Open Access

How Do Apparently Successful Trading Strategies Really Work?

Erhard Reschenhofer*[,1] and Christian Holzmann[2]

[1]*Department of Statistics & Decision Support Systems, University of Vienna, A-1010 Vienna, Austria*

[2]*Department of Economics, University of Vienna, A-1010 Vienna, Austria*

Abstract: This paper investigates a common approach to forecast stock returns. The forecasts are obtained in three steps. First a base set of potential forecasting variables is determined. Then a subset of forecasting variables is selected at each time period. Finally, a regression is run on the selected subset and the estimated regression parameters are used to forecast the return of the next time period. While this approach appears to have high forecasting power, a closer look reveals that none of the three steps contributes significantly to its performance. Moreover, we show that its high forecasting power is simply due to the fact that it mimics a very primitive technical trading strategy, which is based only on the signs of past returns.

Keywords: Forecasting stock returns, automatic model selection, technical trading strategies.

1. INTRODUCTION

There exists a large body of literature on strategies to outperform the market by first predicting future returns on different assets and then switching between these assets accordingly. In the simplest case, there are only two assets, a risky asset (e.g., an individual stock or a stock market index) and a risk-free asset (e.g., treasury bills or just cash). If the predicted return on the risky asset is greater than that on the risk-free asset, then the risky asset is held. Vice versa, if it is less, then the risk-free asset is held. While it is by now well established that both technical trading rules (see, e.g., [1]) and more orthodox econometric methods (see, e.g., [2]) have some predictive power, potential profits quickly vanish when transaction costs are taken into account. Moreover, methods that worked in the past are often useless in the present.

In this paper, we take a closer look at a standard technical approach (Section 2) and a standard econometric approach (Sections 3 and 4) to quantitative trading. However, since neither approach appears to be working today, we examine simple modifications (Section 5), which prove to be effective in the present circumstances and also shed some light on the partial success of the conventional methods.

2. THE TECHNICAL APPROACH

Brock *et al*. [1] explored two of the most popular technical trading rules and found that the buy and sell signals generated by these rules can indeed forecast future returns. In the first case, a buy (or sell) signal is generated when a short-period moving average rises above (or falls below) a long-period moving average. In the second case, a buy (or sell) signal is generated when the asset price moves above (or below) a local maximum (or minimum). For illustration, we apply a common moving-average rule to decide between cash and the Dow Jones Industrial Index. Here the short period is one day and the long period is 200 days. Daily closing prices from October 1, 1928 to January 10, 2010 were downloaded from Yahoo! Finance. Fig. (**1**) shows the performance of this switching strategy in the absence of transaction costs and under 0.25% transaction costs, respectively. In general, the switching strategy outperforms the buy-and-hold strategy only in periods of faltering or declining prices. Its decent performance over the whole observation period is to a certain extent due to its cautious investment behavior during the time of the Great Depression, which started with the stock market crash of October 29, 1929 (Black Tuesday). Of course, we should not attach too much weight to such a singular event, which occurred more than eighty years ago. But even if it cannot beat the benchmark in the long run, the switching strategy can still make sense. A strategy that regularly stays out of the market for extended periods of time and still under-performs only slightly relative to the buy-and-hold strategy must have some predictive power, which could be used to determine opportunities to increase the rate of return by leveraging. For example, Brock *et al*. [1] explored a strategy which borrows and doubles the investment upon a buy signal.

Another possibility to improve the performance of a strategy is to fine-tune it. Fine-tuning is also risky, though not in the same way as leveraging. The danger here is to find spurious patterns by the excessive use of data-mining techniques. Under transaction costs, an obvious starting point is to reduce trading frequency. In the case of moving averages, this can, for example, be achieved by generating no signal when the short moving average is inside a narrow band around the long moving average or when the last signal occurred only a short time ago. Fig. (**2**) shows the results obtained by fiddling about with the four parameters (short period, long period, band, holding period). It seems that in our case with only four parameters and a long series of daily data, the potential data-snooping bias cannot be so large, because there is hardly any improvement over the simple strategy based only on the 200-day moving averages.

*Address correspondence to this author at the Department of Statistics and Decision Support Systems, University of Vienna, Universitätsstraÿe 5, A-1010 Vienna, Austria; Tel: +43 1 4277 38646; Fax: +43 1 4277 38659; E-mail: erhard.reschenhofer@univie.ac.at

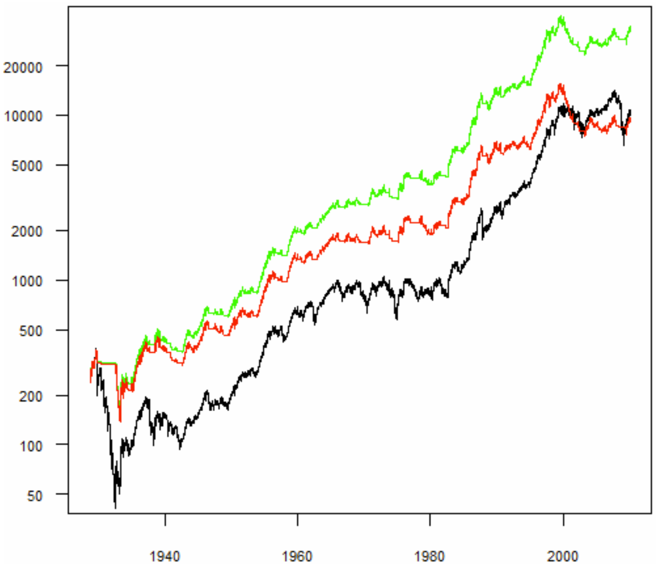

Fig. (1). Performance of a technical trading strategy based on 200-day moving averages in the absence of transaction costs (green line) and under 0.25% transaction costs (red line). The buy-and-hold strategy based on the Dow Jones Industrial Index (black line) serves as a benchmark.

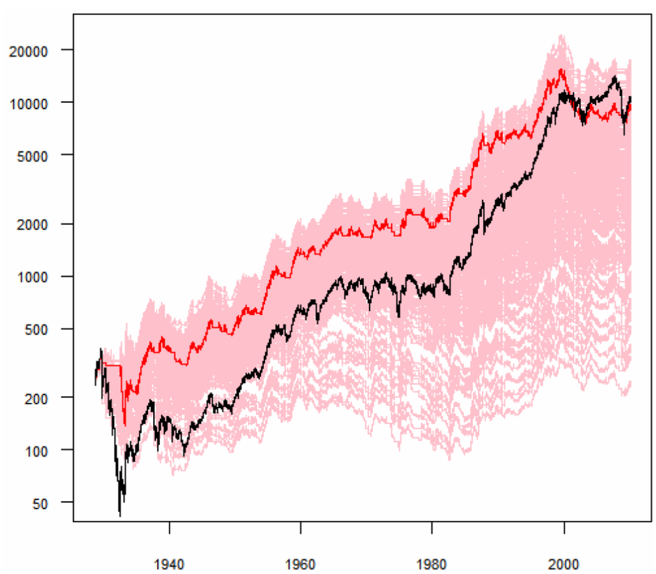

Fig. (2). Performance of various technical trading strategies (pink lines) based on all combinations of 1, 2, 5, 10, and 20-day moving averages (short period), 50, 100, 150, 200, 250, and 500-days moving averages (long period), fixed 0, 5, 10, and 20-day holding periods (after switching), and 0%, 1%, and 2% bands around the moving averages under 0.25% transaction costs. The buy-and-hold strategy based on the Dow Jones Industrial Index (black line) and the simple strategy based only on the 200-day moving average (red line) serve as benchmarks.

3. THE ECONOMETRIC APPROACH

Given the only moderate success of the technical approach described above, we now turn to more sophisticated methods. Pesaran and Timmermann [2] used standard econometric techniques to forecast the monthly excess return on a stock market index over short term treasury bills. First they identified a base set of potentially important forecasting variables. This set consisted of a constant, which was always included in the model, and nine macroeconomic and financial variables (change in industrial output, inflation, earnings-price ratio, dividend yield, bond rates, etc.). Then they used statistical model selection criteria to choose the most promising subset of variables for each point in time. Based on data up to period t they finally forecasted the excess return at time t+1 with a linear regression model containing the chosen subset of forecasting variables. The switching strategies based on these forecasts outperformed the buy and hold strategy only in the 1970s (under transaction costs of 0.5% for the index). A serious limitation of this approach is the combination of an extremely large number of candidate variables for the base set, a relatively large number of parameters (conventional parameters in the regression models plus meta-parameters in the model selection criteria), and a relatively small data set (monthly data), which immensely increases the danger of data snooping. Moreover, because of the lag associated with the publication of macroeconomic data, some variables can only be included with a 2-month time lag. Edwards [3] therefore used in his thesis a closely related approach which is based on daily data. His base set of forecasting variables was given by

(i) the excess return of the Toronto Stock Exchange 300 (TSE 300) Total Return Index over the return on 30-day Canadian treasury bills,

(ii) the excess return on the S&P 500 index over the return on 30-day US treasury bills,

(iii) the return on 90-day Canadian treasury bills,

(iv) the return on 90-day US treasury bills,

(v) the spread between the 12-month Canadian government bond rate and the 30-day Canadian treasury bill rate,

(vi) the spread between the 12-month US government bond rate and the 30-day US treasury bill rate,

(vii) the first difference of the logged Canada-US exchange rate.

The inclusion of these variables was justified by standard economic arguments (close relationship between the US and Canadian economies, inflationary expectations, monetary policies, etc.). For each time period, the optimal subset of variables for the forecasting of the excess return on the TSE 300 index was determined with model selection criteria like AIC [4] and BIC [5]. Since there are seven variables available, there are 2^7=128 different subsets of variables which have to be compared. The forecasts based on the best subsets were then used to switch between the TSE 300 index and 30-day Canadian treasury bills. In the absence of transaction costs, this switching strategy dramatically outperformed the TSE 300 index during the evaluation period from January 2, 1990 to May 28, 1999 (the first year, 1989, was only used for estimation). When transaction costs were implemented, not only the sign of the forecast but also its size was taken into account in order to reduce the trading frequency. More precisely, it was required that the forecasted gain from switching exceeds the transaction costs. Remarkably, the strategy remained profitable under low transactions costs (0.25% two-way). We checked Edwards' [3] strategy again using an updated data set (extending to

October 30, 2009). For technical reasons (availability of the data) we replaced the 30-day US treasury bills in the definitions of the variables (ii) and (vi) by the 90-day US treasury bills. Despite these changes we obtained practically the same results as Edwards [3] for the period until May 28, 1999 (see Fig. **3**). Minor discrepancies may be due to differences in the calculation of the daily returns on the 30-day treasury bills, which are needed to obtain the excess returns. For our calculations we assumed that there are 252 trading days per year.

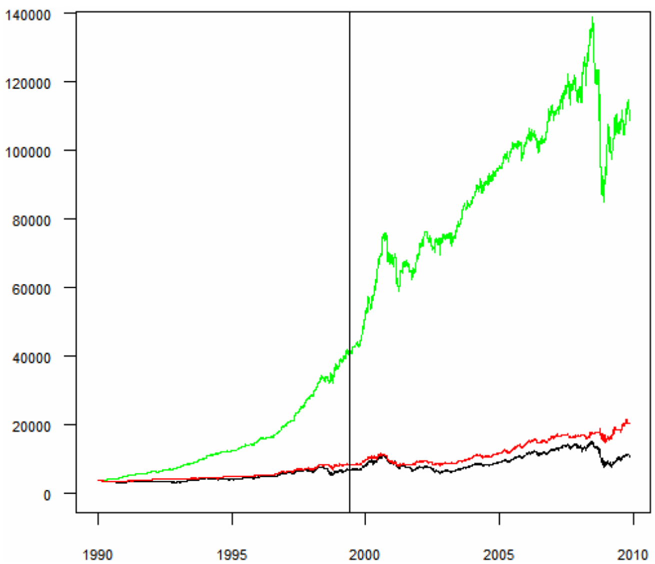

Fig. (3). Performance of Edwards' [3] econometric trading strategy in the absence of transaction costs (green line) and under 0.25% transaction costs (red line). AIC was used for model selection. The vertical line indicates the end of his investigation period. The buy-and-hold strategy based on the Toronto Stock Exchange 300 Total Return Index (black line) serves as a benchmark.

Of course, selling the index will only pay off if the combined effect of the possible loss avoided by not being in the market and the gain from treasury bills exceeds the transaction costs. In the case of frequent trading under non-negligible transaction costs, the net effect of investing in treasury bills will almost certainly be negative. Overall, the switching strategy will only be profitable if significant losses can be avoided. We might therefore consider to switch between the index and cash rather than between the index and treasury bills. In this way, it would be possible to cut transaction costs in half. Using cash instead of treasury bills and returns instead of excess returns has little effect on the performance of our econometric trading strategy (see Fig. **4**). Under low transaction costs, treasury bills and cash yield nearly the same results. That this is also the case in the absence of transaction costs may be surprising, but it is certainly not impossible, because the two strategies differ in more than just the choice of the risk-free asset. Indeed, we have focused on the signs of the forecasted returns in the case of cash and on the signs of the forecasted excess returns in the case of treasury bills. While these choices may seem natural, there is a myriad of alternative options. For example, we could try to determine appropriate threshold values for the assessment of the significance of a forecasted return.

The forecasting power of Edward's [3] econometric approach appears to have decreased after the millennium crash (see Fig. **4**). Of course, there are numerous ways in which it could be developed further, e.g., by using nonlinear models instead of linear regression models or by using model selection criteria that have been designed specially for subset selection instead of AIC and BIC. Trying to improve the forecasting procedure we took a closer look at its major elements. Much to our surprise we found that neither the selection of the base set nor the selection of the subsets plays an important role. These findings will be presented in detail in the next section. Moreover, we will argue in Section 4 that the actual reason for the forecasting power of the strategy is that it mimics a very primitive technical trading rule. We believe that our findings may be relevant also for other apparently successful trading strategies that are based on sophisticated econometric techniques.

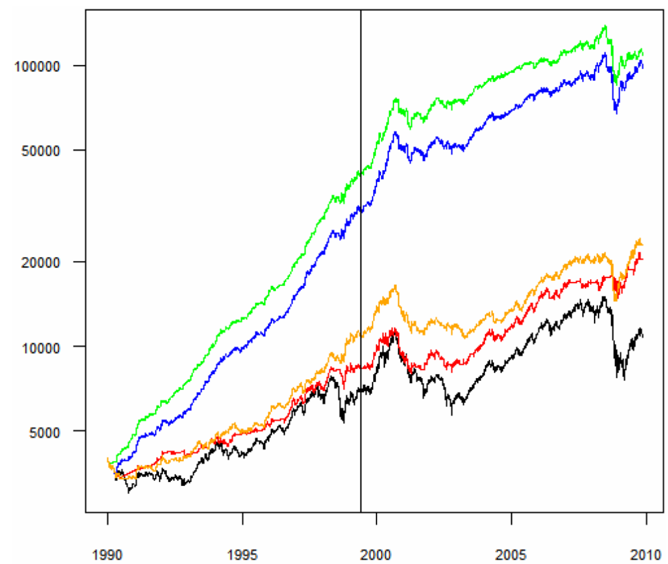

Fig. (4). Performance of Edwards' [3] econometric trading strategy for two different risk-free assets (treasury bills and cash) in the absence of transaction costs (green and blue line) and under 0.25% transaction costs (red and orange line). AIC was used for model selection. The vertical line indicates the end of his investigation period. The buy-and-hold strategy based on the Toronto Stock Exchange 300 Total Return Index (black line) serves as a benchmark.

4. EVALUATING THE INDIVIDUAL COMPONENTS OF THE ECONOMETRIC APPROACH

The first step in forecasting an economic variable is to search for a base set of potential determinants of that variable. Basically there are two types of errors we can make. The first is to include an irrelevant variable and the second is to omit a relevant variable. We do not intend to look for omitted variables. There are simply too many candidate variables. So we focus on examining the base set used in the previous section for forecasting the future return on the TSE 300 index, R_{t+1}. It is quite easy to see that no variables except the current return on the TSE 300 index, R_t, and the current return on the S&P 500 index, r_t, are of any importance. All we have to do is to investigate the forecasting power of all 128 submodels of the full model, which contains all seven variables of the base set. Fig. (**5**) shows that the submodels split into four clearly separated groups of exactly the same size. Remarkably, we can give a striking characterization of each group. The best performing group consists just of those 32 submodels that contain both R_t and r_t. All 32 submodels of the next group contain R_t but not r_t. The third group consists of those 32 submodels that contain r_t but not R_t. Finally, the submodels of the worst performing group contain neither R_t nor r_t. Actually, in the time period before the millennium crash there were only three groups, because the inclusion of r_t in addition to R_t gave no extra advantage.

Next, we check how models selected by model selection criteria perform against the fixed models presented in Fig. (**5**). In this comparison, we also include criteria that have been specially designed for subset selection. In contrast to conventional criteria like AIC and BIC, these special criteria penalize the first regressors entering the model more than the last ones. For a standard linear regression model

$$y = X\beta + u$$

containing k regressors, the criteria AIC and BIC are given by

$$AIC(k) = -2\log(L(y_1,\ldots,y_n;\hat{\beta},\hat{\sigma}^2)) + 2(k+1) \tag{1}$$

and

$$BIC(k) = -2\log(L(y_1,\ldots,y_n;\hat{\beta},\hat{\sigma}^2)) + (k+1)\log(n) \tag{2}$$

respectively, where

$$-2\log(L(y_1,\ldots,y_n;\hat{\beta},\hat{\sigma}^2)) = n + n\log(2\pi) + n\log(\hat{\sigma}^2) = \text{const} + n\log(\hat{\sigma}^2)$$

is minus two times the maximum log likelihood and

$$\hat{\beta} = (X^TX)^{-1}X^Ty,\ \hat{\sigma}^2 = \tfrac{1}{n}(y - X\hat{\beta})^T(y - X\hat{\beta})$$

are the maximum likelihood estimates of the k+1 model parameters. Model selection by minimization of AIC(k) is closely related to model selection by minimization of the final prediction error (FFE) criterion

$$FPE(k) = n\hat{\sigma}^2\,\frac{n+k}{n-k} \tag{3}$$

([6, 7]). Indeed, for large n

$$n\log(\hat{\sigma}^2\,\tfrac{n+k}{n-k}) = n\log(\hat{\sigma}^2) + \log((1+\tfrac{2k}{n-k})^n) \sim n\log(\hat{\sigma}^2) + 2k.$$

If the model is correctly specified, i.e., if $Ey = X\beta$, then the statistic FPE(k) is an unbiased estimator for the mean squared prediction error

$$E(z - X\hat{\beta})^T(z - X\hat{\beta}),$$

where z is an independent sample from the same distribution as y. While the use of the FPE criterion may be arguable for the comparison of two individual models, the comparison of two sets of models is a very different matter. Suppose, for example, that K variables are available for the explanation of the independent variable and we are not sure how many of them should be included in the model. The crucial question for a decision between a particular model M_1 with $k_1 \le K$ regressors and a smaller submodel M_2 with $k_2 < k_1$ regressors is whether the k_2-k_1 additional regressors in the larger model can actually increase the forecasting power. Of course, both the negative maximum log likelihood and the residual sum of squares will always be smaller for M_1 than for M_2. So we need additive (as in the case of AIC and BIC) or multiplicative (as in the case of FPE) penalty terms, which penalize over-parametrization, to allow for a fair comparison. But it is hard to see why a particular penalty term, which has been designed to compare a given small model to a given large model, should also be used to compare the best of few small models to the best of many large models or vice versa (e.g., for K=7 there are 7 models with one regressor, 35 models with 3 regressors, and only one model with 7 regressors). Indeed, it can be shown that under additional assumptions (orthogonal regressors, Ey=0) the alternative statistic

$$FPE_{sub}(k,K) = n\hat{\sigma}^2\,\frac{n+\varsigma_1(k,K)}{n-\varsigma_1(k,K)} \tag{4}$$

is an unbiased estimator for the mean squared prediction error of the (apparently) best model of dimension k [8] (for related criteria see [9-11]). Here $n\hat{\sigma}^2$ is the residual sum of squares of the best model of dimension k and $\zeta_1(k,K)$ is the expected value of the sum of the k largest of K independent $\chi^2(1)$-variables. If k=K, then $\zeta_1(k;K)=k$. In this special case, there is no data-snooping bias from the selection of the best subset, hence FPE_{sub} reduces to FPE. If $k < K$, then $\zeta_1(k;K)>k$. But this does not mean that FPE never can select a smaller model than FPE_{sub}. Once some important regressors have been included, it may be easier for FPE_{sub} than for FPE to clear the hurdle to include further regressors, simply because $\zeta_1(k,K)-\zeta_1(k-1,K)$ will be less than one if k is large. This entails the danger of over-fitting when there are some very important regressors, which are selected anyway. In the case where k_0 of the K variables are certain to be included, the term $\zeta_1(k,K)$ occurring both in the numerator and the denominator of the statistic $FPE_{sub}(k;K)$ must be replaced by $k_0+\zeta_1(k-k_0,K-k_0)$. For example, if a constant is included in each model, then $k_0 \ge 1$. The modified criterion will be denoted by FPE_{k_0}. However, taking only the unconditionally certain variables into account is usually not sufficient to avoid the risk of over-fitting. A more effective approach is to assume for each additional regressor to be included that all k–1 previously included regressors are certain and all K–k+1 remaining regressors are uncertain. Basing our inclusion decision on the expected value of the largest of K–k+1 independent $\chi^2(1)$-variables rather than on the expected value of the k'th-largest of K independent $\chi^2(1)$-variables, we arrive at the criterion

$$FPE_0(k;K) = n\hat{\sigma}^2\,\frac{n+\sum_{j=1}^{k}\varsigma_1(1,K-j+1)}{n-\sum_{j=1}^{k}\varsigma_1(1,K-j+1)} \tag{5}$$

[12], which has to be minimized with respect to k in order to find the best model dimension. This criterion has some similarities to the risk inflation criterion proposed by Foster and George [13]. One difference, among others, is that FPE_0 takes the fact into account that the number of remaining regressors decreases as k increases.

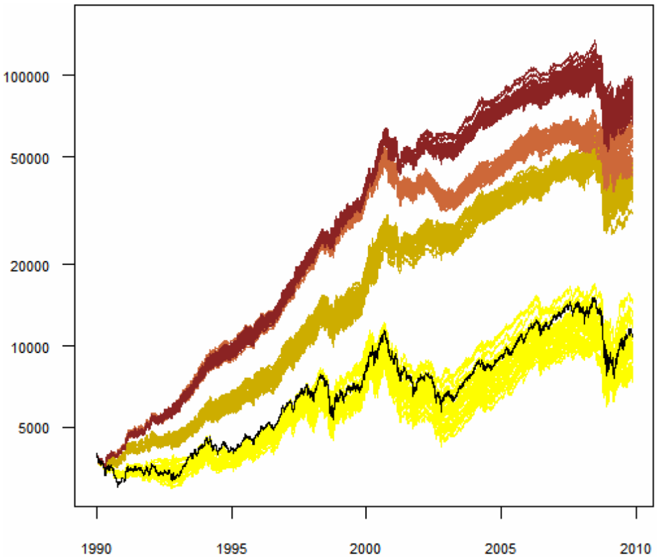

Fig. (5). Performance of strategies (switching between the TSE 300 index and cash) based on the forecasts obtained from 128 fixed regression models with at most seven regressors (plus the constant) in the absence of transaction costs. There are four clearly separated groups of size 32 which differ significantly in their performance. The best performing group (brown lines) consists of those models that contain both the current return on the TSE 300 index, R_t, and the current return on the S&P 500 index, r_t. The models of the second best group (sienna lines) contain R_t but not r_t. The third group (golden lines) consists of those models that contain r_t but not R_t. The models of the worst performing group (yellow lines) contain neither R_t nor r_t. The buy-and-hold strategy based on the TSE 300 index (black line) serves as a benchmark.

We examined the usefulness of a number of model selection criteria for the prediction of the future return on the TSE 300 index. Apart from the constant, which is included in all models, there are seven explanatory variables, hence each model selection criterion must compare 2^7=128 models at each time point. Fig. (**6**) shows how the models selected by the standard criteria AIC and BIC and the subset criteria FPE_{k_0} with k_0 =1 and FPE_0, respectively, performed against the fixed model that contains only the current returns R_t and r_t. Obviously, nothing could be gained from switching between different models. None of the criteria could significantly outperform the fixed model. But there are also hardly any differences between the criteria. Only recently, in a short time period after the housing crash of 2008, AIC was slightly luckier than the other criteria. Of course, trying to select the best model at each time point would certainly be a worthwhile exercise if the parameters of the data generating mechanism changed over time. But just in the case of time-varying parameters, it would make sense to use estimation windows rather than the full sample. Fig. (**7**) shows the results obtained by using at each time point only the last two years for model identification and parameter estimation. While the performance of our simple fixed model containing only the regressors R_t and r_t is now even better than before, the opposite is true for the model selection approach.

As a quick and somewhat naïve (see the discussion in Section 5) check to corroborate our interpretation of Fig. (**7**), we may exemplarily test whether the (relative) increments of different trading strategies differ significantly from each other. As expected, the results obtained with a two-sample t-test (where the variance is estimated separately and the Welch modification to the degrees of freedom is used) indicate no significant difference between the performance of BIC and that of AIC (p-value: 0.1472) and a highly significant difference between the performance of BIC and that of the simple fixed model containing only the current return on the TSE 300 index and the current return on the S&P 500 index (p-value: 0.00227). Our tests are based only on those increments which are not identical under the different strategies. In view of the non-normality of returns, we might think of using more robust tests based only on ranks or even signs (see, e.g., Pesaran and Timmermann [2]). However, such tests are hardly appropriate when our focus is on trading performance. According to a sign test, the absurd strategy of selling an index to avoid several tiny losses and one huge gain would be a wise decision.

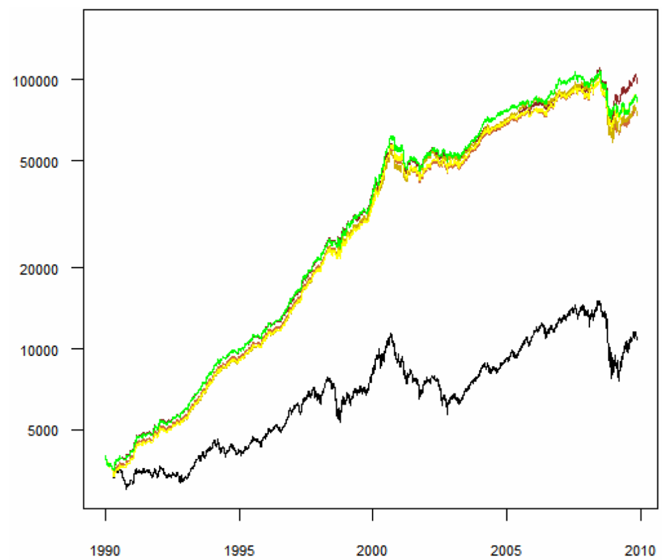

Fig. (6). Performance (in the absence of transaction costs) of strategies (switching between the TSE 300 index and cash) based on the forecasts obtained from regression models selected at each time point by AIC (brown line), BIC (sienna line), FPE_{k_0} (golden line), and FPE_0 (yellow line). The buy-and-hold strategy based on the TSE 300 index (black line) and the simple fixed model containing only the current return on the TSE 300 index and the current return on the S&P 500 index (green line) serve as benchmarks.

5. CONCLUSION

While the predictability of stock returns is indisputable, its practical value depends very much on the level of the transaction costs. This is particularly true for high-frequency trading strategies. Clearly, long-term trading strategies that only try to avoid the major downturns are less affected. The dilemma is that the potential profit increases with the number of trades. Fortunately, transaction costs have fallen steadily over time. Moreover, there is a growing number of stock brokers offering flat fees, so that even the assumption of negligible transaction costs becomes increasingly more plausible. Under this assumption, predictability would equate to profitability and high forecasting power would therefore also be of great economic interest. Figs. (**1**, **3**) suggest that

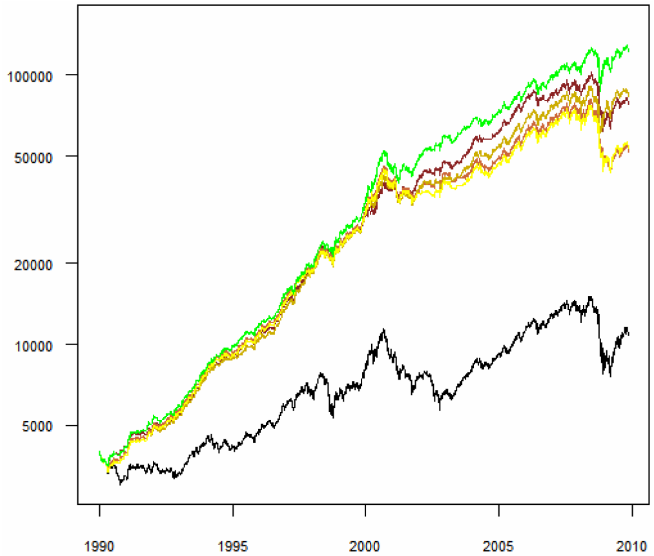

Fig. (7). Performance (in the absence of transaction costs) of strategies (switching between the TSE 300 index and cash) based on the forecasts obtained from regression models selected by AIC (brown line), BIC (sienna line), FPE_{k_0} (golden line), and FPE_0 (yellow line). At each time point only the last two years were used for model identification and parameter estimation. The buy-and-hold strategy based on the TSE 300 index (black line) and the simple fixed model containing only the current return on the TSE 300 index and the current return on the S&P 500 index (green line) serve as benchmarks.

the sophisticated econometric approach significantly outperforms simple technical trading rules. However, the econometric approach has two weaknesses. Firstly, its forecasting power has decreased significantly after the millennium crash. Secondly, its major elements are dubious. Neither economic reasoning (for the selection of the base set of explanatory variables) nor automatic model selection (for the selection of the best subset at each time point) were of any use to improve the forecasting performance. Even during the time of its best performance (before the millennium crash), the econometric approach could not beat the simplest model, which explains tomorrow's return just by today's return. One might argue in favor of the econometric approach that the simplest model is still an econometric model, but Fig. (**8**) shows that practically the same results could have been obtained by just using the sign of today's return instead of running a regression. This does, of course, not mean that buying at the end of a positive day and selling at the end of a negative day would have been a profitable strategy, because transaction costs were much too high in the 1990s (at least for private investors). Now, in the time of flat fees, this strategy does no longer work. But there may be others. Indeed, predictability can be regained (see Fig. **7**) by using an estimation window (two years) and including a second regressor, namely the current return on the S&P 500 index, in addition to the current return on the TSE 300 index. Fig. (**8**) shows that also the performance of this improved model can be matched by a simple trading rule, which is based on the sign of the sum of the two returns. Compared to this rule, the model containing the two returns has a serious disadvantage. It requires an additional tuning parameter, the length of the estimation window.

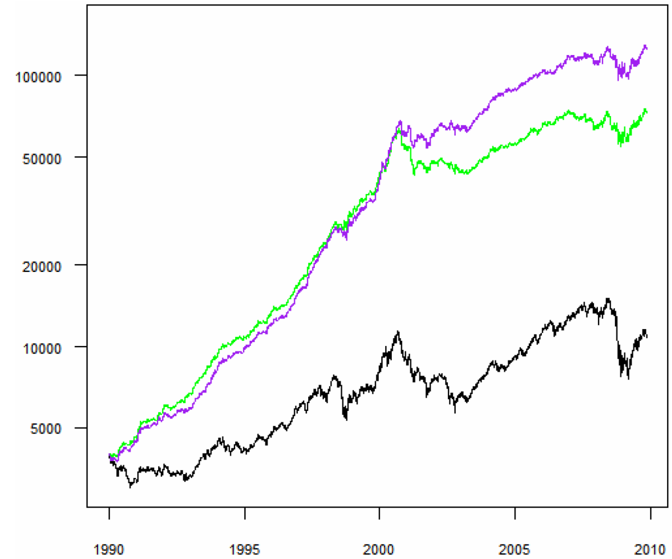

Fig. (8). Performance (in the absence of transaction costs) of simple strategies (switching between the TSE 300 index and cash) that are based on the sign of the current return on the TSE 300 index (green line) and on the sign of the sum of the current returns on the TSE 300 index and the S&P 500 index (purple line), respectively. The buy-and-hold strategy based on the TSE 300 index (black line) serves as benchmark.

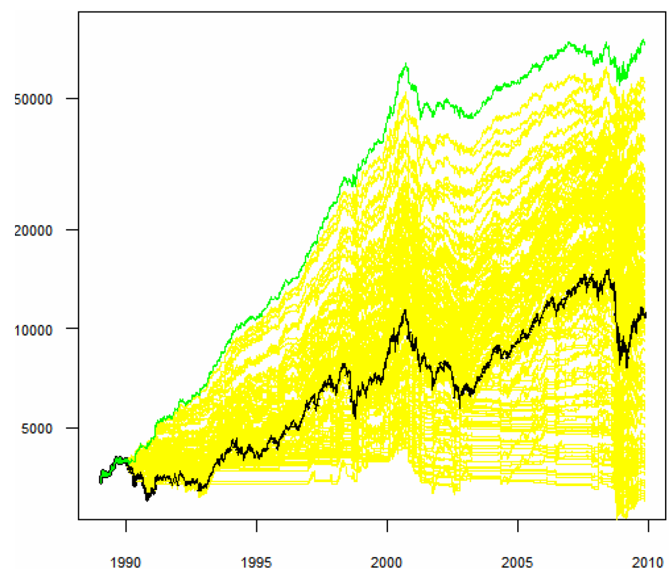

Fig. (9). Performance (in the absence of transaction costs) of simple strategies (switching between the TSE 300 index and cash) that compare the current return on the TSE 300 index to various threshold values (0.000, 0.002,..., 0.020 for buy signals and 0.000, -0.002,..., -0.020 for sell signals; yellow lines). The green line represents the trivial case (0.000 and 0.000), where only the sign of the return is used. The buy-and-hold strategy based on the TSE 300 index (black line) serves as benchmark.

Throughout the paper, we used time paths of cumulative returns to compare different trading strategies. The data provide evidence that one strategy outperforms another strategy, if the gap between the two time paths increases steadily over time. In our opinion, a comparison based only on the final wealth is not sound. Two similar strategies that move parallel to each other most of the time still may differ considerably in their final wealth, just because one strategy luckily avoids a major crash. Also a good performance that is only due to a certain subperiod is not very convincing, particularly when the subperiod is not the most recent one. Of course, we could also use statistical tests to analyze the performance of the forecasting method on which a trading strategy is based. But even a high correlation between forecasted returns and actual returns is no guarantee that the strategy will be profitable. Moreover, there is hardly any test that can cope with all characteristics of financial data (autocorrelation, conditional heteroskedasticity, periodicities, multiple breaks in trend and volatility, etc.). An important aspect of the interpretation of trading performance is the possible presence of data-snooping biases. Figs. (**2**, **5**, **7**, **9**) suggest that in simple settings like the ones investigated in this paper, the normal (non-excessive) use of data-mining techniques cannot produce outstanding results. The danger of over-fitting is much greater when monthly data and/or more complex models are used. In general, the model complexity and the number of tuning parameters should therefore always be kept as low as possible.

ACKNOWLEDGEMENT

We would like to thank two anonymous reviewers for helpful suggestions and comments.

REFERENCES

[1] Brock W, Lakonishok J, LeBaron B. Simple technical trading rules and the stochastic properties of stock returns. J Financ 1992; 47: 1731-64.

[2] Pesaran M, Timmermann A. Predictability of stock returns: Robustness and economic significance. J Financ 1995; 50: 1201-28.

[3] Edwards K. A recursive modelling for market timing of the TSE 300 index. Master thesis, University of Calgary 2000.

[4] Akaike H. Information theory and an extension of the maximum likeliihood principle. In: Petrov B, Csaki F, Eds. Second international symposium on information theory. Akademia Kiado: Budapest 1973; pp. 267-81.

[5] Schwarz G. Estimating the dimension of a model. Ann Stat 1978; 6: 461-4.

[6] Rothman D. Letter to the editor. Technometrics 1968; 10: 432.

[7] Akaike H. Fitting autoregressive models for prediction. Ann Inst Stat Math 1969; 21: 243-7.

[8] Reschenhofer E. On subset selection and beyond. Adv Appl Stat 2004; 4: 265-86.

[9] Tibshirani R, Knight K. The covariance inflation criterion for adaptive model selection. J R Stat Soc Ser B (Statistical Methodology) 1999; 61: 529-46.

[10] George E, Foster D. Calibration and empirical Bayes variable selection. Biometrika 2000; 87: 731-47.

[11] Reschenhofer E. Discriminating between nonnested models. Far East J Theor Stat 2010; 31: 117-33.

[12] Reschenhofer E, Schilde M, Oberecker E, Payr E, Tandogan H, Wakolbinger L. Identifying the determinants of foreign direct investment: a data-specific model selection approach. Working paper, University of Vienna 2009.

[13] Foster D, George E. The risk inflation criterion for multiple regression. Ann Stat 1994; 22: 1947-75.

Received: May 11, 2010 Revised: September 23, 2010 Accepted: October 5, 2010