Other evidences of the predictive power of technical analysis: the moving averages rules on European indexes

P.-J. Detry,

CeReFiM, University of Namur, Belgium

Philippe Grégoire

Université Catholique de Louvain, CeReFiM.

IAG, Place des Doyens 1

1348 Louvain-La-Neuve

Belgium

Tel : 32-(0)10-47.84.46

e-mail : gregoire@fin.ucl.ac.be

Philippe Grégoire will attend the conference.

Other evidences of the predictive power of technical analysis: the moving averages rules on European indexes

Abstract: Many authors discovered that simple forms of technical analysis possessed significant forecast power on various market indexes. We show that these results can be replicated on formally selected European indexes, which almost completely eliminates any influence from data-snooping. Implications of these results in terms of market efficiency are also discussed.

I. Introduction

Technical analysis uses past prices in order to predict future prices. It tries to detect some predefined "patterns" in price series, and claims it is capable of exploiting the trends that it discovers.

Although the vast majority of the professional traders use technical analysis, most academics, until recently, had not recognized the validity of these methods. They prefer the much more theoretical fundamental analysis. In fact, the difference between technical and fundamental analysis is the following: the first uses variables that the economic theory considers as relevant for the estimation of future dividends and of the rate at which they have to be discounted, while the second uses past prices, a set of variables that the efficient market hypothesis, it is widely believed, has shown useless for the prediction of future movements.

However, since the article of Brock, Lackonishok and LeBaron (BLL thereafter) (1992), showing that simple forms of technical analysis can significantly predict daily price movements of the Dow-Jones index, many academics have begun to realize that technical analysis might have some value. Many other were skeptical, mainly because of the huge potential impact of data-snooping when working with an index as much studied as the Dow-Jones. Indeed, when hundreds of researchers try to find predictable patterns on the same sample, they are bound to find one, by pure chance, even if the series follows a random walk. To mitigate the problem, BLL insist on the fact that the technical rules they tested had existed for a very long time (early in the sample). But many other technical rules existed as well, so that the rules they tested may have been filtered from the whole Universe of trading rules precisely by using the Dow-Jones. However, because the Dow-Jones is such a famous index, the sample itself was less likely to have been filtered out of the whole Universe of samples on which trading rules have been tested, another form of data-snooping. In short, BLL solved (part of) the data-snooping bias which was related to the choice of the sample, but they did not solve the problem that the tested rules may have been tested earlier on the same series.

This is why different articles have tried to replicate BLL's results on other samples. For instance, Hudson, Dempsey and Keasey (1996) studied about the same set of rules on the Footsie 30 index. Results were also positive, but, again, using such a famous index made it likely that the Footsie, along with the Dow-Jones, was used to establish, or rather to filter, those rules. On the other hand, Bessembinder and Chan (1995) replicate with success earlier results on a few Asian indexes. Those indexes are less likely to have been used to establish the technical rules. But, again, the skeptical reader may wonder whether the samples have not been selected by intensive testing in earlier unpublished studies.

Instead of studying out of sample results to check the effect of data-snooping, another method is to use the reality check bootstrap, as Sullivan, Timmerman and White (1999). Their conclusion is that data-snooping did not plague the results of BLL. However, this method requires the determination of the Universe of rules from which the most successful rule may have been drawn. We think this method can serve as a "first check", but cannot prove at all that data-snooping did not plague the results. Indeed, determining the Universe of rules from which the best one has been drawn is a rather difficult task. Was the Universe of Sullivan et al. large enough, not only in terms of the number of rules, but also in terms of the space they spanned (inversely proportional to the dependencies across the rules making up the Universe)? If this was not the case, and it is always possible to say that this was not the case, then their conclusion that data-snooping was not the only source of success may appear too hasty...all the more that the set of rules tested by BLL until 1986 does not anymore outperform the market from 1986 to 1999 (LeBaron, 1999).

In this article, we present further evidence of the forecast power of technical rules. By using formal selection procedures for choosing the rules and the samples, we hope to convince even the most skeptical reader that the forecast power of technical rules is (or was) real. Complete resolution of this debate is not possible, however.[1]

The remainder of this article is organized as follows:

In section 2, we briefly present the selection methodology we used to choose the samples on which we tested the technical rules, and the technical rules we tested on the samples. The samples and rules are also presented. Section 3 displays and explains the basic empirical results. Section 4 tries to check whether the basic results are robust to

---

1 One form of data-snooping that could remain is the publication bias: it is a well known fact that studies presenting unusual results are more likely to be published than the studies that just confirm a well known theory.

nonsynchronous trading and other problems regarding the data and rules we used. Section 4 gives an interpretation in terms of market efficiency and underlines the importance of the results for the relevancy of technical methods in an efficient market. Finally, section 5 summarizes and concludes.

II. Technical rules and data

1) *Technical rules*

We chose to evaluate the 10 VMA (variable length moving averages) rules of BLL. Those rules consist of comparing a short moving average of the price with a long moving average. When the short moving average (over 1, 2 or 5 days) is above the long moving average (over 50, 150 or 200 days) plus a certain percentage band, the next day is considered as a buy day. Conversely, when the short average is below the long average minus the band, the next day is classified as a sell day.

Following BLL, we evaluate the following VMA: (1,50,0), (1,150,0), (5,150,0), (1,200,0), (2,200,0), (1,50,0.01), (1,150,0.01), (5,150,0.01), (1,200,0.01), (2,200,0.01), where the first and second figure represent the number of days over which the short and long moving averages are computed, respectively, and where the third figure is the value of the band.

The reason why we chose to evaluate only VMA rules (and not the FMA, fixed length moving averages, nor TRB, trading range break-out) is because the results obtained by VMA rules in BLL were much more significant (see the p-values related to the VMA, FMA and TRB on page 1750). This fact has been confirmed by other studies, so we have good a priori reason to prefer VMA rules.

2) Data

We chose to evaluate the forecast power of technical rules on European indexes, using daily data, to replicate BLL.

Instead of arbitrarily choosing one index, as Hudson et al. (1996), we preferred testing our 10 rules on all 15 countries of the European Union. Again, this reduces the potential data-snooping bias as we avoid presenting to the reader the results on only one index. Once this decision taken, we had to choose the indexes. One possibility would have been to choose " market " indexes, as BLL with the Dow-Jones. The problem with this approach is that there are many " market " indexes in each countries. For instance, in Britain, do we have to choose the Footsie 100, or the Footsie 30 ? Another problem is that market indexes are the ones that are the most studied by researchers, so they may have been used to filter the most ex-post effective technical rules. For these reasons, we chose to select the indexes according to the following formal procedure:

- In datastream, on the index page, we selected all the indexes beginning with the name of the country (in English).
- Then, we selected the index for which the records begin the most early.
- When several indexes start at the same date, we simply chose the first in alphabetical order.

This procedure ensures that only one index is selected per country[2]. The criterion of the starting date is important, because, as Hudson et al. (1996) notice, long periods need to be considered for the results to be statistically significant[3].

---

[2] For the index of Luxemburg, we encountered a download problem, so we simply chose the next index in alphabetical order.

In table 1, we present the 15 selected indexes with some of their characteristics. When the indexes are still "alive" now, we only downloaded the series up to the 1/1/1999.

A few more remarks should be made: we chose to work on “raw” indexes, i.e. as they were available on datastream. The databases are such that even on days where the index is not quoted (except on the week-ends), the price of the previous day is used (for instance, on 25 December or on 1 January of each year). We do not expect this problem to bias our results, so we did not attempt to filter the data one way or the other. On the other hand, for some indexes, the price changes infrequently at the beginning of the sample (it sometimes remains the same for more than 20 days). We guess that the data were not always available, and that "holes" have been filled with the last data available. The only indexes for which the problem exists are the indexes for Austria, Denmark, France and United Kingdom; this problem is more serious and could potentially affect our results. To remedy this, we will split the samples in two equal subperiods, and see if our results remain the same.

At last, it should be noted that the selected indexes are often aeronautic or banking indexes, as those sectors are often first in alphabetical order. We do not see any reason why this should bias our results.

---

[3] Although we did not take into account the fact that some indexes were "dead", which means that they do not exist anymore.

Table 1: The selected Indexes

index name	start and end date	number of daily returns	mean daily returns	Standard deviation of daily returns
AUSTRIA GZ ALLSHARE 'DEAD'	15/12/72-29/11/94	5727	0.000186	0.0076408
BELGIUM-DS BANKS	1/1/73-1/1/99	6784	0.0004043	0.0092351
DENMARK-DS AIRLINES&A'PORTS	1/1/73-1/1/99	6784	0.0006755	0.0197708
FINLAND-DS PHARMACEUTICALS	4/1/93-1/1/99	1564	0.0003833	0.0154014
FRANCE-DS AEROSPACE	1/1/73-1/1/99	6784	0.0004109	0.0241857
GERMANY-DS AIRLINES&A'PORTS	1/1/73-1/1/99	6784	0.0002832	0.0196528
GREECE-DS BANKS	4/1/88-1/1/99	2869	0.0010073	0.0209374
IRELAND-DS BANKS	1/1/73-1/1/99	6784	0.0005426	0.0142020
ITALY-DS AEROSPACE	1/1/73-1/1/99	6784	0.0002962	0.0272278
LUXEMBURG-DS BANKS	2/1/92-1/1/99	1826	0.0009623	0.0094234
PORTUGAL-DS MARKET $	2/1/90-1/1/99	2348	0.0002662	0.0107094
NETHERLAND-DS AIRLINES&A'PORTS	1/1/73-1/1/99	6784	0.0001522	0.0217987
SPAIN- DS MARKET $	2/3/87-26/2/91	1041	0.0003326	0.0129249
SWEDEN-DS MARKET $	4/1/82-1/1/99	4434	0.0005385	0.0134476
UK-DS AEROSPACE	1/1/65-16/3/78	3217	0.0001256	0.0159296

III. Basic results

We used standard t-tests to evaluate whether returns following buy signals are higher than returns following sell signals, and whether those buy (sell) returns are different from the unconditional return. However, to save space, we only present results using the bootstrap methodology. The bootstrap simply consists of creating a large number (we used 1000) of artificial series by randomly rearranging (or as we did drawing with replacement) the returns from the original series, then retesting technical rules on each simulated series. If results are more favorable (for instance if returns following buy signals are higher or less volatile) in x% of the random series than on the original index, then we have a simulated "p-value" of x%. As BLL noticed, there are three advantages to this methodology. Firstly, it allows to compute a test of significance across the 10 rules, taking all their dependencies into account. Secondly, the bootstrap methodology does not require the lognormality of the price series. Thirdly, it is capable of offering a reliable test of significance for the volatility results, showing for instance whether the volatility after buy or sell signals was statistically different to the unconditional volatility. The main disadvantage of the bootstrap methodology is that it takes much more computing time.

In table 2, we present "basic" results for each index. As we work with 15 samples, we economize on space by only presenting results aggregated across all 10 trading rules, as the bootstrap methodology allows us to do so[4]. Detailed results for each rule are available on request.

[4] To obtain such a test, we simply computed, first in the original index then on the 1000 simulations, an average across the 10 rules of the parameter we wanted to test . The aggregated p-value is then the proportion of the 1000 simulations for which the same average (across the 10 rules) of the parameter is more favorable than in the original index.

We begin our description of the results with the first line of table 2, i.e. the results for the index from Austria. The first column, labeled "buy", represents the mean daily return following a day classified as "buy". In reality, this mean buy return is an average of the 10 mean buy returns across the 10 trading rules. This average mean buy return is 0.086%, which is about 24% at an annual rate. This compares to the unconditional daily return of 0.0186% (in table 1), or 4.8% at an annual rate. The p-value in brackets indicates that none of the 1000 simulations gave such a high average mean buy return. The average mean buy return can also be compared to the average mean sell return of -0.03678% (-8.85% at an annual rate). This is done in column 3: the difference between the buy and the sell return is 0.12%, more than 36% at an annual rate ! Again, none of the simulations showed such a high difference between buy and sell returns. Maybe this difference can be explained by increased volatility during buy periods. In column 4, the p-value of 0.943 indicates that in most of the simulations (943 out of 1000) the buy volatility was higher than the buy volatility observed in the original index. And column 7 shows that in only 4 simulations the mean buy volatility was lower than the unconditional volatility. Thus, it appears that volatility was higher during buy periods, and this could explain the higher return observed during buy periods. Column 8 shows that the sell volatility was higher than the unconditional volatility (what we call a favorable outcome, because the technical trader is not in the market) in 603 simulations. So the sell volatility is higher than the unconditional volatility. Although this result is not statistically significant, it is sufficient to say that the lower (more negative) return observed during sell periods cannot be explained by lower risk. In fact, column 9 shows that risk (or rather its proxy, the standard deviation) cannot explain the difference between buy return and unconditional return. Indeed, in no one of the 1000 simulations was the ratio[5] 'buy return/buy standard deviation' as high as in the original series. Thus, the higher buy volatility cannot fully explain the higher returns observed after buy signals. These results for the Austrian support the study of BLL as far as the conditional returns are concerned. But the volatility results are different: on the Dow-Jones index, the buy volatility was lower than the unconditional (and sell) volatility (BLL, 1992).

The last column shows the one way break-even trading costs associated with the 10 VMA rules. These costs (C_p) simply represent the level of transaction costs that would just have eliminated all technical profit in excess of the buy and hold strategy. They are computed as in Bessembinder and Chan (1998):

$$C_p = \pi_p / 2N_p ,$$

where - π_p is the average across the 10 rules of the returns obtained during buy day and during sell day. Thus $\pi_p = \pi_p^b + \pi_p^s$, where π_p^b and π_p^s are the average across the 10 rules of the mean returns obtained during buy and sell periods respectively;

- N_p is the average across the 10 rules of the number of buy and sell signals emitted (the reversal of position after a signal ceased to emit is not computed in N_p).[6]

[5] Using these sorts of Sharpe ratios (column 9 and 10) in conjunction with the bootstrap allows us to easily correct for risk.

[6] For more details, see Bessembinder and Chan (1998).

For the Austrian index, a trader that would have had to pay less than 1.78% of trading costs per transaction over the period would have been better off using the set of trading rules than simply buying and holding the index.

We let the reader examine the results for the other indexes. It can be noted that the results are similar for most of them. The buy-sell return is always positive, except for France and Spain, where this difference is negative (but not significantly), meaning that the returns following sell signals were higher than returns following buy signals. Except for these two indexes, the technical rules seem to perform well as far as the prediction of returns is concerned. And although the results for Finland and Germany may be due to chance, in all of the 11 remaining cases they are significant at the 95% confidence level. Adjusting for risk (columns 9 and 10) only changes the results for Greece: because the rules classify as "buy" the most risky periods, the ratios 'buy returns/buy standard deviation' and 'sell returns/sell standard deviation' do not appear statistically different in the simulations and in the index. For Portugal and Luxemburg, one of those ratios is not anymore significantly different in the simulations to what it is in the original series, while the other remains significant. In fact, we can note that the rules predict the volatility correctly in 11 cases out of 15 (considering column 6). It predicts the volatility incorrectly in the remaining 4 cases, and only significantly for Greece.

Looking at the last column, it can be noted that break-even costs are only negative for France and Spain. This is logical, since on those two indexes the 10 rules were wrong in predicting the buy and sell periods. The traders should have been paid about 0.3% per trade for the set of rules to outperform the buy and hold strategy. Apart from these two cases, break-even costs appear to be positive. And, except for Germany (0.3%), they lie between 0.63% (Portugal) and 3.34% (Luxemburg). We don't know precisely about the real trading costs in these different countries during those periods, but they probably also lied in this range for most traders.

Table 2: results of the 10 VMA rules

The numbers represent the average across the 10 rules of the parameters conditional on buy and sell signals. Numbers in brackets are bootstrap p-values testing whether the parameters are more favorable in the original index than in the simulations.

Index country	buy	sell	buy-sell	buy st dev	sell st dev	st dev buy/sell	st dev buy/uncond	st dev sell/uncond	Buy ret/buy st dev	Sell ret/ sell st dev	one way break-even trading costs
Austria	0.0008618 (0.000)	-0.0003678 (0.000)	0.0012295 (0.000)	0.0084303 (0.943)	0.0078523 (0.318)	1.0792649 (0.753)	1.1033268 (0.996)	1.0276843 (0.307)	0.1018639 (0.000)	-0.0461949 (0.000)	0.0178153
Belgium	0.0007399 (0.012)	-0.0001801 (0.003)	0.0009200 (0.000)	0.0092193 (0.590)	0.0096810 (0.216)	0.9525588 (0.277)	0.9982806 (0.589)	1.0482769 (0.189)	0.0801469 (0.005)	-0.0185854 (0.001)	0.0109342
Denmark	0.0013380 (0.012)	-0.0003299 (0.000)	0.0016679 (0.000)	0.0208350 (0.862)	0.0179261 (0.952)	1.1639799 (0.978)	1.0538260 (0.972)	0.9066953 (0.975)	0.0642332 (0.019)	-0.0185113 (0.001)	0.0252695
Finland	0.0006715 (0.275)	-0.0003631 (0.079)	0.0010346 (0.070)	0.0145078 (0.070)	0.0169862 (0.033)	0.8546941 (0.008)	0.9419768 (0.016)	1.1028985 (0.013)	-0.0057358 (0.227)	0.0459760 (0.081)	0.0088290
France	0.0001456 (0.751)	0.0007732 (0.798)	-0.0006276 (0.868)	0.0225306 (0.008)	0.0264494 (0.009)	0.8542156 (0.001)	0.9315650 (0.001)	1.0935957 (0.000)	0.0062456 (0.741)	0.0292208 (0.714)	-0.0031859
Germany	0.0004706 (0.259)	0.0001195 (0.279)	0.0003511 (0.166)	0.0184426 (0.001)	0.0214927 (0.000)	0.8581465 (0.000)	0.9384190 (0.000)	1.0936190 (0.000)	0.0255131 (0.234)	0.0055633 (0.278)	0.0032200
Greece	0.0017252 (0.046)	0.0001480 (0.06)	0.0015772 (0.009)	0.0219960 (0.897)	0.0196873 (0.853)	1.1195399 (0.958)	1.0505615 (0.981)	0.9402925 (0.906)	0.0781466 (0.082)	0.0076685 (0.072)	0.0160226
Ireland	0.0010571 (0.002)	-0.0002488 (0.000)	0.0013059 (0.000)	0.0132511 (0.004)	0.0161404 (0.001)	0.8227606 (0.000)	0.9330429 (0.000)	1.1364886 (0.000)	0.0797130 (0.000)	-0.0153914 (0.002)	0.0158343
Italy	0.0010295 (0.033)	-0.0005061 (0.037)	0.0015356 (0.005)	0.0250426 (0.262)	0.0297248 (0.172)	0.8456501 (0.073)	0.9197460 (0.077)	1.0917091 (0.071)	0.0411107 (0.021)	-0.0170724 (0.004)	0.0135069

Luxemburg	0.0013813 (0.022)	0.0001175 (0.018)	0.0012639 (0.004)	0.0099655 (0.847)	0.0094366 (0.429)	1.0631610 (0.562)	1.0575235 (0.990)	1.0013950 (0.446)	0.1386542 (0.057)	0.0109526 (0.016)	0.0334
Portugal	0.0006597 (0.058)	-0.0000266 (0.166)	0.0006863 (0.040)	0.0095691 (0.001)	0.0124738 (0.000)	0.7677967 (0.000)	0.8935269 (0.000)	1.1647567 (0.000)	0.0688856 (0.029)	-0.0019497 (0.167)	0.0063483
Netherlands	0.0008612 (0.020)	-0.0005859 (0.020)	0.0014471 (0.002)	0.0209891 (0.052)	0.0225135 (0.097)	0.9323195 (0.018)	0.9628614 (0.001)	1.0327909 (0.034)	0.0410175 (0.009)	-0.0260193 (0.026)	0.0125356
Spain	0.0001079 (0.632)	0.0011333 (0.808)	-0.0010254 (0.810)	0.0099791 (0.000)	0.0157581 (0.020)	0.6350019 (0.001)	0.7720888 (0.001)	1.2192107 (0.014)	0.0110802 (0.614)	0.0742851 (0.739)	-0.0030480
Sweden	0.0008212 (0.099)	0.0000418 (0.054)	0.0007794 (0.018)	0.0118983 (0.000)	0.0168657 (0.000)	0.7060640 (0.000)	0.8847881 (0.000)	1.2541759 (0.000)	0.0690081 (0.043)	0.0021939 (0.046)	0.0102039
UK	0.0006397 (0.06)	-0.0005815 (0.028)	0.0012212 (0.010)	0.0144356 (0.001)	0.0185757 (0.000)	0.7779761 (0.000)	0.9062140 (0.001)	1.1661135 (0.000)	0.0431441 (0.044)	-0.0282760 (0.046)	0.0127336

IV. Sensitivity tests

1) Nonsynchronous trading

The predictive power observed is undoubtedly strong. But it could be fallacious. Indeed, the technical rules that are tested rely on the continuations of movements: the crossing of the two moving averages is supposed to generate some trend that is likely to continue until the short moving average crosses the long average from the opposite sense. The problem when we try to test the predictive power of such rules on a portfolio (or an index) is called "nonsynchronous trading": Scholes and Williams (1977) show that if the securities making up the portfolio are not traded simultaneously, then positive serial correlation in the portfolio returns are likely to be observed. Indeed, news that potentially affects all the stocks of the portfolio may affect the stocks making up the portfolio at different dates, depending on the first trade after the news is released.

To check for the effect of nonsynchronous trading, we have to introduce a lag of a certain period between the moment when the signal is generated and the moment when the position is actually taken. If the stocks making up our indexes trade every day, then a one-day lag would be enough. Using this correction on the Dow-Jones Index, where stocks are heavily traded, Bessembinder and Chan (1995, 1998) show that the predictive power of technical rules diminishes[7], but is still largely significant.

The problems with this approach in our case is that we do not have enough information on the composition of our indexes to know whether the stocks were traded every day. In

[7] However, it is difficult to know whether this decrease reflects the influence of nonsynchronous trading on original results or if some exploitable predictive power has been lost due to the correction.

other words, we cannot be sure that a one day lag will be long enough to correct for the effects of nonsynchronous trading. This is why we chose to introduce a longer lag. In fact, we will use a lag ranging from 0 to 10 days for each index (it seems to be a reasonable assumption that each stock has been traded at least every 10 days). We will then assess the performance of our 10 rules on the base of the break-even trading costs. The higher those costs, the better the predictive power. So, if we see that the break-even costs diminish as we introduce lags, and appear to be equal to 0 (or less) after one, two or more days, we can conclude that nonsynchronous trading may have played an important role in the predictive power of the 10 VMA rules.

In table 3, we present the results for each index. Graph 1 illustrates the figures for the 3 indexes that have the highest break-even costs before any correction for nonsynchronous trading.

Table 3: Introducing lags

The table shows the break-even trading costs associated with the set of 10 rules for each index. Lag x means that a trading lag of x days has been introduced between the signal date and the actual trading date.

Index	Lag 0	Lag 1	Lag 2	Lag 3	Lag 4	Lag 5	Lag 6	Lag 7	Lag 8	Lag 9	Lag 10
Austria	0.0178153	0.0169810	0.0165870	0.0159263	0.0160160	0.0153421	0.0143517	0.0140800	0.0139121	0.0134488	0.0129778
Belgium	0.0109342	0.0102335	0.0097562	0.0099121	0.0101939	0.0100348	0.0096530	0.0095310	0.0084491	0.0079497	0.0079705
Denmark	0.0252695	0.0235344	0.0246815	0.0234234	0.0223848	0.0224507	0.0231708	0.0224784	0.0213565	0.0210477	0.0203970
Finland	0.0088290	0.0077143	0.0070495	0.0067512	0.0059224	0.0070773	0.0069678	0.0069290	0.0076294	0.0067619	0.0058634
France	-0.0031859	-0.0008376	-0.0006553	-0.0006665	-0.0009811	-0.0016127	-0.0012489	-0.0010359	-0.0012199	-0.0015279	-0.0016873
Germany	0.0032200	0.0043705	0.0042193	0.0046130	0.0045788	0.0042426	0.0041981	0.0033977	0.0025960	0.0018840	0.0020369
Greece	0.0160226	0.0141850	0.0135226	0.0121974	0.0110559	0.0101339	0.0082904	0.0075585	0.0064952	0.0070303	0.0068247
Ireland	0.0158343	0.0139805	0.0131275	0.0128808	0.0120406	0.0115671	0.0104629	0.0104186	0.0108505	0.0103251	0.0097517
Italy	0.0135069	0.0155306	0.0161942	0.0156284	0.0143604	0.0133615	0.0113327	0.0117357	0.0120759	0.0105872	0.0088941
Luxemburg	0.0334263	0.0348017	0.0349484	0.0335043	0.0347125	0.0335693	0.0330605	0.0333720	0.0332252	0.0325522	0.0313007
Netherlands	0.0125356	0.0107696	0.0113531	0.0127799	0.0123346	0.0112227	0.0114189	0.0111353	0.0098170	0.0089325	0.0080943
Portugal	0.0063483	0.0055039	0.0046237	0.0046987	0.0053736	0.0059036	0.0063772	0.0060922	0.0057692	0.0053283	0.0045434
Spain	-0.0030480	-0.0044537	-0.0060826	-0.0070599	-0.0061599	-0.0059552	-0.0063141	-0.0071364	-0.0067713	-0.0074478	-0.0077386
Sweden	0.0102039	0.0089167	0.0083733	0.0080252	0.0071334	0.0078406	0.0084123	0.0071698	0.0065275	0.0053936	0.0052447
UK	0.0127336	0.0114953	0.0106932	0.0108553	0.0106489	0.0102172	0.0105622	0.0109828	0.0112738	0.0095969	0.0072598

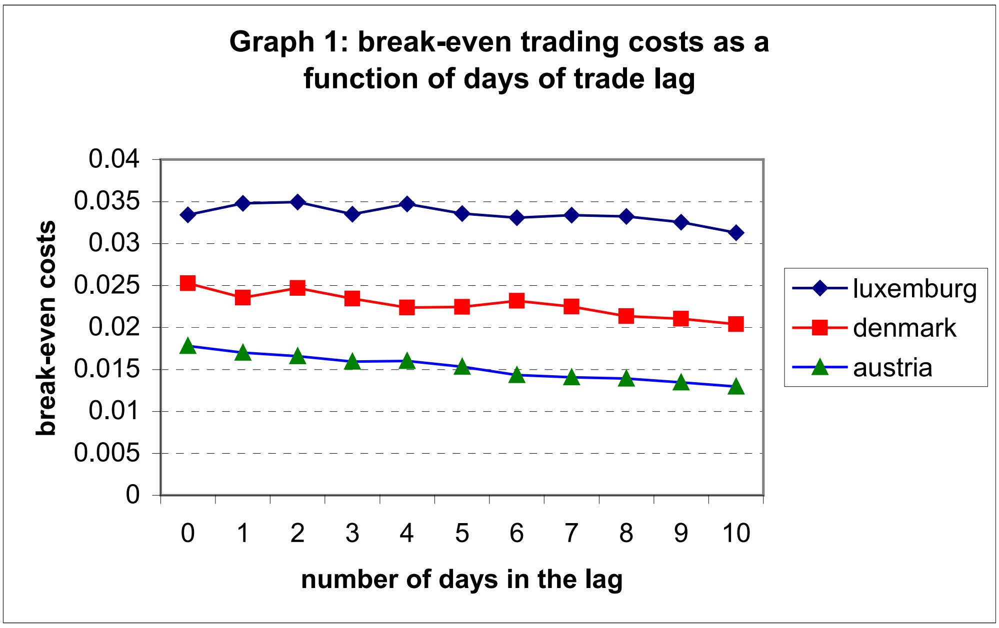

As can be readily seen, the predictive power of the 10 VMA's indeed tend to decrease as we add days in the trade lag. Nonsynchronous trading may thus be responsible for part of it. But two remarks should be made. Firstly, the decrease is not uniform: for some countries (Luxemburg, Germany and Italy) the predictive power first tend to increase with the number of lag days (from 0 to 1 day). Although this result is probably not significant, it is in sharp contrast with the non synchronous trading explanation, that does not seem to hold for those indexes. Most importantly, the decrease in the predictive power is generally small, and even after 10 days, the VMA's can still predict the indexes movements. For instance, the break-even transaction costs remain above 3% for on the index from Luxemburg. Thus, even if nonsynchronous trading has played a role, it is certainly not all the story.

2) Missing data

The second problem we want to address concerns the indexes from Austria, Denmark, France and United Kingdom. As outlined earlier, the data for those indexes seem incomplete: in the first part of the samples, the same price sometimes appear for more that 10 days, as if the "wholes" had been filled with the preceeding data. The autocorrelation induced may have played in favor of the predictive power of the VMA, that exploit some kind of nonlinear autocorrelation.

To check this, we simply splitted the samples in two equal subperiods. Then we tested the 10 rules on those separate samples. We only present the criterion of the break-even costs, for reason of simplicity.

Table 4 shows the results obtained:

Table 4: Break-even costs for subperiods

	Austria period 1	Austria period 2		Denmark period 1	Denmark period 2		France period 1	France period 2		UK period 1	UK period 2
Break-even costs	0.004163	0.029263		0.024619	0.026604		-0.000625	-0.003730		0.002434	0.017840

It is easy to see that, except for France that had negative results in the first and second subperiods (worse in the second subperiods), the 10 VMA rules perform better on the second subperiod, indicating that the autocorrelation present in the first subperiod (due to the lack of data) does not artificially improve our results. It makes them worse.

V. Interpretation

Briefly stated, our results strongly support the conclusions of BLL as for the predictive ability of moving averages rules. Having selected the samples and the rules formally, we can say fairly confidently that this predictive power is real.

But what are the implications in terms of market efficiency ? On most markets, break-even costs do not appear sufficiently high to ensure even to professional traders an excess return. This fact confirms the conclusions of other studies (see for instance Hudson et al. (1996) or Bessembinder and Chan (1995,1998)). The interesting fact is that the order of magnitude of those break-even costs (often from one to two percent) is probably not far from real trading costs encountered by professional traders. This could be easily explained if one considers that traders use technical rules to the point where they are not profitable any more, due to trading costs. This would not only imply a very high (probably too high to be realistic) degree of rationality of technical traders, but also a positive infra-marginal contribution of technical analysis to market efficiency. Stated alternatively, a market where only fundamental analysis would be used would allow the first technical analyst to earn substantial excess returns. In fact, technical analysis may exploit all elements of predictability present in stock price series, leaving only those that cannot be exploited due to transaction costs.

The idea of positive infra-marginal contribution of technical analysis to efficiency will look obvious to some reader. But no robust theory clearly explains the contribution of technical trading to efficiency[8], at least when the market is efficient in the sense that information spreads instantly among the trader's community[9].

8 We can go further: the theory of rationality, implying that the price of an asset is always equal to its expected fundamental value (i.e. the market never over or under reacts) implies that price changes only reflect new fundamental information, random by definition. So price changes should also be random and technical analysis should not have any value if markets were rational.

9 In fact, as Brown and Jennings (1989) explain in the abstract of their article, the use of past prices may be useful when information acquisition is costly, or when information does not spread instantly among the traders community. But even if the market is perfectly efficient and even if information spreads instantly,

In more practical terms, the contribution of technical analysis to efficiency implies that important financial institutions, especially in small markets, must carry on using technical analysis even if the market is efficient. Indeed, if one such institution were to lay off all their teams of technical analysts, the market could become less efficient, a fact from which competitors could profit. In fact, the use of financial analysis in an efficient market could be seen as the result of a strategic equilibrium in an attrition war type game. Modeling this game would be an interesting study, and we leave it for further research.

VI. Summary and conclusion

Our study tried to check whether BLL's results could be replicated on a series of formally selected European indexes. The technical rules we chose to test are the VMA rules presented in BLL, that seemed to perform particularly well on the Dow-Jones Industrial Average, on the Footsie 30 and on a variety of Asian indexes.

We find that in 13 cases out of 15, the VMA rules possess some predictive ability in the sense that the returns following buy signals are higher than returns following sell signals. Only for the indexes from France and Spain is this not the case. In 11 cases this predictive power is statistically significant, and in 10 cases this result is robust to risk adjustment. As far as the volatility results are concerned, our study tends to confirm BLL's results, although not as strongly: in 9 cases, technical rules can significantly select less risky periods (column 6 in table 2). But in four cases, buy signals are followed by

---

technical analysis may still have some infra marginal value, by detecting some systematic over or underreaction to information.

riskier periods than sell signals (although this result is statistically significant only for the Dennish and Greek indexes).

Interestingly, we find that break-even transaction costs, i.e. the level of transaction costs that would just have eliminated all excess profit, are often of the same magnitude as actual transaction costs encountered by professional traders: in 12 cases out of 15, one way break even trading costs lie between 0.5 and 4 percent. If these figures are more or less in accordance to the efficient market hypothesis of Fama (1970), we think they reflect the infra-marginal contribution of technical analysis to market efficiency: In fact, all happens as if chartists exploit predictability in stock prices up to the point where trading costs renders this activity not profitable anymore. This could explain why, in stock prices, we find predictable, but no profitable patterns.

The question that remains is: what is the cause of this predictive ability, of the tendency of stock prices to behave predictably? Two paths may be explored: the gradual diffusion of information, or some irrational tendency from the market[10] to under or overreact to information in some systematic manner. In our study, testing technical rules on daily series, we think the second explanation is the most likely, whereas the first explanation would be more relevant in intra-day.

This is why we believe that the cause of the predictive power of simple forms of technical analysis has to be found in the area of behavioral finance, market psychology, and all branches that do not assume the perfect rationality of markets.

[10] Or rather of an imaginary market where there would not be any technical analyst.

References

Bessembinder H., Chan K. (1995), « The Profitability of Technical Trading Rules in the Asian Stock Markets », *Pacific-Basin Finance Journal*, Vol. 3, No. 257-284.

Bessembinder H., Chan K. (1998), « Market Efficiency and the Returns to Technical Analysis », *Financial Management*, Vol. 27, No. 2, 5-17.

Blume L., Easley D., O'Hara M. (1994), « Market Statistics and Technical Analysis : The Role of Volume », *Journal of Finance*, Vol. XLIX, No. 1, 153-181.

Brock W., Lakonishok J. and LeBaron B. (1992), « Simple Technical Trading Rules and the Stochastic Properties of Stock Returns », *Journal of Finance*, Vol. XLVII, No. 5, 1731-1764.

Brown, D.P. , Jennings R.H. (1989), « On Technical Analysis », *The Review of Financial Studies*, Vol. 2, No. 4, 527-551.

Detry P.-J., (1999), « L'analyse technique, un outil fondamental ? », *FUNDP, Faculté des sciences économiques, sociales et de gestion*, Namur.

Elton E. J., Gruber M. J. (1995), *Modern Portfolio Theory and Investment Analysis*, Chapter 17, Fifth Edition, Jonh Wiley & Sons, INC.

Fama, E. (1970), « Efficient Capital Markets: A Review of Theory and Empirical Work », *Journal of Finance* 25, 383-417.

Fama, E. (1991), « Efficient Capital Markets: II », *Journal of Finance* , Vol. XLVI, No. 5, 1575-1617.

Hudson R., Dempsey M., Keasey K (1996), « A Note on the Weak Form Efficiency of Capital Markets : The Application of Simple Technical Trading Rules to UK Stock Prices-1935 to 1994 », *Journal of Banking & Finance*, 20, 1121-1132.

LeBaron B. (1999), "The Stability of Moving Average Technical Trading Rules on the Dow Jones Index", *National Bureau of Economic Research*, August 1999.

Levich R.M., Thomas L.R. (1991), "The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market : A Bootstrap Approach", *National Bureau of Economic Research*, Working Paper 3818.

Scholes M., Williams J. (1977), « Estimating Betas from Nonsynchronous Data », *Journal of Financial Economics,* 309-327.

Sullivan R., Timmermann A., White H. (1997), « Data-Snooping, Technical Trading Rule Performance, and the Bootstrap », *University of California, San Diego*, Discussion paper 97-31.

Sullivan R., Timmermann A., White H. (1998), « Dangers of Data-Driven Inference : The Case of Calendear Effects in Stock Returns », *University of California, San Diego*, Discussion paper 97-31.