	6026
NOK	0.73	0.10	0.08	0.02	0.20	18.20	6514
PLN	1.72	0.24	0.12	0.03	0.26	19.30	3917
TRY	1.75	-0.80	-0.29	-0.10	0.33	19.79	2768
PEN	0.37	-0.23	-0.36	-0.09	0.25	11.20	4251
CLP	6.07	5.14	3.97	0.90	0.21	15.27	4358
ILS	3.13	1.96	1.80	0.45	0.25	14.11	3749
TWD	1.40	0.77	1.14	0.27	0.26	12.58	3604
CHF_GBP	1.53	1.02	1.19	0.19	0.16	16.94	9168
AUD_GBP	1.43	0.86	0.79	0.13	0.16	18.28	8919
CAD_GBP	2.13	1.62	1.78	0.28	0.16	17.66	9216
JPY_GBP	2.63	2.11	1.84	0.30	0.16	17.97	8981
EUR_GBP	1.86	1.40	1.96	0.32	0.16	15.62	9186
AUD_CHF	1.59	0.98	0.73	0.12	0.17	19.11	8847
CAD_CHF	2.64	2.08	1.83	0.30	0.16	18.55	9149
JPY_CHF	2.66	2.10	2.00	0.33	0.16	17.83	9337
EUR_CHF	0.83	0.39	0.80	0.12	0.16	13.93	9601
CAD_AUD	-0.10	-0.64	-0.67	-0.12	0.18	17.43	8893
JPY_AUD	3.21	2.62	1.73	0.30	0.16	18.81	8632
EUR_AUD	2.35	1.79	1.51	0.25	0.16	17.81	8860
JPY_CAD	3.35	2.83	2.42	0.40	0.16	17.71	8967
EUR_CAD	3.41	2.90	3.09	0.51	0.16	17.23	9157
JPY_EUR	3.75	3.24	3.16	0.53	0.16	16.99	9346
NZD_AUD	-0.89	-1.59	-1.75	-0.37	0.21	17.31	5942
HUF_CHF	-2.49	-4.43	-2.95	-0.82	0.29	17.83	4164
ILS_EUR	-2.91	-4.55	-2.80	-0.94	0.34	18.70	2551
JPY_MXN	1.57	0.59	0.24	0.06	0.26	20.67	3886
Mean	1.90	1.10	1.06	0.17	0.20	17.04	6912

Notes: The table presents the annual gross and net (of transaction costs) excess return and Sharpe ratio averaged across all 17 trading rules for each currency over the full data sample. Sample periods differ by currency.

Table 3

Top 10 ex ante portfolio results

Portfolio #	Gross AR	Net AR	Net AR t-statistic	Sharpe	Sharpe (s.e.)	Trades per year
1	10.08	9.40	4.81	0.78	0.17	13.76
2	3.31	2.41	1.33	0.21	0.16	19.10
3	5.68	4.79	2.61	0.44	0.17	19.25
4	5.58	4.66	2.69	0.44	0.16	19.74
5	8.90	8.00	4.79	0.78	0.17	20.36
6	5.45	4.43	2.62	0.40	0.16	22.96
7	6.18	5.06	3.05	0.50	0.17	25.77
8	5.87	4.77	2.77	0.46	0.17	24.77
9	6.73	5.65	3.12	0.52	0.17	25.30
10	6.24	5.19	2.91	0.49	0.17	24.70
Mean	6.40	5.44	3.07	0.50	0.17	21.57

Notes: The table presents the gross annual excess return (Gross AR) and annual excess return net of transaction costs (Net AR) for the top 10 ranked ex ante portfolio strategies. The sample for the ex ante portfolios is April 1975 to December 2012.

Table 4

Portfolios of technical trading strategies and equity: Sharpe ratios

Weight on each of N FX rules	Weight on equity	Name	Top 10 ex ante rules				Top 50 ex ante rules			
			1975-2012	1975-1987	1988-1999	2000-2012	1975-2012	1975-1987	1988-1999	2000-2012
1/(2N+1)	1/(2N+1)	NE	0.86 (0.16)	1.37 (0.27)	0.40 (0.28)	0.79 (0.27)	0.94 (0.16)	1.47 (0.28)	0.55 (0.28)	0.77 (0.27)
1/(2N)	1/2	NH	0.62 (0.17)	0.76 (0.31)	0.90 (0.30)	0.30 (0.29)	0.61 (0.17)	0.68 (0.31)	1.03 (0.30)	0.24 (0.29)
1/N	zero	NZ	0.80 (0.16)	1.36 (0.27)	0.23 (0.28)	0.79 (0.26)	0.92 (0.16)	1.46 (0.28)	0.49 (0.28)	0.75 (0.26)
optimal	optimal	OE	0.72 (0.16)	1.28 (0.26)	0.28 (0.28)	0.56 (0.25)	0.89 (0.15)	1.29 (0.27)	0.56 (0.28)	0.81 (0.25)
1/2 optimal	1/2	OH	0.59 (0.17)	0.75 (0.31)	0.84 (0.30)	0.24 (0.29)	0.60 (0.17)	0.63 (0.31)	1.05 (0.30)	0.27 (0.29)
optimal	zero	OZ	0.66 (0.16)	1.32 (0.26)	0.10 (0.28)	0.54 (0.25)	0.89 (0.15)	1.31 (0.27)	0.54 (0.28)	0.81 (0.24)
	S&P 500		0.37 (0.17)	0.27 (0.30)	0.94 (0.32)	0.09 (0.29)				

Notes: The table reports Sharpe ratios with standard errors in parentheses. The trading rule portfolios consist of the top 10 and top 50 ranked strategies, respectively, in the left-hand and right-hand panels. The rows provide Sharpe ratios and their standard errors on portfolios constructed from the 10- and 50- rule currency trading strategies and a long position in the S&P 500. The six types of portfolios are constructed as follows: 1/(N+1) weights on N currency strategies and S&P the 500 (NE); 1/(2N) weights on N currency strategies and 1/2 weight on the S&P 500 (NH); 1/N weights on N currency strategies and 0 weight on the S&P 500 (NZ); optimal weights on N currency strategies and the S&P 500 (OE); 1/2 optimal weights on N currency strategies and 1/2 weight on the S&P 500 (OH); optimal weights on N currency strategies and 0 weight on the S&P 500 (OZ). The bottom panel displays the Sharpe ratio to a buy-and-hold position in the S&P 500 over various samples.

Table 5

Portfolios of technical trading strategies and equity using only non-emerging exchange rates: Sharpe ratios

Weight on each of N FX rules	Weight on equity	Name	Top 10 ex ante rules				Top 50 ex ante rules			
			1975-2012	1975-1987	1988-1999	2000-2012	1975-2012	1975-1987	1988-1999	2000-2012
1/(2N+1)	1/(2N+1)	NE	0.67 (0.16)	1.37 (0.27)	0.45 (0.28)	0.13 (0.28)	0.56 (0.16)	1.47 (0.28)	0.41 (0.28)	-0.30 (0.29)
1/(2N)	1/2	NH	0.58 (0.17)	0.76 (0.31)	0.95 (0.30)	0.13 (0.29)	0.52 (0.17)	0.68 (0.31)	1.00 (0.30)	0.02 (0.29)
1/N	zero	NZ	0.59 (0.16)	1.36 (0.27)	0.28 (0.29)	0.09 (0.27)	0.53 (0.16)	1.46 (0.28)	0.36 (0.28)	-0.31 (0.29)
optimal	optimal	OE	0.55 (0.16)	1.28 (0.26)	0.41 (0.28)	-0.03 (0.27)	0.51 (0.16)	1.29 (0.27)	0.53 (0.28)	-0.27 (0.29)
1/2 optimal	1/2	OH	0.55 (0.17)	0.75 (0.31)	0.91 (0.30)	0.09 (0.29)	0.52 (0.17)	0.63 (0.31)	1.07 (0.31)	0.03 (0.29)
optimal	zero	OZ	0.50 (0.16)	1.32 (0.26)	0.22 (0.28)	0.00 (0.27)	0.52 (0.16)	1.31 (0.27)	0.52 (0.28)	-0.24 (0.29)
	S&P 500		0.37 (0.17)	0.27 (0.30)	0.94 (0.32)	0.09 (0.29)				

Notes: See the notes to Table 4.

Table 6

Rule prevalence over the full sample

1			2			3			4			5		
FX rate	rule	% used	FX rate	rule	% used	FX rate	rule	% used	FX rate	rule	% used	FX rate	rule	% used
TRY	Carry Trade	14.3	EUR	ch(10)	14.3	CLP	ch(20)	15.7	EUR	vma(5,20)	10.4	EUR	vma(5,20)	7.0
CLP	Carry Trade	11.2	CLP	Carry Trade	12.6	EUR	ch(10)	9.1	CLP	ch(20)	5.8	JPY_CAD	ch(5)	5.6
EUR	ch(10)	9.5	CLP	ch(20)	7.5	EUR	mom(20)	6.4	EUR	mom(20)	4.6	EUR	mom(20)	5.4
GBP	ch(10)	9.3	CLP	filter .03	5.8	EUR	vma(5,20)	5.6	CLP	mom(20)	4.3	CLP	mom(20)	4.8
EUR	vma(5,20)	7.0	EUR	vma(5,20)	5.4	EUR_CAD	ch(10)	5.4	JPY	vma(5,20)	4.3	JPY_CAD	mom(5)	4.1
RUB	filter .005	6.8	GBP	ch(10)	4.6	CLP	Carry Trade	4.6	JPY_CAD	ch(5)	4.3	HUF_CHF	Carry Trade	3.5
CAD_GBP	ch(20)	6.8	JPY	vma(5,20)	3.9	CLP	ch(5)	3.7	CLP	filter .02	3.7	JPY	vma(5,20)	3.3
EUR_CAD	ch(10)	4.3	CLP	ch(5)	3.9	JPY_CAD	ch(5)	3.7	HUF_CHF	Carry Trade	3.1	EUR	ch(10)	3.1
CAD_GBP	mom(20)	3.7	RUB	vma(1,5)	3.9	CLP	filter .03	3.1	EUR_CAD	ch(10)	3.1	EUR	ch(5)	3.1
JPY	vma(5,20)	3.1	CAD_GBP	ch(20)	3.1	JPY	vma(5,20)	3.1	TWD	vma(5,20)	2.7	TRY	Carry Trade	2.9

Notes: The table reports the largest 10 trading rule frequencies for the top 5 ranked ex ante portfolios over the full sample, 1975-2012. Thus the left-most columns indicate that for the strategy using the top ranked rule, carry trade applied to TRY appeared 14.3 percent of the time, the carry trade applied to the CLP appeared 11.2 percent of the time, and so on.

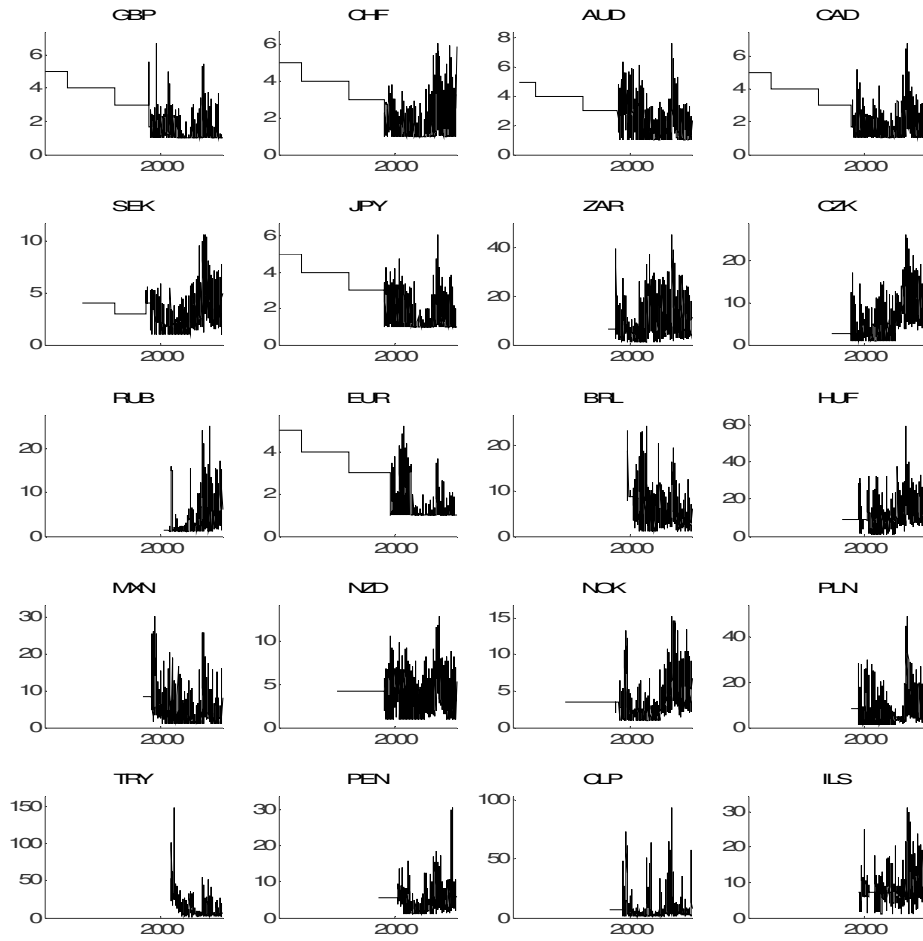
Table 7

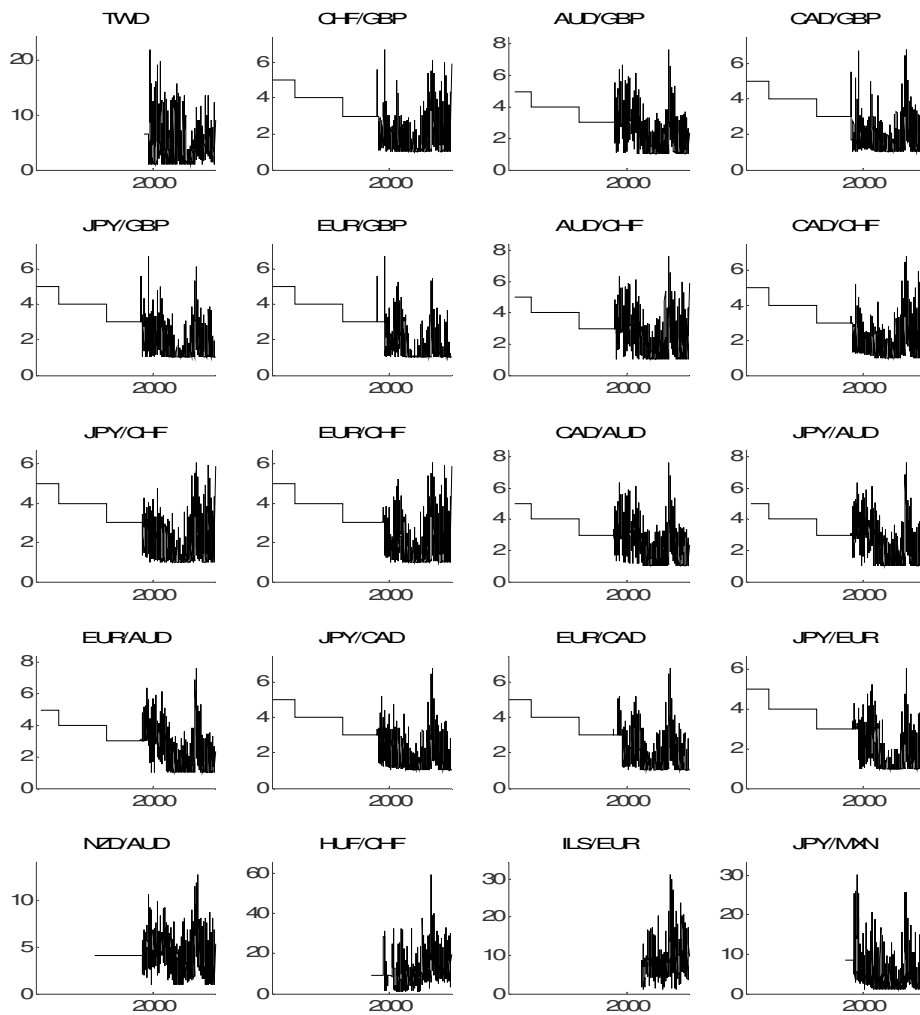
Rule prevalence over subsamples

1973-1982			1983-1992			1993-2002			2003-2012		
FX rate	rule	% used	FX rate	rule	% used	FX rate	rule	% used	FX rate	rule	% used
GBP	ch(10)	45.9	CAD_GBP	ch(20)	26.4	CLP	Carry Trade	41.2	TRY	Carry Trade	53.5
CAD_GBP	mom(20)	18.4	EUR	ch(10)	24.8	EUR_CAD	ch(10)	14.5	RUB	filter .005	25.6
EUR	ch(10)	15.3	EUR	vma(5,20)	24.0	CLP	filter .02	10.7	RUB	vma(1,5)	4.7
EUR	mom(20)	7.1	NOK	vma(1,200)	6.4	JPY	vma(5,20)	10.7	CLP	ch(20)	4.7
CAD_GBP	ch(10)	3.1	NOK	Carry Trade	4.8	MXN	Carry Trade	5.3	RUB	mom(20)	3.1

Notes: The table reports the largest 5 trading rule frequencies for the top ranked portfolio over different sample subperiods. Thus the top row entry in the left panel indicates that for the strategy using the top ranked ex ante rule in the 1973-1982 subsample, the ch(10) applied to the GBP appeared 45.9 percent of the time in the top rule and so on.

Figure 1
Transaction costs






Notes: The figure displays the time series of transaction costs used for each exchange rate in basis points.

Figure 2


Sharpe ratios from the top 662 strategies



Notes: The figure displays the Sharpe ratios for the top 662 ex ante portfolio strategies along with a trendline.

Figure 3

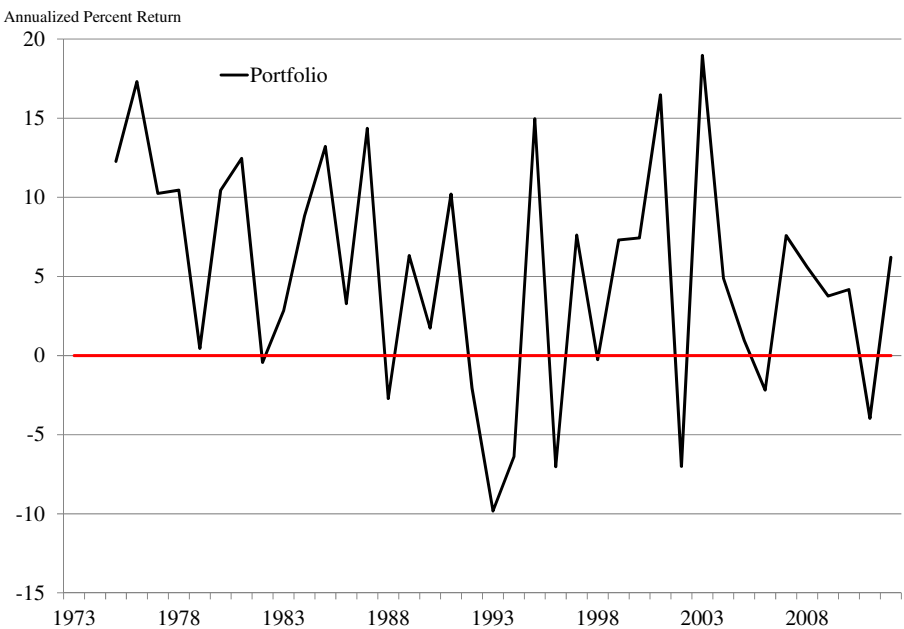

Trades per year



Notes: The panel displays the average number of annual trades for the top 662 ex ante portfolio strategies.

Figure 4

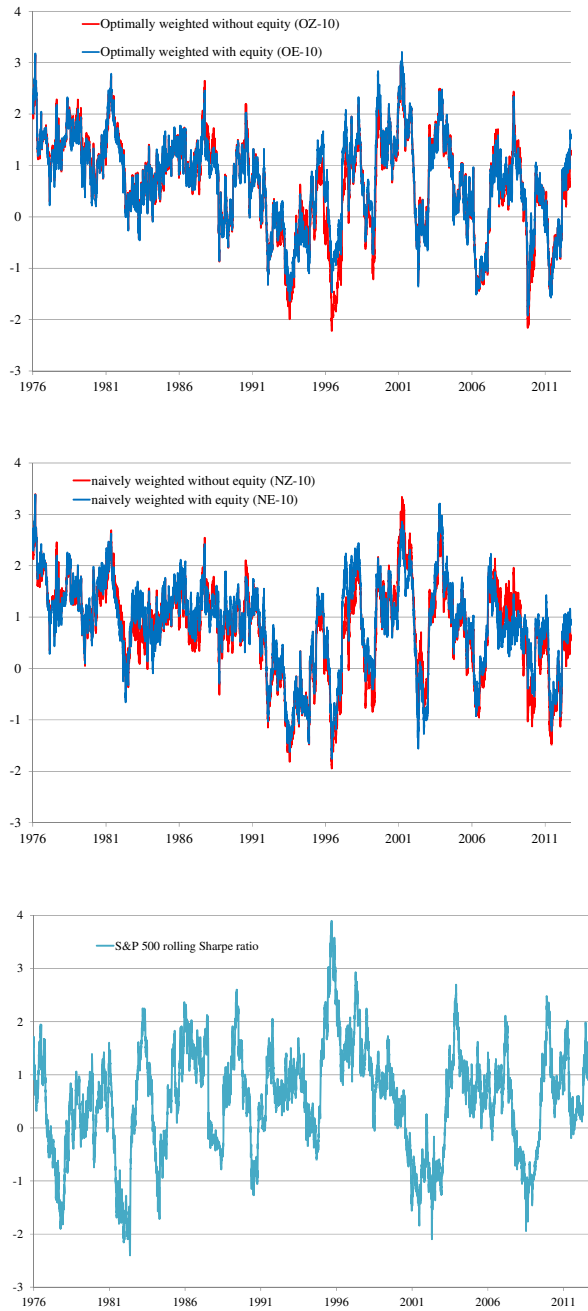
Net returns for the top 5 ranked strategies



Notes: The top panel displays the net annual returns for the top 5 ex ante portfolio strategies, along with the net annual return of the corresponding 1/N portfolio. The bottom panel displays the net annual return of the 1/N portfolio from the top 5 strategies for clarity.

Figure 5

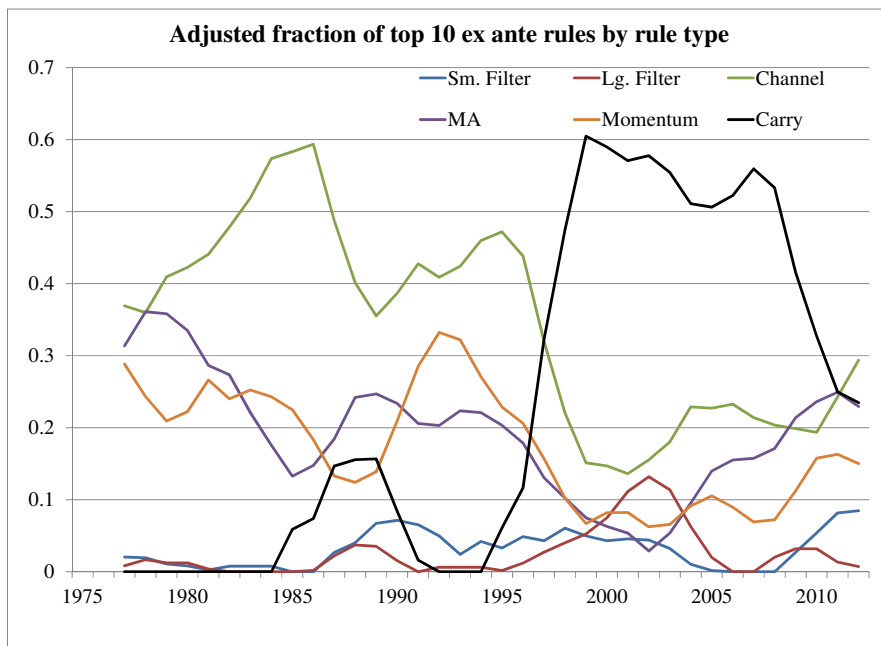

1-year Rolling Sharpe ratios from 1976 for the top 10 strategy portfolios and the S&P 500



Notes: The top (center) panel displays 1-year rolling Sharpe ratios from the OZ-10 and OE-10 (NZ-10 and NE-10) portfolios, from 1976 to 2012. The bottom panel displays the 1-year rolling Sharpe ratios to the S&P 500.

Figure 6

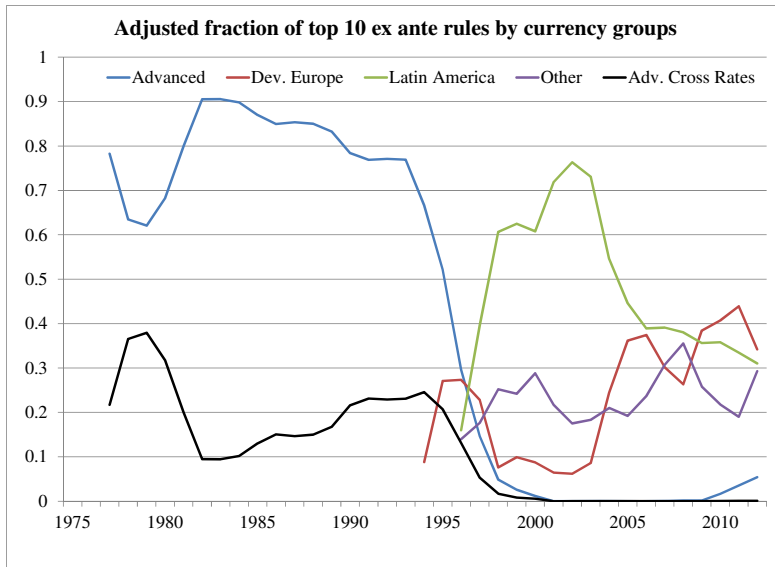

Trading rule prevalence over time



Notes: The panels denote the 3-year moving average prevalence of types of trading rules in the top10 ex ante trading rule strategies. The panel on the top denotes the raw frequency of the rule groups, whereas those on the bottom adjust for group size (see equation (6)). Small filters are those less than or equal to 0.02; large filters are those greater than 0.02.

Figure 7

Exchange rate prevalence over time in the top 10 trading strategies



Notes: The panels denote the 3-year moving average prevalence of currency groups in the best 10 ex ante trading rule strategies. The top panel illustrates the raw prevalence of each group, whereas those on the bottom adjust for group size (see equation (6)). The advanced market exchange rates consist of the AUD, CAD, CHF, EUR, GBP, JPY, NOK, NZD, and SEK; developing Europe consists of the CZK, HUF, PLN and RUB, TRY and HUF/CHF; the Latin American group consists of BRL, CLP, MXN, PEN and JPY/MXN; the Other group consists of ILS, TWD, ZAR and ILS/EUR; and the advanced cross rates group consists of all cross rates between two advanced countries.