Predictability and 'Good Deals' in Currency Markets

Richard M. Levich[a], Valerio Potì[b],

1st Draft: 21st October 2006

This version: 28th March 2008

Comments welcomed

PRELIMINARY AND INCOMPLETE

(Please do not quote without the authors' permission)

This paper studies predictability of currency returns over the period 1971-2006. We examine whether predictability is significant both from an economic and a statistical point of view. To assess economic significance, we construct upper bounds on the explanatory power of predictive regressions. The upper bounds are motivated by no-good deal restrictions. To assess the statistical significance of violations of these bounds, we bootstrap the coefficient of determination R^2 in predictive regressions and we generate its posterior distribution using a Gibbs sampling technique. The evidence on economic and statistical significance is mixed across periods but it is somewhat weaker in the final part of the sample period. Moreover, strategies that attempt to exploit excess-predictability are very sensitive to transaction costs, rendering large scale risk-adjusted profits very unlikely.

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: 3531-7005328, Fax: 3531- 7005446; Email: valerio.poti@dcu.ie.

The author wishes to thank Chris Neely and participants to the 2007 INFINITI conference for helpful comments and suggestions.

Predictability and 'Good Deals' in Currency Markets

This paper studies predictability of currency returns over the period 1971-2006. We examine whether predictability is significant both from an economic and a statistical point of view. To assess economic significance, we construct upper bounds on the explanatory power of predictive regressions. The upper bounds are motivated by no-good deal restrictions. To assess the statistical significance of violations of these bounds, we bootstrap the coefficient of determination R^2 in predictive regressions and we generate its posterior distribution using a Gibbs sampling technique. The evidence on economic and statistical significance is mixed across periods but it is somewhat weaker in the final part of the sample period. Moreover, strategies that attempt to exploit excess-predictability are very sensitive to transaction costs, rendering large scale risk-adjusted profits very unlikely.

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), Olson (2004), 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, Pukthuanthong, Levich and Thomas (2007) find evidence of diminishing profitability of currency trading rules over time. In a comprehensive re-evaluation of the evidence hitherto provided by the extant literature, Neely, Weller and Joshua (2007), also find evidence of declining profitability of technical trading rules.

In this paper, we directly assess whether currency returns are predictable to an extent that implies violation of the efficient market hypothesis (henceforth, EMH) and whether the evidence against the EMH has changed over time. Our analysis entails assessing whether economically significant predictability in currency returns is significant from a statistical point of view. To assess economic significance, we construct a theoretical time-varying upper bound on the explanatory power of predictive regressions. This bound, following Ross (2005), is ultimately a function of the squared maximal Sharpe ratio available in the economy and makes precise the intuitive connection between predictability, risk and reward for risk. Any violation of this bound would imply that, by exploiting predictability, it is possible to generate unduly high Sharpe ratios and thus, in the terminology introduced by Cochrane and Saà Requeio (2000), Cerný and Hodges (2001) and Cochrane (2001), the availability of 'good deals'. To assess the statistical significance of violations of the predictability bound, we bootstrap the coefficient of determination in a simple predictive auto-regression and we generate its posterior distribution using a Gibbs sampling technique. Our approach, therefore, amounts to testing whether, given sensible restrictions on the volatility of either the single (if international financial markets are imperfectly integrated) or the many kernels (if international financial markets are imperfectly integrated) that price the assets, currency return predictability can be exploited to reliably generate 'good deals'.

In a stock market setting, related empirical literature includes the work of Campbell and Thompson (2005) and, with an emphasis on the role of conditioning information, of Stremme, Basu, and Abhyankar (2005). Earlier work that captured the empirical link between predictability and risk (and thus reward for risk) is Pesaran and Timmermann's (1995) study of stock predictability at times of high and low market volatility. While these authors empirically exploit the link between the economy maximal Sharpe ratio and the amount of admissible predictability, they do not directly test for violations of the EMH. This is instead the approach we take here and it represents the main contribution of the paper. As pointed out by Taylor (2005), currency strategies tend to be by far more profitable than strategies that attempt to exploit the predictability of other asset classes. It is therefore rather surprising that this approach has not been previously attempted in a study of the efficiency of the currency market. Empirically, we find evidence of prolonged violations of the EMH, especially in the initial part of the sample period. We find that these violations are often statistically significant. We also find, however, that realistic levels of transaction costs, especially those arising as a result of 'price pressure', can account for much of these violations. Moreover, we find that predictability is less significant from a Bayesian perspective. The reduced significance in a Bayesian setting is due to the extra uncertainty associated to the prior distribution of the parameters, most notably the volatility of the currency returns.

In the next section, we outline the theoretical relation between predictability and time varying expected returns, on one hand, and trading rule profitability, on the other hand. We also introduce Ross' (2005) upper bound on the pricing kernel volatility and we discuss its implications for the maximum amount of explanatory power of predictive regressions of currency returns compatible with foreign exchange market efficiency. In Section 3, we describe our dataset. In Section 4, we describe the simple rolling auto-regressions and autoregressive moving average models (ARMA) that we employ to capture predictability over time and how we construct empirical upper bounds, based on Ross' (2005) theoretical bound, on their coefficient of determination. We also present a bootstrap experiment to test for the significance of violations of the predictability bound. In Section 5, we present back-of-the-envelope calculations of the impact of transaction costs on predictability. In Section 6, we construct the posterior distribution of the auto-regression coefficient of determination to evaluate the impact of Bayesian uncertainty. 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 expected returns, thus representing an equilibrium reward for risk, or from information contained in past prices unexploited by market participants. The former possibility is consistent with the notion that the currency market is efficient, whereas the latter is not. Clearly, being fully able to 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

$$\mu_{t+1} = E(r_{t+1} \mid I_t) = \mu(I_t) \tag{2}$$

Here, I_t is the information set used for pricing at time t and ε_{t+1} is a conditionally zero-mean innovation. Then, following Ross (2005), we can write:

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

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

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

Variation in expected excess returns, in turn, can either come from variation in equilibrium risk premia, consistently with the EMH, or from variation in abnormal expected returns that have not been exploited by the posited rational investor and thus are at odds with the EMH. To discriminate between these two possibilities, one must specify what constitutes the model of the rational expected excess returns, and thus the rational component of σ_μ^2. An equivalent way of representing this fact is to recognize that $\sigma_\mu^2 = E[\mu_{t+1} - E(\mu_{t+1})]^2 \leq E(\mu_{t+1}{}^2)$ and that, as noted by Ross (2005), we can write:

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

The first inequality in (5) is based on an elementary result from descriptive statistics, while the second inequality follows from the fact that, under no-arbitrage and in a friction-less economy, the pricing kernel satisfies $\mu_{t+1} = (1+R_f)Cov(r_{t+1}, m_{t+1})$. Using (5) in (4), we see that predictability is bounded from above by the amount of volatility of the kernel that prices the assets:

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

The pricing kernel, in an EMH setting, is the investors' inter-temporal marginal rate of substitution whereas, if the market is inefficient, a more volatile pricing kernel will be required to price the assets. Finally, by a familiar Hansen and Jagannathan (1991) result, the maximal Sharpe ratio (*SR*), and thus the maximum amount of profitability from any trading strategy consistent with a given pricing kernel is bounded from above by the volatility of the pricing kernel:

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

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 consideration suggests one way to mitigate the stark alternative between conducting a joint test of market efficiency and of a particular asset pricing model and not being able to discriminate between time-variation in equilibrium return and abnormal profitability. A possible solution is to impose just enough restrictions on preferences and make just enough assumptions about how investors form expectations to be able to restrict the volatility of the pricing kernel. This then yields restrictions on the maximal Sharpe ratio of the economy and on predictability. This way it is possible to draw some implications for return predictability without having to fully specify an equilibrium asset pricing model.

Ross (2005) shows that, if we are willing to accept that there is a sufficiently homogeneous and wealthy group of investors with preferences defined over wealth and constant relative risk aversion, then assets are priced by a representative investor with relative risk aversion constant in market wealth. If we further assume that investors' risk aversion is bounded from above and that investors estimate return volatility correctly, then it is possible to place an upper bound on the volatility of the pricing kernel m_{t+1} that prices the assets:

$$\sigma^2(m_{t+1}) \leq \sigma^2(m_{V,t+1}) = RRA_V^2 \sigma^2(r_{m,t+1}) \tag{8}$$

Here, $m_{V,t+1}$ is the inter-temporal marginal rate of substitution between present and future wealth of a representative investor with relative risk aversion RRA_V, the latter is the relative risk aversion upper bound, and $\sigma^2(r_{m,t+1})$ is the volatility of the market excess return $r_{m,t+1}$. Based on (6), the pricing kernel volatility bound in (8) implies an upper bound on the explanatory power of any predictive regression of asset returns:

$$R^2 \leq (1+R_f)\sigma^2(m_{V,t+1}) \tag{9}$$

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 introspection, seems large. We will also experiment with a lower bound, i.e. $RRA \leq RRA_V = 2.5$, as this is the relative risk aversion 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 since 1926. This bound implies that the marginal investor would be willing to pay up to 5 percent per annum, arguably still a relatively large amount, to avoid a 20 percent volatility of his wealth.

3. Data

Our data comprise daily returns on the exchange rate against the US Dollar of the major industrial country currencies (except those that were replaced by the Euro) for the period 1971-2006 taken by Bloomberg at the close of business in London at 6:00 p.m. GMT. These currencies are the Australian and Canadian Dollar (AUD and CAD, respectively), the Japanese Jen (JPY), the British Pound (GPB), 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). As a proxy 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 is taken from Datastream. 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 similar to the results for the underlying currencies. As a proxy for the return on the market portfolio we use daily data on the S&P500 index last traded price provided Datastream.

4. Economic and Statistical Significance of Predictability

To assess the extent to which the returns on the currencies in our sample are predictable, we estimate simple rolling auto-regressions and auto-regressive moving averages of the currency returns. As shown by Taylor (1994), among others, ARIMA models of exchange rates, and thus ARMA models of the currency returns, can capture substantial predictability. Our estimated models are thus specifications of the general ARMA(p, q) model of currency returns, where p denotes the autoregressive lag order and q denotes the order of the moving average term:

$$y_t = const. + b_1 y_{t-1} + ... + b_p y_{t-p} + c_1 u_{t-1} + ... + c_q u_{t-q} + u_t \qquad (10)$$

We apply versions of (10) to both currency returns and to returns adjusted by the interest differential, i.e. the differential between the rate of the funding cost, in US Dollar, and the return from reinvesting the funds in each one of the currencies considered in this study, after the funds have been converted from US Dollars into these currencies. 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 duplications of indistinguishable predictability estimates, we only tabulate those based on return data.

In Figure 1, we plot the sequence of coefficients of determination R^2 of rolling 5-lag auto-regression, i.e. (10) with $p = 5$ and $q = 0$, of the returns on the currencies in our sample. The rolling auto-regressions are estimated over non-overlapping 1-year periods. Visual inspection of Figure 1 suggests a burst of predictability at the beginning of the sample period in the early 1970s. Subsequently, however, no systematic pattern is immediately distinguishable.

To assess the economic significance of predictability, we construct empirical counterparts of the upper bound on the coefficient of determination, i.e. the predictability bound, and we compare the coefficient of determination of various specifications of (10) with the constructed bound. We first estimate a simple ARMA(1,0), i.e. an AR(1) model of the daily returns on the currencies in our sample over a rolling 1-year window. The predictability bound is constructed using, in (8) and (9), yearly averages of GARCH(1,1) market volatility[1] estimates and an upper bound on relative risk aversion equal to 5, i.e. $RRA_V = 5$. In Figure 2, we plot the sequence of coefficients of determination R^2 of the rolling 1-lag auto-regression against the predictability bound. The R^2 of the estimated auto-regressive model often exceeds the bound, thus suggesting the presence of predictability in excess of the threshold that can be explained in terms of variation in risk premia. As shown in Table 1, the bound on daily predictability, i.e. on the R^2 of predictive regressions estimated using daily data, is about 0.25 percent on average over the sample period. The R^2 of the estimated auto-regressive AR(1) model is often much larger. Its average over the sample period ranges between 0.7 and 1 percent and takes values as

[1] This assumption is not very strong as volatilities are relatively easier to estimate than means.

high as 7.5 percent for all currencies except the Japanese Yen. The extent by which the estimated R^2 exceeds the bound is, obviously, even more pronounced in rolling auto-regressions (not reported) with higher order autoregressive and moving average terms.

To assess the importance of sampling error in the estimation of the coefficient of determination, we bootstrap 2-tailed confidence intervals for the coefficient of determination of ARMA(p, q) models of each currency return. To select the number of auto-regressive and moving-average terms in the ARMA(p, q) model, following a Box and Jenkin's (1976) type procedure, we search for a parsimonious specification for which the null of serial correlation in the residuals can be rejected for at least 20 lags (or about one month of trading). We test for residual serial correlation using a Ljung-Box (1978) Q-statistic. The results of this test are reported in Table 2. The auto-regressive model with five lags, ARMA(5,0) or AR(5), is reasonably successful in capturing the serial correlation of currency returns for all periods except the first one, i.e. 1971-1983 (data for the Australian Dollar were not available with sufficient continuity to conduct the test over this period). In tests not tabulated, we find that the null of serial correlation, over the period 1971-1983, can be rejected for lags up to the 9th order for the Japanese Yen and the British Pound and up to the 10th order for the Swiss Franc. As shown in Table 2, an ARMA(5,1) model of the Canadian Dollar returns manages to produce serially uncorrelated residuals up to the 20th order lag (actually, up to any order lag, as shown by tests not tabulated) but, again, its performance is poorer in the case of the Japanese Yen, the British Pound and the Swiss Franc. This is also the case with the ARMA(1,2) specification suggested by

Taylor (1994). In fact, its performance does not materially improve relative to the performance of the more parsimonious ARMA(5,1), i.e., the null of serial correlation can still be rejected only for lags up to the 9th order for the Japanese Yen and the British Pound and up to the 10th order for the Swiss Franc. We thus opt for parsimony and model all the series as an ARMA(5,0) model, except the returns on the Canadian Dollar over the period 1971-1983, for which we use an ARMA(5,1).

To conduct our bootstrapping experiment, we re-sample with replacement 1,000 times the residuals of the estimated auto-regressive model. By this process, we generate 1,000 separate bootstrapped currency return series, for which we then estimate the auto-regressive model and record the coefficient of determination. This generates a bootstrapped distribution based on 1,000 coefficients of determination estimates. In Table 3, we report the predictability upper bound and the bootstrapped confidence intervals for the coefficient of determination of the chosen model for each currency, estimated over the whole sample period and over three sub-periods of roughly equal length, 1971-1983, 1984-1995, 1996-2006. Under the 2.5 upper bound on relative risk aversion, i.e. under RRA_V = 2.5, we can reject the null that the estimated predictability does not violate the bound at conventional significance levels for almost all currencies in the sample. Under a less conservative upper bound on relative risk aversion, i.e. RRA_V = 5, we can reject this null only in 5 out of 17 cases. While the statistical evidence that the estimated predictability violates market efficiency is considerably weaker under the less conservative risk aversion bound, it should be kept in mind that the predictability bound itself might be 'loose', i.e. too little conservative. If the returns on the strategies that exploit predictability are less than perfectly correlated with the pricing kernel, i.e. if these strategies are a diversification opportunity, they would command a lower Sharpe ratio than the maximal Sharpe ratio of the economy and thus the upper bound would be less than the maximal ratio for the economy.

In the case of the GBP, the critical values of the estimated confidence intervals decrease over time in a clear fashion. In all other cases, the lack of a clear declining pattern in the critical values of the estimated confidence intervals contrasts with emerging evidence (in Pukthuanthong, Levich and Thomas (2006) and Neely, Weller and Joshua (2007)) that currency markets have over time become less susceptible to technical trading profits. This contrast can be resolved by noting that confidence intervals depend not only on the mean but also on the variability of a random variable. Therefore, while there might be a downward trend in the coefficient of determination and thus in predictability, these might have become more variable across samples and thus there might be more sampling error that makes them more difficult to be estimated with precision.

5. Transaction Costs and Predictability

The evidence reported in Table 3 suggests that the predictability bound has been violated, often to a statistical extent, for prolonged period of time. Part of the estimated predictability, however, might be un-exploitable due to transaction costs. To formally model the impact of transaction costs on predictability, it is useful to consider the strategies that would have to be implemented in order to optimally exploit it. To this end, we use an 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 is (4) as follows:

$$R^2 = \frac{\sigma_\mu^2}{\sigma_r^2} = \frac{1}{\sigma_r^2}\left(\frac{\mu'\mu}{T} - \overline{\mu}^2\right) = \frac{1}{T}\mu'\left(\frac{DD'}{T}\right)^{-1}\mu - \frac{\overline{\mu}^2}{\sigma_r^2} = \mu'(DD')^{-1}\mu - \frac{\overline{\mu}^2}{\sigma_r^2} \quad (11)$$

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, ...*T*, $\overline{\mu}$ is the unconditional mean return 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 aimed at exploiting predictability as a portfolio made up of as many positions as data points in the sample period, each with its own conditional Sharpe ratio. Recognising that, in daily and higher frequency data, the second term on the far right-hand side of (11) is negligible as it is the square of a typically small percentage number, we can approximate the coefficient of determination as follows,

$$R^2 \cong \mu'(DD')^{-1}\mu \quad (12)$$

Interestingly, if one neglects the possible temporal interdependencies across conditional volatilities, i.e. if one neglects GARCH effects, (12) can be interpreted as the squared maximal Sharpe ratio available by forming ‘portfolios’, i.e. strategies, of one-period positions in the currency under consideration. The weights with which each one-period position enters such strategy are then

$$W = (DD')^{-1}\mu \qquad (13)$$

In the context of our ARMA(p,q), the mean vector equals the conditional mean of (10), 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. $\sigma^2(y_t)I_{T\times T}$. In Figure 3, to illustrate, we report the time-varying weights, calculated using (13) and normalized to add up to unity over the time horizon of the strategy that exploits the predictability of the Canadian Dollar, based on an ARMA(5,1) specification over the period 1996-2006. The corresponding plots for the other currencies are not reported to save space. In all cases, there is substantial variation in the weights of the daily positions, as a result of the conditional time-variation on the mean of the return process, given the chosen ARMA(p,q) specification. We use these weights to calculate the returns of the strategies aimed at exploiting predictability. 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. In Table 4, we report the Sharpe ratios of these strategies. For all the currencies under consideration, except the Swiss Franc, transaction costs of 2 basis point are enough to lower the Sharpe ratio below the level that would correspond to violation of the tightest predictability bounds. With transaction costs of 3 basis points, the Sharpe ratios for the strategies based on the Australian Dollar, the Canadian Dollar, the Japanese Yen and the ECU/Euro become negative. With transaction costs of five basis points, all Sharpe ratios are negative. Nonetheless, the evidence that, for some currencies such as the Swiss Franc and to a lesser degree the ECU/EUR and the British Pound, the Sharpe ratio is relatively high even after accounting for transaction costs of two basis points would seem to point out that there are predictability-based strategies that professional currency traders might find attractive. There is, however, substantial evidence that 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, Chernenko, Howorka and Write (2006), at least for trades executed over a daily horizon. In light of these considerations, the speed at which Sharpe ratios decrease as a function of transaction costs implies that it is in fact rather unlikely that any opportunity to earn sizable abnormal risk-adjusted returns exists.

As shown in Appendix A using a logic similar to Roll (1984), the bid-ask bounce induces an amount of predictability that depends on the relative magnitude of the bid-ask spread and exchange rate variability. In particular, (A5) quantifies the approximate impact of the bid-ask bounce on the auto-regression coefficient of determination. We use this result to assess the level of transaction costs required to explain the observed excess-predictability over the 1972-2006 period. As shown in Table 5, bid-ask transaction costs in the region of 20 to 30 basis points are sufficient to reduce exploitable predictability below the level corresponding to the tightest predictability bound, i.e. the bound corresponding to $RRA = 2.5$. As implied by the evidence provided by the literature on price pressure, i.e. Evans and Lyons (2002) and Berger, Chernenko, Howorka and Write (2006), these levels of transaction costs are to be expected for large transactions.

6. Bayesian Uncertainty and Predictability

To further assess the importance of sampling error in the estimation of the coefficient of determination, we generate the posterior distribution of the currency returns variance and of the coefficient of determination using a Gibbs sampling methodology, as in Gelfand and Smith (1990). Gibbs sampling is applicable when the joint distribution of two or more variables is not known explicitly, but the conditional distribution of each variable is known. We assume that the conditional distribution of the predictive regression residuals is normal, which implies a multivariate normal conditional distribution of the regression coefficient estimates. More formally, consider the AR(5) model, i.e. (10) with p = 5 and q = 0. For notational convenience, we write this model in matrix form, as $Y = \beta X + u$, where $Y = \{y_t\}$, $X = \{1, y_{t-1}, ...y_{t-5}\}$, , $\beta = [const., b_1......b_5]'$ and $u = \{u_t\}$. Letting T denote the sample period length, denoting by I_T a $T \times T$ identity matrix and assuming that $u \mid x \sim N(0, h^{-1}I)$ with prior $\beta \sim N(\beta_p, H^{-1})$ and $v\sigma^2 h \sim \chi^2(v)$, where v denotes the degrees of freedom of a Chi-squared distributed random variable, the posterior density function is proportional to

$$h^{\frac{T+v}{2}-1} e^{-\frac{h}{2}(Y-X\beta)'(Y-X\beta)} e^{-\frac{1}{2}(\beta-\beta_p)' H(\beta-\beta_p)} e^{\frac{-1}{2}v\sigma^2 h}$$

Then, conditional on h, Y and X, β is normal with precision (or variance) $H^* = H\beta_p + hX'X\beta_{OLS}$, where β_{OLS} is the OLS estimate and, conditional on β, Y and X, the precision is chi-squared with v degrees of freedom, i.e. $\left[v\sigma^2 + (Y - X\beta)'(Y - X\beta)\right] h \sim \chi^2(v)$. Drawing in turn from these two conditional densities, the Gibbs sampling algorithm generates an instance from each density conditional on the current values of the other variables. It can be shown, see for example Gelman, Carlin, Stern, and Rubin (1995), that the sequence of samples comprises a Markov chain, and the stationary distribution of this Markov chain is just the required joint distribution of the precision and the regression coefficients.

In our implementation of the Gibb's algorithm, our initial precision and coefficients come from OLS estimates[2]. In Table 6, we compare the predictability bound under relative risk aversion bounds equal to 2.5 and 5 with the 5th and 95th percentiles of the posterior distribution of the coefficient of determination of the autoregressive model with 5 lags, over the whole sample period and over 3 sub-periods of roughly equal length. Only for the Canadian dollar, in the period 1971-1983, we can reject the null that the bound is not violated when relative risk aversion of the marginal investor is as large as 5. For the GBP, however, we can reject this null under the lower 2.5 upper bound on risk aversion for all sample periods but the last.

[2] RATS code to generate the posterior distribution of the model parameters and of its coefficient of determination is available on the corresponding author's website, www.valeriopoti.com.

7. Conclusions and Future Work

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 loose upper bound on relative risk aversion, predictability often violates a theoretically motivated upper bound. Taken at face value, this evidence implies violation of the EMH under a broad class of asset pricing models and for conservative to realistic values of the marginal investor's relative risk aversion. A closer scrutiny reveals, however, that the performance of strategies that attempt to optimally exploit predictability is very sensitive to the level of transaction costs, to the point that much of the observed predictability is un-exploitable given realistic assumptions about transaction costs.

Crucially, our analysis implies that, while it is relatively easy to find strategies that appear profitable, much or all the abnormal risk-adjusted profitability is soaked up by transaction costs, at least for large size transactions. Our findings allow to rationalize both the frequent occurrence of studies that find abnormally profitable strategies, before transaction costs, and the persistence of this apparent excess-profitability. We also find that there is substantial sampling error, especially of a Bayesian type, in the estimation of the coefficient of determination of predictive regressions, and that this sampling error might have increased in recent years. This makes it difficult to reach firm conclusions about whether predictability has genuinely decreased over time.

A possible avenue of future research is a more formal investigation of whether the estimated R^2 series contains a time trend and one or more structural breaks. Considering cross-rates and a wider sample of countries 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 of possible sources of cross-sectional variation in the predictability pf currency returns). Another obvious extension is to consider emerging economies currencies. These extensions would make it possible to better address the important question of whether there is predictability in excess of a level that can be judged consistent with the EMH, even after transaction costs, and whether, if this is the case, this becomes milder over time as a result of learning by economic agents.

Appendix A

We assume that transaction costs are fully reflected in the bid-ask spread bounce S_t and that traded prices are the sum of gross prices P_t^* and the bid or ask spread from gross prices, which takes a positive or a negative sign depending on whether the transaction is the result of a purchase or a sale order, i.e.

$$P_t = P_t^* + S_t \ .$$

We can then write the covariance between two successive price changes as $Cov(\Delta P_t, \Delta P_{t+1}) = Cov(\Delta P_t^* + \Delta S_t, \Delta P_{t+1}^* + \Delta S_{t+1})$. This can be written out as

$$Cov(\Delta P_t, \Delta P_{t+1}) = Cov(\Delta P_t^*, \Delta P_{t+1}^*) + 2Cov(\Delta P_t^*, \Delta S_{t+1}) + Cov(\Delta S_t, \Delta S_{t+1}) \quad \text{(A1)}$$

Following Roll (1984), assume that the bid-ask bounce S_t can take, with equal probability, values *s* or -*s*. This implies that $S_t = \pm s = \frac{\mathbf{ask_t - bid_t}}{\mathbf{2}}$ and that the gross price P_t^* coincides, by construction, with the mid price. Also assume that, in successive draws, the probability of each outcome is the same regardless of the outcome in previous draws. The number of occurrences of *s* or –*s* in *T* trials is then distributed as a binomial random variable with mean $\frac{1}{2}T$ and variance $T\frac{1}{2}\left(1-\frac{1}{2}\right) = \frac{1}{4}T$. The variable S_t, thus, takes values *s* or –*s* with equal probability in repeated draws and it is identically and independently distributed over time with zero mean and variance equal to $\frac{s^2}{4}$. The covariance in (A1) between two successive price changes, then, simplifies to the following

$$Cov(\Delta P_t \ ,\Delta P_{t+1}) = Cov(\Delta P_t^*,\Delta P_{t+1}^*) - Cov(S_t, S_{t+1}) = Cov(\Delta P_t^*,\Delta P_{t+1}^*) - \frac{s^2}{4} \quad \text{(A2)}$$

The above result is analogous to the relation between the size of the bid-ask spread and price auto-covariance derived by Roll (1984), except for the $Cov(\Delta P_t^*,\Delta P_{t+1}^*)$ term that allows for serial correlation in the mid-price and thus for serial correlation that does not depend on the bid-ask bounce. Using (A2), we can decompose the coefficient of determination in auto-regressions as follows,

$$\begin{aligned} R^2 &= \left[\frac{Cov(\Delta P_t \ ,\Delta P_{t+1})}{Var(\Delta P_t \ )}\right]^2 = \left[\frac{Cov(\Delta P_t^*,\Delta P_{t+1}^*)}{Var(\Delta P_t \ )} - \frac{s^2}{4Var(\Delta P_t \ )}\right]^2 \\ &= \left[\frac{Cov(\Delta P_t^*,\Delta P_{t+1}^*)}{Var(\Delta P_t \ )}\right]^2 + \frac{s^4}{16Var(\Delta P_t \ )^2} - \frac{Cov(\Delta P_t^*,\Delta P_{t+1}^*)}{Var(\Delta P_t \ )}\frac{s^2}{Var(\Delta P_t \ )}\frac{1}{2} \end{aligned}$$

Denoting by R^{2*} the coefficient of determination in predictive regressions of the mid-price and recognizing that $R^{2*} \cong \left[\frac{Cov(\Delta P_t^*,\Delta P_{t+1}^*)}{Var(\Delta P_t \ )}\right]^2$,

$$\begin{aligned} R^2 &= R^{2*} + \frac{s^4}{16Var(\Delta P_t \ )^2} - R^{2*}\frac{s^2}{2Var(\Delta P_t \ )} \\ &= R^{2*}\left[1 - \frac{s^2}{2Var(\Delta P_t \ )}\right] + \frac{s^4}{16Var(\Delta P_t \ )^2} \end{aligned} \quad \text{(A3)}$$

Thus, solving for R^{2*}

$$R^{2*} \cong \frac{R^2 - \dfrac{s^4}{16Var(\Delta P_t)^2}}{1 - \dfrac{s^2}{2Var(\Delta P_t)}} \tag{A4}$$

The bid-ask bounce is typically a small fraction of currency volatility and thus s^2 should be negligible relative to $2Var(\Delta P_t)$ and therefore $1 - \dfrac{s^2}{2Var(\Delta P_t)} \cong 1$.

Letting $R^2_{spread} \equiv \dfrac{s^4}{16Var(\Delta P_t)^2}$, we can thus re-write (A4) as follows

$$R^{2*} \cong R^2 - R^2_{spread} \tag{A5}$$

While R^{2*} can be interpreted as the portion of predictability that can be exploited by trading at the mid price, R^2_{spread} can be interpreted as the approximate impact on predictability of the bid-ask bounce and thus the portion of predictability of the traded price ("the last price") that would be costly to exploit. Thus, (A5) provide a way of adjusting predictability estimates based on time series of returns calculated using the "last price" to take transaction costs into account.

Figure 1
Predictability of Currency Returns
(Non-Overlapping Annual Periods)

Panel A
(1972-2006)

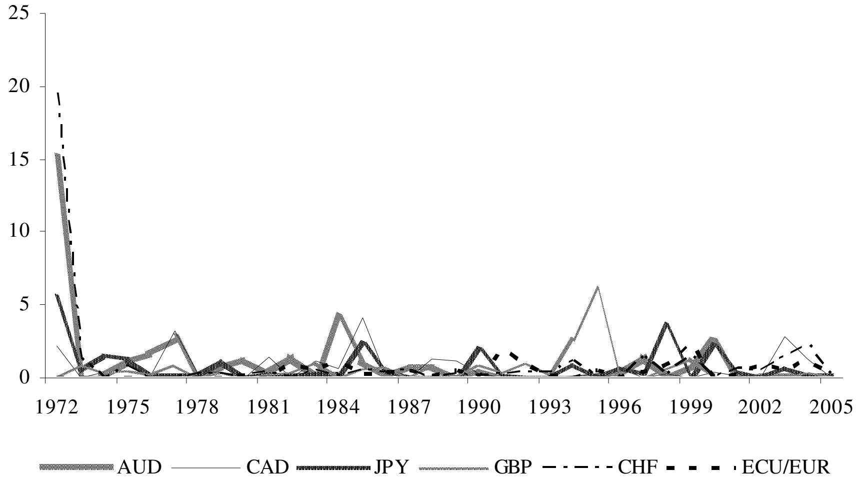

Panel B
(1975-2006)

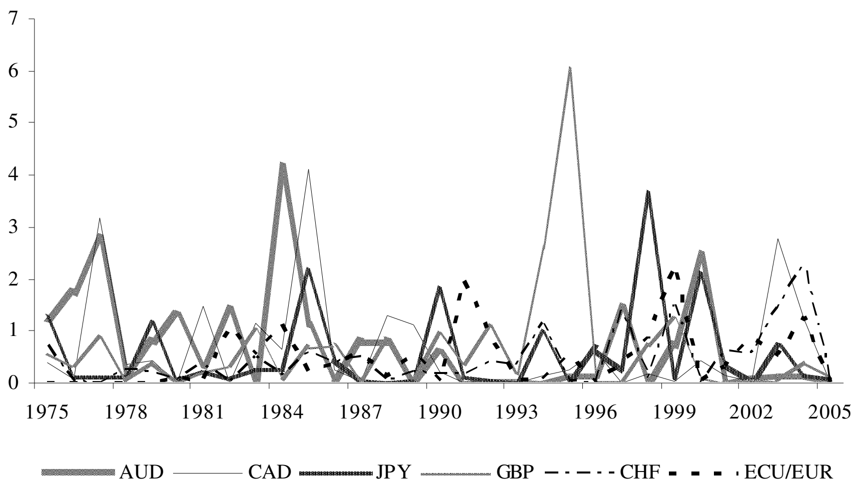

Notes. These figures plot the series of coefficients of determination of auto-regressions of each currency return on its own 5 lags and a constant, i.e. AR(5), estimated over sequential non-overlapping 1-year periods.

Figure 2
Economic Significance of Currency Return Predictability

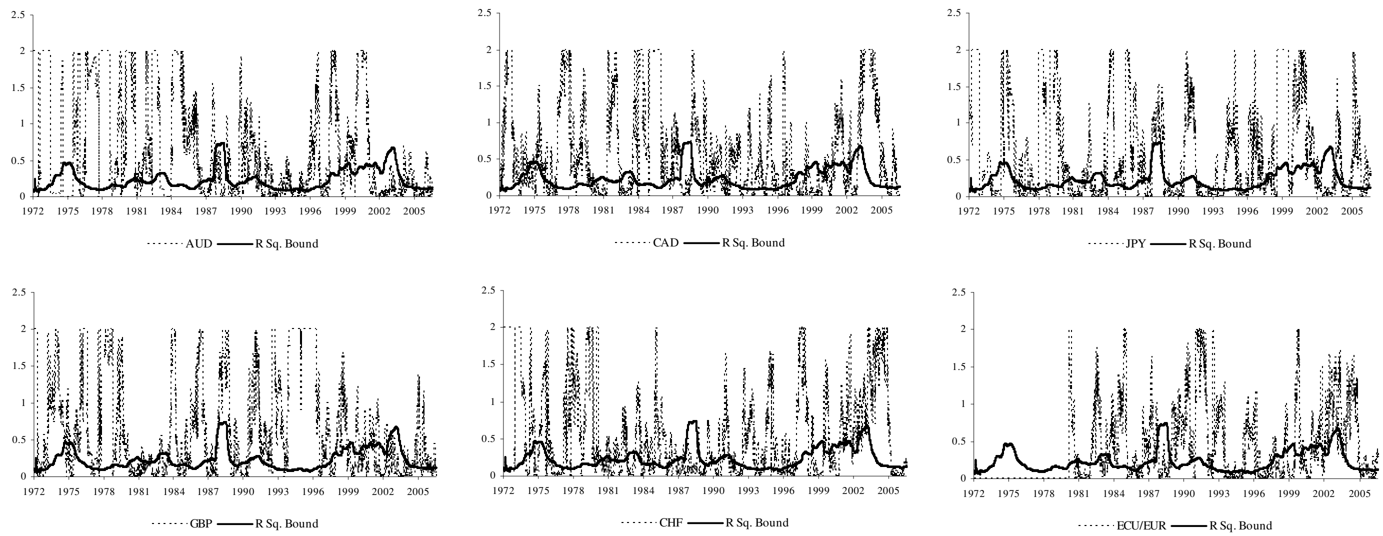

Notes. These figures plot the sequences of the percentage coefficients of determinations (shown by the dotted line) of rolling auto-regressions for each major currency in our sample against their conditional upper bound (shown by the solid line). The latter is computed under a relative risk aversion upper bound of 5. The autoregressive model includes 1 lag. The estimation window of each auto-regression is one year and the sample period is 1971-2006. The values of all the series have been cut off at 7.5 to improve visual clarity.

Table 1
Descriptive Statistics for Estimated R^2 in Rolling Regressions

	AUD	CAD	JPY	GBP	CHF	ECU	ECU/ EUR	R Sq. Bound
Average	0.97	0.64	0.68	0.85	0.75	0.75	0.78	0.24
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06
Max	7.50	5.65	7.50	7.50	7.50	7.50	7.50	2.50

Notes. This table reports descriptive statistics, average, minimum and maximum, for the estimated coefficient of determination of an AR(1) model estimated over a rolling 1-year window of daily data. The sample period is 1972-2006 and the reported figures are in percentage. The last column reports the same statistics for the predictability upper bound computed under a relative risk aversion upper bound set to 5.

Table 2
ARMA(p,q) Model Selection
Ljung Box Tests

	1971-2006	1971-1983	1984-1995	1996-2006	1971-1983	1971-1983
		ARMA(5,0) Q(20-5)			ARMA(5,1) Q(20-6)	ARMA(1,2) Q(20-3)
AUD			22.16	21.28		
			(0.103)	(0.130))		
CAD	*23.14	*23.07	16.00	16.3	13.11	19.14
	(0.081)	(0.082)	(0.380)	(0.351)	(0.517)	(0.321)
JPY	**30.40	**41.46	*22.52	14.69	**32.27	**32.16
	(0.010)	(0.000)	(0.094)	(0.472)	(0.003)	(0.014)
GBP	20.60	**39.65	14.33	13.31	**37.17	**57.23
	(0.149)	(0.000)	(0.500)	(0.575)	(0.000)	(0.000)
CHF	14.17	**29.65	9.10	10.91	**27.16	**29.26
	(0.512))	(0.002)	(0.871)	(0.758)	(0.018)	(0.032)
ECU/EUR	15.66	18.26	14.83	7.39	18.22	21.83
	(0.404)	(0.249)	(0.463)	(0.945)	(0.196)	(0.191)

Notes. This table reports Ljung-Box test statistics and associated significance levels for the entire sample period 1971-2006 and three sub-samples of about equal length, 1971-1983, 1984-1995, 1996-2006. For added visual clarity, we use one and two asterisks to draw attention to test statistics significant at the 10 and 5 percent level, respectively.

Table 3
Statistical Significance of Predictability
Bootstrapped Percentage 2-Tailed Confidence Intervals

	1971-2006	1971-1983	1984-1995	1996-2006
$\text{Bound}_{RRA=2.5}$	0.06	0.05	0.06	0.07
$\text{Bound}_{RRA=5}$	0.24	0.20	0.22	0.26
		AR(5)		
AUD				
5.0%			*0.14 0.79	*0.12 0.53
2.5%			*0.10 0.94	*0.09 0.64
CAD				
5.0%	*0.08 0.29	**0.49 1.28	*0.10 0.50	**0.44 1.24
2.5%	*0.06 0.34	**0.41 1.42	*0.08 0.60	**0.37 1.42
JPY				
5.0%	0.05 0.22	*0.12 0.57	*0.09 0.46	*0.10 0.50
2.5%	0.03 0.25	*0.09 0.66	*0.06 0.55	*0.07 0.56
GBP				
5.0%	**0.56 1.57	**0.47 1.48	*0.17 0.70	*0.07 0.40
2.5%	**0.43 1.73	**0.38 1.65	*0.12 0.80	0.05 0.48
CHF				
5.0%	0.02 0.12	*0.08 0.40	*0.09 0.46	**0.28 0.97
2.5%	0.01 0.15	*0.06 0.48	*0.07 0.57	*0.22 1.07
ECU/EUR				
5.0%			*0.14 0.60	*0.19 0.77
2.5%			*0.10 0.70	*0.14 0.86
		ARMA(5,1)		
CAD				
5.0%		**0.75 1.75		
2.5%		**0.62 1.92		

Notes. The first two rows of this table report percentage unconditional upper bounds on the explanatory power of predictive regressions of each major currency in our sample under a relative risk aversion upper bound equal to 2.5 and 5, respectively. The other rows report, for each currency, 90 and 95 two-tailed confidence intervals for the coefficient of determination (in percentage) of 5-lag auto-regressions AR(5) for all currency and ARMA(5,1) for CAD. The sample period is 1971-2006 and three sub-samples of about equal length, 1971-1983, 1984-1995, 1996-2006. In the table, one and two asterisks denote when the upper bound is violated at the significance level corresponding to the value reported in the left-most column under a RRA bound of 2.5 and 5, respectively.

Figure 3
Canadian Dollar Predictability

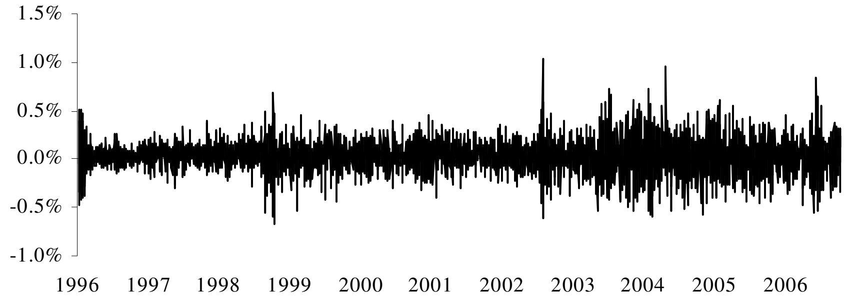

Notes. This Figure plots the time-varying weights of the strategy that optimally exploits the Canadian Dollar return predictability, based on estimates from an ARMA(5,1) model. The weights are rescaled in such a way that they add up to 1 over the 1996-2006 sample period.

Table 4
Percentage Sharpe Ratios of Predictability-Based Strategies

t.c. (bps)	0	1	2	3	5
AUD	57.3	37.2	17.4	-2.3	-41.7
CAD	128.5	81.45	34.5	-12.4	-106.3
JPY	49.1	19.03	-11.0	-41.1	-101.2
GBP	48.4	34.2	20.0	5.7	-22.7
CHF	104.3	75.9	47.7	19.6	-36.7
ECU/EUR	82.1	52.4	22.7	-7.1	-66.6

Notes. This Table reports percentage annualized Sharpe ratios of strategies that optimally exploit estimated predictability, as a function of various levels of transaction costs (in basis points in the tow row). The sample period is 1996-2006. The annualized maximal SR bound under a RRA upper bound equal to 2.5 and 5 is 44 and 85 percent, respectively.

Table 5
Predictability and Transaction Costs

	AUD	CAD	JPY	GBP	CHF	ECU/EUR
bound	0.06	0.06	0.06	0.06	0.06	0.06
R^2	0.09	0.37	0.07	0.05	0.22	0.14
Excess R^2	0.03	0.31	0.01	-0.01	0.16	0.08
$16\sigma^2(\Delta P)$	0.12	0.03	2081.63	0.02	0.29	0.18
Implied s	0.23	0.20	22.81		0.54	0.36
Implied s/P	0.19	0.16	0.13		0.29	0.41

Notes. This Table reports the percentage bid-ask spread, denoted by s, in dollar terms and as a fraction of the average value of the exchange rate under consideration, denoted by s/P, that is required to justify the excess of estimated predictability over the bound corresponding to RRA = 2.5. The excess R^2 denotes the excess of the coefficient of determination in predictive regressions over the predictability bound. The sample period for the calculations is 1972-2006.

Table 6
Statistical Significance of Predictability
Gibbs Sampling

	1971-2006	1971-1983	1984-1995	1996-2006
$Bound_{RRA=2.5}$	0.06	0.05	0.06	0.07
$Bound_{RRA=5}$	0.24	0.20	0.22	0.26
AUD				
R^2	0.11	0.14	0.31	0.13
5^{th} percent.	-0.02	-0.30	-0.02	-0.23
CAD				
R^2	0.13	0.74	0.16	0.18
5^{th} percent.	0.00	**0.28	-0.01	-0.02
JPY				
R^2	0.07	0.19	0.13	0.06
5^{th} percent.	-0.04	-0.12	-0.03	-0.12
GPB				
R^2	0.21	0.65	0.34	0.12
5^{th} percent.	*0.06	*0.11	*0.11	-0.07
CHF				
R^2	0.01	0.08	0.07	0.03
5^{th} percent.	-0.10	-0.26	-0.10	-0.13
ECU/EUR				
R^2	0.17	0.92	0.29	0.20
5^{th} percent.	-0.01	-0.17	-0.00	0.00

Notes. The first two rows of this table report percentage unconditional upper bounds on the explanatory power of predictive regressions of each major currency in our sample under a relative risk aversion upper bound equal to 2.5 and 5, respectively. The other rows report for each currency the percentage coefficient of determination of 5-lag auto-regressions and the 5^{th} percentile of its posterior distribution. The posterior probability is estimated using the Gibbs sampling methodology. One and two asterisks denote cases in which we can reject at the 95 percent level the null that the upper bound is not violated under relative risk aversion bounds of 2.5 and 5, respectively.

Bibliography

Andersen, T.G., Bollersev, T., Diebold, F.X. and H. Ebens, 2001, The Distribution of Stock Return Volatility, Journal of Financial Economics 61, 43-76.

Andersen, T.G., Bollerslev, T., Diebold, F.X. and J. Wu, 2004, Realized Beta: Persistence and Predictability, PIER Working Paper No. 04-018.

Berger, D.W., Chaboud, A.P., Chernenko, S.V., Howorka E. and J.H. Write, 2006, Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data, Board of Governors of the Federal Reserve System.

Box, G.E.P. and G.M. Jenkins (1976), Time Series Analysis, Forecasting and Control, San Francisco: Holden Day.

Campbell, J.Y. and S.B. Thompson, 2005, Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?, unpublished manuscript.

Cerný, A. and S.D. Hodges, 2000, The Theory of Good-Deal Pricing in Financial Markets, Tanaka Business School Discussion Paper TBS/DP04/14.

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: Spring, 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., 2001, Asset Pricing, Princeton University Press, Princeton.

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.

Dooley, M.P. and J. Shafer. Analysis of Short Run Exchange Rate Behavior: March 1973 – September 1975, 1976, 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, Martin D. D., and R.K. Lyons, 2002, Order flow and exchange rate dynamics, Journal of Political Economy 110, 170–180.

Gelfand, A.E. and A.F.M. Smith, 1990, Sampling Based Approaches to Calculating Marginal Densities, JASA 85, 398-409.

Gelman, A., Carlin, J.B., Stern, H.S. and Donald B. Rubin, 1995, Bayesian Data Analysis. London: Chapman and Hall.

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

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

Hansen, L.P. and R. Jagannathan, 1991, Implications of Security Market Data for Models of Dynamic Economies, Journal of Political Economy 99, 225-262.

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.

Ljung, G.M. and G.E.P. Box (1978), On a Measure of Lack of Fit in Time Series Models, Biometrika 67, 297-303.

Neely, C., Weller, P., and R. Dittmar, 1997, Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach, Journal of Financial & Quantitative Analysis 32, 405-426.

Neely, C.J., Weller, P.A. and U. Joshua, 2007, The Adaptive Markets Hypothesis: Evidence from the Foreign Exchange Market, FRB of St. Louis Working Paper No. 2006-046B. Available at SSRN: http://ssrn.com/abstract=922345

Pesaran, M.H. and A. Timmermann, Predictability of Stock Returns: Robustness and Economic Significance, Journal of Finance 50, 1201-1228.

Olson, D., 2004, Have Trading Rule Profits in the Foreign Exchange Markets declined through time?, Journal of Banking and Finance 28, 85-105.

Pukthuanthong, K., Levich, R. M. and L.R. Thomas III, 2007, Do Foreign Exchange Markets Still Trend?, Journal of Portfolio Management, forthcoming.

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., 1976, Arbitrage Theory of Capital Asset Pricing, Journal of Economic Theory 13, 341-60.

Schulmeister, S., 2006, The Interaction Between Technical Currency Trading and Exchange Rate Fluctuations, Finance Research Letters 3, 212-233.

Stremme, A., Basu, D. and A. Abhyankar, 2005, The Optimal Use of Return Predictability: An Empirical Analysis, AFA 2006 Boston Meetings, Paper Available at SSRN: http://ssrn.com/abstract=687083.

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.