DUŠAN ISAKOV AND MARC HOLLISTEIN*

Application of simple technical trading rules to Swiss stock prices: Is it profitable?

1. Introduction

Technical analysis is a generic term which includes many different techniques whose goal is to predict the future evolution of asset prices from the observation of past prices. There are two approaches to technical analysis. The first is purely graphical as it looks for patterns in past data. The second approach derives some trading rules on the basis of filters applied to past data. These techniques were introduced a long time before modern financial theory was born and have therefore no theoretical foundation. This is one of the reasons why academics have looked at these techniques with contempt. Several other facts have contributed to this situation. The main reason is that technical analysis violates one of the basic principles of financial theory: the efficient market hypothesis, which claims that it is impossible to predict future prices from the observation of past prices. Another reason is that a major part of these techniques cannot be tested as they are purely graphical and they do not have precise rules. Finally, early tests of technical trading rules have produced very poor results which reinforced the general feeling of academics towards technical analysis. However, practitioners are still using these techniques to make investment decisions often in conjunction with more traditional tools as fundamental analysis[1]. Recently, some academics have slightly changed their mind towards technical analysis as they found that it is possible to predict future returns with some simple technical trading rules.

Our paper is in line with recent literature on technical trading rules as it tests if these rules are profitable when they are applied to Swiss stock prices. We consider different trading rules which are all based on one of the main tools of technical analysis: moving averages. The idea is that financial prices are volatile but that they follow some trend. Moving averages are supposed to capture trends and leave aside the "noisy" part of the evolution of prices. According to this rule, buy or sell signals are generated by two moving averages of the level of the index: a long period moving average and a short period moving average. The strategy involves buying (being long in) the asset when the short average is above the long moving average and selling (being short in) the asset when the short period moving average is below the long period moving average. We also test these rules

* We thank Karl Keiber, Walter Wasserfallen and Heinz Zimmermann (the referees) and Pierre-André Dumont for helpful comments, suggestions and discussions. Dušan Isakov, HEC-Université de Genève, 102 Boulevard Carl-Vogt, 1211 Genève 4, tél.: ++ 41 - 22 - 705 86 11, email: Dusan.Isakov@hec.unige.ch. Marc Hollistein, Banque Cantonale de Genève, 2 Rue Diday, 1211 Genève 2, tél.: ++41 - 22 - 809 35 38, email: hollis@worldcom.ch.

with the addition of other signals such as oscillators, which are supposed to detect trend reversals. We consider two popular oscillators in this paper: the relative strength index and the stochastic indicator. We also test these rules by adding bands to the moving averages in order to avoid false signals. These strategies are tested on the Swiss Bank Corporation General Index for the period 1969–1997. The results show that a simple buy-and-hold strategy on the SBC index produces a daily average return of 0.025% or 6.25% yearly. The use of technical trading rules produces a daily average return of 0.097% or 24.30 % annualy, which is significantly different and above the buy-and-hold average return. These results are obtained with simple moving averages with a short window of one day and a long period moving average of five days.

The predictability of asset returns could be due to some well-known features of the data as non-normality, serial correlation and time-varying moments. In order to check if these features do not bias the test statistics we conduct some bootstrap tests which assume that returns follow an AR(1) and a GARCH(1,1) processes. The results show that these features are present in our data set but they are not the cause of the profitability of technical trading rules. Finally, we consider if these results still hold for individual stocks and in the presence of transaction costs. The results for individual stocks are similar to those found for the SBC index. When we consider transaction costs, we find that small investors cannot benefit from the profits generated by the trading rules which means that the weak form efficient market hypothesis cannot be rejected for a large fraction of the market. Only some large investors fulfilling certain conditions could possibly get some profit from these techniques.

The paper is organized as follows. Section 2 summarizes the literature on the use of technical trading rules. Section 3 presents the data used in this study and section 4 gives the results for different trading rules based on moving averages and oscillators. Section 5 provides the empirical results obtained with the bootstrap methodology and section 6 considers the real profitability of these strategies by considering the results obtained on individual stocks and also by incorporating the trading costs in the computation of the trading rule results. Section 7 offers some conclusions.

2. Previous research

Technical trading rules investigated in academic literature can be divided in two major areas: filter rules and moving average rules. Early research focused on filter rules. This rule involves buying a security if it had risen by x% on the last period or selling it if its price has decreased by x% on the last period. ALEXANDER (1961) was the first to examine the profitability of this kind of rule on individual stocks and he found that they were profitable. In a second article, ALEXANDER (1964) included transaction costs and found that the profits generated by this strategy vanished. FAMA and BLUME (1966) confirmed this conclusion and this led the academic community to be skeptical about technical analysis not only because it lacked theoretical justification but also because it yielded poor results. SWEENEY (1988) re-examined the results of FAMA and BLUME (1966) for another period and found that, depending on the level of transaction costs, filter rules still yielded profitable results.

At the beginning of the nineties, research focused on moving average rules. BROCK, LAKONISHOK and LEBARON (1992) investigated moving average rules on a century of daily data of the Dow Jones Industrial Index. They found that these rules yielded profitable results and that the signals generated by these rules were able to detect abnormal returns when compared to the average buy-and-hold return. They investigated moving averages of length 1,2 and 5 days for the short period and 50 to 200 days for the long period moving average. They also investigated other rules based on resistance levels showing that they were also generating signals which are able to detect abnormal returns. They conducted some bootstrap tests and showed that the results obtained with these strategies were robust to other specifications of the return generating process. They did not include transaction costs in their tests, so they could not conclude if these strategies were really implementable and profitable. HUDSON, DEMPSEY and KEASEY (1996) who replicate BROCK, LAKONISHOK and LEBARON (1992) tests on the UK stock market for the period 1935 to 1994 considered this issue. They also found profitable results with the moving average strategies but these profits vanished when transaction costs were considered. LEVICH and THOMAS (1993) and KHO (1996) also considered moving average strategies but on another asset: currency futures. Both studies found profitable results for these strategies even by taking account of transaction costs. KHO (1996) showed that these results are partly due to a time-varying risk-premia, a new avenue for future research in this field.

The literature on the use of technical trading rules indicates that it is possible to obtain profitable results by using these strategies. It is interesting that only a few of the possible strategies provided by technical analysis have been investigated so far. However, there is no clear-cut conclusion on the profitability of these strategies when transaction costs are considered.

Table 1: Summary statistics for daily returns of the SBC index

Number of observations	7084
Mean	0.000250
Standard deviation	0.008422
Skewness	–1.164
Kurtosis	19.576
ρ(1)	0.101*
ρ(2)	0.011
ρ(3)	0.004
ρ(4)	0.045*
ρ(5)	0.030*

* Indicate a significant number at the 5% level for a two-tailed test, ρ are the autocorrelation coefficients

3. Data description

Our study examines the profitability of technical trading rules applied to the Swiss Bank Corporation General Index for the period running from the beginning of January 1969 to the end of December 1997. We have chosen this index because it is the only broadly based Swiss index for which data is available on a long period of time. The SBC index was created in 1963 to reflect as closely as possible the evolution of the Swiss stock market (SBC (1963)). Although the index was recomputed until December 1958, we have only obtained daily data on the index from Datastream International since January 1969. This index is a large-scale, value-weighted index, as it includes all the available securities on the market. It is computed with the Laspeyres formula.

The SBC index is a price index and therefore does not include dividends. In order to truly reflect the evolution of stock prices we should use a performance index (which includes dividends) to test the trading rules. The omission of dividends in the index could lead to wrong signals from the trading rule as the dividend payment induces a drop in stock prices[2]. Unfortunately a performance index over the considered period is not available for the Swiss stock market. However as the dividend yield is relatively low in Switzerland (on average it is never higher than 5% over the period) and as we use a value-weighted index we expect that the effects of individual dividend payments on the index are diluted and therefore that the results for the trading rules should be close for both type of indexes[3]. The composition of the index is adjusted twice a year to reflect the fact that some new securities are available and that others disappear from the market. Table 1 presents summary statistics for the daily returns of this index, which

are computed as the first difference of price logarithms.
The figures in Table 1 show that the return series is asymmetric as indicated by the negative skewness coefficient and that it is leptokurtic, i.e. it has fatter tails than the normal distribution. There is also some positive autocorrelation in returns which is a common phenomenon in indexes. As the SBC index contains the majority of Swiss stocks, a non-negligible fraction of them is relatively illiquid and therefore stale prices (due to stocks which are not traded every day) could explain the large first-order autocorrelation. Despite the fact that the fourth- and fifth-order autocorrelation coefficients are statistically significant, it is likely that they are spurious as they do not make much sense economically. The mean daily buy-and-hold return for the index is 0.025% or, assuming 250 working days in a year, an average of 6.25% yearly. We also apply technical trading rules to individual stocks. Data for the stocks is also obtained from Datastream. In certain cases, some trading rules give a neutral signal, i.e. neither a buy or sell signal. We assume throughout our study that when the investor receives a neutral signal he invests its assets in a risk-free asset. It is therefore necessary to choose an appropriate rate to reflect the yield on the risk-free asset. We use a one-day money market rate called the "tomorrow next" rate which was either set in Zürich or on the Euromarket. This rate is obtained in various issues of the monthly bulletin of the Swiss National Bank.

4. Empirical results

4.1 Moving averages

One of the simplest, oldest and most widely used technical trading rule is the moving average rule. According to this rule, buy or sell signals are generated by two moving averages of the level of the index: a long period moving average and a short period moving average. The strategy involves buying (being long in) the asset when the short average is above the long moving average and selling (being short in) the asset when the short period moving average is below the long period moving average. The use of moving average rules is based on the fact that financial time-series are volatile and on the belief there exist some underlying trends in these series. When a short period moving average cuts a long period moving average, a trend is supposed to be initiated. The most popular moving average rule used is (1,200), where the short period is one day (in fact it is the index itself) and the long period is 200 days (almost a year). The academic literature has shown that the best results were obtained when the short average is one day but has not reached any distinct conclusion on the length of the long period. The different lengths considered in these papers were 200, 150, 100 and 50 days. Our paper also investigates shorter lengths for the long period as 30, 10 and 5 days. Numerous variations on moving average rules exist. They basically add other signals to the relative positions of moving average to detect trend reversals or other phenomena. We consider two popular kinds of variations: bands and oscillators. Bands are used to eliminate „noisy" signals. A band of 1% around the long-term moving average is often used in practice. This means that if the difference between the long-term and short-term moving average is less than 1% of the value of the long-term average, there is no clear signal and the investor is neutral. In those situations, he should be out of the market and invest in the risk-free asset.
The results for various lengths of moving averages are presented in Table 2. The moving average rule is used to divide the full sample in either buy or sell periods. The strategy investigated here is the following: when the investor observes a buy signal he holds a long position in the index and when he observes a sell signal the investor holds a short position in the index. The first column of Table 2 indicates the length of the moving averages for the trading rules. The next two columns indicate the number of days when the investor is long or short.

The figures in brackets indicate the proportion of right signals i.e. the percentage of positive returns observed after buy signals and the percentage of negative returns observed after sell signals. The next two columns report the average daily return obtained in long or short positions. For these figures, we compute t-statistics which test if the average return obtained in long (or short) positions is significantly different from the average return obtained by the buy-and-hold strategy. According to BROCK, LAKONISHOK and LEBARON (1992) these t-statistics are computed in the following way:

$$\frac{\mu_z - \mu}{\left(\frac{\sigma^2}{N_z} + \frac{\sigma^2}{N}\right)^{1/2}} \qquad (1)$$

where μ_z is either the buy or sell period mean return and N_z is the number of observations in these periods. μ and N are respectively the unconditional mean of the series and the total number of observations. σ^2 is the estimated variance for the entire sample. The figures in parentheses under the average returns are the standard deviations of the different periods. What really matters to the investor systematically following these rules is to know whether or not the return obtained from these strategies earn him a return which is superior to that obtained by the buy-and-hold strategy. The mean return of this global strategy is[4]:

$$\mu_{str} = \frac{1}{T}\sum_{t=1}^{T} r_t d_t \qquad (2)$$

where r_t is the return on the index at time t, d_t is a variable which equals 1 if the signal is buy and equals –1 if the signal is sell. The test which checks whether the average return obtained through the global strategy based on the trading rule is different from the buy-and-hold return is performed with the following t-statistics:

$$\frac{\mu_{str} - \mu}{\left(\frac{\sigma^2_{str}}{N_{str}} + \frac{\sigma^2}{N}\right)^{1/2}} \qquad (3)$$

where μ_{str} is as defined in equation (2) and σ^2_{str} is the variance of the return of this strategy. The last column of the table presents the number of trades generated by the strategy. This is the number of times it is necessary to buy or sell the index according to the signals of the trading rule. This figure is important when we consider transaction costs. We consider different combinations for computing moving averages. It appears that results obtained with short period moving averages of 2 or 5 days are all dominated by those obtained with a one-day moving average. Therefore, we only present the results obtained with a one-day moving average in Table 2.

The results of Table 2 are striking. They show that the signals produced by the trading rule based on moving averages are able to clearly identify positive and negative returns on the index. Moreover, in the majority of cases, the average returns are significantly different from the buy-and-hold return. The fact that the rule is appropriate is confirmed by the number of buy positions, which is superior to the number of sell positions. This is consistent with an upward-sloping trend. We also observe that buy signals are more accurate than sell signals as shown by the larger fraction of right signals obtained in buy periods. The second more remarkable result is that the global strategy consisting of being long in the market after buy signals and short after sell signals produces an average daily return which is above and significantly different from the buy-and-hold return. For instance, the average return of the strategy (1,5) is 0.0972%, which in annual terms amounts to 24.30%, an impressive average compared to the 6.25% obtained with the buy-and-hold strategy. Notice that the average return of these strategies increases monotonically when the length of the long period moving average decreases. Strategies with a long average below 50 days all yield returns which are significantly different from the buy-and-hold return.

Moreover, the results are all economically significant as the returns from every trading rule are above the buy-and-hold return. As expected, the number of trades increases with the reduction of the window of the long moving average, because trading rules are more sensitive to variation in the index. In terms of volatility, buy periods have a lower standard deviation than sell periods. This is consistent with a well-known feature of asset returns, called the leverage effect and initially documented by BLACK (1976), where the volatility associated to negative returns is larger than the volatility associated to positive returns. What is more puzzling is that the results of the global strategies have a higher average return than a buy-and-hold policy but a similar standard deviation of about 0.0084.

Another interesting question is whether these results hold on subperiods. We can report that the results are fairly stable as the same type of results is obtained for subperiods of fifteen and ten years[5]. We also find that the highest returns are obtained with the (1,5) rule and that they are systematically above the buy-and-hold returns.

Finally, it could be argued that the results of Table 2 are not feasible on the Swiss market as short positions could be relatively difficult to build over the period. Another way to achieve similar results is to use the following strategy: when he observes a buy signal the investor borrows and doubles his investment in the index. This yields twice the market return less the risk-free rate. When the investor observes a sell signal, he sells the index and invests all his money in the risk-free asset. If there is an equal number of buy and sell signals and if the borrowing and lending rates are close, then such a strategy would yield similar results to the long-short strategy investigated in Table 2. When we implement this alternative strategy on our sample we find that the results of such

Table 2: Results for moving average rules

Trading rule	N(buy)	N(sell)	μ(buy)	μ(sell)	μ(strategy)	N(trades)
(1,200)	4466	2618	0.000519	–0.000210*	0.000405	144
	[52.4]	[49.3]	(0.00747)	(0.00982)	(0.00842)	
(1,100)	4331	2753	0.000545	–0.000214*	0.000417	220
	[52.7]	[48.6]	(0.00736)	(0.00984)	(0.00842)	
(1,50)	4220	2864	0.000662*	–0.000358*	0.000539*	366
	[52.6]	[48.0]	(0.00720)	(0.00992)	(0.00841)	
(1,30)	4182	2902	0.000774*	–0.000506*	0.000664*	478
	[53.0]	[48.8]	(0.00705)	(0.01003)	(0.00840)	
(1,10)	3971	3113	0.000923*	–0.000608*	0.000785*	962
	[53.3]	[48.7]	(0.00723)	(0.00967)	(0.00839)	
(1,5)	3863	3221	0.001120*	–0.000794*	0.000972*	1500
	[54.6]	[49.3]	(0.00718)	(0.00960)	(0.00837)	

The column "Trading rule" gives the length of the moving averages for the trading rules. N(buy) and N(sell) indicate the number of days when the investor is long or short. Figures in brackets indicate the proportion of right signals, e.g. having a positive (negative) return after a buy (sell) signal. μ(buy) and μ(sell) report the average daily return obtained in long or short positions. μ(strategy) gives the average daily return obtained with the strategy over the whole period. Figures in parentheses are the standard deviations of the strategies. N(trades) is the number of times it is necessary to change the position according to the trading rule.* indicates that the average return is significantly different from the average buy-and-hold return at the 5% level for a two-tailed test.

an investment policy yield very similar results[5]. confirming therefore that the results of Table 2 are relevant and feasible.

4.2 Moving averages and bands

Let us now turn to some refinements of the basic moving average rule. Table 3 presents the results of the use of a 1% band with moving averages. The idea behind the use of bands is to avoid "noisy" signals or in other words to be sure that a trend is really initiated. The principle is the following: when the distance between the short moving average and the long moving average is less than 1% of the long moving average, it is considered that the relative positions of moving averages cannot give reliable indications regarding the existence of a trend in stock prices. If such a situation happens the individual should not invest in the market and should hold the risk-free asset. The main difference with the previous strategy based on the crossing of moving averages alone is that there is not only a possibility of being either short or long but also neutral, that is out of the market and holding the risk-free asset. In the case of bands, the computation of the mean return of the global strategy implied by the technical trading rule is the following:

$$\mu_{str} = \frac{1}{T}\sum_{t=1}^{T}\left(r_t d_{1t} + f_t d_{2t}\right) \tag{4}$$

where r_t is the daily return on the index at time t, f_t is the one-day risk-free return, d_{1t} is a variable which equals 1 if the signal is buy, –1 if the signal is sell and 0 if the signal is neutral, d_{2t} is a variable which equals one if the signal is neutral and 0 if the signal is either buy or sell. The computation of the t-statistic is identical to equation (3).
Table 3 shows that the number of long and short positions decreases with respect to Table 2. Ex-

Table 3: Results for moving average rules with 1% band

Trading rule	N(buy)	N(sell)	μ(buy)	μ(sell)	μ(strategy)	N(trades)
(1,200)	4201	2290	0.000519	–0.000115	0.000355	285
	[52.6]	[49.2]	(0.00750)	(0.0101)	(0.00816)	
(1,100)	3878	2315	0.000634*	–0.000301*	0.000459	420
	[53.4]	[49.4]	(0.00745)	(0.01024)	(0.00804)	
(1,50)	3378	2155	0.000788*	–0.000389*	0.000518*	678
	[53.8]	[48.3]	(0.00708)	(0.01072)	(0.00768)	
(1,30)	2999	1947	0.000939*	–0.000529*	0.000577*	857
	[54.5]	[48.9]	(0.00713)	(0.01087)	(0.00735)	
(1,10)	1822	1269	0.001391*	–0.001114*	0.000624*	1429
	[55.5]	[51.1]	(0.00782)	(0.01279)	(0.00673)	
(1,5)	836	741	0.001762*	–0.001032*	0.000410	1368
	[55.2]	[49.7]	(0.00931)	(0.01459)	(0.00573)	

The column "Trading rule" gives the length of the moving averages for the trading rules. N(buy) and N(sell) indicate the number of days when the investor is long or short. Figures in brackets indicate the proportion of right signals, e.g. having a positive (negative) return after a buy (sell) signal. μ(buy) and μ(sell) report the average daily return obtained in long or short positions. μ(strategy) gives the average daily return obtained with the strategy over the whole period. Figures in parentheses are the standard deviations of the strategies. N(trades) is the number of times it is necessary to change the position according to the trading rule.* indicates that the average return is significantly different from the average buy-and-hold return at the 5% level for a two-tailed test.

cept for the first strategy, the mean returns of buy and sell periods are all higher in absolute value with this rule than without it, showing that the introduction of a band has removed days with poorer performance and it permits to sort out the most extreme returns. This is confirmed by the proportion of right signals which have increased with respect to those obtained in Table 2. However, the average return of the global strategy is inferior to those obtained without a band in Table 2. This is due to the fact that the use of bands induces neutral positions which yield much less return (the risk-free rate) than days when the investors is in the market. This is particularly true for the (1,5) strategy, where the rule induces only 1577 days where the investor is in the market and 5507 days with neutral position. Despite the use of bands permits to identify higher buy or sell returns, the global profits are lower than strategies without bands because of the neutral positions.
As before we notice that the volatility of short periods is higher than those of buy periods. We also notice that the volatility of the global strategy is always smaller than the volatility of the buy-and-hold return. This is due to the fact that this strategy has a reduced risk when the investor has neutral position (and holds the risk-free asset). Finally it is of interest that except for the (1,5) strategy the number of trades is larger in Table 3 than the number of trades incurred by the strategy without bands in Table 2. This phenomenon can be explained in the following way: when the short moving average crosses the 1% band without crossing the long moving average a trade is generated when the rule with bands is used. If the moving average rule without bands is used no trade would have been generated in this situation. As the number of trades has increased significantly from Table 2 to Table 3 it can be deduced that the short moving average crosses more often the 1% band than the long moving average.

4.3 Moving averages and oscillators

Bands were used to get more clear-cut signals on the beginning of a trend. Another aim of technical analysis is the prediction of trend reversals. This is typically what is achieved by tools called oscillators. These indicators are complementary to moving averages as they are supposed to give appropriate signals to neutralize (to step out of the market) a short or a long position. Oscillators try to detect if an asset is overbought or oversold in which case they give the signal to neutralize the position. According to PRING (1991) and BÉCHU and BERTRAND (1998) the two most popular oscillators are the relative strength index (RSI) and the stochastic indicator (SI). Our paper considers both of them. The RSI has been proposed by WILDER (1978). It is defined as:

$$RSI_t = 100 - \frac{100}{1 + RS_t} \text{ with}$$

$$RS_t = \frac{\sum_{n=t-d}^{t} \max(0, p_n - p_{n-1})}{\sum_{n=t-d}^{t} \left|\min(0, p_n - p_{n-1})\right|} \quad (5)$$

where d is the number of days on which the RSI is computed and p_n is the price at time n. Intuitively, the RSI compares the magnitude of increases in the price level of an asset with the magnitude of decreases over a given period. A high ratio means that the rises in prices have been more frequent and larger than decreases in prices. This situation is considered as overbought, and the asset under consideration should be sold as a return reversal is expected in the near future. On the other hand a low ratio means that the rises in prices have been less frequent and lower than decreases in prices. This situation is considered as oversold, and the position in the asset under consideration should be neutralized as a return reversal is expected in the near future.
The number of days d and the level of neutralization must be determined before applying the RSI.

As for moving averages these parameters do not correspond to some theory but are rather determined by practice. According to BÉCHU and BERTRAND (1998), popular levels for the number of lags used is 5, 14 and 21 days (which represent roughly 1, 3 and 4 weeks of trading) and levels of neutralization are 90 for overbought situations and 10 for oversold situations (the magnitude of rise (decreases) has been 9 times larger than decreases (rises)). This means that when the RSI is over 90 or under 10, the position is neutralized. As the results for different lags d and for different levels of neutralization are very close, we only present the results for d equals to 21 days and neutralization levels 90/10 in Table 4.

The results are close to those obtained without oscillators in Table 2. This is due to the fact that the number of days where the position is neutralized by the RSI is small. The limits of 90 and 10 are only crossed a few times by the RSI. The only time when it happens, it does not remove enough returns to significantly improve the performance of the strategy.

The other popular oscillator is the stochastic indicator (SI). It is defined as follows:

$$SI_t = 100\left\{\frac{1}{y}\sum_{t=1-y}^{0}\left[\frac{p_{t-1} - L_{t-1,t-1-x}}{H_{t-1,t-1-x} - L_{t-1,t-1-x}}\right]\right\} \qquad (6)$$

where $x > y$, p_t is the price at time t, $L_{t-1,t-1-x}$ is the lowest price between time t–1 and t–1–x, $H_{t-1,t-1-x}$ is the highest price between time t–1 and t–1–x. The SI is another way of depicting overbought or oversold situations. Instead of focusing on a series of variations, it focuses on the distance between the last quoted price and the high/low of a price on a certain window.

More precisely, the SI compares the distance between the last price of an asset and the lowest price in a period of x days before with the distance between the highest and lowest price on the same

Table 4: Results for moving average rules with 21-day relative strength index

Trading rule	N(buy)	N(sell)	μ(buy)	μ(sell)	μ(strategy)	N(trades)
(1,200)	4385 [52.4]	2607 [49.3]	0.000501	–0.000200*	0.000384	181
(1,100)	4248 [52.8]	2742 [48.6]	0.000534	–0.000204*	0.000401	262
(1,50)	4137 [52.6]	2853 [47.9]	0.000653*	–0.000349*	0.000523	408
(1,30)	4099 [53.0]	2891 [48.7]	0.000767*	–0.000497*	0.000648*	520
(1,10)	3895 [53.3]	3102 [48.6]	0.000926*	–0.000601*	0.000773*	999
(1,5)	3789 [54.6]	3210 [49.2]	0.001124*	–0.000772*	0.000959*	1522

The column "Trading rule" gives the length of the moving averages for the trading rules. N(buy) and N(sell) indicate the number of days when the investor is long or short. Figures in brackets indicate the proportion of right signals, e.g. having a positive (negative) return after a buy (sell) signal. μ(buy) and μ(sell) report the average daily return obtained in long or short positions. μ(strategy) gives the average daily return obtained with the strategy over the whole period. N(trades) is the number of times it is necessary to change the position according to the trading rule.* indicates that the average return is significantly different from the average buy-and-hold return at the 5% level for a two-tailed test.

Table 5: Results for moving average rules with 10/20 days stochastic indicator

Trading rule	N(buy)	N(sell)	μ(buy)	μ(sell)	μ(strategy)	N(trades)
(1,200)	2974 [51.0]	2118 [49.6]	0.000258	–0.000267*	0.000223	459
(1,100)	2748 [51.3]	2171 [49.1]	0.000267	–0.000311*	0.000236	574
(1,50)	2575 [51.1]	2260 [48.4]	0.000450	–0.000511*	0.000365	743
(1,30)	2534 [51.9]	2293 [49.4]	0.000643*	–0.000690*	0.000493	844
(1,10)	2623 [52.3]	2640 [48.9]	0.000859*	–0.000674*	0.000600*	1159
(1,5)	2746 [53.5]	2852 [49.5]	0.001046*	–0.000827*	0.000764*	1664

The column "Trading rule" gives the length of the moving averages for the trading rules. N(buy) and N(sell) indicate the number of days when the investor is long or short. Figures in brackets indicate the proportion of right signals, e.g. having a positive (negative) return after a buy (sell) signal. μ(buy) and μ(sell) report the average daily return obtained in long or short positions. μ(strategy) gives the average daily return obtained with the strategy over the whole period. N(trades) is the number of times it is necessary to change the position according to the trading rule.* indicates that the average return is significantly different from the average buy-and-hold return at the 5% level for a two-tailed test.

period of x days. This ratio is then recomputed for y preceding periods and averaged over time. This means that if the latest price is systematically close to the lowest observed price in a certain period, the asset is considered as oversold and that prices are expected to rise in the near future. On the other hand if the last price is systematically close to the highest price observed then the SI indicates that the prices are going to drop soon. Intuitively, this means that if the latest observed price is systematically close to the highest/lowest price during several days, a trend reversal is expected.

Again, the user of this indicator must choose some levels and no theoretical arguments are available for this choice. Popular number of lags for x and y include 5/10 days (1 and 2 trading weeks) and 10/20 days (2 and 4 trading weeks). Levels of neutralization are also 90 and 10. Again we compute the results for 5/10 and 10/20 lags and for different levels of neutralization. As we observe very similar results we only present the results for the 10/20 SI and 90/10 neutralization level in Table 5. Compared to the RSI, the SI neutralizes more frequently long or short positions as can be seen from Table 5. However the days which are neutralized are not the worst days as the average returns are again very close to those obtained with simple moving averages. For the global strategy the lower average return can also be attributed to the presence of the neutral signals.

The conclusion of tests performed with various bands or oscillators is that they do not lead to a great improvement over results obtained with technical trading rules using only double moving averages. The use of bands permits to identify higher returns but this is compensated by the fact that it also induces neutral positions which reduce the average return of the global strategy. Oscillators do not really permit to identify higher returns than simple moving averages. In the rest of the paper, we only analyze results obtained with simple moving averages.

5. Empirical results of bootstrap tests

Statistical tests performed in the previous section assume that returns are normally distributed, that the observations are independent and that the distribution does not change trough time. Since FAMA (1965) a number of studies have documented that asset returns have special statistical properties. In particular return distributions are known to be non-normal, returns present some degree of dependance and they have time-varying moments. For instance, if the returns are normally distributed they should have a skewness of 0 and a kurtosis of 3. Table 1 shows that it is not the case and that the return distribution of the SBC index has fatter tails than the normal and is asymmetric. Moreover, the returns are not independent as is witnessed by the significant $\rho(1)$, $\rho(4)$ and $\rho(5)$ coefficients. The apparent predictability of the returns could simply be due to these features. To be more precise the problem is that as usual statistical tests used in the previous section do not take into account these deviations, they could indicate that the average return obtained with the trading strategies are statistically significantly different from the average return that is obtained with the buy-and-hold strategy, but in reality they can not be considered as being different. In order to check if these features of asset returns modify the distribution of test statistics, we use the bootstrapping method. The basic idea is to simulate the empirical distribution of returns and compute p-values with respect to these simulated distributions instead of the theoretical normal distribution. There are three steps in this approach: estimate the alternative model (AR and/or GARCH model), simulate the empirical distribution and finally compare the parameters with the empirical p-value.

The application of the bootstrap methodology in this context has been proposed by BROCK, LAKONISHOK and LEBARON (1992).[6] Empirical distribution of the parameters (in our case, the mean return produced by the trading rule) are simulated under various null hypothesis for the return generating process. The null model is estimated on the original series, e.g. a random walk with drift. Then the residuals of the model are resampled; i.e. they are randomly drawn with replacement from the original residual series, to form a new simulated series of returns and prices for the index. Technical trading rules are then applied to this simulated series and an average return is obtained for the buy and sell periods as well as for the global strategy. This step (resampling) is repeated 500 times[7] to get the empirical distributions of the mean return of the trading rule under the null model. The return obtained with the original series is then compared to this distribution as it would have been done with a theoretical distribution. In our paper, we consider several null models which all coincide with observed features of the data. The first null model is a random walk with drift, which simply assumes independence and identical distribution for returns but does not assume a normal distribution. The second specificiation that we consider is an AR(1) model because of the strong first order autocorrelation in returns documented in Table 1. This model is the following:

$$r_t = a_0 + a_1 r_{t-1} + \varepsilon_t \text{ with } |a_1| < 1 \qquad (7)$$

This model assumes time-varying expected return and it is estimated on the SBC index returns using ordinary least squares. The two other models considered for null hypothesis are ARCH-type models. These models explicitly take account of the fact that financial series exhibit a time-varying conditional variances[8] or more precisely that the conditional variance of returns changes through time according to an AR process. DUBOIS and DURINI (1995) found that these models capture properly the dynamics of Swiss stock returns. We first estimate an AR(1)-GARCH(1,1) model on the returns SBC index:

$$\begin{aligned} r_t &= a_0 + a_1 r_{t-1} + \varepsilon_t \\ \sigma_t^2 &= \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \end{aligned} \qquad (8)$$

Where ε_t is a normally distributed residual with mean zero and variance σ_t^2. This model assumes that both conditional mean and variance are time-varying. Different ARCH models exist and they differ by the specification of the variance equation. Another common feature of stock returns we have already mentioned is the leverage effect, where there are asymmetric responses to past shocks on variance, i.e. past negative and positive shocks have different impacts on the present conditional volatility. Different models have been proposed to capture this feature of the data as the EGARCH model of NELSON (1991) or the Asymmetric GARCH of GLOSTEN, JAGANNATHAN and RUNKLE (1993). We only present the results obtained with the latter (GJR) specification as they are close to those obtained with the EGARCH model. The GJR model can be written as:

$$\begin{aligned} r_t &= a_0 + a_1 r_{t-1} + \varepsilon_t \\ \sigma_t^2 &= \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 d_{t-1} \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \end{aligned} \quad (9)$$

where d_t is a dummy which takes the value of 1 if ε_t is negative and zero otherwise and ε_t is a normally distributed residual with mean zero and variance σ_t^2. The results of the various estimations are presented in Table 6. GARCH models have been estimated with the maximum likelihood method and t-statistics are computed with the robust BOLLERSLEV-WOOLDRIDGE (1992) standard errors.

For every estimated model the residuals are resampled with replacement and trading rules are applied to the new series. This operation is repeated 500 times and then we compute the fraction of simulated returns which have a greater mean (or standard deviation) than the mean return (or standard deviation of returns) obtained from the original SBC index. These fractions can be interpreted as p-values and are shown in Table 7. The results in the μ columns are the results for the average return and those in the σ columns are the results for standard deviations. Each row represents a different null model. The number in the column μ(buy) and row RW for the strategy (1,200) is 0.01 and it means that only 1% of the simulations generated by the random walk model yields an average buy return higher than the average buy return obtained on the original series. It means that the high level of average buy return cannot be explained by the fact that the series follows a random walk. For the σ(buy), we see that 100% of the standard deviations generated by the random walk are greater than the random walk of the original SBC index. This means that the random walk model cannot explain the low level of volatility observed in buy periods. For the sell periods, we see that 100% of the mean returns generated by the random walk are higher than the

Table 6: Estimation results for various null models

Panel A: AR(1) model			
a_0: 0.000251*	a_1: 0.100772*		
Panel B: AR(1)-GARCH(1,1) model			
a_0: 0.000534*	a_1: 0.168754*		
α_0:0.00000555*	α_1: 0.166264*	β_1: 0.758429*	
Panel C: AR(1)-GJR(1,1) model			
a_0: 0.000320*	a_1: 0.177412*		
α_0:0.00000576*	α_1: 0.082022*	α_2: 0.173189*	β_1: 0.751590*

Numbers marked with a * are significant at the 5% level for a two-tailed test average return of the original series and that all the volatilities generated by the sell signals on the random walk model are lower than the standard deviation of the sell periods of the original series. Again the random walk model can neither explain the level of the mean or the standard deviation of returns observed on the original series. Finally, we see that there is 0.2% of the simulated average returns which are larger than the average return of the global strategy applied to the original series. In this case, the standard deviations are similar. These results indicate that the random walk with drift model cannot explain the various results obtained by applying trading rules to the original series. We repeated these simulations for the various strategies and for the various null models as can be seen from Table 7. Globally, the results are very similar to those obtained for the random walk model for the (1,200) strategy. The only exception is for the GARCH and GJR model which can produce slightly larger average buy returns and also slightly higher standard deviations in sell periods. But this is not enough to explain the results obtained with the trading strategy on the original SBC index. The conclusion of the bootstrap simulations is that the predictability and profits obtained by applying trading rules on the original SBC index are not the result of the omission of one of the well-known features of asset returns as non-normality, autocorrelation, or time-varying mean or variance.

Table 7: Results of bootstrap simulations

Trading rule	Null model	μ(buy)	σ(buy)	μ(sell)	σ(sell)	μ(strat.)	σ(strat.)
(1,200)	RW	0.010	1.000	1.000	0.000	0.002	0.604
	AR(1)	0.034	1.000	0.992	0.000	0.010	0.574
	GARCH(1,1)	0.084	0.974	0.976	0.058	0.028	0.460
	GJR(1,1)	0.108	0.992	0.978	0.098	0.026	0.578
(1,100)	RW	0.006	1.000	1.000	0.002	0.000	0.582
	AR(1)	0.046	1.000	0.982	0.000	0.014	0.542
	GARCH(1,1)	0.110	0.990	0.964	0.086	0.030	0.478
	GJR(1,1)	0.166	0.998	0.972	0.092	0.042	0.572
(1,50)	RW	0.004	1.000	1.000	0.000	0.000	0.554
	AR(1)	0.018	1.000	1.000	0.000	0.002	0.608
	GARCH(1,1)	0.048	1.000	0.982	0.064	0.000	0.548
	GJR(1,1)	0.096	1.000	0.990	0.076	0.008	0.584
(1,30)	RW	0.000	1.000	1.000	0.000	0.000	0.596
	AR(1)	0.002	1.000	0.998	0.000	0.000	0.592
	GARCH(1,1)	0.020	1.000	1.000	0.050	0.002	0.496
	GJR(1,1)	0.068	1.000	0.998	0.082	0.004	0.630
(1,10)	RW	0.000	1.000	1.000	0.000	0.000	0.610
	AR(1)	0.006	1.000	1.000	0.000	0.000	0.604
	GARCH(1,1)	0.048	1.000	0.988	0.066	0.000	0.470
	GJR(1,1)	0.094	0.998	0.986	0.088	0.006	0.602
(1,5)	RW	0.000	1.000	1.000	0.000	0.000	0.650
	AR(1)	0.000	1.000	1.000	0.002	0.000	0.656
	GARCH(1,1)	0.034	1.000	0.996	0.078	0.000	0.470
	GJR(1,1)	0.072	1.000	0.984	0.100	0.010	0.588

RW indicates that the null model used in the simulations is the random walk with drift, AR(1) stands for the AR(1) model, GARCH(1,1) for the AR(1)-GARCH(1,1) model and GJR(1,1) for the AR(1)-GJR(1,1) model.

6. Profitability of technical trading rules

All the results presented so far have been obtained on the SBC General Index which is not an easily replicable index and we have not included trading costs in our tests, which makes the results unrealistic. This section checks if these results can also be obtained on more easily holdable securities as individual stocks. It also analyzes the previous results by including trading costs to see if technical trading rules are really profitable.

Table 8 shows the results obtained for 5 individual stocks which are chosen from the main industrial sectors of the Swiss market. We only present results for the (1,5) rule as this is the trading rule which gives the best results on the index. All the considered securities are bearer shares. The test periods are different as some stocks were merged into other categories before the end of 1997 and some data is missing before 1980. Again, the average returns based on the trading strategies are all higher than the return of the buy-and-hold strategy. Only two of them are not statistically significantly different from the buy-and-hold return. As for the SBC Index, volatilities for buy-and-hold returns and for returns of the global strategies are very close and cannot explain the difference in average returns. The results for individual stocks are similar to those obtained for the index which shows that technical trading strategies are also operational with individual stocks.

Finally, we consider the effect of trading costs on the profits generated by the trading rules. As has been emphasized among others by SWEENEY (1988), the level of transaction costs charged to investors depends largely on the type of investor considered. Let us consider from this point of view the level of fees faced by Swiss investors. BRUAND and GIBSON-ASNER (1998) estimate that trading costs are between 0.3% and 1.6% depending on the magnitude of the order and if the investor has direct access to the market or not. Clearly, 0.3% would be the fee charged to an important financial institution which has direct access to the market and 1.6% would be the fee charged to the individual investor. Discussions with practitioners have shown that 0.3% is a very conservative figure for big institutions and that they probably face lower transaction costs. In or-

Table 8: Results of moving averages trading rule (1,5) for individual stocks

Stock	Period	Nb obs.	μ(buy & hold)	μ(strategies)	N(trades)
UBS	1.1.80–31.12.97	4320	0.000323 (0.01199)	0.001324* (0.01192)	969
ABB	1.1.80–31.12.97	4320	0.000421 (0.01599)	0.001526* (0.01599)	1011
Nestlé	1.1.80–28.05.93	3162	0.000415 (0.01184)	0.000905 (0.01181)	762
Ciba-Geigy	1.1.80–29.11.96	4049	0.000515 (0.01572)	0.000731 (0.01571)	1005
Zürich	1.1.80–30.06.95	3691	0.000249 (0.01396)	0.001152* (0.01391)	857

Nb obs. gives the number of observations. μ(buy & hold) reports the average daily return obtained with a buy-and-hold strategy. μ(strategy) gives the average daily return obtained with the strategy over the whole period. N(trades) is the number of times it is necessary to change the position according to the trading rule.* indicates that the average return is significantly different from the average buy-and-hold return at the 5% level for a two-tailed test.

der to have an idea of the impact of these costs on the previous profits we compute the resulting returns when costs are included. The average costs to be deduced from the average daily return of the strategy depends on the number of trades an investors makes. In the case of the (1,5) trading rule for the SBC index there are 1500 trades out of 7084 holding days. If the trading cost is 0.3%, the average trading costs is 0.003(1500/7084) = 0.000635. As this trading rule yields an average 0.000972 and the average buy-and-hold period is 0.000250, there remains an excess return of using this trading strategy of 0.0087% or 2.18% yearly. This means that a large investor can still earn 2.18% over the buy-and-hold strategy after deduction of fees. Table 9 provides the average excess returns from the trading strategy over the buy-and-hold return after deduction of fees. We have also computed the level of fees for which the trading strategy based on moving averages yields the same return as the buy-and-hold strategy.

Table 9: Trading rules and transaction costs

Asset-Rule	0.3%	1.6%	Max. fees
SBC-(1,200)	0.0094%	–0.0170%	0.764%
SBC-(1,100)	0.0074%	–0.0330%	0.535%
SBC-(1,50)	0.0134%	–0.0537%	0.560%
SBS-(1,30)	0.0212%	–0.0666%	0.621%
SBC-(1,10)	0.0128%	–0.1638%	0.395%
SBC-(1,5)	0.0087%	–0.2666%	0.343%
UBS-(1,5)	0.0328%	–0.0251%	0.445%
ABB-(1,5)	0.0403%	–0.0264%	0.469%
Nestlé-(1,5)	–0.0023%	–0.3366%	0.205%
Ciba-Geigy-(1,5)	–0.0053%	–0.3755%	0.088%
Zürich-(1,5)	0.0206%	–0.2812%	0.389%

The 0.3% and 1.6% columns indicate the daily average excess returns of the buy-and-hold strategy in presence of these transaction costs. The max. fees column gives the maximum amount of fees an investor can face in order to get a higher return than the buy-and-hold strategy.

Table 9 shows that for a 0.3% transaction cost, all trading rules are still profitable for the SBC index, and that some of them are also profitable for individual stocks. On the other hand, for a 1.6% trading cost, no technical trading rule yields profitable results anymore. This means that an individual investor cannot gain anything with these technical trading rules. This is confirmed by the maximum level of fees which is far below 1.6%. Does this mean that large investors can benefit from these trading rules and that markets are inefficient? An additional condition has to be fulfilled before this can be seen as true, i.e. that these investors should be able to trade these assets at the closing prices used in this study. The profitability of these simple technical trading rules depend critically on the fulfillment of these conditions.

In order to illustrate the impact of transaction costs on the profitability of the strategies we have plotted in figure 1 the evolution of the wealth invested according to a strategy based on moving averages of one and thirty days[9]. Note that the scale of the graph is logarithmic. We use this type of graph because the magnitudes of the series are very different (especially at the end of our period). This means that in a graph with arithmetic scale the 1.6% line would be difficult to distinguish from the x-axis and it would be difficult to compare the evolution of the series. Figure 1 shows the evolution of 100.– CHF invested on January 1, 1969. This would have yielded 457.17 CHF at the end of 1997 for an investor who would have invested in the SBC index (bold line). An investor using the (1,30) strategy during all the period and who would have 0.3% fees per transaction would have ended with with 2058.07 CHF at the end of 1997. On the other hand a small investor who would have followed the same strategy but with 1.6% fees per transaction would have ended with only 3.92 CHF at the end of 1997. Clearly even if the strategy seems very profitable without transaction costs, their inclusion changes the picture and shows that only investors with very low transaction costs could gain some profits from these strategies. For a small investor the level of trans-

Figure 1: Evolution of wealth according to the (1,30) rule and including transaction costs

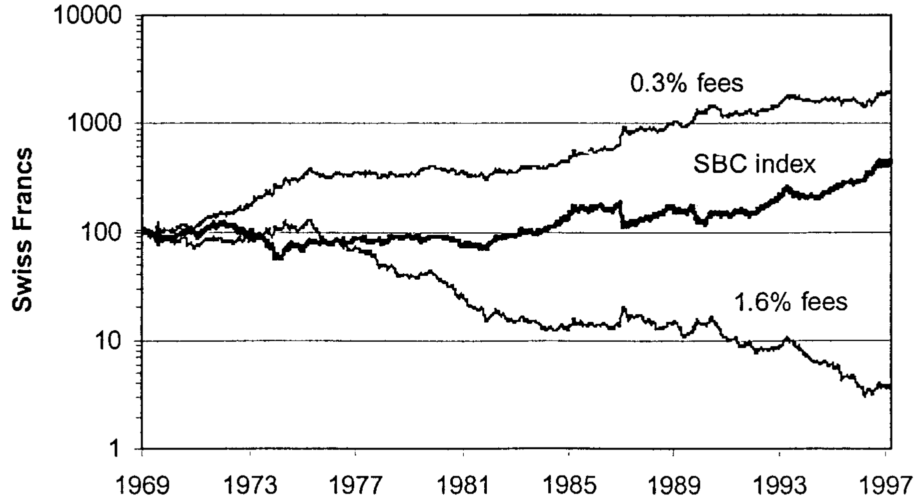

action costs he faces largely offsets the potential profits generated by the trading rule. If this type of investors is rational it will not use such investment strategies.

7. Conclusions

This paper tests if simple technical trading rules are profitable on Swiss stock prices. It considers different trading rules as simple moving averages or moving averages with bands and oscillators. Tests of the various rules are performed on daily prices of the Swiss Bank Corporation Index for the period January 1969 to December 1997. The most profitable rule appears to be a double moving average with respective windows of one and five days. This technique yields an annual average return on the SBC Index of 24.30% compared to a buy-and-hold return on the same index of 6.25%. These results are statistically significantly different from each other. The trading rules permit to identify clearly periods with positive returns from periods with negative returns. We also find that buy periods are characterized by lower volatility and sell periods are characterized by high volatility which is consistent with a leverage effect. We find that the use of bands permits to isolate even higher returns for buy and sell periods. However, this strategy does not lead to higher returns for the investor as he faces neutral periods where he must stay out of the market. We introduce and perform tests with oscillators, such as the relative strength index and the stochastic indicator, which aims at detecting trend reversals. The results show that they do not allow to improve significantly the performance obtained with simple moving averages.

As asset returns are known to present a certain number of features as non-normality, serial correlation and time-varying conditional moments, we perform bootstrap simulations to check if previous results are not due to one of these characteristics. We find that it is not the case and that an AR(1) and GARCH(1,1) component although present in the data is not responsible for the documented profits. Finally, we investigate whether the results are feasible from an investor's point of view as the SBC index is not easily replicable and does not include transaction costs. Tests are repeated on individual assets and the same kind of profitable results are obtained. When transaction costs are considered we find that the results of the trading strategies only yield profitable results for large investors who must fulfill two conditions: pay transaction costs of 0.3% or lower and trade the index at closing prices. As small investors cannot achieve such conditions, they cannot get any profits from these simple technical rules. This could also explain why the profit opportunities associated with these strategies have not disappeared as a large fraction of the market participants could not get any profit because of the presence of transaction costs. If these investors are rational, it is very likely they have used other investment strategies during this period. This also means that the hypothesis of weak form efficiency of the market cannot be rejected for a large fraction of market participants.

Footnotes

[1] This is documented in the survey by TAYLOR and ALLEN (1992).

[2] We are grateful to Karl Keiber for pointing out this problem.

[3] In order to test for this hypothesis we compare the results of trading rules on the subperiod May 1993–December 1998 where we have data on the SPI index ex-dividend with that of the SPI index (which includes dividends). We find close results for both indexes. The results are available upon request to the authors.

[4] In the literature, this type of return has not been considered yet. Usually, only the difference between the average return of buy and sell periods is investigated. Unfortunately, this figure does not indicate to the investor the potential profit of using trading rules.

[5] These results are available upon request to the authors.

[6] The interested reader can find the details and demonstrations relative to the application of the bootstrap methodology in our framework in BROCK, LAKONISHOK and LEBARON (1992). MADDALA and LI (1996) discuss more generally the application of bootstrap tests in financial models.

[7] BROCK, LAKONISHOK and LEBARON (1992) have shown that the results obtained with 500 simulations are reliable.

[8] The interested reader can find the details on GARCH models and their application to financial data in the survey of BOLLERSLEV, CHOU and KRONER (1992).

[9] As the SBC index is a price index, the evolution of wealth depicted in figure 1 does not include the effects of dividends payments.

References

ALEXANDER, S. S. (1961): "Price Movements in Speculative Markets: Trends or Random Walks?", Industrial Management Review 2, pp. 7–26.

ALEXANDER, S. S. (1964): "Price movements in speculative markets: Trends or random walks? No 2", in P. Cootner (ed.): The Random Character of Stock Market Prices, Cambridge: MIT Press.

BÉCHU, T. and E. BERTRAND (1998): L'Analyse Technique: Pratiques et Méthodes, (3rd ed.), Paris: Economica.

BLACK, F. (1976): "Studies of Stock Price Volatility Changes", Proceedings of the American Statistical Association, pp. 177–181.

BOLLERSLEV, T. and J. WOOLDRIDGE (1992): "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances", Econometric Reviews 11, pp. 143–172.

BOLLERSLEV, T., R. CHOU and K. KRONER (1992): "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence", Journal of Econometrics 52, pp. 5–59.

BROCK, W., J. LAKONISHOK and B. LEBARON (1992): "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns", Journal of Finance 47, pp. 1731–1764.

BRUAND, M. and R. GIBSON-ASNER (1998): "The Effects of Newly Listed Derivatives in a Thinly Traded Stock Market", Review of Derivatives Research, forthcoming.

DUBOIS, M. and J. M. DURINI (1995): "La Dynamique des Cours Boursiers: Une Etude Empirique sur Quatre Marchés Européens", Swiss Journal of Economics and Statistics 131, pp. 85–119.

FAMA, E. F. (1965): "The Behavior of Stock Market Prices", Journal of Business 58, pp. 34–105.

FAMA, E. F. and M. BLUME (1966): "Filter Rules and Stock Market Trading Profits", Journal of Business 39, pp. 226–241.

GLOSTEN, L., R. JAGANNATHAN and D. E. RUNKLE (1993): "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Returns on Stocks", Journal of Finance 48, pp. 1779–1801.

HUDSON, R., M. DEMPSEY and K. KEASEY (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 and Finance 20, pp. 1121–1132.

KHO, B.-C. (1996): "Time-Varying Risk Premia, Volatility, and Technical Trading Rule Profits: Evidence from Foreign Currency Futures Markets", Journal of Financial Economics 41, pp. 249–290.

LEVICH, R. M. and L. R. THOMAS (1993): "The Significance of Technical Trading Rules Profits in Foreign Exchange Markets: A Bootstrap Approach", Journal of International Money and Finance 12, pp. 451–474.

MADDALA, G. S. and H. LI (1996): "Bootstrap based tests in financial models", in G. S. Maddala and C. R. Rao (eds.): Handbook of Statistics, (Volume 14), Amsterdam: North-Holland.

NELSON, D. (1991): "Conditional Heteroskedasticity in Asset Returns: A New Approach", Econometrica 59, pp. 347–370.

PRING, M. J. (1991): "Technical Analysis Explained", (3rd ed.), New York: McGraw-Hill.

SWEENEY, R. (1988): "Some New Filter Rule Tests: Methods and Results", Journal of Financial and Quantitative Analysis 23, pp. 285–300.

SWISS BANK CORPORATION (1963): "Bulletin", no 3.

TAYLOR, M. P. and H. ALLEN (1992): "The Use of Technical Analysis in Foreign Exchange Markets", Journal of International Money and Finance 11, pp. 304–314.

WILDER, J. N. (1978): "New Concepts in Technical Trading Systems", Greensboro: Friend Research.