

Do Moving Average Trading Rule Results Imply Nonlinearities in Foreign Exchange Markets?

Blake LeBaron
Department of Economics
University of Wisconsin - Madison
1180 Observatory Drive
Madison, WI 53706
(608) 263-2516
BLAKEL@macc.wisc.edu

March 1992
Revised: October 1992

ABSTRACT

This paper tests whether fitted linear models can replicate results from moment tests inspired by moving average technical trading rules for weekly foreign exchange series. Estimation is performed using standard OLS and maximum likelihood methods, along with a simulated method of moments technique which incorporates the trading rule moments into the estimation procedure. Results show that linear models are capable of replicating the trading rule moments along with the small autocorrelations observed in these series. This result holds for parameter values estimated using SMM and GARCH disturbances, but not for parameters estimated using maximum likelihood. The estimated models are simulated to examine the amount of predictability over long horizons.

Acknowledgments

The author is grateful to Rochelle Antoniewicz, W. A. Brock, Wen-Ling Lin, and Ken West for helpful comments on an earlier draft. The author also is grateful to the National Science Foundation (SES-9109671) and the University of Wisconsin Graduate School for support.

I. Introduction

When work on nonlinearities in financial time series began it was seen by many technical traders as justification for their work. They saw a clear connection between technical trading rules and nonlinearities.¹ For some rules this connection is indisputable. Rules that look for general patterns such as “head and shoulders” and other more complicated figures are clearly attempting to find some kind of nonlinearity in these series. For other rules this link is not so obvious, and it is not clear whether rules that are simple trend followers, such as the moving average or oscillator rules, are connected to nonlinearities. This paper will test the strength of some of these linkages.

Several recent papers have presented conflicting evidence for the presence of nonlinearities. Hsieh(1989) finds most of the evidence for nonlinearities for daily exchange rates to be coming from changing conditional variances. Diebold and Nason(1990), and Meese and Rose(1990) found few improvements in out of sample forecasts using nonparametric techniques. However, Kim(1989), and LeBaron(1992), show small out of sample forecast improvements using models which combine forecasts of the conditional mean with information contained in conditional variances.

Results using technical trading rules have been more consistent. Dooley and Shafer(1983) presented some of the earliest evidence suggesting that technical trading rules might be detecting changes in conditional mean returns in foreign exchange rate series. Sweeney(1986) also finds results supportive of the profitability of similar rules. Also, studies by Schulmeister(1987), Curcio and Goodhart(1992), and Taylor(1992a,b) find similar evidence for even more extensive sets of rules and data series.² Recent papers by LeBaron(1991), and Levich and Thomas(1991) have followed the paper by Brock et al.(1992) on stock returns by using bootstrap simulations to demonstrate the statistical significance of the moving average technical trading rule results against certain alternatives.

The results of these two lines of research, nonlinear forecasting, and moving average technical trading rules, are interesting, but are they detecting the same characteristics in these series? The moving average types of trading rules appear well designed to pick out long range persistent trends rather than the more complicated dynamics generated by certain nonlinear models. Antoniewicz(1992) contains some of the first power tests done on moving average rules. She found that the rules had little power against some simple nonlinear models. This suggests that the profitability of trend following trading rules may be more closely

¹ For a recent example of this see, “Technical Analysis: Tilting at Chaos”, in *The Economist* (August 15, 1992).

² Taylor and Allen(1992) report on survey results which find that 64 percent of dealers surveyed used moving averages, or some other trend following system, as a part of their forecasting system.

related to models with persistent trends such as those used in Engel and Hamilton(1990) and Taylor(1980), or long range forecasts such those in Mark(1992).³

This paper approaches this problem by searching the space of traditional linear models. New estimation techniques are used which incorporate the technical trading rules in the estimation procedure. Parameters are estimated using simulated method of moments to match both trading rule results, and low order autocorrelations. These models are then simulated and compared with the original series. The objective is to see if there do exist linear models which can match the trading rule predictability while at the same time being relatively uncorrelated. The results indicate that the new estimation procedures change the estimated values for some of the parameters, and that simulations using these parameters show a good match between the simulated linear models and the actual series on both these dimensions. Tests are also performed on linear models with GARCH disturbances, and these do a good job of matching the original series. In all cases the parameters of the estimated models are often close to being nonstationary. This is consistent with some of the evidence for persistence of returns to forward speculation presented in Evans and Lewis(1992a).

These results should not be taken as indicating that foreign exchange series are necessarily linear. Earlier papers have documented both changes in conditional variances, and some amount of nonlinear forecastability. What this paper does suggest is that the connections between moving average trading rules and nonlinearities may be somewhat weak in foreign exchange series. The fact that linear models do a good job of replicating the trading rules also strengthens the case for models containing slowly changing risk premiums either do to consumption risk, or learning about policy changes.⁴ Models with different trader types such as Brock(1992), Frankel and Froot(1990), Goodhart(1988), and Tsibouris(1992), are also consistent with these results when they generate the necessary persistence in the observed series.

Some of the contributions of this paper are methodological. Using the technical trading rules as part of a simulated method of moments estimation procedure is a new approach. In future work this estimation strategy will be extended to theoretical models of exchange rate behavior.⁵ These results also ask important questions about estimation and the parametric bootstrap.⁶ When a model specification is rejected using a set of estimated parameters how confident can we be in rejecting that entire class of models? This question

³ In a recent paper Neftci(1991) proves that trading rules only give useful information beyond that from linear forecasting models when nonlinearities are present. However, this assumes that the true parameters for the linear model are known.

⁴ See Lewis(1989), and Evans and Lewis(1992b) for models where agents are learning about policy changes.

⁵ Canova and Marrinan(1990) is an early application of simulated method of moments for structural exchange rate model estimation.

⁶ Tsay(1992) gives a good description of the parametric bootstrap approach.

is intimately related to what time series characteristics we are interested in replicating. Unfortunately, the bootstrap does not give a lot of information about how well models simulated at other parameter values do in replicating certain features of interest. Results in this paper find that this may be a question that needs to be considered more seriously when using parametric bootstrap style specification tests.

Finally, the forecastability of the fitted models is estimated using a long simulated time series. Several interesting results emerge when this analysis is performed. First, the amount of forecastability, in terms of Sharpe ratios, is very sensitive to the ARMA parameters. Second, the moving average trading rules perform very well in comparison to traditional time series forecasts. This is very surprising since the traditional models are given the true parameter values to make their forecasts.⁷

Section II provides a description the technical trading rules used, the SMM estimation procedure, and the linear models that will be fit. Section III presents the empirical results. Section IV looks at simulations of the actual fitted models to test the amount of predictability of different models, and compares moving average prediction methods with more traditional time series forecasts. Section V presents conclusions and suggestions for future research.

II. Technical Trading Rules and Exchange Rate Models

A. Moving Average Rules

This section outlines the technical rules used in this paper which are closely related to those used by actual traders. All the rules are of the moving average or oscillator type. Signals are generated based on the relative levels of the price series and a moving average of past prices,

$$m_t = (1/L) \sum_{i=0}^{L-1} p_{t-i}.$$

One possible trading rule based on the moving average generates a buy signal when the current price level is above the moving average and a sell signal when it is below.⁸

Some of the estimation procedures used here will be based on moving average trading rules. Define the return at time t as

$$r_t = \log(p_t) - \log(p_{t-1}).$$

⁷ These results are similar to some results reported in Taylor(1992a,b) where the trading rules perform well in relation to rules based on traditional time series forecasts.

⁸ There are many variations of this simple rule in use. One is to replace the price series with another moving average. A second modification is to only generate signals when the price differs from the moving average by a certain percentage. Many other modifications are discussed in Schulmeister(1987), Sweeney(1986), and Taylor(1992a).

Now define a trading rule moment as,

$$E\left\{S\left(\frac{p_{t-1}}{m_{t-1}}\right)r_t\right\}, \quad (2.1)$$

where $S(x) = 1$ if $x \geq 1$ and $S(x) = -1$ if $x < 1$. This is close to the return for an investor taking the appropriate long and short positions in the two currencies. This still ignores interest differentials and transactions costs and cannot be viewed as a true test of the usefulness of this trading rule.⁹ This paper uses this moment as a useful feature for estimation purposes.

Results from the original foreign exchange series are compared with those from simulated null models for foreign exchange movements using these trading rule moments. Also, more standard aspects of the time series will be used such as autocorrelations. These tests can be extended to higher order moments such as variance, skewness, and kurtosis. Tests of these moments will not be done here, but they are reported in LeBaron(1991).

B. Null Models for Foreign Exchange

The trading rule specification presents a very specific challenge for various candidate models for the movements in foreign exchange series. However, the usefulness of results from these tests also depends on the choice of null models for simulation and comparison. Since the important point of this paper is to test linear specifications, only linear models will be used. Also, higher moments will be ignored, focusing attention on modeling conditional means. This is done because results in LeBaron(1991) show that conditional variances in foreign exchange series do not depend on moving average trading rules. In other words, the variance is the same during buy and sell periods.¹⁰

Most of the series studied here show some amount of autocorrelation, so short range autoregressive models(AR) of the form,

$$r_t = a + b_1 r_{t-1} + b_2 r_{t-2} + \epsilon_t,$$

will be used. These models test whether the small amount of autocorrelation seen in the series along with the trading rule tests can be modeled using a simple short range AR. Later sections of the paper will include some longer range AR's.

⁹ See LeBaron(1991) for some tests accounting for transactions costs.

¹⁰ This differs from results on stock returns in Brock et al.(1991). Also, this does not say that conditional variances are constant, but that they do not depend on the trading rule buy or sell signals. Papers by Hsieh(1988) and Baillie and Bollerslev(1988) present evidence showing that conditional variances are changing. Finally, the conditional variances may depend on the price-moving average ratios in a more complicated fashion than through the coarse buy-sell signal.

A more interesting process for the foreign exchange series includes a stochastic trend. Models of this form have been suggested by many authors including Taylor(1986,1992a) and Hodrick and Srivastava(1987). If the trend at time t , μ_t , follows,

$$\mu_t = \rho\mu_{t-1} + \eta_t, \quad (2.2)$$

and the return at time t is,

$$r_t = \mu_t + \epsilon_t, \quad (2.3)$$

then the returns process r_t follows an ARMA(1,1) process. Let σ_η^2 be the variance of the iid. process η_t , and σ_ϵ^2 be the variance of iid. process for ϵ_t . The ARMA process for the returns series will be,

$$r_t = \rho r_{t-1} - \rho \left(\frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_\eta^2} \right)^{(1/2)} e_{t-1} + e_t, \quad (2.4)$$

where the variance of the new independent disturbance e_t is $\sigma_\epsilon^2 + \sigma_\eta^2$. This puts certain constraints on the parameters of the estimated ARMA model. First, the signs of the MA and AR components should be reversed. Second, the MA term must be smaller than the AR term (in absolute value). The size of this reduction is determined by the ratio of the variances of η_t and ϵ_t

C. Simulated Method of Moments Estimators

Estimation of the linear models will be done using standard OLS and maximum likelihood procedures augmented by a technique which explicitly takes the trading rule moments into consideration. Simulated method of moments will be used to estimate parameters for stochastic processes that match both autocorrelations and the trading rule results. This technique was developed for cross sectional data by McFadden(1989) and Pakes and Pollard(1989). It is extended to time series cases in Duffie and Singleton(1989) and Ingram and Lee(1991). This section will briefly outline the necessary assumptions and the techniques used in this paper. It looks at the conditions set forth in Duffie and Singleton(1989) in the context of the problems addressed here.

The SMM technology is a method for estimating the parameters for a transition function for a stochastic process Y_t in R^N ,

$$Y_{t+1} = H(Y_t, \epsilon_{t+1}, \beta_0),$$

where β_0 is the parameter vector, and ϵ_{t+1} is a vector i.i.d. sequence. Estimation of β_0 uses moments defined by the vector function $f_t^* = f(Z_t, \beta_0)$, $E\{f_t^*\}$, where $Z_t = [Y_t, Y_{t-1}, \dots, Y_{t-l-1}]$, and $f : R^{Nl} \times \Theta \rightarrow R^M$. Sometimes these moments are known analytically in which case this estimation procedure is standard GMM.

¹¹ However, in many cases analytic moment conditions are not available. The simulated method of moments procedure assumes that the econometrician has access to a sequence of random variables $\hat{\epsilon}_t$ which is iid. and independent of ϵ_t from which to construct simulated series for $\beta \in \Theta$,

$$Y_{t+1}^\beta = H(Y_t^\beta, \hat{\epsilon}_{t+1}, \beta).$$

The simulated moments are now constructed as time averages from $f_t^\beta = f(Z_t^\beta, \beta)$ where $Z_t^\beta = [Y_t^\beta, Y_{t-1}^\beta, \dots, Y_{t-l-1}^\beta]$. For any β let,

$$G_T(\beta) = \frac{1}{T} \sum_{t=1}^T f_t^* - \frac{1}{nT} \sum_{s=1}^{nT} f_s^\beta,$$

where n defines how long the simulation length will be relative to the sample length T . This can be a general increasing function of T as long as some limiting properties hold.¹²

The estimation procedure searches $\beta \in \Theta$ to get the vector $G_T(\beta)$ “close” to zero in some metric. Formally, the SMM estimator, b_T , is defined as,

$$b_T = \underset{\beta \in \Theta}{\operatorname{argmin}} G_T(\beta)' W_T G_T(\beta),$$

where W_T is a consistent sample estimate of

$$W_0 = (\Sigma_0)^{-1} = \left[\sum_{j=-\infty}^{j=\infty} E([f_t^* - E(f_t^*)][f_{t-j}^* - E(f_{t-j}^*)]) \right]^{-1}.$$

The SMM procedures requires several assumptions for consistency and asymptotic normality. These assumptions are:

1. $\|f_t^\beta\|_{2+\delta}$ is bounded for some $\delta > 0$. f_t^β must be continuously differentiable and $E f_t^\beta(Z_t^\beta, \beta)$ is a continuous function of β .¹³ This required smoothness on the moment conditions will easily hold for the standard moment conditions, but will require some modifications to the trading rule moments used here.

¹¹ A good example of this would be Euler equations restrictions used in finance. These theoretical models, under certain parametric forms, give very specific analytic moment conditions to be met by the series and the true parameter values.

¹² As $T \rightarrow \infty$ the ratio of simulations to sample should go to some constant. This obviously holds for nT simulation lengths.

¹³

$$\|x\|_q = [E\|x\|^q]^{1/q}$$

2. The process Y_t^β is ergodic for all $\beta \in \Theta$ and the state process Y_t is geometrically ergodic. This condition requires convergence of the process to the unconditional probability measure to occur at a certain rate. For all of the linear processes and parameters used in this paper this can be easily verified using methods outlined in Duffie and Singleton(1989).
3. $W_T \rightarrow W_0$ almost surely.
4. Let $C(\beta) = G_\infty(\beta)'W_0G_\infty(\beta)$. $C(\beta_0) < C(\beta)$ for all $\beta \in \Theta$.
5. β_0 and b_T are in the interior of Θ .

To get the asymptotic distribution for the parameter estimates the following conditions are necessary.

6. $D_0 = E(\partial f_t^{\beta_0} / \partial \beta)$ exists, is finite, and has full rank.

Since the simulations are used to estimate D_0 suitable smoothness is needed there too. In other words:

7. $D_\beta f_t^\beta$ satisfies a Lipschitz condition given in Duffie and Singleton(1989) for all $\beta \in \Theta$. Also, $E(|D_\beta f_\infty^\beta|) < \infty$, and $E(D_\beta f_\infty^\beta)$ is a continuous function of β .

Under these assumptions,

$$\sqrt{T}(b_T - \beta_0) \xrightarrow{d} N(0, (1 + (1/n))(D_0'W_0D_0)) \quad T \rightarrow \infty.$$

Linear processes are fit to the data using a set of moment conditions that includes modified versions of the conditions given in section 2A. Equation (2.1) gave a trading rule moment condition,

$$E\left\{S\left(\frac{p_{t-1}}{ma_{t-1}}\right)r_t\right\}$$

where $S(x) = 1$ if $x \geq 1$ and $S(x) = -1$ if $x < 1$. This cannot be used for simulated method of moments since the first derivatives will not necessarily be continuous in the parameters of the process r_t . The condition must be replaced with a smooth substitute. The hyperbolic tangent does a good job of accomplishing this.¹⁴ Replace the above condition with

$$E\left\{\tanh\left((1/\mu)\left(\frac{p_{t-1}}{ma_{t-1}} - 1\right)\right)r_t\right\}.$$

This condition can now be added to a more standard set of moment conditions.¹⁵

¹⁴

$$\tanh(x) = \frac{-e^{-x} + e^x}{e^{-x} + e^x}$$

¹⁵ The parameter μ will be set to constant according to the variability of the price-moving average ratio.

When using any method of moments estimator, choosing the moment conditions to use is not always a trivial procedure. Here, the choice of moments will follow the goal of trying to see whether a linear model does a good job of replicating some properties of the data (autocorrelations) as well as the trading rule results, and to see how taking the trading rules into account affects the parameter estimates. The actual data will be aligned to simulated data using the mean, variance, the first five lagged autocovariances, and three trading rule moments. This gives a total of ten moment conditions. For the trading rule moment condition the 20, 30, and 50 week moving averages are used.

There are two final details left for estimation. The variance covariance matrix is estimated using the Newey-West(1987) weighting using 15 lags. The procedure in Newey-West(1992) was used to determine the optimal lag length. It recommended lag lengths in the range of 10-20. No important changes are observed in estimated parameters or standard errors when the lag length is extended past 15. This is important for this case since the moving average may generate very long range dependence in the estimated moments. Lastly, the number of simulations is set to 50 times the sample size, $n = 50$.

III. Empirical Results

A. Data Summary

The data used in this paper are all from the EHRA macro data tape from the Federal Reserve Bank. Weekly exchange rates for the British Pound (BP), German Mark (DM), and Japanese Yen (JY) are sampled every Wednesday from January 1974 through July 1992 at noon New York time. If Wednesday is a holiday Thursday is used instead.

Returns are created using log first differences of these weekly exchange rates quoted in dollars/fx. Table 1 presents some summary statistics for these return series. All three series show little evidence of skewness and are slightly leptokurtic. These properties are common for many high frequency asset returns series. The first 10 autocorrelations are given in the rows labeled, ρ_n . The Bartlett asymptotic standard error for these series is 0.032. The BP shows little evidence of any autocorrelation except for lags 4 and 8, while the DM shows some weak evidence of autocorrelation, and the JY shows strong evidence for some autocorrelation at shorter horizons. The Ljung-Box-Pierce statistics are shown in the last row. These are calculated for 10 lags and are distributed $\chi^2(10)$ under the null of iid. The p-values are given in parenthesis. Only the JY series gives a strong rejection of the null hypothesis.

These three series are adjusted for the interest rate differential using weekly eurorates. The rates are the one week bid prices from the London Financial Times. The interest rate series extends from January 1979 through November 25th, 1991. The adjusted series is then defined using,

$$r_t = \log(p_t) - \log(p_{t-1}) - (i_{t-1} - i_{t-1}^*). \quad (3.1)$$

Assuming the continuously compounded form of covered parity this is the return to forward speculation,

$$r_t = \log(p_t) - \log(f_{t-1}),$$

where f_{t-1} is a one week forward rate. Under risk neutrality the returns to this zero cost investment should follow a martingale. However, it is well known that this series is forecastable using past data, suggesting either a time varying risk premium, or market inefficiencies.¹⁶

Table 2 presents summary statistics for these series. The series show similar results to table 1. They all show some amount of leptokurtosis and some weak autocorrelations. The autocorrelation patterns are stronger for the DM and JY series. These results are very close to those from table 1. These series will be referred to as, BPIA, DMIA, and JYIA.

B. Trading Rule Summaries

This section compares results for the moving average technical trading rules on the original series with some simple stochastic processes. Some of the basic properties of the rules are examined along with the important puzzles these present for empirical work.

Figure 1 shows the British pound exchange rate in dollars/pound for the entire sample along with a 30 day moving average. The basic trading rule test would define all periods where the price is above the moving average as buy periods, and where the price is below the moving average as sell periods. The figure appears to show long run persistent trends which are captured by the moving average trading rules generating buy signals during persistent upswings and sell signals during persistent downswings.

One of the questions that figure 1 asks is whether these apparent trends and the deviations from the moving averages are real, or statistical artifacts. This is answered in table 3. This table looks at the trading rule moments defined as,

$$E(S_{t-1}r_t),$$

where S_t is 1 if the price at time t is above a moving average of past prices, and -1 if it is below. The time averages are estimated from the observed series,

$$(1/T) \sum_{t=1}^T (S_{t-1}r_t). \quad (3.2)$$

These averages are then compared with the distribution of the same statistic generated by a simulated random walk for each series. The random walk is generated by scrambling the actual returns series for each exchange rate, and rebuilding a price series from these scrambled returns. Table 3 presents the fraction of

¹⁶ See Hodrick(1987) for a summary of this evidence.

1000 simulated random walks with a trading rule moment greater than that from the original series. Results are given for 3 different moving averages, the 20, 30, and 50 week. Finally, an average over the three rules is also given. This would correspond to,

$$(1/T) \sum_{t=1}^T (1/3)(S_{t-1}^{20} + S_{t-1}^{30} + S_{t-1}^{50})r_t.$$

It presents a summary over the three tests. This value is also compared to the appropriate bootstrapped distributions which correctly accounts for the dependence across the individual tests.

The results in table 3 dramatically reject the random walk for all three exchange rates with and without interest differential adjustment. The largest fraction (or simulated p-value) given is less than 3 percent for the 20 week moving average on the BPIA series, indicating that very few of the simulated random walks generated values close to those from the original series. Other p-values are very small with many of these well below 1 percent. This demonstrates that the moving average moments are clearly detecting some pattern in these series which is statistically significant.

One potential problem here is that of data snooping and the choice of the moving average test. After observing a series it will always be possible to find some kind of trading “rule” which will find a feature in the series which cannot be replicated. This problem is difficult to eliminate entirely, but some tests can be run to measure its impact. Figure 2 shows the sensitivity of the trading rule moment defined in (3.2) to the choice of the moving average length. It shows that the actual moving averages used in the simulations were not chosen to maximize the in sample trading rule profits.¹⁷ The figure actually shows that two of the rules, the 50 and 20 week moving averages, may be two of the worst rules to use from the standpoint of trading profits. The trading rules used here are based on those used by actual traders, but they are much simpler. Data snooping is an important reason for this simplification. Adding further complexity to the rules will add more parameters that need to be optimized in some way, increasing the chances for discovering spurious results. Even though these simple rules may not be optimal from a trading point of view, they are better for the statistical studies performed here.

Table 3 showed that the trading rules clearly rejected a random walk for all the series. This leaves us with the question of what type of process is capable of generating these results. All of the series showed some weak autocorrelations in tables 1 and 2. Could this be enough to generate the results from table 3? Table 4 approaches this question and strengthens the puzzle put forth by the trading rules.

¹⁷ The MA lengths were chosen to follow common trader practice. Technical traders may often follow several different MA’s. The 30 week (150 day) moving average is very common. The other rules were included as two nearby rules of longer and shorter duration.

This table compares the trading rule results for the DM series with those from a simulated AR(1),

$$r_t = \rho r_{t-1} + \epsilon_t.$$

This table presents results for both the trading rule moments and the first three autocorrelations along with the Ljung-Box-Pierce statistic for 10 lags. Means from the series and the simulations are given along with the fraction of simulations larger than the original series. The AR(1)'s are simulated using a normal random number generator.¹⁸ When the parameter is zero, results similar to table 3 are observed. The random walk is unable to replicate the trading rule results. The simulated autocorrelations are also a little low with a LBP simulated p-value of (0.119). As the AR parameter is increased the simulated trading rule profits begin to increase. However, this increase appears to be quite slow. The simulations do not match the actual series for the trading rule tests until the parameter value is increased to 0.3. At this level the autocorrelations in the simulated series are clearly much larger than those from the original series. The simulated p-value for the LBP statistic is (0.999) for an AR parameter of 0.3 indicating that the autocorrelation properties of the simulated series are far from those of the DM series.

It is now clear that linear processes can show large trading rule profits using the moving average rules. A simple AR(1) is quite capable of doing this. However, the results in table 4 show the difficulty of using this model. It appears to be hard to find an AR(1) that can match both the autocorrelations in the series, along with the trading rule profits. This strong tension makes it look doubtful that linear models in general will be able to pass this test, and that nonlinear specifications may be necessary. This question will be explored in the following sections.

C. AR(2) Processes

This section presents results for AR(2) models fit to the foreign exchange series. Parameters will be estimated using the procedures from section II, and then the processes will be simulated using the bootstrap procedure of using the scrambled estimated residuals to generate replications of the estimated parametric model.

The results in table 1 show some weak correlations for the each of the exchange rate series. However, results in table 4 suggest that a simple autoregressive model will not be able to replicate the trading rule results without introducing a large amount of correlation into the series. The usefulness of a simple short autoregressive structure is tested more thoroughly in this section.

¹⁸ This is done since these forced processes will, by design, be simulated far from the estimated parameters, and the estimated residuals will therefore not be appropriate.

Autoregressive models are estimated using both simple OLS, and the simulated method of moments procedure described earlier. The simulated method of moments estimation procedure will match moments from simulations of the fitted model with those from data. For each estimation procedure 10 moments will be used. The first and second moments of the series are matched along with the first five covariances without mean adjustment,

$$(1/T) \sum_t r_t r_{t-j} \quad j = 1, 2, 3, 4, 5.$$

Finally, three trading rule moments are used. As mentioned earlier, the hyperbolic tangent is added to smooth the trading rule moment. These take the form

$$(1/T) \sum_t \tanh((1/\mu)(\frac{p_{t-1}}{ma_{t-1}} - 1)) r_t,$$

where ma_t is the price moving average of length, 20, 30, and 50. The scaling parameter μ is fixed at 0.0005 which is about (1/100) of the standard deviation of the price moving average ratio. The Newey-West(1987) variance covariance estimator is used with 15 lags, and the simulation length is set to 50 times the length of the original series.

Parameter estimates are given in table 5. The rows beginning with OLS are standard OLS estimates, and those beginning with SMM are the simulated method of moments estimates. The estimated parameters show very little change for the two different estimation techniques. They are generally very small with only the JY having values that are significantly different from zero. This reflects the small amount of autocorrelation seen in each of the series. The goodness of fit test for the SMM estimates rejects the overidentifying restrictions for the BP series, but not for the other two series.¹⁹

The conditional heteroskedasticity of these series is well known.²⁰ In these simulations this will be accounted for using a GARCH framework used by the previously mentioned papers and many others. Table 6 presents estimates for a fitted GARCH(1,1)-AR(2) model. It shows a strong persistence in the conditional variance for all the series given by the estimates for β . Also, the estimated autoregressive parameters have increased in every case except the second lag for the JY. This change could be due to the better efficiency of this estimator, or it could be related to some of the results on connections between volatility and autocorrelation.²¹

¹⁹ In LeBaron(1991) estimation is done using three autocovariances, and only one trading rule moment. Using these conditions the AR(2) specification was rejected for the BP and JY series. For this paper precision of the parameter estimates is the most important criteria, so more moment conditions are used.

²⁰ See Baillie and Bollerslev(1989), Hsieh(1988), or the survey by Bollerslev et al.(1990).

²¹ Bilson(1989), Kim(1989), and LeBaron(1992), all find an inverse relation between the level of the conditional variance and the autocorrelations in foreign exchange series.

The AR(2) specification will now be simulated using the different parameter estimates to see if it can replicate the trading rule results seen in actual series while also staying close to the autocorrelations in each series. The results are presented in table 7. The row labeled “original” gives the estimated value of each statistic from the original series. The rows labeled “OLS” and “SMM” show the fraction of simulations with a value greater than the original series for the OLS and SMM estimated parameters respectively. The columns labeled MAn refer to the trading rule moment condition using a moving average of length n weeks, and AVE refers to the average of of the three trading rule moments. The columns ρ_i represent the autocorrelations at lag i , and LBP is the Ljung-Box-Pierce statistic.

The AR(2) model is simulated 1000 times for each exchange rate using scrambled estimated residuals redrawn with replacement. For each of the three series these simulations do a good job of matching the autocorrelations, but fail for the trading rule tests. For example, for the BP, the MA30 trading rule moment is 0.0017 and only 0.3 percent of the simulations give a value this large for both the OLS and SMM parameters. Looking at the AVE column shows no simulated p-values larger than 5 percent. These results do not differ across the two estimation methods, or the simulated GARCH models. This shows that for all three exchange rate series the AR(2) was not able to match the trading rule results for each series either in a homoskedastic framework, or using the GARCH model.

The AR(2) does much better in matching up with the autocorrelations in the series. This is seen on the right side of table 7. For all the series the autocorrelations from the simulations do a good job of matching those in the actual series. Most of the estimated autocorrelations are within the central parts of the simulated distributions. This should be expected since an autoregressive model should be able to faithfully replicate the autocorrelations seen in the series.²²

C. ARMA(1,1) Processes

In this section the ARMA(1,1) described earlier is estimated and simulated. This process was derived from a model with a time varying drift following an AR(1) process. This is an important test of the possibility of a slow moving risk premia as an explanation for these results. Table 8 gives the estimated parameters for the ARMA(1,1) process using the two different estimation techniques, maximum likelihood(ML), and simulated method of moments (SMM). The parameter estimates are all in the range consistent with the stochastic drift model, $b_1 > b_2$. Also, for all the series the MA and AR parameters are very close. This is due to the fact that they must fit the very low autocorrelation structure seen for each of the series.

²² Most of the results of the AR(2) seen here have been replicated on the interest rate adjusted series with little change.

For all the series the estimated parameters change when moving from ML estimation to SMM estimation. Both the AR and MA parameters increase. This increase is quite dramatic for the unadjusted series. For example, for the DM series, the estimated AR coefficient increases from 0.53 to 0.96. For the BP the estimated AR coefficient is close to 1 with a value of 0.98. The change is smaller for the interest rate adjusted series, BPIA, DMIA, and JYIA. Unlike the AR(2) estimates, the trading rule moments have had a strong impact on the values of the estimated parameters. In all cases the goodness of fit statistic is not large. The largest is for the DM with a p-value of 0.297.

The ARMA(1,1) is also estimated with a GARCH processes for the disturbance term. The parameter estimates are shown in table 9. The variance persistence parameters, α_1 and β are very close to those estimated in table 6 for the AR(2) model. This suggests some amount of independence of the estimated variance process to the type of linear returns process estimated. This is probably due to the fact that in either case the means process is a small fraction of the variability of the series. The estimated values for the ARMA parameters are close to those estimated in the homoskedastic case for the interest adjusted series. However, they increased for the nonadjusted series, BP, DM, and JY.

The important question is again whether the estimated ARMA model is able to capture both the trading rule moments along with the weak autocorrelations. Results of these tests are given in table 10. In this table the series are compared with results from 1000 simulations of the ARMA(1,1) generated using scrambled estimated residuals drawn with replacement. The right hand side of table 10 again shows this model fitting the autocorrelations of the series very well. There are a few exceptions where the ML estimates did not do a good job matching certain specific correlations, but the overall matchup on the first ten correlations given by the LBP statistic does well. The most unusual p-value is 0.913 for the ARMA-GARCH model for the BP LBP statistic.

Turning to the trading rule results for the ML parameter estimates a familiar pattern is observed. This test strongly rejects this model with this set of estimated parameters. For the AVE test the simulated p-values are 0.004, 0.014, and 0.015 for the BP, DM, and JY respectively. Turning to the SMM and GARCH simulations a different picture is revealed. For both these cases the p-values indicate that the simulated processes are doing a good job of replicating the trading rule results from the foreign exchange series. For example, the p-values for the SMM estimated parameters are 0.803, 0.391, and 0.419 for the BP, DM, and JY respectively. This model, in these two cases, has done a generally good job of matching the two important properties being tested here. ²³

²³ The LBP tests used here cover only 10 lags. Simulations have been performed at lags up to 50 to see if unusual correlations get detected. There is no evidence of unusually large correlation in the simulations using the SMM parameters.

The second part of table 10 presents results from simulating the interest adjusted series. These simulations are a little trickier since they involve simulating the interest adjusted series, and then adding the interest differential back to get a simulated version of the actual series. The simulated series is constructed from,

$$\hat{r}_t = \hat{e}_t + (i_{t-1} - i_{t-1}^*),$$

where \hat{e}_t is the bootstrapped simulation of the appropriate ARMA process, and i_t and i_t^* are the actual interest rate series. This rebuilds the simulated drift process adding the interest rate differential back in to get a simulated exchange rate returns series.

Findings in the bottom of table 10 indicate a similar, but weaker pattern to those from the original series. All the estimated parameters do a good job of replicating the overall autocorrelations seen in the p-values for the LBP statistic on the right hand column of the table. The ML parameter estimates show marginal rejections of the trading rule test for each of the series with simulated p-values of 0.120, 0.094, and 0.142, for the BPIA, DMIA, and JYIA respectively. However, examining the results for either the SMM parameters, or the GARCH model, gives a different story. These are again able to do a good job of replicating the trading rule results. For example, the simulated p-values for the AVE test on the ARMA(1,1) using the SMM parameters are 0.832, 0.262, and 0.380 for the BPIA, DMIA, and JYIA respectively.

The results for SMM and GARCH ARMA models are quite striking. One interesting question is how stable these results are over time. Table 11 uses the SMM parameter estimates from the entire sample to generate 1000 bootstrapped simulations for two subsamples. The series is divided exactly in half with the cutoff occurring in March of 1983. Simulated subsamples are generated using scrambled residuals from each subsample. The table presents both the trading rule moments for each series subsample along with the simulated p-values. For the trading rule tests the simulated stochastic ARMA models does a good job of replicating the results for each subsample. For the correlations the model replicates the actual series well with a few exceptions. One important exception is the LBP statistic for the BP2 series. The simulated p-value here is 0.940 suggesting a large amount of autocorrelation from the simulated series relative to the actual series.

D. Long AR's

This section explores the possibility for replicating these results using longer range autoregressive models estimated using OLS. The results for two models, an AR(10) and an AR(20), are given in table 12. The models are estimated using OLS and simulated using the estimated residuals. They once again do a very good job of matching up with autocorrelations in the series, as they should. However, matching up with the trading rule moment results is again difficult as shown by the left side of table 12. For the all three

series the simulated p-value for the average test is less than 10 percent for all the simulations except the JY using an AR(20) where it is 0.18. It is either the case that these models are still unable to capture the long range persistence seen in the series, or the results might be better replicated if parameters from the SMM estimation techniques were used.²⁴

E. Diagnostics and Discussion

The proceeding sections show that it is possible to find linear models which replicate both the trading rule moments and the autocorrelations in these series. ARMA(1,1) models, which are consistent with slowly changing risk premia are the only linear models of those tested here which are capable of doing this. Also, the ARMA(1,1) specification only matches the results for the actual series using the SMM estimated parameters, or the GARCH disturbance process, but not using maximum likelihood estimated parameters.

Many of the results in the previous section showed a dramatic change when moving from the ML estimated models to parameters estimated using SMM. This section checks to see if anything is going seriously wrong in the SMM estimation procedure. All the analysis here centers on the ARMA(1,1) model estimation for the DM series. This series was chosen since it displayed the largest changes in moving from the ML estimates to the SMM estimates.

Figure 3 displays the sum of squared residuals for the ARMA(1,1) for different parameter values. The objective function is set to zero when these parameters are out of the range that is consistent with the stochastic trends model which constrains the MA parameter to be less than the AR parameter. The figure shows a generally well behaved objective which looks symmetric. Also, it looks like it would be hard to discern between any of the possible parameter combinations which are on or near the diagonal. Figure 4 replaces the objective with the SMM objective including the trading rule moments. This plot is clipped at the top for some of the very large values. The picture shows some evidence of asymmetry in that the drop off near the diagonal appears much steeper for the larger AR values than for the smaller values.

The actual change in the parameter estimates is shown in figure 5 which compares the objective function on a straight line through the minimum parameters found by the SMM estimation. The line is constructed to be parallel to the diagonal, $AR=MA$. For example, for the DM the SMM estimated parameters (AR,MA) are, (0.95, 0.92). Figure 5 plots the objective for the parameter pairs $(x, x - 0.03)$ with x on the x-axis. This line passes very close to the ML optimum at (0.53, 0.48). This figure shows the big difference in the estimated AR coefficient using the two different methods. For the usual sum of squares objective the graph shows a minimum near 0.5. The objective is quite flat in the neighborhood of the minimum. For the

²⁴ Unfortunately, with this many parameters SMM becomes infeasible for small computer estimation since calculation of the gradient using numerical techniques becomes extremely costly.

SMM objective the minimum is closer to 0.95 with a very sharp jump down and then a quick and abrupt increase. These figures do not show anything seriously wrong or strange about either objective function, and the different parameter estimates appear to be very dramatic. It remains to be seen just what is causing these differences.

IV. Simulated Sharpe Ratios

This section uses the parameter estimates from the previous section to test some properties of the forecastability in these series. Predictability here is measured using conditional Sharpe ratios from a strategy of forward market speculation. This is an essentially zero cost strategy (ignoring margins) of going long or short in the forward market conditional on the trading rule result.

For this one security case this can be viewed as a standard deviation bound on the IMRS needed to explain the trading rule returns. This is closely related to the technique of scaling returns at time $t + 1$ with time t information suggested in Hansen and Jagannathan(1991) and used in Bekaert and Hodrick(1992). Dynamic asset pricing models such as Lucas(1982) give the result that returns on a zero cost strategy must be orthogonal to $u'(c_{t+1})$ and therefore,

$$0 = E_t(m_{t,t+1}(p_{t+1} - f_t)),$$

where $m_{t,t+1}$ is the IMRS from t to $t+1$. Adding time t conditioning information in the form of the trading rule, S_t , gives,

$$0 = E_t(m_{t,t+1}S_t(p_{t+1} - f_t)).$$

Writing the above expectation as the covariance plus the expected values and using the Cauchy-Schwarz inequality gives,

$$\frac{\sigma_m}{E(m)} \geq \frac{|E(S_t(p_{t+1} - f_t))|}{\sigma(S_t(p_{t+1} - f_t))}$$

The purpose of table 13 is to find out whether this measurement of predictability is different for some of the different parameter values and models fit in the previous sections. This is done by generating extremely long time series of 10,000 weeks which amounts to almost 200 years of data. Another comparison is made by using two different types of trading strategies. One uses the 30 day moving average already described. When the price is above this level a buy is generated, and when it is below a sell is generated. For this case a pseudo price series is generated by directly exponentiating the simulated series, $s_{t+1} - f_t$. The second strategy uses the ARMA(1,1) forecast using the true parameter values. If the forecast is positive, a buy is generated, and if negative, a sell is generated. This is the theoretical optimal forecast for this process. Results will also be compared to the numbers from the original series. However, these numbers should be viewed with caution since they are sampled from much shorter series. The Sharpe ratios are given at

the monthly horizon for comparison with previous results. However, the portfolio positions are allowed to change at the weekly frequency.

The first row of table 13 shows the results for the BPIA series. The numbers for the original series are somewhat smaller than those in Bekaert and Hodrick(1992). They are 0.27 while Bekaert and Hodrick obtain values in the range of 0.3 to 0.4.²⁵ The row labeled ML gives the mean value for the ARMA(1,1) simulated at the ML estimated parameters. The Sharpe ratio for the ARMA forecast is 0.131, and the ratio for the moving average trading rule is 0.123. This result suggests only a moderate drop off in the Sharpe ratio when shifting from the optimal ARMA forecast using the true parameters to the trading rule. Given that in real life the true parameters are unknown, or may be changing over time, these results give an interesting explanation for why trading rules might be preferred to more traditional time series analysis.

The second important fact to note for the BPIA series is that the Sharpe ratio changes dramatically when changing from the ML parameters to the SMM parameters. For both forecasting techniques it almost doubles. This suggests that in terms of predictability there is a big difference between the two different sets of parameters. From table 8, the AR parameter is 0.922 for ML, and 0.992 for SMM. The GARCH model produces a Sharpe ratio between the SMM and ML.

The entries in table 13 corresponding to the DMIA series show some similarities and differences to the BPIA series. First, the differences between the ARMA trading rule and moving average rule remain small. For the ML parameters the difference is about 0.009 and similar differences are seen for the SMM and GARCH simulations. An important difference for this series is that the changes across different models is much smaller than for the BPIA. There is an increase for both trading rules from the ML to SMM parameters of about 30 percent, but from the standard errors presented this does not appear to be a significantly large increase.

The results for the JYIA in the bottom of table 13 are slightly different. The yen series displays large differences between the ARMA and moving average forecasts for two of the three simulations, the ML, and the GARCH. Also, the Sharpe ratio increase is not consistent across the different types of trading rules. For the SMM parameters, the ARMA, and moving average rules give very similar Sharpe ratios of 0.206, 0.181, respectively. However, for the other processes the differences across the two rules are much larger.

These results demonstrate several interesting features of the trading rules, and the fitted stochastic processes. First, for many of the simulations the differences in Sharpe ratios between the moving average rules and ARMA forecasts were not great. This is very surprising since the ARMA forecasts are given a

²⁵ Bekaert and Hodrick use both the scaled and unscaled returns in the Hansen-Jagannathon bounds calculation. This is equivalent constructing Sharpe ratios from the optimal portfolio composition. The trading rule used in this paper simply moves back and forth between longs and shorts and no attempt is made to find the optimal portfolio composition.

tremendous in advantage in using the true parameters. It is therefore likely that in real life situations where the parameters need to be estimated, or are slowly changing over time the moving average forecast might dominate.²⁶ The second interesting result is that for some of the series the amount of predictability, as measured by the Sharpe ratio, changes dramatically as the parameters are changed. This shows that the different estimated parameters are give very different economic conclusions about the amount of predictability in these series.

These differences emphasize the need to further diagnose the different models and parameters estimated here. Also, this work looks much more closely at the fitted stochastic models than the economic significance of the estimated predictability. Future work will try to address this second question more thoroughly considering both transactions costs and other securities.

V. Conclusions

The results presented in this paper show that it is possible for a linear model to replicate both trading rule moments as well as traditional autocorrelations using both raw and interest rate adjusted foreign exchange returns. This implies that the trading rule results themselves are not necessarily indicative of nonlinearities in foreign exchange series. However, these trading rule moments may prove useful in estimating models with long highly persistent trend components.

These findings do not imply that nonlinearities in foreign exchange series are not important. The trading rules used were a very specific type of trend following rule, and there is a much larger space of trading rules available. Results such as those in Bilson(1989), Kim(1989) and LeBaron(1992) still show evidence for nonlinearities. Finally, the results should not be taken as holding for all asset markets. It is still not clear what may be found for other markets and frequencies.

An important issue that this paper does not address is the connection between the stochastic trend models estimated here with results from the forward and futures markets. Papers such as Hansen and Hodrick(1980,1983) and Bilson(1981) document the bias in the predictions of forward rates for future spot rates. Results here are generally supportive of models with slowly changing risk premia, but this risk premia is not connected to any specific type of economic risk. The results are also supportive of evidence that movements in the risk premia are larger than movements in expected exchange rate changes. This is shown by the relatively small change in the parameter estimates when shifting from the raw series to the interest adjusted series. This agrees with results in Fama(1984) and Hodrick and Srivastava(1986).

²⁶ These results are consistent with some results in Taylor(1992a,b) which show simple trading rules doing well in comparison to more sophisticated time series techniques.

Forecasting experiments show that for most of the estimated parameters the trading rules perform quite favorably when compared to the optimal time series forecasting models. This result is consistent with some of the findings in Taylor(1992a,b) and gives some support for the use of oscillators and channels as opposed to more traditional time series models. In a world where parameters need to be estimated, or where there is some nonstationarity in the parameters themselves it is quite possible that technical rules may be the optimal forecasts.²⁷

This paper indicates some delinking of technical trading results and nonlinearities for foreign exchange series. The trading rule results can be replicated by linear models exhibiting persistent trends. Far from being useless, the technical trading rules themselves helped to estimate these models. It remains to be seen whether they will have further uses in estimating persistent trend models. Finally, further nonlinear effects and the connections between volatility, trading rules and risk premia need to be further explored.

²⁷ Granger and Newbold(1986), page 174, report on cases where a close relative of the moving average rules, exponential smoothers, will be an optimal forecaster. One case where this is true is where the state variable follows a random walk and then is obscured by additive noise. This is close to the stochastic trend models fitted here, since they are nearly nonstationary at the estimated parameters.

References

- Antoniewicz, Rochelle. 1992. An Empirical Analysis of Stock Returns and Volume. Ph. D. Thesis, University of Wisconsin, Madison, WI.
- Baillie, Richard, and Tim Bollerslev. 1989. The Message in Daily Exchange Rates: A Conditional-Variance Tale. *Journal of Business and Economic Statistics* 7 (3) : 297-305.
- Bekaert, G. and Hodrick, R. J. 1992. Characterizing Predictable Components in Excess Returns on Equity and Foreign Exchange Markets. *Journal of Finance* 47 (2): 467-510.
- Bilson, J. F. 1989. "Technical" Currency Trading. The Chicago Corporation. Chicago, IL.
- Bollerslev, T., R. Y. Chou, N. Jayaraman, and K. F. Kroner. 1990. ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence. *Journal of Econometrics* 52 (1) : 5-60.
- Brock, W. A., J. Lakonishok, and B. LeBaron. 1992 forthcoming. Simple Technical Trading Rules and the Stochastic Properties of Stock Returns. *Journal of Finance*.
- Cochrane, J. H. and Hansen, L. P. 1992. Asset Pricing Explorations for Macroeconomists. NBER wp 4088, Cambridge, MA.
- Canova, F. and Marrinan, J. 1990. Nominal Profits, Risk and Uncertainty in Foreign Exchange Markets. Boston College, Dept. of Economics, Chestnut Hill, MA.
- Curcio, R. and C.A.E. Goodhart. 1992. When Support/Resistance Levels are Broken, Can Profits be Made? Evidence from the Foreign Exchange Market. LSE Discussion paper 142, London.
- Diebold, Francis X., and James M. Nason. 1990. Nonparametric Exchange Rate Prediction? *Journal of International Economics* 28 : 315-332.
- Dominguez, Kathryn M. 1986. Are Foreign Exchange Forecasts Rational? New Evidence from Survey Data. *Economics Letters* 21 : 277-281.
- Domowitz, I., and C. S. Hakkio. 1985. Conditional Variance and the Risk Premium in the Foreign Exchange Market. *Journal of International Economics* 19 : 47-66.
- Dooley, Michael P., and Jeffrey Shafer. 1983. Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981. In *Exchange Rate and Trade Instability: Causes, Consequences, and Remedies*. Edited by D. Bigman and T. Taya. 43-72. Cambridge, MA: Ballinger.
- Duffie, D. and K. Singleton. 1988. Simulated Moments Estimation of Markov Models of Asset Prices. Graduate School of Business, Stanford University, Stanford, CA.
- Efron, B. 1979. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* 7 (1) : 1-26.
- Engel, Charles , and James D. Hamilton. 1990. Long Swings in the Dollar: Are They in the Data and Do Markets Know It? *American Economic Review* 80 (4) : 689-713.
- Evans, M. D. D. and K. K. Lewis. 1992a. Peso Problems and Heterogeneous Trading: Evidence From Excess Returns in Foreign Exchange and Euromarkets. *National Bureau of Economic Research Working Paper 4003*.
- Evans, M. D. D. and K. K. Lewis. 1992b. Trends in Expected Returns in Currency and Bond Markets. *National Bureau of Economic Research Working Paper 4116*.
- Frankel, Jeffrey A., and Kenneth A. Froot. 1990. Chartists, Fundamentalists, and Trading in the Foreign Exchange Market. *American Economic Review* 80 (2) : 181-185.
- Goodhart, Charles. 1988. The Foreign Exchange Market: A Random Walk with a Dragging Anchor. *Economica* 55 : 437-60.
- Hodrick, Robert J. 1987. *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. New York: Harwood Academic Publishers.
- Hsieh, David. 1988. The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983. *Journal of International Economics* 24 : 129-45.

- Hsieh, David. 1989. Testing For Nonlinear Dependence in Daily Foreign Exchange Rates. *Journal of Business* 62 (3) : 339-68.
- Kim, Chong Man. 1989. Volatility Effect on Time Series Behavior of Exchange Rate Changes. Korea Institute for International Economic Policy,
- LeBaron, Blake. 1991. Technical Trading Rules and Regime Shifts in Foreign Exchange. University of Wisconsin - Madison.
- LeBaron, Blake. 1992 forthcoming. Forecast Improvements Using a Volatility Index. *Journal of Applied Econometrics*.
- Lee, Bong-Soo, and Beth F. Ingram. 1991. Simulation Estimation of Time-Series Models. *Journal of Econometrics* 47 : 197-205.
- Levich, R. M., and L. R. Thomas. 1991. The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach. New York University, Graduate School of Business, New York.
- Lewis, Karen. 1989. Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange. *American Economic Review* 79 : 621-36.
- Mark, N. C. 1992. Exchange Rates and Fundamentals: Evidence on Long-horizon Predictability and Overshooting. Dept. of Economics, Ohio State University, Columbus, Ohio.
- McFadden, D. 1989. A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica* 57 (5) : 995-1026.
- Meese, R., and K. Rogoff. 1983. Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample? *Journal of International Economics* 14 : 3-24.
- Meese, R. A., and A. K. Rose. 1990. Nonlinear, Nonparametric, Nonessential Exchange Rate Estimation. *American Economic Review* 80 (2) : 192-196.
- Neftci, S. N. 1991. Naive Trading Rules in Financial Markets and Wiener-Kolmogorov Prediction Theory: A Study of "Technical Analysis". *Journal of Business* 64 (4) : 549-571.
- Newey, W. and K. D. West. 1987. A Simple, Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55 : 703-708.
- Newey, W. and K. D. West. 1992. Automatic Lag Selection in Covariance Matrix Estimation. University of Wisconsin, Madison, WI.
- Pakes, A. and D. Pollard. 1989. Simulation and the Asymptotics of Optimization Estimators. *Econometrica* 57 (5) : 1027-1058.
- Schulmeister, Stephan. 1987. An Essay on Exchange Rate Dynamics. Wissenschaftszentrum Berlin für Sozialforschung, Berlin. IIM/LMP 87-8.
- Sweeney, Richard J. 1986. Beating the Foreign Exchange Market. *Journal of Finance* 41 (1) : 163-182.
- Sweeney, R. J., and P. Surarjara. 1989. The Stability of Speculative Profits in the Foreign Exchanges. In A Reappraisal of the Efficiency of Financial Markets. Edited by R. Guimaraes, B. Kingsman and S. J. Taylor. Heidelberg: Springer-Verlag.
- Taylor, M. P. and H. Allen. 1992. The Use of Technical Analysis in the Foreign Exchange Market. *Journal of International Money and Finance* 11: 304-314.
- Taylor, Stephen J. 1980. Conjectured Models for Trends in Financial Prices, Tests and Forecasts. *Journal of the Royal Statistical Society A* 143 (3) : 338-362.
- Taylor, Stephen J. 1986. *Modelling Financial Time Series*. Chichester: John Wiley.
- Taylor, Stephen J. 1992a. Rewards Available to Currency Futures Speculators: Compensation For Risk or Evidence of Inefficient Pricing? *Economic Record* 68 (Supplement): 105-116.
- Taylor, Stephen J. 1992b. Trading Futures Using the Channel Rule: A Study of the Predictive Power of Technical Analysis with Currency Examples. The Management School, Lancaster University, UK.
- Technical Analysis: Tilting at Chaos. 1992. *The Economist* 7772, August 15th, 1992.

- Tsay, R. S. 1992. Model Checking via Parametric Bootstraps in Time Series Analysis. *Applied Statistics* 41: 1-15.
- Tsibouris, G. C. 1992. Essays on Nonlinear Models of Foreign Exchange. Ph.D. Thesis. University of Wisconsin, Madison, WI.




Figure 1: \$/British Pound and 30 week moving average.




Figure 2: Moving average trading profit (Mean (Buy+(-1)Sell) Returns, British Pound. Moving average is varied from 5-50.




Figure 3: Sum of Squared errors, ARMA(1,1), DM. AR varies from zero to 1, MA varies from 0 to -1. If $-(MA) > AR$ the objective is set to zero.




Figure 4: SMM Objective, ARMA(1,1), DM. AR varies from zero to 1, MA varies from 0 to -1. If $-(MA) > AR$ the objective is set to zero.




Figure 5: Simulated Method of Moments versus ML objectives for ARMA(1,1). MA Parameter = -(AR Parameter - 0.033).

Table 1
Summary Statistics
Weekly Exchange Rates : Log First Difference

Description	BP	DM	JY
Sample Size	965	965	965
Mean*100	-0.019	0.062	0.084
Std*100	1.465	1.473	1.389
Skewness	0.041	0.158	0.376
Kurtosis	5.569	4.529	5.118
Max	0.074	0.081	0.065
Min	-0.065	-0.068	-0.066
ρ_1	0.036	0.038	0.088
ρ_2	-0.009	0.077	0.095
ρ_3	0.022	-0.010	0.065
ρ_4	0.081	0.041	0.039
ρ_5	0.021	-0.013	0.026
ρ_6	-0.012	-0.042	-0.005
ρ_7	0.019	0.005	-0.036
ρ_8	0.063	0.051	-0.003
ρ_9	-0.024	0.014	0.000
ρ_{10}	0.001	0.046	-0.085
Bartlett	0.032	0.032	0.032
LBP(10)	13.530	15.552	30.749
(p-value) $\chi^2(10)$	0.804	0.887	0.999

Summary statistics for BP (British Pound), DM (German Mark), JY (Japanese Yen) weekly exchange rates from 1974-July 1992.

Table 2
Summary Statistics
Weekly Exchange Rates : Interest Rate Adjusted Log First Difference

Description	BPIA	DMIA	JYIA
Sample Size	659	659	659
Mean*100	0.027	-0.033	0.014
Std*100	1.614	1.585	1.505
Skewness	0.235	0.371	0.512
Kurtosis	4.658	3.934	4.235
Max	0.075	0.080	0.065
Min	-0.052	-0.047	-0.050
ρ_1	0.047	0.047	0.101
ρ_2	-0.009	0.073	0.087
ρ_3	0.035	-0.002	0.052
ρ_4	0.085	0.064	0.042
ρ_5	0.031	-0.004	0.058
ρ_6	-0.022	-0.060	-0.003
ρ_7	0.011	-0.003	-0.033
ρ_8	0.092	0.079	0.014
ρ_9	-0.019	0.007	0.007
ρ_{10}	0.020	0.060	-0.077
Bartlett	0.039	0.039	0.039
LBP(10)	14.356	16.743	21.831
p-values $\chi^2(10)$	0.843	0.920	0.984

Summary statistics for BP (British Pound), DM (German Mark), JY (Japanese Yen) weekly exchange rates. The log first differences are adjusted using the interest differentials for the appropriate two countries. The series cover the period from January 1979 through November 1991.

Table 3
Random Walk Bootstraps

Series	MA20	MA30	MA50	AVE
BP	0.006	0.001	0.002	0.002
DM	0.001	0.001	0.011	0.003
JY	0.001	0.001	0.017	0.000
BPIA	0.028	0.001	0.012	0.004
DMIA	0.004	0.003	0.009	0.002
JYIA	0.001	0.006	0.017	0.003

Fraction of 1000 simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 4
DM AR(1) Experiment

AR(1) Parameter	Description	MA20	MA30	MA50	AVE	ρ_1	ρ_2	ρ_3	LBP
	Original	0.0017	0.0015	0.0013	0.0015	0.0382	0.0767	-0.0099	15.55
0.0	Mean	0.0000	0.0000	0.0000	0.0000	-0.0014	-0.0031	-0.0022	10.14
	p-value	(0.000)	(0.000)	(0.002)	(0.000)	(0.110)	(0.006)	(0.602)	(0.119)
0.1	Mean	0.0005	0.0004	0.0003	0.0004	0.0998	0.0097	-0.0009	20.17
	p-value	(0.008)	(0.008)	(0.028)	(0.003)	(0.974)	(0.018)	(0.606)	(0.683)
0.2	Mean	0.0009	0.0008	0.0006	0.0008	0.1972	0.0381	0.0091	50.09
	p-value	(0.085)	(0.065)	(0.066)	(0.049)	(0.999)	(0.130)	(0.706)	(0.998)
0.3	Mean	0.0015	0.0012	0.0009	0.0012	0.2976	0.0871	0.0233	105.16
	p-value	(0.345)	(0.273)	(0.225)	(0.256)	(0.999)	(0.608)	(0.824)	(0.999)
0.4	Mean	0.0021	0.0018	0.0014	0.0018	0.3986	0.1583	0.0609	195.56
	p-value	(0.790)	(0.669)	(0.556)	(0.680)	(0.999)	(0.988)	(0.981)	(0.999)

Results of simulated AR(1)'s compared with the original DM series for various AR parameters. Original refers to the estimated trading rule moments and autocorrelations from the actual series. Means are the means from 1000 simulated autoregressive models. P-values (in parenthesis) are the fraction of simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 5
AR(2) Parameter Estimates

$$r_t = a + b_1 r_{t-1} + b_2 r_{t-2} + \epsilon_t$$

Series	Estimation	a	b_1	b_2	χ^2
BP	OLS	-0.0001 (0.0004)	0.034 (0.041)	-0.001 (0.035)	
	SMM	-0.0001 (0.0006)	0.038 (0.032)	0.009 (0.041)	14.18 (0.028)
DM	OLS	0.0006 (0.0004)	0.036 (0.036)	0.077 (0.037)	
	SMM	0.0006 (0.0005)	0.038 (0.030)	0.061 (0.042)	8.71 (0.191)
JY	OLS	0.0008 (0.0004)	0.081 (0.042)	0.088 (0.035)	
	SMM	0.0014 (0.0005)	0.090 (0.035)	0.084 (0.035)	7.97 (0.240)

Estimated parameters using ordinary least squares (OLS) and simulated method of moments (SMM). Numbers in parenthesis for the parameter estimates are OLS and SMM standard errors. χ^2 is the goodness of fit test for the method of moments objective. It is asymptotically distributed $\chi^2(6)$ under the null model specification. Numbers in parenthesis below the χ^2 value are asymptotic p-values. Moment conditions include mean, standard deviation, the first five autocovariances, and three technical trading rules using, 20, 30, and 50 week moving averages. The SMM weighting matrix and standard errors are estimated using a Newey-West autocovariance matrix with 15 lags.

Table 6
GARCH(1,1)-AR(2)

$$r_t = a + b_1 r_{t-1} + b_2 r_{t-2} + \epsilon_t$$

$$\epsilon_t = h_t^{1/2} z_t$$

$$h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta h_{t-1}$$

$$z_t \sim N(0, 1)$$

Series	a	b_1	b_2	$\alpha_0 * 10^5$	α_1	β
BP	-0.0003 (0.0004)	0.075 (0.038)	0.041 (0.038)	1.896 (0.308)	0.153 (0.028)	0.766 (0.033)
DM	0.0006 (0.0004)	0.056 (0.037)	0.095 (0.034)	1.227 (0.306)	0.179 (0.027)	0.776 (0.029)
JY	0.0007 (0.0004)	0.105 (0.037)	0.077 (0.036)	1.901 (0.316)	0.207 (0.033)	0.710 (0.037)

Estimated using maximum likelihood. Numbers in parenthesis are asymptotic standard errors.

Table 7
AR(2) Simulations

Series	Estimation	MA20	MA30	MA50	AVE	ρ_1	ρ_2	ρ_3	LBP
BP	Original	0.0012	0.0017	0.0014	0.0014	0.0355	-0.0090	0.0223	13.53
	OLS	(0.019)	(0.003)	(0.007)	(0.004)	(0.444)	(0.605)	(0.235)	(0.285)
	SMM	(0.016)	(0.000)	(0.005)	(0.001)	(0.495)	(0.711)	(0.264)	(0.318)
	GARCH	(0.064)	(0.011)	(0.022)	(0.022)	(0.788)	(0.877)	(0.346)	(0.836)
DM	Original	0.0017	0.0015	0.0013	0.0015	0.0382	0.0767	-0.0099	15.5519
	OLS	(0.011)	(0.021)	(0.042)	(0.015)	(0.477)	(0.484)	(0.663)	(0.546)
	SMM	(0.009)	(0.015)	(0.039)	(0.012)	(0.502)	(0.321)	(0.659)	(0.437)
	GARCH	(0.043)	(0.066)	(0.074)	(0.046)	(0.643)	(0.658)	(0.683)	(0.898)
JY	Original	0.0019	0.0017	0.0013	0.0017	0.0877	0.0950	0.0650	30.7486
	OLS	(0.011)	(0.030)	(0.084)	(0.024)	(0.483)	(0.489)	(0.065)	(0.309)
	SMM	(0.012)	(0.027)	(0.101)	(0.023)	(0.601)	(0.495)	(0.047)	(0.379)
	GARCH	(0.027)	(0.034)	(0.103)	(0.036)	(0.647)	(0.434)	(0.168)	(0.637)

Results of 1000 simulated AR(2)'s compared with the original series using estimated parameters. Original refers to the estimated trading rule moments and autocorrelations from the actual series. P-values (in parenthesis) are the fraction of simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 8
ARMA(1,1) Parameter Estimates

$$r_t = a + b_1 r_{t-1} - b_2 \epsilon_{t-1} + \epsilon_t$$

Series	Estimation	a	b_1	b_2	χ^2
BP	ML	-0.0001 (0.0003)	0.442 (0.260)	0.408 (0.265)	
	SMM	0.0001 (0.0005)	0.983 (0.016)	0.957 (0.026)	6.311 (0.389)
DM	ML	0.0003 (0.0003)	0.538 (0.247)	0.482 (0.259)	
	SMM	0.0004 (0.0004)	0.956 (0.068)	0.923 (0.085)	7.261 (0.297)
JY	ML	0.0006 (0.0003)	0.368 (0.106)	0.270 (0.112)	
	SMM	0.0010 (0.0005)	0.914 (0.053)	0.856 (0.070)	2.495 (0.869)
BPIA	ML	0.0000 (0.0000)	0.922 (0.055)	0.890 (0.089)	
	SMM	-0.0002 (0.0008)	0.992 (0.030)	0.972 (0.070)	4.525 (0.606)
DMIA	ML	0.0002 (0.00001)	0.892 (0.091)	0.852 (0.105)	
	SMM	-0.0003 (0.0007)	0.967 (0.041)	0.940 (0.054)	5.886 (0.436)
JYIA	ML	0.0000 (0.0003)	0.682 (0.154)	0.583 (0.167)	
	SMM	0.0003 (0.0007)	0.898 (0.080)	0.838 (0.100)	2.105 (0.910)

Estimated parameters using maximum likelihood and simulated method of moments (SMM). Numbers in parenthesis for the parameter estimates are ML and SMM standard errors. χ^2 is the goodness of fit test for the method of moments objective. It is asymptotically distributed $\chi^2(6)$ under the null model specification. Numbers in parenthesis below the χ^2 value are asymptotic p-values. Moment conditions include mean, standard deviation, the first five autocovariances, and three technical trading rules using, 20, 30, and 50 week moving averages. The SMM weighting matrix and standard errors are estimated using a Newey-West autocovariance matrix with 15 lags.

Table 9
GARCH(1,1)-ARMA(1,1)

$$r_t = a + b_1 r_{t-1} - b_2 \epsilon_{t-1} + \epsilon_t$$

$$\epsilon_t = h_t^{1/2} z_t$$

$$h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta h_{t-1}$$

$$z_t \sim N(0, 1)$$

Series	a	b_1	b_2	$\alpha_0 * 10^5$	α_1	β
BP	-0.0001 (0.0002)	0.747 (0.138)	0.672 (0.158)	1.889 (0.310)	0.152 (0.028)	0.766 (0.034)
DM	0.0003 (0.0003)	0.582 (0.238)	0.506 (0.259)	1.249 (0.310)	0.179 (0.027)	0.774 (0.029)
JY	0.0004 (0.0003)	0.586 (0.175)	0.469 (0.198)	1.823 (0.302)	0.204 (0.033)	0.716 (0.036)
BPIA	0.0000 (0.0001)	0.899 (0.073)	0.846 (0.094)	2.777 (0.026)	0.111 (0.026)	0.782 (0.073)
DMIA	-0.0000 (0.0001)	0.892 (0.079)	0.830 (0.099)	5.420 (1.506)	0.199 (0.046)	0.596 (0.075)
JYIA	-0.0000 (0.0002)	0.683 (0.167)	0.572 (0.196)	7.342 (2.253)	0.114 (0.041)	0.546 (0.122)

Estimated using maximum likelihood. Numbers in parenthesis are asymptotic standard errors.

Table 10
ARMA(1,1) Simulations

Series	Estimation	MA20	MA30	MA50	AVE	ρ_1	ρ_2	ρ_3	LBP
BP	Original	0.0012	0.0017	0.0014	0.0014	0.0355	-0.0090	0.0223	13.53
	ML	(0.021)	(0.001)	(0.007)	(0.004)	(0.478)	(0.780)	(0.302)	(0.306)
	SMM	(0.812)	(0.720)	(0.833)	(0.803)	(0.691)	(0.928)	(0.772)	(0.838)
	GARCH	(0.524)	(0.212)	(0.268)	(0.321)	(0.814)	(0.925)	(0.669)	(0.913)
DM	Original	0.0017	0.0015	0.0013	0.0015	0.0382	0.0767	-0.0099	15.55
	ML	(0.012)	(0.019)	(0.044)	(0.014)	(0.700)	(0.088)	(0.777)	(0.374)
	SMM	(0.310)	(0.406)	(0.511)	(0.391)	(0.685)	(0.283)	(0.947)	(0.815)
	GARCH	(0.151)	(0.167)	(0.217)	(0.159)	(0.779)	(0.258)	(0.752)	(0.845)
JY	Original	0.0019	0.0017	0.0013	0.0017	0.0877	0.0950	0.0650	30.7486
	ML	(0.006)	(0.017)	(0.074)	(0.015)	(0.639)	(0.037)	(0.054)	(0.139)
	SMM	(0.306)	(0.409)	(0.559)	(0.419)	(0.422)	(0.279)	(0.546)	(0.659)
	GARCH	(0.163)	(0.207)	(0.339)	(0.212)	(0.763)	(0.334)	(0.338)	(0.690)
BPIA	Original	0.0013	0.0019	0.0015	0.0016	0.0333	-0.0225	0.0233	11.2974
	ML	(0.271)	(0.056)	(0.153)	(0.120)	(0.523)	(0.925)	(0.582)	(0.658)
	SMM	(0.858)	(0.770)	(0.876)	(0.832)	(0.804)	(0.974)	(0.862)	(0.893)
	GARCH	(0.815)	(0.633)	(0.638)	(0.710)	(0.708)	(0.942)	(0.695)	(0.893)
DMIA	Original	0.0020	0.0019	0.0015	0.0018	0.0403	0.0675	-0.0078	15.3778
	ML	(0.099)	(0.088)	(0.156)	(0.094)	(0.604)	(0.276)	(0.867)	(0.510)
	SMM	(0.195)	(0.236)	(0.394)	(0.262)	(0.588)	(0.371)	(0.879)	(0.632)
	GARCH	(0.744)	(0.724)	(0.727)	(0.748)	(0.722)	(0.504)	(0.897)	(0.863)
JYIA	Original	0.0018	0.0017	0.0013	0.0016	0.0897	0.0757	0.0407	19.1501
	ML	(0.145)	(0.137)	(0.215)	(0.142)	(0.693)	(0.476)	(0.592)	(0.653)
	SMM	(0.343)	(0.346)	(0.466)	(0.380)	(0.355)	(0.432)	(0.682)	(0.692)
	GARCH	(0.344)	(0.265)	(0.339)	(0.307)	(0.724)	(0.571)	(0.617)	(0.761)

Results of 1000 simulated ARMA(1,1)'s compared with the original series using estimated parameters. Original refers to the estimated trading rule moments and autocorrelations from the actual series. P-values (in parenthesis) are the fraction of simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 11
ARMA(1,1) Simulations: Subsamples

Series	Estimation	MA20	MA30	MA50	AVE	ρ_1	ρ_2	ρ_3	LBP
BP1	Original	0.0018	0.0020	0.0022	0.0020	0.0855	0.0073	0.0723	15.0102
	smm	(0.653)	(0.624)	(0.626)	(0.631)	(0.352)	(0.811)	(0.451)	(0.644)
BP2	Original	0.0008	0.0017	0.0006	0.0010	0.0067	-0.0089	-0.0118	6.4583
	smm	(0.807)	(0.611)	(0.885)	(0.778)	(0.827)	(0.864)	(0.896)	(0.940)
DM1	Original	0.0020	0.0022	0.0023	0.0021	0.0722	0.0677	0.0676	52.3467
	smm	(0.350)	(0.621)	(0.600)	(0.520)	(0.324)	(0.251)	(0.546)	(0.558)
DM2	Original	0.0020	0.0022	0.0018	0.0020	0.0108	0.0725	-0.0535	8.4280
	smm	(0.410)	(0.376)	(0.488)	(0.412)	(0.830)	(0.404)	(0.986)	(0.881)
JY1	Original	0.0020	0.0021	0.0014	0.0019	0.1135	0.1432	0.0801	29.3061
	smm	(0.233)	(0.174)	(0.435)	(0.247)	(0.300)	(0.123)	(0.435)	(0.412)
JY2	Original	0.0019	0.0015	0.0015	0.0016	0.0706	0.0633	0.0499	13.8232
	smm	(0.417)	(0.602)	(0.545)	(0.519)	(0.613)	(0.607)	(0.643)	(0.803)

Results of simulated ARMA(1,1)'s compared with the original series using estimated parameters. Original refers to the estimated trading rule moments and autocorrelations from the actual series. P-values (in parenthesis) are the fraction of simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 12
Long AR Simulations

Series	Estimation	MA20	MA30	MA50	AVE	ρ_1	ρ_2	ρ_3	LBP
BP	AR(10)	(0.183)	(0.024)	(0.051)	(0.045)	(0.454)	(0.521)	(0.480)	(0.879)
BP	AR(20)	(0.217)	(0.045)	(0.075)	(0.083)	(0.477)	(0.579)	(0.480)	(0.867)
DM	AR(10)	(0.038)	(0.076)	(0.108)	(0.063)	(0.529)	(0.499)	(0.501)	(0.877)
DM	AR(20)	(0.054)	(0.074)	(0.109)	(0.059)	(0.474)	(0.487)	(0.497)	(0.890)
JY	AR(10)	(0.012)	(0.020)	(0.070)	(0.016)	(0.493)	(0.425)	(0.575)	(0.712)
JY	AR(20)	(0.093)	(0.176)	(0.360)	(0.180)	(0.468)	(0.415)	(0.537)	(0.729)

Results of simulated AR's compared with the original series using estimated parameters. Original refers to the estimated trading rule moments and autocorrelations from the actual series. P-values (in parenthesis) are the fraction of simulated random walks with trading rule moments greater than the original series. The trading rule moment is defined as $\sum_t S_{t-1} r_t$, where S_{t-1} is 1 if a buy period is indicated, and -1 for a sell period.

Table 13
Conditional Sharpe Ratios

Series	Estimation	ARMA	TR
BPIA	Original	0.278	0.206
	ML	0.131	0.123
		(0.023)	(0.023)
	SMM	0.254	0.227
(0.036)		(0.034)	
DMIA	GARCH	0.172	0.155
	ML	0.131	0.122
		(0.021)	(0.021)
	SMM	0.167	0.162
(0.025)		(0.024)	
JYIA	GARCH	0.193	0.171
	ML	0.203	0.119
		(0.021)	(0.019)
	SMM	0.206	0.181
(0.024)		(0.022)	
GARCH	0.221	0.128	
	(0.020)	(0.020)	

Expected return divided by standard deviation for dynamic forward strategies. The return is $S_t(s_{t+1} - f_t)$, where S_t is 1 for a buy signal and -1 for a sell. Simulated Sharpe ratios are the ratios of the conditional mean to standard deviation using 10,000 weeks of simulated data. Aggregation is done by adding the dynamic weekly returns to obtain monthly returns. Trading is still allowed at weekly frequency. Numbers in parenthesis are the standard deviations of the estimated Sharpe ratios from 100 simulations of the 10,000 week samples.