

Predictability and ‘Good Deals’ in Currency Markets

Richard M. Levich^a

Valerio Potì^b

Initial Draft: 21 October 2006

This version: 04 August 2009

This paper studies predictability of currency returns over the period 1971-2006. To assess the economic significance of currency predictability, we construct an upper bound on the explanatory power of predictive regressions. The upper bound is motivated by “no good-deal” restrictions that rule out unduly attractive investment opportunities. We find that predictability exceeds this bound during recurring albeit relatively short-lived episodes. Excess-predictability is highest in the 1970s and tends to decrease over time, but it is still present in the final part of the sample period. Moreover, periods of high and low predictability tend to alternate. These stylized facts pose a challenge to Fama’s (1970) Efficient Market Hypothesis but are consistent with Lo’s (2004) Adaptive Market Hypothesis, coupled with slow convergence towards efficient markets. Transaction costs can explain much of the daily excess-predictability but not monthly excess-predictability. Moreover, maximal Sharpe strategies that rationally exploit monthly predictability are not spanned by either the Fama and French (1993) equity-based factors or the AFX Currency Management Index.

Key words: Foreign Exchange, Predictability, Filter Rules, Market Efficiency

JEL Classification: F31

Contact details:

^a Richard M. Levich, New York University Stern School of Business, 44 West 4th Street, New York, NY 10012-1126, USA. Tel: 212-998-0422, Fax: 212-995-4256; E-mail: rlevich@stern.nyu.edu.

^b Valerio Potì, Room Q233, Dublin City University Business School, Glasnevin, Dublin 9, Ireland; Tel: 3531-7005823, Fax: 3531- 7005446; Email: valerio.poti@dcu.ie.

The authors wish to thank Chris Neely (Federal Reserve Bank of St. Louis), Ming-Yuan Leon Li (National Cheng Kung University, Taiwan), Stephen Taylor (Lancaster University), Devraj Basu (EDHEC), and discussants and participants to the INFINITI 2007 conference, the EFM 2008 Symposium on Risk and Asset Management, the EFA 2008 meeting and seminars at Lancaster University and Reading University ICMA Centre for helpful comments and suggestions. Any remaining errors are the authors’ sole responsibility.

Predictability and ‘Good Deals’ in Currency Markets

This paper studies predictability of currency returns over the period 1971-2006. To assess the economic significance of currency predictability, we construct an upper bound on the explanatory power of predictive regressions. The upper bound is motivated by “no good-deal” restrictions that rule out unduly attractive investment opportunities. We find that predictability exceeds this bound during recurring albeit relatively short-lived episodes. Excess-predictability is highest in the 1970s and tends to decrease over time, but it is still present in the final part of the sample period. Moreover, periods of high and low predictability tend to alternate. These stylized facts pose a challenge to Fama’s (1970) Efficient Market Hypothesis but are consistent with Lo’s (2004) Adaptive Market Hypothesis, coupled with slow convergence towards efficient markets. Transaction costs can explain much of the daily excess-predictability but not monthly excess-predictability. Moreover, maximal Sharpe strategies that rationally exploit monthly predictability are not spanned by either the Fama and French (1993) equity-based factors or the AFX Currency Management Index.

1. Introduction

In a literature that spans more than thirty years, various studies have reported that filter rules, moving average crossover rules, and other technical trading rules often result in statistically significant trading profits in currency markets. Beginning with Dooley and Shafer (1976, 1984) and continuing with Sweeney (1986), Levich and Thomas (1993), Neely, Weller and Dittmar (1997), Chang and Osler (1999), Gencay (1999), LeBaron (1999), and Schulmeister (2006), among others, this evidence casts doubts on the simple efficient market hypothesis, even though it is not incompatible with efficient markets under time-varying risk premia and predictability induced by time-varying expected returns. More recently, however, and contrary to the bulk of these earlier findings, a number of authors, including Olson (2004) and Pukthuanthong, Levich and Thomas (2007), find evidence of diminishing profitability of currency trading rules over time. In a comprehensive re-evaluation of

the evidence, Neely, Weller and Ulrich (2009), also find evidence of declining profitability of technical trading rules.

Based on these recent studies, it is tempting to conclude that currency markets have become more efficient. This conclusion, however, rests on an implicit assumption that the set of trading rules examined by the studies span all the strategies that currency market participants could deploy to earn excess-profits. In other words, to infer that currency markets have become more efficient from evidence of diminishing profitability of specific trading rules, one must assume that the econometrician can identify and test all trading strategies that rational investors could have possibly devised, based on available information (that the econometrician, realistically, only partially observes), to exploit predictability of currency returns. Because the econometrician necessarily works with only a subset of available information, and hence can identify at best only a subset of possible trading strategies, findings in support of efficiency based on this methodology may be suspect.

In this paper, we overcome this inherent shortcoming of prior studies by directly assessing whether currency returns are predictable to an extent that implies violation of the efficient market hypothesis (henceforth, EMH). To this end, we test whether an investor endowed with rational expectations (RE), as defined by Muth (1961), can exploit currency return predictability to generate “good deals.” Following the terminology introduced by Cochrane and Saà Requeio (2000), Cerný and Hodges (2001) and Cochrane (2005), “good deals” are defined as investment opportunities that offer unduly high Sharpe ratios. In doing this, we examine how predictability and excess-predictability, i.e. predictability in excess of the “good deals” threshold,

have varied over time and we compare and contrast their patterns with historical patterns in the profitability of technical trading rules considered by the extant literature.

To check on the availability of “good deals”, we follow two methods. First, we compare the explanatory power of predictive regressions, after adjusting for sampling error, with a theoretically motivated “no good deal” bound. This bound is a function of the volatility of the kernel that prices the assets traded in the economy. Following Ross (2005), if prices are set by RE investors, the bound ultimately depends on the volatility of discount factors. Second, we check whether predictability-based strategies expand the investment opportunity set of an investor endowed with RE. To do so, we follow Gibbons, Ross and Shanken (1989) and test whether the maximal Sharpe ratio attainable by strategies that rationally exploit predictability is matched by the maximal Sharpe ratio attainable by portfolios that mimic known asset pricing factors. Importantly, in applying both these methods, we follow Kirby (1998) and explicitly model the implications of realistic levels of transaction costs for the ability of an investor endowed with RE to exploit predictability.

It should be emphasized that we base our inferences on in-sample predictability. This is not, however, a limitation of our analysis. To the contrary, as long as in-sample moments provide consistent estimates of ‘population’ moments, it allows us to check two specific implications of the EMH, namely that the ‘population’ R^2 should not exceed the theoretically motivated predictability upper bound and the ‘population’ maximal Sharpe ratio attainable exploiting predictability should not exceed the Sharpe ratio of strategies spanned by known asset pricing factors. Borrowing from

terminology popularized by Cochrane (2005), the first implication has to do with the “absolute pricing” of observed currency predictability, i.e. whether it is compatible with a model of how expected returns are determined in equilibrium, whereas the second implication is a statement about the “relative pricing” of currency predictability, i.e. whether it is compatible with observed returns on other assets and asset classes. An additional benefit of our reliance on in-sample predictability is that, as demonstrated by Inoue and Kilian (2004), in-sample tests have greater power against the null of no-predictability compared with out-of-sample tests, for a given size of the test, especially in the presence of possible un-modelled structural change in the parameters of the predictive regression under consideration. In many circumstances, as emphasized by Inoue and Kilian (2004), this greater power implies that genuine predictability picked up by in-sample tests is not accompanied by significant out-of-sample predictability and, therefore, results of in-sample tests of predictability are to be seen, perhaps against conventional wisdom, as more credible than results of out-of-sample tests.

A short-cut to understanding our approach can be had by drawing an analogy to studies of excess volatility in equity markets, where researchers compare the volatility of share returns (in-sample) to the volatility of dividends, earnings, and discount factors (also in-sample). Along the same lines, but in the context of currency markets, Brennan and Xia (2006) relate the volatility of exchange rates to the volatility of the economy pricing kernel and, ultimately, to the volatility of discount factors. As noted by Cochrane (2005, p. 396), “‘excess volatility’ is *exactly* the same thing as return predictability” (quotes and italics in the original text). In the same

vein, our study of foreign exchange markets examines whether in-sample predictability is too high relative to the admissible variability of discount factors.

In this respect, our approach can be seen as building upon the intuition developed in Kirby's (1998) seminal article on rational asset pricing and predictability. Kirby (1998) offers a formal analysis of the restrictions that rational asset pricing models place on the coefficient of determination of predictive regressions as well as on the intercept and slope coefficients of such regressions, and uses in-sample moments to make inferences about whether specific asset pricing models can account for observed in-sample predictability of CRSP stock deciles. In a stock market setting, related empirical literature includes the work of Campbell and Thompson (2008) and, with an emphasis on the role of conditioning information, Basu, Abhyankar and Stremme (2005), Stremme, Abhyankar and Basu (2007) and Lettau and van Nieuwerburgh (2008). Pesaran and Timmermann (1995) also study the empirical link between predictability and risk (and thus reward for risk) by examining stock predictability at times of high and low market volatility.

While some of these authors, especially Stremme, Abhyankar and Basu (2007) and Campbell and Thompson (2008), empirically exploit the link between the economy's maximal Sharpe ratio and the amount of admissible predictability, they do not directly test for violations of the EMH. This is the task we undertake here and, in applying and extending Kirby's (1998) approach in a currency market setting, it represents the main empirical contribution of the paper. As pointed out by Taylor (2005), currency strategies have tended to be, by far, more profitable than strategies that attempt to exploit the predictability of other asset classes. A natural question is

then whether such profitability stems from inefficient pricing. It is, therefore, rather surprising that this approach has not been previously attempted in a study of currency market efficiency. The empirical results in our study offer evidence of violations of the upper bound on predictability and, in that sense, violations of the EMH. While such violations are especially severe in the initial part of our sample, excess-predictability of currency returns has not disappeared from the mid-1990s onwards. Our results thus contrast with the vanishing profitability of many popular technical trading rules reported in several recent studies referenced earlier. Importantly, we find that predictability varies over time in a roughly cyclical manner with recurring albeit relatively short-lived episodes during which it exceeds the no good-deal upper bound.

While our evidence is in contrast with the EMH, it is consistent with implications of Lo's (2004) Adaptive Market Hypothesis (AMH), in that bursts of predictability would appear to occur when shifts in market conditions require market participants to re-learn how to make efficient forecasts. Although realistic levels of transaction costs, especially those arising as a result of 'price pressure', e.g. Evans and Lyons (2002), can account for part of these violations and daily predictability is difficult to exploit because it requires frequent trading, monthly predictability can be exploited by strategies that are much less sensitive to transaction costs and expand the investment opportunity set, thus rationalizing market participants' enduring tendency to engage in technical analysis and other active currency management practices.

In the next section, we outline the relation between predictability and time-varying expected returns, on the one hand, and trading rule profitability, on the other hand. In

Section 3, we introduce Ross' (2005) upper bound on pricing kernel volatility and discuss its implications for the maximum amount of predictability compatible with market efficiency. In Section 4, we describe our dataset. In Section 5, we present preliminary empirical results on the predictability of the currencies in our sample. In Section 6, we consider strategies that exploit predictability and evaluate the impact of transaction costs on their profitability. In Section 7, in the spirit of White's (2000) reality checks, we assess the possible impact of sampling error on our inferences. In Section 8, we adopt an explicit multi-factor asset pricing perspective to assess to what extent strategies that exploit predictability, net of both transaction costs and sampling error at conventional confidence levels, expand the investment opportunity set of an investor endowed with RE. In the final Section, we summarize our main findings and offer conclusions.

2. Predictability, Time-Varying Expected Returns and Pricing Kernel Volatility

Trading rule profitability implies that returns are to some extent predictable. This predictability, in turn, can stem either from time-varying discount rates, thus representing changing equilibrium reward for risk, or from information contained in past prices unexploited by market participants. The former possibility is consistent with the notion of market efficiency, whereas the latter is not. Clearly, being able to fully discriminate between these two possibilities requires an equilibrium asset pricing model. More formally, consider the following model of excess returns:

$$r_{t+1} = \mu_{t+1} + \varepsilon_{t+1} \tag{1}$$

where

$$\mu_{t+1} \equiv E(r_{t+1} | I_t) \equiv \mu(I_t) \quad (2)$$

Here, I_t is the information set at time t and ε_{t+1} is a conditionally zero-mean innovation. We can write the unconditional variance of both sides of (1) as follows:

$$\sigma^2(r_{t+1}) = \sigma_\mu^2 + \sigma^2(\varepsilon_{t+1}) \quad (3)$$

Here, $\sigma_\mu^2 = \sigma^2[\mu(I_t)]$. Dividing both sides by $\sigma^2(r_{t+1})$ and rearranging, we see that predictability, as originally formalized by Kirby (1998)¹, is related to variation σ_μ^2 in mean excess returns,

$$R^2 = 1 - \frac{\sigma^2(\varepsilon_{t+1})}{\sigma^2(r_{t+1})} = \frac{\sigma_\mu^2}{\sigma^2(r_{t+1})} \quad (4)$$

Variation in mean returns (conditional expected returns), in turn, can either come from variation in discount rates, consistent with the EMH, or from variation in abnormal mean returns that has not been exploited by the posited rational investor and thus is at odds with the EMH.² To discriminate between these two possibilities, one must identify the component of σ_μ^2 arising from variability of equilibrium discount rates. One way of doing this, as in Kirby (1998), is to specify the model that determines the discount rate component of μ_{t+1} , but this approach entails the difficult task of specifying a full model of equilibrium rational asset pricing, an ongoing

¹ See Equation (14) in Kirby (1998).

endeavour that has hitherto not enjoyed unchallenged theoretical and empirical success. An alternative approach, and the one that we pursue here, is to model σ_μ^2 directly. To do this, we start by recognizing, as already noted by Ross (2005), that we can write:

$$\sigma_\mu^2 \leq E(\mu_{t+1}^2) \leq (1 + R_f)^2 \sigma^2(r_{t+1}) \sigma^2(m_{t+1}) \quad (5)$$

Here, R_f denotes the unconditional risk-free rate and m_{t+1} is a kernel that prices the traded assets. The first inequality in (5), $\sigma_\mu^2 = E[\mu_{t+1} - E(\mu_{t+1})]^2 \leq E(\mu_{t+1}^2)$, is based on an elementary result from descriptive statistics. The second inequality follows from the fact that, under no-arbitrage and in a friction-less economy, the pricing kernel satisfies $\mu_{t+1} = (1 + R_f) \text{Cov}(r_{t+1}, m_{t+1} | I_t)$, while the correlation between the kernel and the asset excess-return is bounded from above, in absolute value, by one. Notably, the inequality in (5) must hold for every kernel that prices the assets, including the one exhibiting minimum variance. Henceforth, unless otherwise specified, m_{t+1} will denote the minimum variance pricing kernel and the term pricing kernel itself will be used to indicate this variable. Using (5) in (4), we see that predictability is bounded from above by a quantity that depends on the amount of volatility of the kernel:

$$R^2 \leq (1 + R_f)^2 \sigma^2(m_{t+1}) \quad (6)$$

² Or time-varying (conditional) ‘alphas’ in the terminology of the empirical asset pricing literature, the literature on performance attribution and event studies.

Notably, the restriction in (6) holds unconditionally and thus, if the data generating processes of the asset returns and the kernel are sufficiently ergodic, the restriction holds also for the *in-sample* coefficient of determination of any consistently estimated predictive regression. By a familiar Hansen and Jagannathan (1991) result, the economy maximal Sharpe ratio for any trading strategy, and thus the maximum amount of profitability per unit of risk, is bounded from above by the volatility of the pricing kernel:

$$\frac{E(\mu_{t+1})}{\sigma(r_{t+1})} = SR \leq (1 + R_f)\sigma(m_{t+1}) \quad (7)$$

Thus, from (6) and (7), it is clear that the volatility of the pricing kernel places an upper bound on both predictability and the maximal SR of the economy. This circumstance is a consequence of the fact that the coefficient of determination of predictive regressions can itself be interpreted as a squared maximal Sharpe ratio. To show this, we use the elementary statistical result that relates the variance of a random variable to its second moment and the square of its mean, and re-write the coefficient of determination in (4) as follows:

$$R^2 = \frac{\sigma_\mu^2}{\sigma_r^2} = \frac{1}{\sigma_r^2} \left(\frac{\mu' \mu}{T} - \bar{\mu}^2 \right) = \frac{1}{T} \mu' \left(\frac{DD'}{T} \right)^{-1} \mu - \frac{\bar{\mu}^2}{\sigma_r^2} = \mu' (DD')^{-1} \mu - \frac{\bar{\mu}^2}{\sigma_r^2} \quad (8)$$

Here, μ is the $T \times 1$ vector that stacks the conditional means of the currency return at each point in time t , $t = 1, \dots, T$, $\bar{\mu}$ is the unconditional mean return (a scalar) and D denotes a $T \times T$ diagonal matrix with elements along the main diagonal that

contain the conditional standard deviation of the currency return at each point in time t . In using this notation, we are essentially interpreting a strategy consisting of a long position in one unit of the currency under consideration as a portfolio made up of as many positions as data points in the sample period, each with its own ‘conditional’ SR. In monthly and higher frequency data, the second term on the far right-hand side of (8) is negligible, as it is the square of a typically small percentage number.³ In light of this, we can approximate the coefficient of determination as follows,

$$R^2 \cong \mu'(DD')^{-1}\mu \quad (9)$$

Interestingly, if one neglects the possible temporal interdependencies across conditional volatilities, i.e. if one neglects GARCH effects, (9) implies that the coefficient of determination in predictive regressions can be interpreted as the squared maximal SR attainable by forming ‘portfolios,’ i.e. strategies, of one-period positions in the currency under consideration. Up to a constant of proportionality, the weights with which each one-period position enters such strategy are then

$$W = (DD')^{-1}\mu \quad (10)$$

Intuitively, a trading strategy based on the above inter-temporal weights amounts to using a predictive model that combines a directional signal, the conditional mean μ_t , with a volatility filter, i.e. the elements σ_t of D .

³ Currencies have typically a relatively small drift rate.

3. EMH and Predictability

To draw testable implications of the EMH for return predictability, we proceed by assuming that there exist a group of greedy, risk-averse investors endowed with Muth's (1961) RE and that such investors are sufficiently homogeneous and wealthy to act as the marginal investor. Ross (2005) argues that the volatility of the intertemporal marginal rate of substitution (IMRS) of such marginal investor provides an upper bound to the volatility of the (minimum variance) pricing kernel. Thus, if we restrict the curvature of this investor's utility function by imposing a relative risk aversion upper bound RRA_V , we can place the following upper bound on the volatility of the (minimum variance) kernel m_{t+1} that prices the assets:

$$\sigma^2(m_{t+1}) \leq \sigma^2(m_{V,t+1}) \cong RRA_V^2 \sigma^2(r_{m,t+1}) \quad (11)$$

Here, $m_{V,t+1}$ is the IMRS between present and future wealth of an investor with relative risk aversion RRA_V , the latter is the relative risk aversion upper bound, and $\sigma(r_{m,t+1})$ is the volatility of the market excess-return $r_{m,t+1}$. Given the latter, the postulated relative risk aversion (RRA) upper bound restricts the curvature of the investors' utility function and thus, under RE, the volatility of discount rates. Combined with (6), then, the pricing kernel volatility bound in (11) implies the following restriction on the explanatory power of any predictive regression of asset returns:

$$R^2 \leq (1 + R_f)^2 RRA_V^2 \sigma^2(r_{m,t+1}) \cong RRA_V^2 \sigma^2(r_{m,t+1}) \quad (12)$$

The right-hand side of (12), which we denote in short as $\phi \equiv RRA_V^2 \sigma^2(r_{m,t+1})$, provides an upper bound on predictability. Violations of this predictability bound amount to excessive volatility of (conditional) expected return, picked up by a high R^2 in predictive regressions, relative to the admissible volatility of discount rates. Equivalently, an excessive R^2 in predictive regressions implies that the conditional expected return is more volatile than the discount rate and thus, ultimately, that the return demanded by investors to hold the asset differs from the return that, based on a forecast that rationally uses all available information, can be expected to earn. Crucially, investors endowed with RE know the economy data generating process (DGP) and thus would be able to exploit and eliminate this discrepancy⁴. In other words, in the language of Cochrane and Saà Requeio (2000) and Cochrane (2005), violations of the predictability bound amount to unexploited “good deal” opportunities and, as such, they violate the EMH.

To operationalize the predictability bound in (12), we need to specify the RRA upper bound RRA_V . Ross (2005) suggests imposing an upper bound of 5 on the relative risk aversion of the marginal investor, i.e. $RRA_V \leq 5$. Among the motivations advanced by Ross (2005) to do so, the one that most easily applies to a world with possibly non-normally distributed returns and non-quadratic utility is the simple observation that a relative risk aversion higher than 5 implies that the marginal investor would be willing to pay more than 10 percent per annum to avoid a 20 percent volatility of his wealth (i.e., about the unconditional volatility of the S&P from 1926) which, by

⁴ This is because, due to the RE assumption, they know the conditional expected return at each point in time and, of course, they know the return they demand.

introspection, seems large. A study by Meyer and Meyer (2005) has recently provided a comprehensive re-evaluation of the empirical evidence on investors' risk aversion. They show that RRA estimates reported by the extant literature are less heterogeneous and extreme if one takes into account measurement issues and the outcome variable with respect to which each study defines risk aversion. Using returns on stock investments as the outcome variable, calculations by Meyer and Meyer (2005) show that the RRA coefficient in the classical Friend and Blume's (1975) study of household asset allocation choices ranges between 6.4 and 2.0, and decreases in investors' wealth. Using returns on the investors' overall wealth, including real estate and a measure of human capital, the RRA estimate ranges between 3.0 and 2.4. The same calculations show that the RRA implied by Barsky et al. (1997) experiment ranges between 0.8 and 1.6.⁵ Importantly, these estimates, at least for the wealthiest cohorts of investors, are always considerably lower than 5, i.e. the upper bound suggested by Ross (2005).

In light of the evidence reviewed by Meyer and Meyer (2005), alongside the RRA upper bound of 5 suggested by Ross (2005), we will also experiment with a lower value for the RRA upper bound, i.e. $RRA_v = 2.5$. This value is just above the RRA of the marginal investor in the stock market, if we assume that this investor's preferences are described by a power utility function and we estimate the mean and volatility of the stock market using the historical average and standard deviation of the returns on the S&P index since 1926. This RRA bound implies that the marginal

⁵ When Meyer and Meyer (2005) consider estimates provided by studies based on asset pricing data, e.g. studies of the equity premium puzzle, they calculate somewhat higher values. Since in these studies the estimates of risk aversion are backed out parametrically from estimates of a particular asset pricing model, often based on a narrow definition of the market portfolio, they are of no interest for the purpose of computing the SDF volatility bound. Their use would imply a circular argument.

investor would be willing to pay up to 5 percent per annum, again an arguably large amount, to avoid a 20 percent volatility of his wealth.

It might be questioned whether a risk aversion upper bound of 2.5 or even 5 is large enough in light of evidence, provided by the empirical literature on the “equity premium puzzle”, that points to much larger values. For example, Mehra and Prescott’s (1985) seminal study suggests that, given the real rate of return on risk free assets, risk aversion in excess of 50 is needed to explain the US equity premium in a model with frictionless capital markets and standard preference assumptions. As formally shown by Ross (2005), however, requiring the upper bound on RRA to exceed such empirical lower bounds provided by the literature on the “equity premium puzzle” would neglect the important circumstance that it is the marginal investors’ risk aversion that matters, not the average investor’s one, and in the former the risk aversion of wealthier and less risk averse investors takes a much larger weight than in the latter. Related arguments offered by Ross (2005) draw on the reasonable idea that aggregate consumption is much less volatile than the portfolio realistically held by the marginal investors, i.e. by investors who have the capacity of influencing prices.

4. Data

In this study, we take the perspective of the US marginal investor. The numeraire currency is thus the US Dollar (USD). Our dataset includes the spot exchange rates against the USD of the 6 major currencies (except those that were replaced by the Euro) over the period 1971-2006. The latter include the Australian and Canadian

Dollar (AUD and CAD, respectively), the Japanese Yen (JPY), the British Pound (GBP), the Swiss Franc (CHF) and the Euro (denoted as ECU/EUR because we combine data on the ECU before the introduction of the Euro in 1999 and on the latter after its launch) provided by Bloomberg at the close of business in London at 6:00 p.m. GMT.⁶ The exchange rate is expressed in terms of units of USD for a unit of the currency under consideration. We construct daily and monthly currency return series by calculating the daily and monthly rate of change of the corresponding exchange rates. To remain consistent with the perspective of the American marginal investor, we use daily and monthly returns on the S&P500 index, constructed from last traded price and dividend data provided by Datastream, to proxy for the return on the market portfolio.

5. Preliminary Results on Currency Predictability

As shown by Taylor (1994), among others, ARIMA models of exchange rates, and thus ARMA models of currency returns, capture substantial predictability. Our estimated predictive models are thus specifications of the general ARMA(p,q) model, where p denotes the autoregressive (AR) lag order and q denotes the order of the moving average (MA) term:

$$y_t = \text{const.} + b_1 y_{t-1} + \dots + b_p y_{t-p} + c_1 u_{t-1} + \dots + c_q u_{t-q} + u_t \quad (13)$$

⁶ We also use daily data, provided by Bloomberg, on the front month futures contract on the exchange rate of each of the above currencies against the US Dollar traded on the Chicago Mercantile Exchange (CME), but the results are not reported because they are qualitatively indistinguishable from, and quantitatively very similar to, the results for the underlying currencies.

We fit versions of (13) to both raw returns and returns adjusted by the interest differential (i.e. the differential between the funding cost in US Dollars and the return from reinvesting the funds in each one of the foreign currencies). While data on a suitable risk-free asset proxy is not available throughout the sample period, we find that adjusting returns for the interest differential has virtually no impact on estimated predictability. This is because the volatility of the interest differential is negligible relative to currency returns volatility. Thus, to avoid duplication of indistinguishable predictability estimates and problems that arise from the incomplete availability of data, we report results for currency return data only.⁷

The predictive regressions in (13) are intended as reduced-form representations of the DGP of currency returns. As prescribed by (6), their in-sample coefficient of determination represents our measure of predictability R^2 , while the difference between the coefficient of determination and the predictability bound captures excess-predictability. To compute the predictability bound ϕ , i.e. the quantity in the right-hand side of (12), we use the standard deviation of the excess-return on the market portfolio proxy over the predictive regression estimation window as an estimator of the unconditional volatility of the market portfolio, and multiply by the square of RRA_V . The latter is set to either 2.5 or 5. To help fix ideas, we refer to the difference between our measure of predictability and the predictability bound as the

⁷ As a proxy for the risk-free rate on assets denominated in the currencies included in our dataset, we use daily middle rate data on Australian Dollar and German Mark inter-bank ‘call money’ deposits, on Canadian Dollar and Swiss Franc Euro-market short-term deposits (provided by the Financial Times/ICAP), on inter-bank overnight deposits in GBP and the middle rate implied by Japan’s Gensaki T-Bill overnight contracts (a sort of repo contract used by arbitrageurs in Japan to finance forward positions). The rate on German Mark deposits is used as a proxy for the rate at which it is possible to invest funds denominated in ECU, while the overnight Euribor is used as a proxy for the rate at which it is possible to invest Euro denominated funds. As a proxy for the US risk-free rate, we use daily data on 1 month T-Bills (yields implied by the mid-price at the close of the secondary market). The interest rate data are taken from Datastream.

‘boundary violation index’ (henceforth BVI), i.e. we define $BVI = R^2 - \phi$. The inequality in (6) implies that BVI should be non-positive for all predictive regressions of the returns on all traded assets priced by the kernel m_{t+1} . In presenting our results, to facilitate their interpretation, we exploit the insight offered by (8) and express excess-predictability in annualized Sharpe ratio units. To do so, we take the square root of the annualized value of BVI ,

$$\gamma_{i,t_0 \rightarrow t_1} \equiv \sqrt{BVI_{i,t_0 \rightarrow t_1} \frac{l}{years}} \quad (14)$$

Here, t_0 and t_1 are the beginning and end points, respectively, of the estimation window, i.e. the sample period over which the predictive regression is estimated and the predictability bound constructed. The quantity under the square root is multiplied by the ratio of the number of observations l to the number of years in the estimation window length, i.e. $\frac{l}{years}$, to annualize (e.g., $\frac{l}{years} = \frac{60}{5} = 12$ if working with 5-year windows of monthly data). The quantity $\gamma_{i,t_0 \rightarrow t_1}$ has an appealing economic interpretation. Based on (8), it can be seen as the annualized excess-SR that can be earned by a RE investor by exploiting observed predictability, assuming that such investor trades at the indicated frequency (i.e. daily or monthly). The value taken by $\gamma_{i,t_0 \rightarrow t_1}$, therefore, can be seen as a measure of “good deal” availability.⁸ Clearly, $\gamma_{i,t_0 \rightarrow t_1}$ is defined only when $BVI_{i,t_0 \rightarrow t_1} \geq 0$.

⁸ We acknowledge that the empirical results are in-sample and may appear less convincing as evidence of “good deals” availability and thus EMH violations. However, our methodology is consistent with the essence of equation (6) and not a limitation. In addition, the literature cited earlier finds that many

The predictive models of daily returns are ARMA(5,0) or, equivalently, AR(5) for all sample periods and currencies. The chosen AR(5) model captures reasonably well the daily predictability of currencies in our sample, as suggested by Ljung-Box (1978) tests of the null of residual serial correlation, rejected for lags of up to the 36th order.⁹ In Table 1, we report the coefficients of determination of the AR(5) models of daily returns estimated over the full sample period as well as six consecutive non-overlapping 5-year windows. The Table also reports the computed values of $\gamma_{i,t_0 \rightarrow t_1}$. These excess predictability measures are in many cases positive and economically sizable, suggesting the attainability of excess-SRs often larger than 100 per annum. Excess-predictability is especially high in the initial part of the sample periods, i.e. in 1971-1976. In subsequent periods, the computed values of $\gamma_{i,t_0 \rightarrow t_1}$ are generally lower, except for a surge in the final 5-year period. In the case of ECU/EUR, there is a burst of predictability between 1989 and 1994, possibly in relation to market adjustments leading to the adoption of the Euro.

Finding a single ARMA(p,q) specification capable of capturing monthly predictability for all currencies over all estimation windows proves harder than in the daily case. We thus conduct a full search over alternative versions of the ARMA(p,q) model in (13) allowing for AR and MA terms p and q of up to the fifth order and using the Akaike Information Criterion (AIC), as recommended by Pesaran and

currency speculators employed technical rules similar to (9) and captured returns in line with our calculations.

⁹ The results of these tests are not tabulated to save space, but they are available from the authors upon request. Moreover, the Akaike Information Criterion (AIC) suggests that the AR(5) model provides a better fit than lower order specifications.

Timmermann (1995), as a model-selection criterion.¹⁰ Table 2 reports the AR and MA orders as well as the coefficients of determination and related excess-predictability measures of the chosen ARMA(p,q) specifications, estimated over the full sample period and three sub-sample periods of roughly equal length. Predictability of monthly returns exceeds the bound corresponding to $RRA_V = 2.5$ and, at times, even the bound corresponding to $RRA_V = 5$. Violations are most severe in the first two sub-sample periods, with excess-SRs larger than 100 per cent per annum in the case of CAD and JPY. They are not absent, however, in the 1995-2006 period. In particular, the coefficient of determination of the predictive regressions estimated over this period exceeds the bound corresponding to $RRA_V = 2.5$ for all currencies except CAD and, in the case of JPY, it also exceeds the tighter bound corresponding to $RRA_V = 5$. These results are consistent with the evidence of profitability of momentum strategies during the 1990s reported by Okunev and White (2003) and Harris and Yilmaz (2009).¹¹

To gain visual perspective on how predictability has evolved over time, we estimate specifications of (13) over rolling windows of daily and monthly data for all currency

¹⁰ Unlike the Bayesian Information Criterion (BIC), the Akaike Information Criterion (AIC) is not consistent (just like the popular R^2 criterion), in the sense of selecting the ‘true’ model as the sample size increases without bounds. Pesaran and Timmermann (1995), however, note that, in the context of forecasting asset returns, where the ‘true’ model or the ‘correct’ list of regressors is clearly unknown and may be changing over time, the consistency property is not as important as it may appear at first. In such a context, they suggest that the ability to select a forecasting equation that could be viewed at the time as being a reasonable approximation to the DGP is of greater importance. The AIC, although statistically inconsistent, displays such ability, in that it has the property of yielding an approximate model and, as shown by Shibata (1976), strikes a good balance between giving biased estimates when the order of the model is too low and the risk of increasing the variance when too many regressors are included.

¹¹ In a similar vein, Della Corte, Sarno and Tsiakas (2008), using combined forecasts and Bayesian model averaging, find that the predictive ability of the forward currency premium has substantial economic value and that a mean-variance investor would pay large fees to reward a portfolio manager able of delivering a comparable performance. It is conceivable that our ARMA(p,q) reduced model specifications pick up some of the predictive ability of the forward currency premium, and vice versa. We leave the investigation of this intriguing possibility for future research.

returns in our sample, and record their coefficients of determination. Each predictive autoregression is estimated over a window that runs between $t-l$ and t , where window length l equals 252 for daily data and 60 for monthly data. For each currency i , this yields daily and monthly series of coefficients of determination $R_{i,(t-l) \rightarrow t}^2$ estimated over rolling estimation window of returns from $t-l$ to t . The predictive models for daily returns are, as before, ARMA(5,0). For monthly returns, however, instead of the specification selected by the AIC, we use a simple ARMA(1,1) to facilitate comparison across currencies and over time. We also estimate a time-varying predictability bound ϕ_t at the daily (monthly) frequency. To do so, we proxy for the variance of the market return between $t-l$ and t , i.e. $\sigma_{(t-l) \rightarrow t}^2(r_{S\&P,t-l+s})$ with $s \in [1, 2, \dots, l]$, as the average, over rolling windows of 1 year (5 years) of daily (monthly) GARCH(1,1) S&P500 returns $r_{S\&P,t-l+s}$. To compute a rolling predictability bound $\phi_{(t-l) \rightarrow t}$, we then multiply $\sigma_{(t-l) \rightarrow t}^2(r_{S\&P,t-l+s})$, as prescribed by (12), by the square of the chosen RRA upper bound, i.e. by the square of RRA_V .

In Figure 1, we plot a simple equally-weighted average across currencies of the daily

BVI series, i.e. $BVI_{avg,t} = \frac{1}{6} \sum_{i=1}^6 BVI_{i,t}$, where $BVI_{i,t} = R_{i,(t-l) \rightarrow t}^2 - \phi_{(t-l) \rightarrow t}$ are daily

excess-predictability series constructed, for each currency, as the difference between the daily $R_{i,(t-l) \rightarrow t}^2$ and the corresponding predictability bound, under $RRA_V = 5$.

While $BVI_{avg,t}$ is always positive throughout the sample period, suggesting excess-predictability for the average currency, it appears to be considerably higher at the beginning of the 30-year period under consideration, especially in the 70s, and

subsequently decreasing over time. The monthly time-series of the rolling coefficients of determination $R_{i,(t-l)\rightarrow t}^2$ for each currency are plotted in Figure 2 against the time series of the rolling predictability bound $\phi_{(t-l)\rightarrow t}$, computed by setting $RRA_V = 5$. Visual inspection of the Figure suggests the occurrence of bursts of excess-predictability at various points, especially in the 1970s, around times of market turmoil and shifting economic and institutional/geopolitical circumstances, such as the so-called Asian crisis and the events surrounding the European monetary integration process in the second part of the 1990s. The rolling coefficients of determination series for all currencies are, however, within the bound in the final part of the 30-year period under consideration. This circumstance, coupled with the declining behaviour of daily excess-predictability for the average currency emphasized earlier by Figure 1, is in agreement with the diminishing abnormal profitability of technical trading rules reported by Olson (2004) and Neely, Weller and Ulrich (2009), in that decreasing excess-predictability presumably makes it more difficult for such rules to spot profitable trends.¹²

Overall, our results partially support the view that excess-predictability has declined over time but also suggest great caution before concluding that excess-predictability has disappeared from the early 1990s onwards. The instances of violation of the predictability bound that occur in the latter part of the sample period contrast with this conclusion. To reconcile the results reported in Table 1 and Table 2 with the findings of disappearing profitability of technical trading rules reported by Olson

¹² Olson (2004) applies double moving-average rules to GBP, CHF, JPY and the German Mark exchange rate against the US dollar and finds evidence that they would have generated abnormal profitability over the periods 1976-1980 and 1986-1990 but also that excess-profitability disappeared

(2004) and Neely, Weller and Ulrich (2009), one must posit that the rules considered by these authors, as well as the simple ARMA(1,1) on which the rolling estimate R^2 reported in Figure 2 are based, do not capture all predictability.

One possibility is therefore that markets have learned to exploit the predictability captured by simple predictive models, e.g. our ARMA(1,1), and relatively traditional trading strategies, e.g. the moving average and trading rules examined by Olson (2004) and Neely, Weller and Ulrich (2009), but not all excess-predictability, e.g. the predictability captured by the more flexible ARMA(p,q) specifications considered in Table 2 and the momentum strategies studied by Okunev and White (2003) and Harris and Yilmaz (2009). Evidence provided by Pukthuanthong, Levich and Thomas (2007) support this conjecture. These authors report that once profitable trend-following rules now lose money, whereas the corresponding counter-trending rules, i.e. rules that do exactly the opposite, are increasingly profitable. Our measure of excess-predictability, in principle, captures the sources of profitability of both rules.

On balance, our findings offer *prima facie* evidence that there is excess-predictability in currency markets and that this excess-predictability, in recent years, has declined from its 1970s peaks without disappearing entirely. This implies that there might be good reasons why currency traders, in their pursuit of profitability and against academic advice, have long engaged in technical analysis and other practices aimed at exploiting predictable patterns in currency returns. Taken at face value, these results represent potential evidence against the EMH. There is the possibility, however, that high transaction costs might have to be incurred to rationally exploit

after 1991. Neely, Weller and Ulrich (2009) examine a more comprehensive set of trading rules and

the estimated predictability and that our estimates of the coefficient of determination R^2 may be inflated by sampling error. We now investigate these important possibilities in turn.

6. The Impact of Transaction Costs

We use the ‘weights’ in (10) to construct predictability-based ‘maximal SRs strategies’ and calculate the returns on such strategies after transaction costs. In the context of our ARMA(p,q), the mean vector in (10) equals the conditional mean of (13), i.e. $\mu_t = y_t - u_t$, while DD' collapses to the currency return sample variance times a $T \times T$ identity matrix, i.e. $\sigma^2(y_{t+1})I_{T \times T}$. The strategies under consideration replicate the trading decisions of an investor endowed with RE who aims to generate the maximal SRs attainable by trading in each currency. Much of the extant literature considers transaction costs of about 0.05 percent, or 5 basis points, realistic for a typical round trip trade between professional counterparts, see Levich and Thomas (1993) and Neely, Weller and Dittmar (1997). This corresponds to about 2-3 basis points on each one way, i.e. buy or sell, transaction. In calculating the return to these strategies, therefore, we allow for transaction costs of up to 5 basis points. For comparison, we also experiment with transaction costs of 25 basis points.

In Figure 3, to illustrate the potential impact of transaction costs, we plot the time-varying weights, calculated using (10) and normalized to add up to unity, of the maximal SR strategies that rationally exploit the daily and monthly predictability of the Canadian Dollar, based on AR(5) and ARMA(5,2) specifications, respectively,

estimated over the period 1995-2006. The corresponding plots for other currencies, predictive models and sample periods are not reported to save space. In all cases, there is substantial variation in the weights of the (daily) positions entailed by the maximal SR strategies that optimally exploit daily predictability, as a result of the conditional time-variation of the mean of the return process. There is much less variation in the weights of the (monthly) positions required to exploit monthly predictability. This means that strategies that exploit daily predictability are rebalanced more frequently than those that exploit monthly predictability and therefore transaction costs are likely to have a greater impact on the former than on the latter. Notably, in classic filter and moving-average strategies, trading positions change relatively infrequently.¹³ This is because such strategies often exploit predictability at low frequency and thus avoid the burden of high transaction costs.

In Table 3, we report the SRs offered by maximal SR strategies that rationally exploit daily and monthly predictability. As before, the predictive model for daily returns is ARMA(5,0). For monthly returns, to ease our computational burden, we use a single ARMA(5,2) model for all currencies and sample periods.¹⁴ Table 3 shows that, for all the currencies under consideration, except CHF, transaction costs of 3 basis points are enough to lower the SRs of the daily strategies below the level that corresponds to the tightest predictability bound, and the maximal SRs of the strategies based on the daily predictability of AUD, JPY and ECU/Euro become negative. With transaction costs of five basis points, the maximal SRs of daily strategies are negative

¹³ For example, Levich and Thomas (1993) report that over their 15 year sample period of major currencies, the 5 day / 20 day moving average rule traded 13 times per year.

¹⁴ Because of the relatively large number of AR and MA terms, our chosen ARMA(5,2) specification produces a volatile prediction compared to most specifications considered in Table 2. Thus, such a specification will not underestimate the impact of transaction costs. If anything, it will overestimate it.

for all currencies. The strategies that exploit monthly predictability, however, are much less sensitive to transaction costs. In all sub-sample periods, the SRs for the maximal SR monthly strategies are positive even with transaction costs of 5 basis points.

Overall, these findings suggest that, while a RE investor cannot exploit daily predictability because of high transaction costs, monthly predictability might be amenable to generate high SRs for such an investor, because trading frequency and transaction costs are reduced. In turn, this suggests that trading strategies based on low-frequency currency predictability might be attractive for professional investors, at least for those who can use available information better than the representative investor and face moderate yet realistic levels of transaction costs.

A word of caution is in order at this point with respect to the likely magnitude of any available “good deal.” There is substantial evidence that transaction costs depend on the size of the transaction and, more specifically, on “price pressure.” For example, Evans and Lyons (2002) estimate that a buy order of 1 million US dollars increases the execution exchange rate against the Deutsche Mark and the Japanese Yen by as much as 0.54 percent, or 54 basis points. Similar figures are provided by Berger et al. (2008), at least for trades executed over a daily horizon. As shown in Table 4, transaction costs of 25 basis points are enough, with few exceptions, to lower SRs below the threshold that corresponds to the wider predictability bound, i.e. the bound corresponding to $RRA_V = 5$, and often below the level implied by the tighter predictability bound, i.e. the bound corresponding to $RRA_V = 2.5$. Similar or even

higher levels of transaction costs, as implied by the evidence provided by the literature on “price pressure,” are to be expected for large transactions.

In a similar vein, high transaction costs by themselves have been shown to generate apparent predictability. Roll (1984), for example, shows that the bid-ask bounce induces an amount of predictability that depends on the relative magnitude of the bid-ask spread and exchange rate variability. This predictability is not exploitable because, by construction, any attempt to exploit it would be costly.¹⁵ The evidence of high predictability and these considerations on the impact of transaction costs on the profitability of large-size transaction, taken together, allow one to rationalize, on the one hand, the frequent occurrence of studies that find abnormally profitable strategies and, on the other hand, the persistence of excess-profitability.

7. The Impact of Sampling Error

To make inferences on the statistical significance of our predictability estimates, we use an estimator of sampling error based on the familiar F-distribution of the coefficient of determination of restricted regressions under the null that the restrictions are not binding and the restricted model is correct. More specifically, denote by R_c^2 the coefficient of determination of an ARMA(p,q) model, which includes an intercept, with linear restrictions on s of its parameters, $s \leq k = p + q$. It is

¹⁵ We conjecture that available “good deal” opportunities might persist over time because, though in principle advantageous, they do not attract enough investors or investors with enough risk capital, perhaps due to the presence of a fixed component of transaction costs, e.g. entry costs. We leave, however, a formal investigation of this issue, i.e. the link between transaction costs, transaction size and persistence of profit opportunities, for future research.

well known that, under OLS assumptions¹⁶ and the null that $R^2 - R_c^2 = 0$ and the restricted model is correct,

$$\frac{(T-s-1)(R^2 - R_c^2)}{s(1-R^2)} \sim F_{s, T-s-1} \quad (15)$$

Here $F_{s, T-s-1}$ denotes an F-distribution with s and $T-s-1$ degrees of freedom and T denotes the sample size. Nelson (1976) argues that this distribution approximates reasonably well the distribution of $\frac{(T-s-1)(R^2 - R_c^2)}{s(1-R^2)}$ for ARMA(p,q) models. As illustrated by Kirby (1998), restrictions on predictability can be equivalently expressed either as restrictions on the coefficient of determination R^2 of predictive regressions or as linear restrictions on the parameters of such regressions. In our case, the restricted model is a predictive regression with coefficient of determination that does not exceed the predictability bound ϕ . Therefore, denoting by \hat{R}^2 the coefficient of determination of the estimated predictive regression and letting $R_c^2 \leq \phi$, the following provides an upper bound to the distribution of $\hat{R}^2 - \phi$ under the null that such a quantity does not exceed zero and the predictability bound is not too tight (i.e., the restricted model is correct),

$$\hat{R}^2 - \phi \leq \hat{R}^2 - R_c^2 \sim \frac{p+q}{(T-p-q-1)(1-\hat{R}^2)} F_{p+q, T-p-q-1} \quad (16)$$

¹⁶ A similar and equally well known result, with the Chi-squared distribution in place of the F-distribution, holds asymptotically when the assumption that the regression residuals are i.i.d. normally

The 5th percentile of such distribution, then, provides an upper bound to the 95th percentile of the estimated R^2 sampling error distribution, $s.e._{R^2,95\%}$, under the null of no excess-predictability.¹⁷ We use this sampling error upper bound estimator to construct a modified version of BVI , i.e. $BVI_{adj} = BVI - s.e._{R^2,95\%}$, and ‘translate’ this quantity into annualized SRs units,

$$\gamma_{adj;i,t_0 \rightarrow t_1} \equiv \sqrt{BVI_{adj;i,t_0 \rightarrow t_1} \frac{l}{years}} \quad (17)$$

Clearly, $\gamma_{adj;i,t_0 \rightarrow t_1}$ is defined only when $BVI_{adj,t_0 \rightarrow t_1} \geq 0$ and thus only in the presence of a statistically significant violation of the predictability bound. Because the profitability of strategies that exploit daily predictability is seriously affected by transaction costs, as shown in Table 3 and discussed in the preceding section, we focus on the assessment of the statistical significance of monthly predictability. Table 4 reports estimates of the R^2 sampling error upper bound, based on (16), for the monthly predictive regressions considered earlier and illustrated in Table 2. It also reports (an upper bound to) the F-statistic of the test that the R^2 of such regressions does not exceed the predictability bound, as well as the related excess-predictability measure $\gamma_{adj;i,t_0 \rightarrow t_1}$ constructed as in (17). The p-values of the F-test indicate that we can reject the null of no excess-predictability at conventional levels for a number of currencies and periods. Over the full sample period, the predictability bound under $RRA_V = 2.5$ is violated in the case of AUD (but continuous data for this currency is

distributed is relaxed.

only available from 1972), JPY and CHF, with annualized excess-SRs of about 57 percent, 46 percent and 26 percent, respectively. Over the same period, the predictability bound corresponding to $RRA_V = 5$ is not violated to a statistically significant extent for any currency. Violations of this tighter bound occur, however, over the periods 1971-1982 and 1982-1994 in the case of both CAD and JPY, and in the case of the latter also over the period 1995-2006.

We also construct times series of BVI_{adj} , based on ARMA(p,q) predictive regressions estimated using rolling 5-year windows of monthly data, i.e. letting $l = 60$ and $years = 5$, and plotted in Figure 4. The ARMA(p,q) predictive regressions are estimated by maximum likelihood and, when this method fails to converge, using in sequential order the Broyden, Fletcher, Goldfarb and Shanno method described in Press et al. (1988), the simplex method and a genetic search procedure. The AR and MA orders of the ARMA(p,q) models are selected, for each estimation window and currency, using the AIC. Figure 4 shows that bursts of statistically significant excess-predictability occurred at various points over the sample period, for example in the 1970s and 1980s, around the European Monetary System (EMS) crisis of the early 1990s and at the time of the Asian Financial Crisis in the second half of the 1990s. In the more recent part of the sample period, a number of currencies, especially AUD, JPY, CHF and EUR, also experienced episodes of significant excess-predictability. Table 5 shows that the predictability bound is violated in about 20 to 30 percent of the estimation windows for most currencies.

¹⁷ We define the percentile of the F-distribution in a manner consistent with the popular convention of referring to its critical value at the level α as the integral from the critical value to infinity, i.e. an integral 'over the right tail' of the distribution.

Because of the link between excess-predictability and excess-volatility, this means that, at the beginning of more than a fifth of the 5-year estimation windows and from the perspective of a RE investor, the currencies under consideration were mispriced against the USD. The strategies considered in (10) represent, for this investor, one way to exploit such mispricing by generating, over a 5-year horizon, statistically significant excess-maximal SRs of the magnitude depicted in Figure 4. It should be noted, however, that it is not all bad news for the RE perspective. In fact, Table 5 also shows that the number of average consecutive estimation windows characterized by a statistically positive BVI ranges between 1 and 3 for most currencies and sub-periods and never exceeds 3.4. This means that it takes generally less than just 4 months, and often considerably less, for currency markets to correct mispricing to within a level consistent with ruling out statistically significant “good-deals”.

Overall, as emphasized by the 12-month moving average superimposed to the BVI series, statistically significant excess-predictability displays a roughly cyclical pattern, i.e. periods of high and low predictability alternate over time, consistent with Lo’s (2004) AMH. Suggestively, episodes of excess-predictability appear to present themselves at times of changing economic conditions, only to relatively quickly come to an end as market participants re-learn efficient information processing, i.e. to do so in the context of the new economic conditions.¹⁸

¹⁸ To double-check on our assessment of sampling error and especially as a further robustness check on our inferences about the presence of excess-predictability, we consider Hosking’s (1979) results on the asymptotic distribution of the coefficient of determination of ARMA(p,q) models and also bootstrap 2-tailed confidence intervals for the coefficient of determination of the estimated predictive regressions. To do the bootstrap, we employ a version of the so-called “estimation bootstrap” methodology, which has been introduced by Freedman and Peters (1984) and Peters and Freedman (1984) and has been used by Karolyi and Kho (2004) to test the profitability of momentum strategies. The results of this bootstrapping experiment broadly confirm our inferences about the presence of excess-predictability for the currencies in our sample and over time. These results are not reported to

8. Significance of Predictability net of Sampling Error and Transaction Costs

To further assess to what extent predictability-based strategies offer “good deal” opportunities to an investor endowed with RE, we check whether they expand the investment opportunity set of such an investor. To do so, while continuing to take both transaction costs and sampling error into account, we construct maximal SR strategies for the individual currencies and then combine such strategies into an overall maximal SR strategy. We then compare the performance of these maximal SR strategies to a set of benchmark strategies. The latter include the basket of currency management strategies tracked by the AFX index, introduced by Lequeux and Acar (1998) and designed to track the performance of technical analysis rules commonly followed by active currency managers, as well as known equity market factors. To take a conservative stance on the amount of exploitable or detectable predictability, we consider the maximal SR strategies that exploit the predictability implied by relatively parsimonious ARMA(5,0) models of monthly currency returns. We denote by r^* the excess return on the overall maximal SR strategy. The maximal SR strategies for the individual currencies, constructed using the weights defined by (10), enter the overall maximal SR strategy with weights calculated as follows:

$$w^* = \Sigma^{*-1} \bar{\mu}^* \tag{18}$$

save space but are available in previously circulated versions of this paper, i.e. Levich and Potì (2008), and also from the authors upon request.

Here, Σ^* denotes the variance-covariance matrix of the returns on the individual currencies maximal SR strategies and $\bar{\mu}^*$ denotes the vector of their unconditional expected returns. The hypothesized level of transaction costs is two basis points per transaction implied by the time series of weights defined by (10) and (18). The weights in (18), normalized to sum to one, are reported in the first column of Table 6. As reported in the second column of the Table, the SR of r^* is considerably higher than the SR of the individual currencies maximal SRs strategies.

We first compare our constructed maximal SR strategies with the AFX index. Monthly values of this index are available from 1984 onwards. As shown in Table 6, especially in the more recent sub-sample period, the SR of the overall maximal SR strategy r^* is much higher than the SR of the AFX currency management index. In fact, while the SR of the AFX currency management index is lower in 1996-2006 than in 1986-2006, the SR of r^* is actually much higher in the more recent sub-sample period. As shown in Table 7, the correlation of the AFX index and r^* drops from 52.36 percent in 1985-1990 to just over 41 percent in 2003-2006. During the same time, while the SR of the AFX index becomes negative, the SR of r^* exceeds 106 percent per annum.

In a similar vein, Figure 5 shows that the series of the rolling 12-month SR of the AFX index and r^* move remarkably closely until about 1996, but subsequently their correlation breaks down.¹⁹ As shown in Table 6, this correlation becomes negative in 2003-2006. Taken together, these results suggest that the combination of moving-

¹⁹ Interestingly, this is probably about the time when working paper versions of Lequeux and Acar's (1998) article began to circulate.

average rules and currencies considered by the AFX index does not fully capture the estimated amount of currency predictability, especially in recent times. One way to rationalize this evidence is to postulate that, while the excess-profitability of the specific moving average rules considered by the AFX index has dried up as market participants have employed them in their trading strategies, alternative and not yet fully exploited sources of excess-profitability have emerged and manifest themselves as excess-predictability. Again, this is consistent with the Adaptive Market Hypothesis (AMH) perspective put forth by Lo (2004) and advocated, in a currency market setting, by Neely, Weller and Ulrich (2009).

Next, we take an explicit asset pricing perspective and ask whether maximal SR predictability-based strategies are spanned by the Fama and French (1993) factors, which some studies suggest span the investment opportunity set. To this end, we simply regress the excess return on each one of the individual currencies maximal SR strategies, the overall SR strategy and, for comparison, the AFX currency management index against the Fama and French (1993) factors, i.e. we estimate

$$r_{i,t} = \alpha_i + \beta_{i,m}r_{m,t} + \beta_{i,SMB}SMB_t + \beta_{i,HML}HML_t + \varepsilon_{i,t} \quad (19)$$

Here, $r_{i,t}$ is the excess return on either the overall maximal SR strategy, r_i^* , an individual currency maximal SR strategy or the AFX currency management index, α_i denotes the regression intercept, $r_{m,t}$, SMB_t and HML_t are the excess-returns on the Fama and French (1993) market, size and book-to market factor mimicking portfolios, respectively, and $\beta_{i,m}$, $\beta_{i,SMB}$ and $\beta_{i,HML}$ denote their corresponding

factor loadings, while $\varepsilon_{i,t}$ denotes the regression error term. As shown in Table 7, the maximal SR strategies for a number of individual currencies and the overall maximal SR strategy offer a positive and statistically significant α_i , especially over the period 1984-2006. By contrast, the alpha of the AFX index is statistically significant only in the early part of this period, i.e. the sub-sample period 1984-1996. Perhaps more interestingly, the factor loadings on these strategies are always either very small and statistically insignificant or negative and statistically significant. This implies that the strategies either carry little systematic risk or they act as a hedge against the latter.²⁰

This fact, coupled with the sign and magnitude of the factor loadings and the significance of the ‘alpha’ terms, suggest that the strategies that exploit currency predictability expand the investment opportunity set, i.e. they are not spanned by the Fama and French (1993) factors. To formally test whether this is the case, we use the ‘alphas’ of these strategies to compute a Gibbons, Ross and Shanken (1989) test-statistic, i.e. we form

$$GRS = \frac{T - N - K}{K} \left[1 + E(f)' \Omega^{-1} E(f) \right]^{-1} \alpha' \Sigma^{-1} \alpha \sim F_{K, T-N-K} \quad (20)$$

Here, T is the sample size, N is the number of factors f , $E()$ denotes the unconditional expectation operator, Ω denotes the factor variance-covariance matrix, α is the vector of the intercepts from (19), and Σ is the variance-covariance matrix of the strategies residuals not explained by the factor model, i.e. the error terms ε_t in (19). In our application, there are 6 maximal SR strategies that exploit the

²⁰ A recent study by Burnside, et al. (2007) also shows that currency returns are uncorrelated with

predictability of AUD, CAD, JPY, GBP, CHF and ECU/EUR and thus $K = 6$, while the factors f are the excess-returns on the Fama and French (1993) market, size and book-to market factor mimicking portfolios and thus $N = 3$. The GRS statistics for the periods 1986-2006 and 1996-2006, reported in the last column of Table 7, are both highly statistically significant. In computing (20), we estimate population moments using their sample counterparts. Gibbons, Ross and Shanken (1989) demonstrate that comparing the GRS statistic with the 5 percent critical value of its finite sample distribution (under the null that pricing errors equal zero), i.e. the F distribution with K and $T - N - K$ degrees of freedom, amounts to testing whether the factors are on the ex-post mean-variance frontier. The significance of the GRS statistic in our tests thus implies that the Fama and French (1993) factors do not span the predictability-based strategies and, therefore, that the latter expand the efficient frontier, at least from the point of view of a mean-variance investor endowed with RE.

9. Conclusions and Final Remarks

In this paper, we assess the statistical and, more importantly, economic significance of predictability in currency returns over the period 1971-2006. We find that, even under a relatively wide bound on RRA, predictability at times violates the attendant theoretically motivated upper bound. Closer scrutiny reveals that the performance of strategies that attempt to optimally exploit daily predictability is very sensitive to transaction costs and this limits the extent to which predictability can be relied upon to generate genuine “good deals.” On the other hand, the performance of strategies

other asset classes.

that attempt to optimally exploit monthly predictability is robust to moderate yet realistic levels of transaction costs and, at least over portions of the sample period, sufficiently large for us to conclude that it is not due to sampling error.

Taken at face value, this evidence implies the availability of “good deals” and thus violation of the EMH under a broad class of asset pricing models, with conservatively high values of the marginal investor’s RRA and with realistic levels of transaction costs. Excess-predictability is highest in the 1970s and, for most currencies in our sample, tends to decrease over time without disappearing. In addition, we find that strategies based on monthly predictability expand the investment opportunity set, even after transaction costs. This effect is also present in the latter part of the sample period and, crucially, it does not disappear after the mid-1990s, contrary to the conclusions of several recent studies. Taken together, our findings pose a challenge to the EMH but they are consistent with Lo’s (2004) AMH.²¹

Offering confirmation that technical trading is still alive in the currency domain, Pojarliev and Levich (2008) have shown that currency hedge funds behave as if they follow standard technical trading strategies. Over the 1990-2006 period, the authors show that a technical trading factor was the single most significant explanatory variable of currency hedge fund returns. The returns of currency hedge funds were significantly correlated with the AFX index over the 1990-2006 period. The strength

²¹ On a similar note, Lo (2005) offers, on pp. 35-36, a suggestive discussion of the cyclical behaviour of the first-order autocorrelation of the S&P Composite Index. In particular, on p. 35, Lo (2005) argues: “Rather than the inexorable trend to higher efficiency predicted by the EMH, the AMH implies considerably more complex market dynamics, with cycles as well as trends, and panics,

of this relation declined somewhat over the 2001-06 sub-sample, but remained highly significant. And in the recent 2008 financial crisis, Deutsche Bank (2008) reported that their technical trading currency benchmark earned 8.8% over LIBOR through October 20, similar to magnitudes observed in the 1970s and 1980s.

A possible avenue of future research is a more formal investigation of whether the estimated R^2 series contains a time trend, a cyclical component and one or more structural breaks. Considering cross-rates and a wider sample of countries, possibly including emerging economies, might also allow the estimation of possible time trends and structural breaks, perhaps adopting a panel approach. A random coefficient model, along the lines of Swamy (1970), would appear particularly promising to accommodate the difficulty of modelling possible sources of cross-sectional variation in the predictability of currency returns. These extensions would make it possible to better address the important question of whether predictability in excess of a level that can be judged consistent with the EMH has become milder over time as a result of learning by economic agents, or whether excess-predictability exhibits a persistently cyclical pattern that can be more easily explained by Lo's (2004) AMH.

manias, bubbles, crashes and other phenomena that are routinely witnessed in natural market ecologies. These dynamics provide the motivation for active management.”

Table 1
Daily Predictability vs. Bound
ARMA(5,0)

Sample Period	AUD	CAD	JPY	GBP	CHF	ECU/ EUR	Bound	Bound p.a.
	Predictability (R²)							
1971-2006	0.09	0.13	0.07	0.21	0.01	0.21		
1971-1976	-	1.62	0.98	1.52	0.76	-		
1977-1982	1.06	0.96	0.27	0.48	0.26	-		
1983-1988	0.64	0.79	0.99	1.07	0.36	0.19		
1989-1994	0.56	0.28	0.19	0.38	0.19	0.74		
1995-2000	0.43	0.55	0.12	0.94	0.24	0.19		
2001-2006	0.53	0.81	0.31	0.10	0.89	0.67		
	Annualized Excess-Predictability under RRA_V = 2.5 (γ_{RRA=2.5})							
1971-2006	27.5	42.0	15.9	61.5	-	61.5	0.06	38.9
1971-1976	-	198.3	152.3	191.8	132.8	-	0.06	38.9
1977-1982	159.5	151.4	74.5	104.1	72.7	-	0.05	35.5
1983-1988	117.7	132.8	150.6	157.1	82.5	50.2	0.09	47.6
1989-1994	114.5	77.8	61.5	92.6	61.5	132.8	0.04	31.7
1995-2000	95.2	110.0	35.5	148.1	65.5	55.0	0.07	42.0
2001-2006	107.7	136.6	77.8	27.5	143.7	123.0	0.07	42.0
	Annualized Excess-Predictability under RRA_V = 5 (γ_{RRA=5})							
1971-2006	-	-	-	-	-	-	0.25	79.4
1971-1976	-	187.2	137.5	180.3	115.6	-	0.23	76.1
1977-1982	148.9	140.2	47.6	86.9	44.9	-	0.18	67.3
1983-1988	86.9	106.5	128.0	135.6	22.4	-	0.34	92.6
1989-1994	101.6	57.2	31.7	76.1	31.7	121.9	0.15	61.5
1995-2000	61.5	82.5	-	129.0	-	-	0.28	84.0
2001-2006	85.5	119.8	42.0	-	128.0	104.1	0.24	77.8


Notes. This table reports the percentage coefficient of determination R² for each currency, the predictability upper bounds under RRA upper bounds equal to 2.5 and 5, annualized versions of the bounds (“Bound p.a.”), and a measure of excess-predictability, converted to annualized SR units, under a RRA upper bound set to either 2.5 or 5. The estimation method is maximum likelihood.

Table 2
Monthly Predictability vs. Bound
AIC-Based Model Selection

	AUD	CAD	JPY	GBP	CHF	ECU/ EUR	Bounds with RRA _v = 2.5 RRA = 5.0	Bound p.a.
1971-2006								
p,q	3,2	3,3	4,4	1,0	3,2	2,4		
Algorithm	GN	GN	GN	GN	GN	GFGS		
R ²	5.31	1.82	4.37	0.5	3.13	2.20		
$\gamma_{RRA=2.5}$	70.1	26.8	61.5	-	47.9	34.3	1.22	38.3
$\gamma_{RRA=5.0}$	23.0	-	-	-	-	-	4.87	76.4
1971-1982								
p,q	1,0	5,4	2,5	3,2	1,0	-		
Algorithm	GN	GN	BFGS	GEN	GN	-		
R ²	1.1	15.76	15.56	4.70	0.20	-		
$\gamma_{RRA=2.5}$	-	131.5	130.6	63.4	-	-	1.35	40.2
$\gamma_{RRA=5.0}$	-	111.5	110.4	-	-	-	5.40	80.5
1983-1994								
p,q	2,2	3,4	3,5	1,0	1,0	1,0		
Algorithm	GN	BFGS	BFGS	GN	GN	GN		
R ²	3.30	32.56	36.02	0.84	0.21	0.44		
$\gamma_{RRA=2.5}$	50.3	194.0	204.4	-	-	-	1.19	37.8
$\gamma_{RRA=5.0}$	-	182.6	193.7	-	-	-	4.76	75.6
1995-2006								
p,q	2,3	0,1	3,2	2,0	1,1	0,1		
Algorithm	GN	GN	GN	GN	GN	GN		
R ²	3.6	0.03	8.17	3.01	2.42	1.29		
$\gamma_{RRA=2.5}$	54.3	-	91.8	47.4	39.2	13.4	1.14	37.0
$\gamma_{RRA=5.0}$	-	-	65.8	-	-	-	4.56	74.0


Notes. This table reports, for the full sample period and three sub-sample periods of about equal length, the autoregressive p and moving average q terms order lags selected by the Akaike Information Criterion (AIC), the percentage coefficient of determination R^2 of the chosen ARMA(p,q) predictive regressions, as well as their corresponding monthly and annualized (p.a.) predictability upper bounds under RRA upper bounds equal to 2.5 and 5.00. The table also reports excess predictability measures $\gamma_{RRA=2.5}$ and $\gamma_{RRA=5.0}$ computed under each RRA bound and converted to annualized SR units. The estimation method, as specified in the rows labelled "Algorithm", is a Gauss-Newton (GN) algorithm with numerical derivatives (the default choice in RATSTM) or, when this procedure fails to converge, the Broyden, Fletcher, Goldfarb and Shanno (AFGS) method described in Press et al. (1988) or a genetic search (GEN) algorithm. The data frequency is monthly.

Figure 1
Daily Excess-Predictability



Notes. This figure plots, for each point in our sample period, the average of the percentage BVI over the cross-section of the currencies in our sample. The latter is based, for all currencies, on rolling AR(5) auto-regressions and a RRA upper bound of 5, i.e. $RRA_V = 5$. The estimation window of each auto-regression is one year and the sample period is 1971-2006. The average BVI series has been cut off at 10.0 for improved visual clarity. The solid and dotted lines in bold are a 252-day moving average and a linear interpolation, respectively, of the average BVI series.

Figure 2
Monthly ARMA(1,1) Predictability vs. Predictability Bound



Notes. These figures plot the sequences of the percentage coefficients of determinations (shown by the thin dotted line), as well as their 12-month moving average (the thick dotted line), of rolling ARMA(1,1) models of the monthly return on each currency in our sample against their upper bound (shown by the solid thin and lightly coloured line). The bound is computed under a relative risk aversion upper bound of 5. The estimation window of each predictive regression is 5 years and they run from 1971 to 2006. The estimation method is maximum likelihood.

Figure 3
Weights for the Maximal SR Strategy for the Canadian Dollar

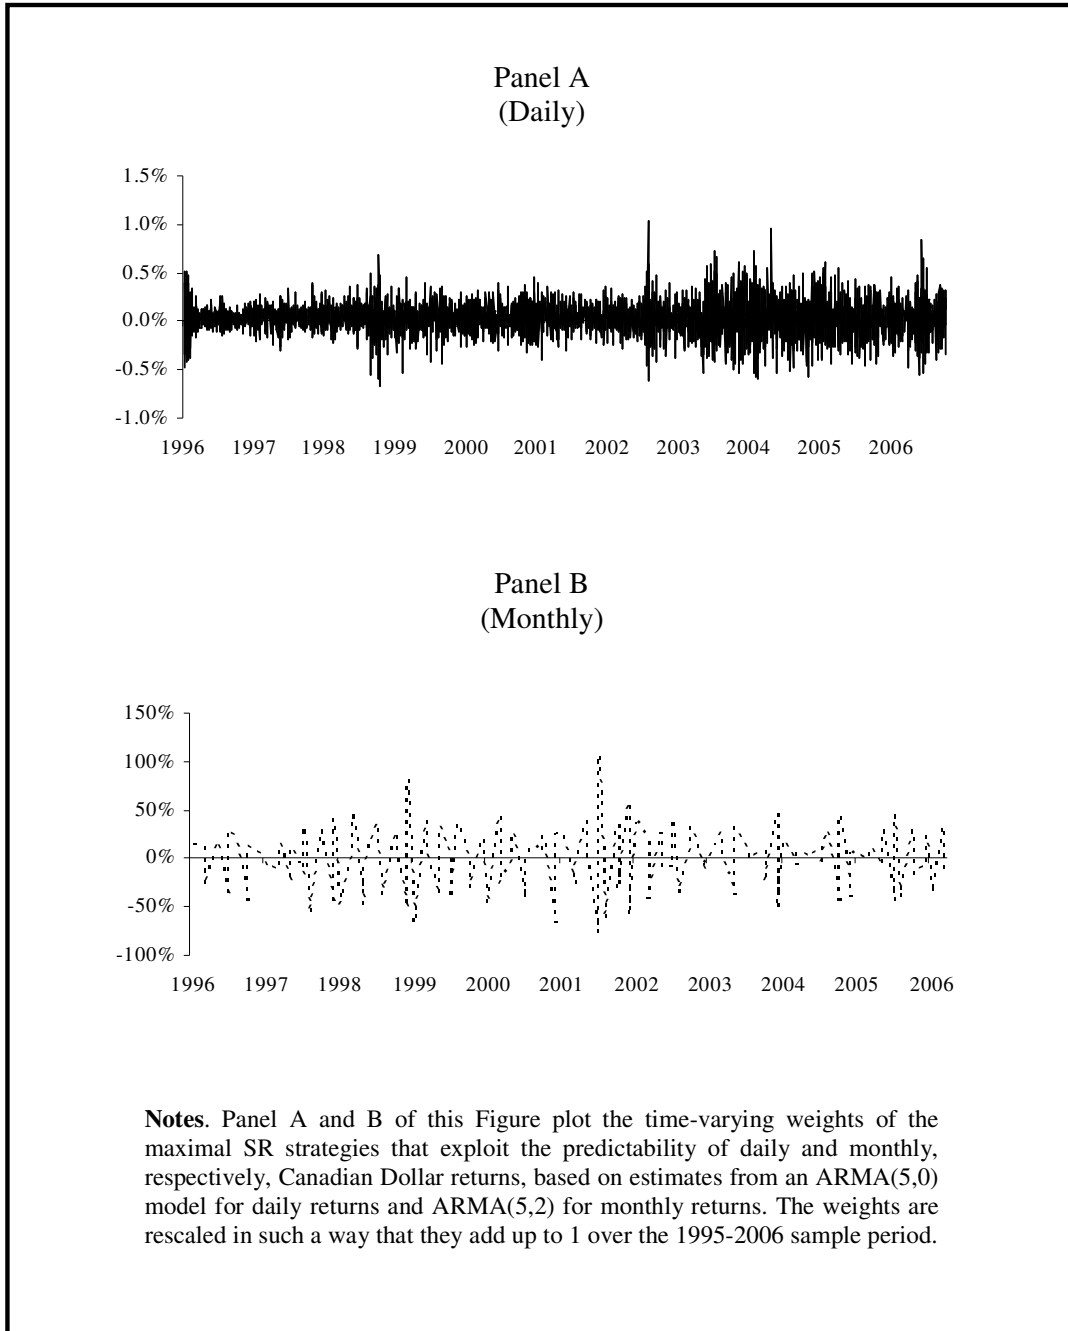


Table 3
**Impact of Transaction Costs on Percentage Maximal SRs of Predictability-
Based Strategies**

Transaction costs (bps)	0	2	3	5	25	Bound RRA _v = 2.5	Bound RRA _v = 5.0
	Daily (1995-2006)					46.0	88.0
AUD	*57.3	17.4	-2.3	-41.7			
CAD	**130.1	*47.1	5.62	-77.4			
JPY	*49.1	-11.0	-41.1	-101.2			
GBP	*48.4	20.0	5.7	-22.7			
CHF	**104.3	*47.7	19.6	-36.7			
ECU/EUR	*82.1	22.7	-7.1	-66.6			
	Monthly (1972-2006)					39.2	78.3
AUD	*43.3	*40.9	*39.8	37.4	14.0		
CAD	*39.5	34.0	31.1	25.6	-30.7		
JPY	*53.0	*51.4	*50.6	*48.9	32.2		
GBP	37.4	35.5	3.5	32.5	12.7		
CHF	*60.2	*58.0	*57.0	*54.8	32.7		
ECU/EUR							
	(1972-1982)					40.1	80.1
AUD	*52.3	*50.6	*49.6	*47.7	27.3		
CAD	**130.9	**123.8	**120.3	**113.0	35.7		
JPY	**128.0	**126.4	**125.7	**124.2	**109.2		
GBP	**181.4	**179.2	**178.1	**175.9	**154.0		
CHF	*76.01	*74.3	*73.4	*71.7	*54.1		
ECU/EUR							
	(1983-1994)					44.0	87.8
AUD	*75.5	*73.7	*72.8	*71.0	*52.8		
CAD	*74.4	*69.8	*67.5	*62.9	16.2		
JPY	**93.4	**91.3	**90.3	**88.2	*66.8		
GBP	*54.6	*52.9	*52.1	*51.3	33.6		
CHF	*63.0	*60.8	*59.7	*57.5	34.8		
ECU/EUR	*85.8	*83.0	*81.5	*78.7	49.4		
	(1995-2006)					38.6	77.1
AUD	*81.5	*79.2	*78.0	*75.7	*51.9		
CAD	*72.0	*68.1	*66.1	*62.0	20.3		
JPY	**106.4	**105.0	**104.4	**103.0	**89.6		
GBP	**78.6	*75.0	*73.3	69.8	34.1		
CHF	*56.0	*54.9	*54.4	*53.4	*43.7		
ECU/EUR	*63.9	*62.2	*61.3	*59.7	*42.8		


Notes. This Table reports percentage annualized maximal Sharpe ratios of strategies that rationally exploit estimated predictability of daily and monthly currency returns, as a function of various levels of transaction costs (in basis points in the top row). The estimated daily predictive regression models are ARMA(5,0) for all currencies. The estimated monthly predictive regression models are ARMA(5,2) for all currencies. The last two columns report the annualized maximal SR bounds under RRA upper bounds equal to 2.5 and 5. The SR bound is computed by taking the square root of the predictability bound and annualizing. One and two asterisks are used to draw attention to SRs in excess of the bound corresponding to RRA = 2.5 and RRA = 5, respectively.

Table 4
Monthly Excess Predictability Adjusted for Sampling Error

		AUD	CAD	JPY	GBP	CHF	ECU/ EUR
1971-2006	BVI _{RRA=2.5}	4.09	0.6	3.15	-	1.91	0.98
	F _{p+q,T-p-q-1}	**9.33	1.32	**7.11	-	**4.26	2.16
	p-value	(0.000)	(0.268)	(0.001)	-	(0.015)	(0.116)
	s.e. 5% level	1.34	1.39	1.35	-	1.37	1.38
	$\gamma_{adj,RRA=2.5}$	57.45	-	46.44	-	25.46	-
	BVI _{RRA=5.0}	0.44	-	-	-	-	-
	F _{p+q,T-p-q-1}	1.00	-	-	-	-	-
	p-value	(0.367)	-	-	-	-	-
	s.e. 5% level	1.39	-	-	-	-	-
	$\gamma_{adj,RRA=5.0}$	-	-	-	-	-	-
1971-1982	BVI _{RRA=2.5}	-	14.54	14.34	3.48	-	-
	F _{p+q,T-p-q-1}	-	**37.28	**36.68	**7.89	-	-
	p-value	-	(0.000)	(0.000)	(0.000)	-	-
	s.e. 5% level	-	1.19	1.20	1.35	-	-
	$\gamma_{adj,RRA=2.5}$	-	126.55	125.59	50.58	-	-
	BVI _{RRA=5.0}	-	10.89	10.69	-	-	-
	F _{p+q,T-p-q-1}	-	**27.92	**27.35	-	-	-
	p-value	-	(0.000)	(0.000)	-	-	-
	s.e. 5% level	-	1.24	1.25	-	-	-
	$\gamma_{adj,RRA=5.0}$	-	107.59	106.45	-	-	-
1983-1994	BVI _{RRA=2.5}	2.08	31.34	34.8	-	-	-
	F _{p+q,T-p-q-1}	**4.65	**100.38	**117.49	-	-	-
	p-value	(0.010)	(0.000)	(0.000)	-	-	-
	s.e. 5% level	1.37	0.96	0.91	-	-	-
	$\gamma_{adj,RRA=2.5}$	29.24	190.94	201.66	-	-	-
	BVI _{RRA=5.0}	-	27.69	31.15	-	-	-
	F _{p+q,T-p-q-1}	-	**88.69	**105.16	-	-	-
	p-value	-	(0.000)	(0.000)	-	-	-
	s.e. 5% level	-	1.01	0.96	-	-	-
	$\gamma_{adj,RRA=5.0}$	-	178.93	190.33	-	-	-
1995-2006	BVI _{RRA=2.5}	2.38	-	6.95	1.79	1.2	0.07
	F _{p+q,T-p-q-1}	**5.33	-	**16.35	**3.99	*2.66	0.15
	p-value	(0.005)	-	(0.000)	(0.019)	(0.071)	(0.858)
	s.e. 5% level	1.36	-	1.30	1.37	1.38	1.40
	$\gamma_{adj,RRA=2.5}$	34.93	-	82.34	22.41	-	-
	BVI _{RRA=5.0}	-	-	3.3	-	-	-
	F _{p+q,T-p-q-1}	-	-	**7.76	-	-	-
	p-value	-	-	(0.000)	-	-	-
	s.e. 5% level	-	-	1.35	-	-	-
	$\gamma_{adj,RRA=5.0}$	-	-	48.37	-	-	-

Notes. This table, based on the ARMA(p,q) estimates illustrated in Table 2, reports the BVIs under $RRA_V = 2.5$ and $RRA_V = 5$, their associated F-statistic and the p-value of the latter under the null that BVI equals zero, and the corresponding percent γ_{adj} , computed as the square root of the annualized BVI $- 1.64 \times \text{s.e.}(R^2)$. One and two asterisks are used to draw attention to when the F-statistic is significant at the 10 and 5 percent level, respectively.

Figure 4
Monthly Excess-Predictability Adjusted for Sampling Error
(Rolling 5-Year Estimation Windows)



Notes. These figures plot, for each point in our sample period and each currency in our sample, the percentage BVI adjusted for sampling error, i.e. $BVI_{i,adj}$, based on ARMA(p,q) predictive regressions, with p and q selected by the AIC, and a RRA upper bound of 5, i.e. $RRA_v = 5$. The estimation window of each predictive regression is 5 years of monthly data from 1971 to 2006. The solid lines formatted in bold are 12-month moving averages. The estimation is conducted by maximum likelihood and, when this method fails to converge, using in a sequential order the Broyden, Fletcher, Goldfarb and Shanno method described in Press et al. (1988), a simplex method or a genetic search algorithm.

Table 5
Descriptive Statistics of Significant
Boundary Violations Occurrences
(Monthly Data)

		AUD	CAD	JPY	GBP	CHF	ECU /EUR
1971-2006	(1)	360	359	359	359	240	355
	(2)	88	70	95	104	84	115
	(3)	24.4	19.5	26.4	28.9	35.0	32.3
	(4)	2.2	2.3	1.5	2.1	2.2	3.4
1971-1983	(1)	86	86	86	85	-	84
	(2)	37	6	26	9	-	17
	(3)	43.0	6.7	30.2	10.5	-	20.3
	(4)	3.0	3.4	1.0	2.6	1.5	-
1984-1995	(1)	83	82	82	83	83	80
	(2)	6	26	15	16	38	25
	(3)	7.2	31.7	18.2	19.2	45.7	31.2
	(4)	2.8	1.8	1.6	2.0	1.5	2.0
1996-2006	(1)	69	69	69	69	69	69
	(2)	18	21	23	27	17	23
	(3)	26.0	30.4	33.3	39.1	24.6	33.3
	(4)	1.5	2.4	1.9	1.9	3.2	2.2

Notes. This table reports, for the rolling predictive regressions considered in Figure 4 for which the estimation procedure converged, (1) the number of windows over which the estimation converged for the full sample period and in each of 3 sub-sample periods, (2) the number and the (3) percentage frequency of positive BVI values, i.e. (2) over (1), as well as (4) the average number of consecutive instances of positive BVI values. BVI is calculated, as explained in the text, under a RRA upper-bound equal to 5.

Table 6
SRs, Correlations and Factor Structure of Maximal-SR Strategies

	w*	SR	Corr. vs. r*	Corr. vs. AFX	alpha	R _m -R _f	SMB	HML	GRS
Panel A (1984-2006)									
AUD	46.1	43.9	58.8	22.5	*0.45 (2.54)	*-5.17 (-2.27)	*-11.78 (-2.44)	*-8.86 (-2.04)	
CAD	15.9	15.4	20.7	4.3	0.16 (0.57)	6.52 (1.09)	5.49 (0.49)	0.85 (0.10)	
JPY	31.6	43.9	58.8	19.1	*0.59 (2.22)	*-12.44 (-2.09)	-5.41 (-0.80)	*-15.19 (-2.01)	
GBP	4.4	24.3	32.5	21.8	0.52 (1.30)	*-14.55 (-3.06)	2.12 (0.19)	-10.91 (-1.45)	
CHF	-15.6	35.3	47.3	41.8	*0.19 (2.71)	*-3.75 (-2.17)	-3.99 (-1.70)	*-6.37 (-2.42)	
ECU-EUR	17.6	39.6	53.1	45.0	*0.85 (2.60)	-6.27 (-0.59)	-10.27 (-0.73)	-17.62 (-1.23)	
r*	100.0	74.7	100.0	46.9	*0.56 (4.11)	*-6.44 (-1.97)	-7.39 (-1.59)	*-11.38 (-2.67)	*2.79 (0.012)
AFX		56.7	46.9	100.0	*0.34 (2.72)	0.52 (0.17)	*-4.71 (-1.11)	-3.60 (-0.92)	
Panel B (1996-2006)									
AUD	6.9	44.2	31.8	10.2	0.34 (1.18)	7.68 (1.11)	-2.22 (-0.18)	5.47 (0.77)	
CAD	12.5	42.7	30.7	11.2	0.38 (1.58)	-0.54 (-0.12)	14.27 (1.70)	-12.52 (-1.35)	
JPY	30.9	74.4	53.5	4.8	0.38 (1.59)	3.19 (1.58)	*-17.93 (-3.23)	2.11 (0.32)	
GBP	22.9	64.6	46.4	-3.0	*0.45 (2.71)	3.15 (0.89)	10.25 (1.90)	1.29 (0.16)	
CHF	21.7	59.3	42.6	23.5	*0.30 (2.07)	2.38 (0.77)	-1.54 (-0.25)	1.84 (0.26)	
ECU-EUR	5.1	53.4	38.3	40.7	*0.86 (2.21)	16.61 (1.19)	-5.42 (-0.20)	-25.72 (-1.55)	
r*	100.0	139.2	100.0	27.9	*0.40 (4.16)	3.53 (1.47)	-2.19 (-0.52)	-1.14 (-0.25)	*2.73 (0.016)
AFX		35.8	27.9	100.0	0.20 (1.40)	1.01 (0.29)	9.52 (1.56)	*-11.65 (-1.77)	

Notes. The first column of this table reports the percentage weights w* with which the maximal Sharpe Ratio predictability-based strategies for each currency enter the overall maximal Sharpe Ratio strategy, denoted by r*. The second column reports percentage annualized Sharpe ratios of these strategies and of the AFX currency management index. The other columns reports the percentage estimated intercept and factor loadings and, in brackets, the associated t-static based on Newy and West (1987) autocorrelation and heteroskedasticity adjusted standard errors. An asterisk denotes significance at the 5 percent level. The last column reports the GRS statistic (and its p-value in brackets) for the 6 individual currencies maximal SR strategies. The hypothesized level of transaction costs is two basis points per each way transaction and the predictive model is ARMA(5,0) for all currencies. The data frequency of the underlying return series is monthly.

Table 7
AFX vs. Maximal SR Strategy

Period	Corr. AFX vs. r^*	SR_{AFX}	SR_{r^*}	Corr. SR_{AFX} vs. SR_{r^*}
1985-1990	52.36	97.17	115.89	52.92
1991-1996	46.29	38.38	58.62	45.64
1997-2002	50.20	49.53	20.00	11.62
2003-2006	41.08	-5.44	106.14	-44.55

Notes. This Table reports, for four 5-year periods between 1985 and 2006, the percentage coefficient of correlation between the AFX Currency Management index return and the maximal SR strategy return r^* , their SR and the correlation between their 12-month moving.

Figure 5
AFX vs. Maximal SR Strategy



Notes. This Figure plots the rolling 12-month SR of the maximal SR strategy r^* that exploits monthly predictability, based on estimates from an ARMA(5,0) model, and of the AFX Currency Management Index. The sample period is 1984-2006.

Bibliography

Barsky, R.B., F.T. Juster, M.S. Kimball, and M.D. Shapiro, 1997, Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study, *The Quarterly Journal of Economics* 112, 537-539.

Berger, D.W., Chaboud, A.P., Chernenko, S.V., Howorka E. and J.H. Write, 2008, Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data, *Journal of International Economics* 75, 93-109.

Brennan, M.J. and Y. Xia, 2006, International Capital Markets and Foreign Exchange Risk, *Review of Financial Studies* 19, 753 - 795.

Burnside, C., M. Eichenbaum, I. Kleschelski, and S. Rebelo, 2007, The Returns to Currency Speculation in Emerging Markets, *American Economic Review Papers and Proceedings* 97, 333-338.

Campbell, J.Y. and S.B. Thompson, 2008, Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies* 21, 1509-1531.

Cerný, A. and S.D. Hodges, 2001, The Theory of Good-Deal Pricing in Incomplete Markets, in *Mathematical Finance-Bachelier Congress 2000*, H. Geman, D. Madan, S. Pliska and T. Vorst, eds. Berlin: Springer, 175-202.

Chang, P. H. K. and C.L. Osler, 1999, Methodical Madness: Technical Analysis and the Irrationality of Exchange-Rate Forecasts, *The Economic Journal* 109, 636-661.

Cochrane, J.H., 2005, *Asset Pricing*, 2nd Ed., Princeton University Press, Princeton.

Cochrane, J.H. and J. Saà-Requejo, 2000, Beyond Arbitrage: Good Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy* 108, 79-119.

Deutsche Bank, 2008. "DB Currency Returns: Weekly Update," October 20.

Della Corte, P., Sarno, L. and I. Tsiakas, 2008, An Economic Evaluation of Empirical Exchange Rate Models, *Review of Financial Studies* (forthcoming), doi:10.1093/rfs/hhn058.

Dooley, M.P. and J. Shafer, 1976, Analysis of Short Run Exchange Rate Behavior: March 1973 – September 1975, *International Finance Discussion Papers*, No. 76. Washington, D.C., Federal Reserve System.

Dooley, M.P. and J. Shafer, 1984, Analysis of Short-Run Exchange Rate Behaviour: March 1973–November 1981, in Bigman, D. and T. Taya (eds.), *Exchange Rate and Trade Instability*, Ballinger, Cambridge, MA.

- Evans, M.D.D., and R.K. Lyons, 2002, Order flow and exchange rate dynamics, *Journal of Political Economy* 110, 170–180.
- Fama, E., 1970, Efficient Capital Markets: a Review of Theory and Empirical Work, *Journal of Finance* 25, 383-417.
- Fama E. and K. French, 1993, Common Risk Factors in the Returns on Stocks and Bonds, *Journal Financial Economics* 33, 3-56.
- Freedman, D.A. and S.C. Peters, 1984, Bootstrapping an Econometric Model: Some Empirical Results, *Journal of Business and Economic Statistics* 2, 150-158.
- Friend, I. and M.E. Blume, 1975, The Demand For Risky Assets, *American Economic Review* 65, 900-922.
- Gencay, R., 1999, Linear, Non-Linear and Essential Foreign Exchange Rate Prediction with Simple Technical Trading Rules, *Journal of International Economics* 47, 91-107.
- Gibbons, M., Ross, S.A., and J. Shanken, 1989, A Test of The Efficiency of A Given Portfolio, *Econometrica* 57, 1121-1152.
- Hansen, L.P. and R. Jagannathan, 1991, Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy* 99, 225-262.
- Harris, R.D.F. and F. Yilmaz, 2009, A Momentum Trading Strategy Based on the Low Frequency Component of the Exchange Rate, *Journal of Banking and Finance* 33, 1575-1585.
- Hosking, J.R.M., 1979, The Asymptotic Distribution of R^2 for Autoregressive-Moving Average Time Series Models When Parameters Are Estimated, *Biometrika* 66, 156-7.
- Inoue, A., and L. Kilian, 2004, In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?, *Econometric Reviews* 23, 371–402.
- Karolyi, G.A. and B.-C. Kho, 2004, Momentum Strategies: Some Bootstrap Tests, *Journal of Empirical Finance* 11, 509-536.
- Kirby, C., 1998, The Restrictions on Predictability Implied by Rational Asset Pricing Models, *Review of Financial Studies* 11, 343–382.
- LeBaron, B., 1999, Technical Trading Rule Profitability and Foreign Exchange Intervention, *Journal of International Economics* 49, 125-143.
- Lequeux, P. and E. Acar, 1998, A Dynamic Index for Managed Currencies Funds Using CME Currency Contracts, *European Journal of Finance* 4, 311-330.

- Lettau, M. and S. van Nieuwerburgh, 2008, Reconciling the Return Predictability Evidence, *Review of Financial Studies* 21, 1607-1652.
- Levich, R. M. and L.R. Thomas III, 1993, The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach, *Journal of International Money and Finance* 12, 451-74.
- Levich, R. M. and V. Poti, 2008, Predictability and ‘Good Deals’ in Currency Markets, NBER Working Paper No w14597.
- Ljung, G.M. and G.E.P. Box (1978), On a Measure of Lack of Fit in Time Series Models, *Biometrika* 67, 297-303.
- Lo, A.W., 2004, The Adaptive Market Hypothesis: Market Efficiency from an Evolutionary Perspective, *Journal of Portfolio Management*, 30th Anniversary Issue, 15-29.
- Lo, A.W., 2005, Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis, *Journal of Investment Consulting* 7, 21-44.
- Mehra, R. and E. Prescott, 1985, The Equity Premium Puzzle, *Journal of Monetary Economics* 15, 145-161.
- Meyer, D.J. and J. Meyer, 2005, Relative Risk Aversion: What Do We Know? *Journal of Risk and Uncertainty* 31, 243-262.
- Muth, J.F., 1961, Rational Expectations and the Theory of Price Movements, *Econometrica* 29, 315-335.
- Neely, C., Weller, P., and R. Dittmar, 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. and J. Ulrich, 2009, The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market, *Journal of Financial and Quantitative Analysis* 44, 467-488.
- Nelson, C.R., 1976, The Interpretation of R^2 in Autoregressive-Moving Average Time Series Models, *The American Statistician* 30, 175-180.
- Newy, W.K. and D.K. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-8.
- 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 Foreign Exchange Markets declined through time? *Journal of Banking and Finance* 28, 85-105.
- Pesaran, M.H. and A. Timmermann, 1995, Predictability of Stock Returns: Robustness and Economic Significance, *Journal of Finance* 50, 1201-1228.
- Peters, S.C. and D.A. Freedman, 1984, Some Notes on the Bootstrap in Regression Problems, *Journal of Business & Economic Statistics* 2, 406-409.
- Pojarliev, M. and R.M. Levich. 2008, Do Professional Currency Managers Beat the Benchmark?, *Financial Analysts Journal* (Sept/Oct.), 18-32.
- Press, W.H., Flannery, B.P., Teukolsky, S.A. and W.T. Vettering, 1988, *Numerical Recipes in C*, Cambridge University Press, New York.
- Pukthuanthong, K., Levich, R. M. and L.R. Thomas III, 2007, Do Foreign Exchange Markets Still Trend? *Journal of Portfolio Management*, vol. 34, no. 1 (Fall): 114-28..
- Roll, R., 1984, A Simple Implicit Measure of the Effective Bid-Ask Spread in Financial Markets, *Journal of Finance* 39, 1127-1139.
- Ross, S. A., 2005, *Neoclassical Finance*, Princeton University Press, Princeton, New Jersey.
- Schulmeister, S., 2006, The Interaction Between Technical Currency Trading and Exchange Rate Fluctuations, *Finance Research Letters* 3, 212-233.
- Shibata, R., 1976, Selection of the order of an autoregressive model by Akaike's information criterion, *Biometrika* 63, 117-126.
- Basu, D., Abhyankar, A. and A. Stremme, 2005, The Optimal Use of Return Predictability: An Empirical Analysis, AFA 2006 Boston Meetings Paper. Available at SSRN: <http://ssrn.com/abstract=687083>.
- Stremme, A., Abhyankar, A and D. Basu, 2007, Portfolio Efficiency and Discount Factor Bounds with Conditioning Information: An Empirical Study, *Journal of Banking and Finance* 31, 419-437.
- Swamy, P.A.V.B, 1970, Efficient Inference in a Random Coefficient Regression Model, *Econometrica* 38, 311-323.
- 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.

Taylor, S.J., 2005, *Asset Price Dynamics, Volatility and Prediction*, Princeton University Press, Princeton, NJ.

White, H., 2000, A Reality Check for Data Snooping, *Econometrica* 68, 1097–1126.