

Predicting a Rank Measure for Portfolio Selection

Thomas Hellström
UMINF 00.18
ISSN-0348-0542
Department of Computing Science
Umeå University
SE-901 87 Umeå, Sweden
thomash@cs.umu.se

August 18, 2000

Abstract

Investigations regarding stock price predictability and market efficiency normally focus on individual time series. In this paper we introduce a rank measure that takes into account a large number of securities and grades them according to the relative returns. It turns out that this rank measure, besides being more related to a real trading situation, is more predictable than the individual returns. The rank is predicted with a linear model and the empirical results show hit rates around 60% for the sign of the 1-day predictions. The predicted ranks are also used as a basis for a portfolio selection algorithm, which significantly outperforms the benchmark when tested on the Swedish stock market over 1993-1997. The result is in strong contradiction to the random walk hypothesis and also a strong indication that market inefficiencies exist and can be exploited with multi-stock predictions such as the presented rank approach.

Key words: Asset allocation, financial series, forecasting, portfolio management, rank, trading simulation, stock prediction

1 Introduction

The returns of individual securities are the primary targets in most researches that deal with the predictability of financial markets [12][2][11][10][8][3][4][9][1]. In this paper we focus on the observation that a real trading situation involves not only attempts to predict the individual returns for a set of interesting securities, but also a comparison and selection among the produced predictions. What an investor really wants to have is not a large number of predictions for individual returns, but rather a grading of the securities in question. Even if this can be achieved by grading the individual predictions of returns, it is not obvious that it will yield an optimal decision based on a limited amount of noisy data. Our approach in this paper is to introduce a relative rank measure for a set of securities, and instead of formulating the prediction problem for individual returns, to do so for this entity.

In Section 2 we introduce a rank measure that takes into account a large number of securities and grades them according to the relative returns. The general statistical properties of the measure are discussed and investigated empirically. In Section 3 a model for prediction of the rank measure is formulated and historical data is used to estimate the parameters in the model. The empirical results are evaluated in Section 4, both as time series predictions and by simulated trading. The astonishingly successful results are discussed and to some extent explained. Section 5 concludes the work by general conclusions and ideas for future research.

2 Defining a Rank Measure

The k -day return $R_k(t)$ for a stock m with close prices $y^m(1), \dots, y^m(t_1)$ is defined for $t \in [k+1, \dots, t_1]$ as

$$R_k^m(t) = \frac{y^m(t) - y^m(t-k)}{y^m(t-k)}. \quad (1)$$

We introduce a rank concept A_k^m , based on the k -day return R_k as follows: The k -day rank A_k^m for a stock s_m in the set $\{s_1, \dots, s_N\}$ is computed by ranking the N stocks in the order of the k -day returns R_k . The ranking orders are then normalized so the stock with the lowest R_k is ranked -0.5 and the stock with the highest R_k is ranked 0.5 . The definition of the k -day rank A_k^m for a stock m belonging to a set of stocks $\{s_1, \dots, s_N\}$, can thus be written as

$$A_k^m(t) = \frac{\#\{R_k^i(t) | R_k^m(t) \geq R_k^i(t), 1 \leq i \leq N\} - 1}{N - 1} - 0.5 \quad (2)$$

where the $\#$ function returns the number of elements in the argument set. This is an integer between 1 and N . R_k^m is the k -day returns computed for stock m . The scaling between -0.5 and $+0.5$ assigns the stock with the median value on R_k the rank 0. A positive rank A_k^m means that stock m performs better than this median stock, and a negative rank means that it performs worse. This new measure gives an indication of how each individual stock has developed relatively to the other stocks, viewed on a time scale set by the value of k .

The scaling around *zero* is convenient when defining a prediction task for the rank. It is clear that an ability to identify, at time t , a stock m , for which $A_h^m(t+h) > 0, h > 0$ means an opportunity to make profit in the same way as identifying a stock, for which $R_h(t+h) > 0$. A method that can identify stocks m and times t with a mean value of $A_h^m(t+h) > 0, h > 0$, can be used as a trading strategy that can do better than the average stock. The hit rate for the predictions can be defined as the fraction of times, for which the sign of the predicted rank $A_h^m(t+h)$ is correct. A value greater than 50% means that true predictions have been achieved. The following advantages compared to predicting returns $R_h(t+h)$ can be noticed:

1. The benchmark for predictions of ranks $A_h^m(t+h)$ performance becomes clearly defined:
 - A hit rate $> 50\%$, for the predictions of the sign of $A_h^m(t+h)$ means that we are doing better than chance. When predicting returns $R_h(t+h)$, the general positive drift in the market causes more than 50% of the returns to be > 0 , which means that it is hard to define a good benchmark.

- A positive mean value for predicted positive ranks $A_h(t+h)$ (and a negative mean value for predicted negative ranks) means that we are doing better than chance. When predicting returns $R_h(t+h)$, the general positive drift in the market causes the returns to have a mean value > 0 . Therefore, a mere positive mean return for predicted positive returns does not imply any useful predicting ability.
2. The rank values $A_k^1(t), \dots, A_k^N(t)$, for time t and a set of stocks $1, \dots, N$ are uniformly distributed between -0.5 and 0.5 provided no return values are equal. Returns R_k^m , on the other hand, are distributed with sparsely populated tails for the extreme low and high values. This makes the statistical analysis of rank predictions safer and easier than predictions of returns.
 3. The effect of global events gets automatically incorporated into the predictor variables. The analysis becomes totally focused on identifying deviations from the average stock, instead of trying to model the global economic situation.

2.1 Serial Correlation in the Ranks

We start by looking at the serial correlation for the rank variables as defined in (2). In Table 1 mean ranks $A_1^m(t+1)$ are tabulated as a function of $A_k^m(t)$ for 207 stocks from the Swedish stock market 1987-1997. Table 2 shows the “*Up fraction*”, i.e. the number of positive ranks $A_1^m(t+1)$ divided by the number of non-zero ranks. Table 3 finally shows the number of observations of $A_1^m(t+1)$ in each table entry. Each row in the tables represents one particular value on k , covering the values 1, 2, 3, 4, 5, 10, 20, 30, 50, 100. The label for each column is the mid-value of a symmetrical interval. For example, the column labeled 0.05 includes points with k -day rank $A_k^m(t)$ in the interval $[0.00, \dots, 0.10[$. The intervals for the outermost columns are open-ended on one side. Note that the stock price time series normally have 5 samples per week, i.e. $k=5$ represents one week of data and $k=20$ represents approximately one month. Example: There are 30548 observations where $-0.40 \leq A_2^m(t) < -0.30$ in the investigated data. In these observations, the 1-day ranks on the following day, $A_1^m(t+1)$, have an average value of 0.017, and an “*Up fraction*” = 52.8%.

Table 1: Mean 1-step ranks for 207 stocks

	k-day rank									
k	-0.45	-0.35	-0.25	-0.15	-0.05	0.05	0.15	0.25	0.35	0.45
1	0.067	0.017	-0.005	-0.011	-0.011	-0.004	-0.005	-0.010	-0.014	-0.033
2	0.060	0.017	0.002	-0.004	-0.010	-0.003	-0.007	-0.015	-0.017	-0.032
3	0.057	0.016	0.003	-0.005	-0.003	-0.008	-0.011	-0.011	-0.015	-0.034
4	0.054	0.018	0.003	-0.003	-0.005	-0.008	-0.011	-0.013	-0.012	-0.032
5	0.051	0.015	0.004	-0.002	-0.004	-0.009	-0.010	-0.009	-0.016	-0.032
10	0.040	0.013	0.005	-0.001	-0.003	-0.006	-0.007	-0.009	-0.012	-0.030
20	0.028	0.008	0.003	-0.003	-0.002	-0.002	-0.006	-0.011	-0.009	-0.019
30	0.021	0.007	0.002	0.004	-0.003	-0.003	-0.006	-0.006	-0.011	-0.015
50	0.014	0.005	0.000	-0.000	-0.001	-0.002	-0.005	-0.004	-0.006	-0.010
100	0.007	0.003	0.001	-0.002	-0.003	-0.004	-0.004	-0.004	-0.003	-0.008

Table 2: Fraction up/(up+down) moves (%)

	k-day rank									
k	-0.45	-0.35	-0.25	-0.15	-0.05	0.05	0.15	0.25	0.35	0.45
1	59.4	52.9	49.1	47.3	48.0	49.6	49.5	48.2	47.8	46.4
2	58.4	52.8	49.7	48.9	48.4	49.7	48.9	47.4	47.6	46.2
3	58.1	52.4	50.3	48.9	49.1	49.0	48.1	48.1	47.8	46.1
4	57.5	52.5	50.4	49.2	49.0	48.7	48.0	47.9	48.6	46.3
5	57.1	52.0	50.4	49.4	49.1	48.5	48.2	48.6	47.7	46.3
10	55.6	51.7	50.4	49.8	49.3	48.8	48.7	48.5	48.2	46.3
20	53.8	51.1	50.2	49.6	49.4	49.5	49.0	48.3	48.7	47.8
30	52.7	50.9	50.3	50.8	49.1	49.2	48.8	48.9	48.5	48.4
50	52.0	50.7	49.6	49.9	49.6	49.6	49.0	49.3	49.1	48.9
100	51.4	50.4	49.9	49.5	49.2	49.2	49.0	49.2	49.6	49.1

Table 3: Number of points

	k-day rank									
k	-0.45	-0.35	-0.25	-0.15	-0.05	0.05	0.15	0.25	0.35	0.45
1	30878	30866	31685	30837	30434	31009	31258	30539	30951	31550
2	30926	30548	31427	30481	30442	31116	31263	30435	30841	31675
3	30922	30440	31202	30404	30350	31146	31061	30449	30814	31697
4	30887	30315	31052	30320	30371	31097	31097	30328	30777	31776
5	30857	30293	30951	30275	30191	31049	31144	30254	30701	31816
10	30755	30004	30648	29958	30004	30875	30889	30155	30571	31775
20	30521	29635	30306	29591	29679	30560	30580	29836	30377	31692
30	30388	29371	30083	29388	29567	30349	30437	29652	30190	31503
50	30117	29006	29728	28979	29306	29876	30109	29236	29927	31159
100	29166	28050	28790	28011	28238	29015	29049	28254	29012	30460

The only clear patterns that can be seen in the table are a slight negative serial correlation: negative ranks are followed by more positive ranks and vice versa. To investigate whether this observation reflects a fundamental property of the process generating the data, and not only idiosyncrasies in the data, the relation between current and future ranks is also presented in graphs, in which one curve represents one year. Figure 1 shows $A_1^m(t+1)$ versus $A_1^m(t)$ in the diagram on the left. I.e.: 1-day ranks on the following day versus 1-day ranks on the current day. The same relation for 100 simulated random-walk stocks is shown in the diagram on the right for comparison.

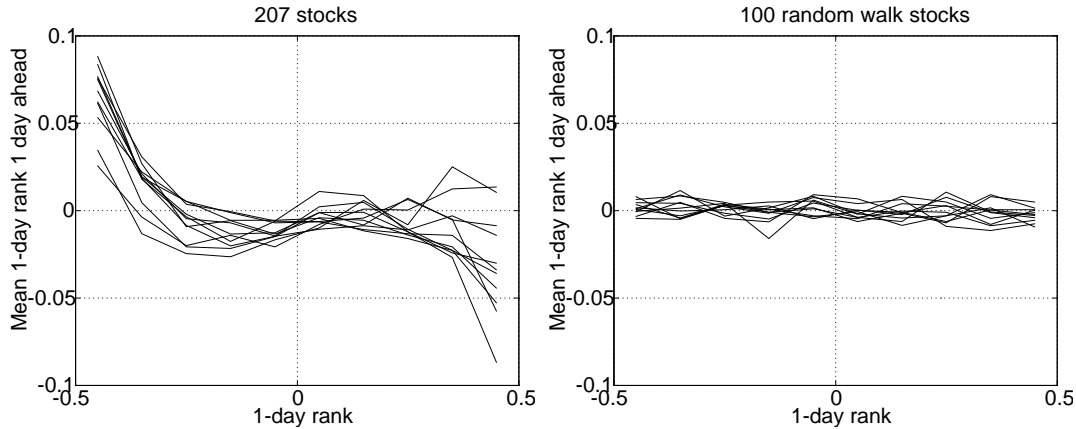


Figure 1: 1-day ranks $A_1^m(t+1)$ versus $A_1^m(t)$. Each curve represents one year between 1987 and 1997. Real data in the diagram on the left and simulated random walk in the diagram on the right.

From Figure 1 we can conclude that the rank measure exhibits a mean reverting behavior, where a strong negative rank in mean is followed by a positive rank. Furthermore, a positive rank on average is followed by a negative rank on the following day. Looking at the “*Up fraction*” in Table 2, the uncertainty in these relations is still very high. A stock m with a rank $A_1^m(t) < -0.4$ has a positive rank $A_1^m(t+1)$ the next day in no more than 59.4% of all cases. However, the general advantages described in the previous section, coupled with the shown correlation between present and future values, do make the rank variables very interesting for further investigations. In [5] the observed mean reverting behavior is exploited in a simple trading system. The rank measure in the next section is used both as input and output in a model for prediction of future ranks.

3 Predicting the Rank

For a stock m , we now attempt to predict the h -day-rank h days ahead by fitting a function g_m so that

$$\hat{A}_h^m(t+h) = g_m(A_1^m(t), A_2^m(t), A_5^m(t), A_{20}^m(t)). \quad (3)$$

The choice of function g_m could be a neural network or a simpler function. Our first attempt is a linear function, i.e. the model is:

$$\hat{A}_h^m(t+h) = p_0 + p_1 A_1^m(t) + p_2 A_2^m(t) + p_3 A_5^m(t) + p_4 A_{20}^m(t) \quad (4)$$

where the parameter vector $(p_0, p_1, p_2, p_3, p_4)$ can be determined by linear regression on historical data. For a market with N stocks, N separate models are built, each one denoted by the index m . The h -day rank $A_h^m(t+h)$ for time $t+h$ is predicted from the 1-day, 2-day, 5-day and 20-day ranks, computed at time t . To facilitate further comparison of the m produced predictions, they are ranked in a similar way as in the definition of the ranks themselves:

$$\hat{A}_h^m(t+h) \leftarrow \frac{\#\{\hat{A}_h^i(t) | \hat{A}_h^m(t) \geq \hat{A}_h^i(t), 1 \leq i \leq N\} - 1}{N - 1} - 0.5. \quad (5)$$

In this way the N predictions $\hat{A}_h^m(t+h)$, $m = 1, \dots, N$, get values uniformly distributed between -0.5 and 0.5 with the lowest prediction having the value -0.5 and the highest prediction having the value 0.5 .

3.1 Data and Experimental Set-Up.

The data that has been used in the study comes from 80 stocks on the Swedish stock market from January 1, 1989 till December 31, 1997. We have used a sliding window technique, where 1000 points are used for training and the following 100 are used for prediction. The window is then moved 100 days ahead and the procedure is repeated until end of data. The sliding window technique is a better alternative than cross validation, since data at time t and at time $t+k$, $k > 0$ is often correlated (consider for example the returns $R_5^m(t)$ and $R_5^m(t+1)$). In such a case, predicting a function value $A_1^m(t_1+1)$ using a model trained with data from time $t > t_1$ is cheating and should obviously be avoided. The sliding window approach means that a prediction $\hat{A}_h^m(t+h)$ is based on close prices $y^m(t-k), \dots, y^m(t)$. Since 1000 points are needed for the modeling, the predictions are produced for the years 1993-1997.

4 Evaluation of the Rank Predictions

The computed models g_m , $m = 1, \dots, N$ at each time step t produce N predictions of the future ranks $A_h^m(t+h)$ for the N stocks. The N predictions \hat{A}_h^m , $m = 1, \dots, N$, are uniformly distributed by transformation 5 in $[-0.5, \dots, 0.5]$. As we shall see in the following section, we can construct a successful trading system utilizing only a few of the N predictions. Furthermore, even viewed as N separate predictions, we have the freedom of rejecting predictions if they are not viewed as reliable or profitable¹. By introducing a cut-off value γ , a selection of predictions can be made. For example, $\gamma = 0.4$ means that we are only considering predictions $\hat{A}_h^m(t+h)$ s.t. $|\hat{A}_h^m(t+h)| > 0.4$.

4.1 Daily predictions

The results for 1-day predictions of 1-day ranks $\hat{A}_1^m(t+1)$ for a $\gamma = 0.0, 0.45$ and 0.49 are presented in Tables 4, 5 and 6. The corresponding 2-day predictions are presented in Tables 7, 8 and 9. Each column in the tables represents performance for one trading year with the rightmost column showing the mean values for the entire time period. The rows in the table contain the following performance measures:

¹As opposed to many other prediction and classification problems, where the performance has to be calculated as the average over the entire test data set.

1. *Hitrate*₊. The fraction of predictions $\hat{A}_h^m(t+h) > \gamma$, with correct sign. A value significantly higher than 50% means that we are able to identify higher-than-average performing stocks better than chance.
2. *Hitrate*₋. The fraction of predictions $\hat{A}_h^m(t+h) < -\gamma$, with correct sign. A value significantly higher than 50% means that we are able to identify lower-than-average performing stocks better than chance.
3. *Meanrank*₊. Mean value of the h-day ranks $A_h^m(t+h)$ for predictions $\hat{A}_h^m(t+h) > \gamma$. Should be significantly higher than 0.
4. *Meanrank*₋. Mean value of the h-day ranks $A_h^m(t+h)$ for predictions $\hat{A}_h^m(t+h) < -\gamma$. Should be significantly lower than 0.
5. *Return*₊. 100*Mean value of the h-day returns $R_h^m(t+h)$ for predictions $\hat{A}_h^m(t+h) > \gamma$. Should be significantly higher than the measure on the next line.
6. *Return*₋. 100*Mean value of the h-day returns $R_h^m(t+h)$ for predictions $\hat{A}_h^m(t+h) < -\gamma$. Should be significantly lower than the measure on the previous line.
7. *Return*_{tot}. 100*Mean value of the h-day returns $R_h^m(t+h)$ for all predictions. Compare measures 5 and 6 to this value.
8. *#Pred*₊. Number of predictions $\hat{A}_h^m(t+h) > \gamma$. Should be as high as possible to increase the statistical significance of measures 1, 3 and 5.
9. *#Pred*₋. Number of predictions $\hat{A}_h^m(t+h) < -\gamma$. Should be as high as possible to increase the statistical significance of measures 2, 4 and 6.
10. *#Pred*. Total number of predictions $\hat{A}_h^m(t+h)$. Compare measures 8 and 9 to this value.

All presented values are average values over time t and over all involved stocks m .

4.1.1 Results for the 1-day predictions $\hat{A}_1(t+1)$.

The performance for the one-day predictions are shown in the Tables 4, 5 and 6. In Table 4 with $\gamma = 0.00$, the hit rates *Hitrate*₊ and *Hitrate*₋ are not significantly different from zero and indicate low predictability. However, the difference between the mean returns for positive and negative rank predictions (*Return*₊ and *Return*₋) shows that the sign of the rank prediction really separates the returns significantly. By increasing the value for the cut-off value γ to $\gamma = 0.45$, the hit rate goes up to 61.2% for predicted positive ranks (Table 5). In Table 6 where $\gamma = 0.49$, this hit rate has increased to 63.0%. Furthermore, the difference between the mean returns for positive and negative rank predictions (*Return*₊ and *Return*₋) is impressive. Positive predictions of ranks are in average followed by a return of 0.827% while a negative rank prediction in average is followed by a return of 0.003%. This is a significant difference since the average unconditional return is 0.142%. The rows *#Pred*₊ and *#Pred*₋ show the number of selected predictions, i.e. the ones greater than γ and the ones less than γ respectively. For $\gamma = 0.49$ (Table 6) these numbers add to about 2.7% of the total number of predictions. This is normally considered insufficient when single securities are predicted, both on statistical grounds and for practical reasons (we want to trade more often than a few times per year). But since the ranking approach produces a uniformly distributed set of predictions each day (in the example 80 predictions) there is always at least one selected prediction for each day, provided $\gamma < 0.5$. Therefore, we can claim that we have a method by which, every day we can pick a stock that goes up the following day with probability 63%. This is by itself a very strong result compared to most published single security predictions [2][11][10].

4.1.2 Results for the 2-day predictions $\hat{A}_2^m(t+2)$.

The performance for the 2-day predictions is shown in the Tables 7, 8 and 9. The performance is much lower than for the 1-day case and can hardly be utilized for trading purposes. However, the difference between the mean returns for positive and negative rank predictions, *Return*₊

Table 4: 1-day predictions of 1-day ranks $|\hat{A}_1(t+1)| > 0.00$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	51.3	53.4	53.1	53.1	53.0	52.8
<i>Hitrate</i> ₋	51.9	53.6	53.1	53.4	53.1	53.0
<i>Meanrank</i> ₊	0.011	0.025	0.022	0.021	0.018	0.020
<i>Meanrank</i> ₋	-0.014	-0.025	-0.021	-0.023	-0.019	-0.020
<i>Return</i> ₊	0.396	0.101	0.139	0.253	0.190	0.217
<i>Return</i> ₋	0.247	-0.171	-0.081	0.042	-0.009	0.006
<i>Return</i> _{tot}	0.391	0.027	0.034	0.143	0.111	0.142
<i>#Pred</i> ₊	7715	8321	8311	8923	8162	41506
<i>#Pred</i> ₋	7788	8343	8342	8942	8171	41664
<i>#Pred</i>	15503	16664	16653	17865	16333	83170

Table 5: 1-day predictions of 1-day ranks $|\hat{A}_1(t+1)| > 0.45$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	56.3	63.7	61.7	64.2	59.8	61.2
<i>Hitrate</i> ₋	54.6	54.4	57.1	57.0	55.7	55.7
<i>Meanrank</i> ₊	0.051	0.092	0.077	0.095	0.065	0.077
<i>Meanrank</i> ₋	-0.028	-0.037	-0.060	-0.053	-0.038	-0.043
<i>Return</i> ₊	0.920	0.611	0.548	0.689	0.500	0.661
<i>Return</i> ₋	0.328	-0.252	-0.211	-0.080	-0.122	-0.063
<i>Return</i> _{tot}	0.391	0.027	0.034	0.143	0.111	0.142
<i>#Pred</i> ₊	846	865	870	925	861	4374
<i>#Pred</i> ₋	868	885	881	939	869	4450
<i>#Pred</i>	15503	16664	16653	17865	16333	83170

Table 6: 1-day predictions of 1-day ranks $|\hat{A}_1(t+1)| > 0.49$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	57.7	67.3	63.3	65.7	60.6	63.0
<i>Hitrate</i> ₋	51.8	57.0	56.4	57.5	55.5	55.7
<i>Meanrank</i> ₊	0.061	0.115	0.088	0.115	0.077	0.092
<i>Meanrank</i> ₋	-0.007	-0.070	-0.069	-0.046	-0.034	-0.045
<i>Return</i> ₊	1.202	0.841	0.618	0.726	0.686	0.827
<i>Return</i> ₋	0.801	-0.490	-0.332	0.073	-0.064	0.003
<i>Return</i> _{tot}	0.391	0.027	0.034	0.143	0.111	0.142
<i>#Pred</i> ₊	215	214	218	230	213	1092
<i>#Pred</i> ₋	220	221	220	233	218	1114
<i>#Pred</i>	15503	16664	16653	17865	16333	83170

Table 7: 2-day predictions of 2-day ranks $|\hat{A}_2(t+2)| > 0.00$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	50.0	52.2	52.0	52.2	51.6	51.6
<i>Hitrate</i> ₋	50.9	52.2	52.0	52.4	51.6	51.9
<i>Meanrank</i> ₊	-0.001	0.014	0.011	0.014	0.009	0.010
<i>Meanrank</i> ₋	-0.005	-0.013	-0.011	-0.015	-0.009	-0.011
<i>Return</i> ₊	0.298	0.022	0.047	0.180	0.094	0.132
<i>Return</i> ₋	0.302	-0.060	0.004	0.095	0.036	0.080
<i>Return</i> _{tot}	0.386	0.008	0.030	0.134	0.084	0.131
<i>#Pred</i> ₊	7687	8356	8263	8889	8158	41544
<i>#Pred</i> ₋	7759	8382	8315	8963	8159	41771
<i>#Pred</i>	15446	16738	16578	17852	16317	83315

Table 8: 2-day predictions of 2-day ranks $|\hat{A}_2(t+2)| > 0.45$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	54.8	56.0	53.5	52.1	51.5	53.5
<i>Hitrate</i> ₋	53.6	53.0	54.3	56.2	53.2	54.1
<i>Meanrank</i> ₊	0.024	0.032	0.024	0.020	0.006	0.020
<i>Meanrank</i> ₋	-0.019	-0.015	-0.028	-0.038	-0.022	-0.024
<i>Return</i> ₊	0.565	0.075	0.164	0.234	0.108	0.229
<i>Return</i> ₋	0.319	0.002	0.029	0.108	-0.042	0.095
<i>Return</i> _{tot}	0.386	0.008	0.030	0.134	0.084	0.131
<i>#Pred</i> ₊	831	873	858	923	868	4372
<i>#Pred</i> ₋	856	891	877	950	872	4466
<i>#Pred</i>	15446	16738	16578	17852	16317	83315

Table 9: 2-day predictions of 2-day ranks $|\hat{A}_2(t+2)| > 0.49$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hitrate</i> ₊	51.7	56.9	59.8	50.2	50.5	53.8
<i>Hitrate</i> ₋	48.8	55.2	59.5	59.2	56.4	55.7
<i>Meanrank</i> ₊	0.000	0.038	0.075	0.008	-0.003	0.023
<i>Meanrank</i> ₋	-0.000	-0.024	-0.068	-0.050	-0.031	-0.033
<i>Return</i> ₊	0.809	0.142	0.372	0.207	0.091	0.317
<i>Return</i> ₋	0.840	0.041	0.052	0.205	-0.038	0.237
<i>Return</i> _{tot}	0.386	0.008	0.030	0.134	0.084	0.131
<i>#Pred</i> ₊	203	216	214	223	214	1075
<i>#Pred</i> ₋	217	223	220	238	218	1121
<i>#Pred</i>	15446	16738	16578	17852	16317	83315

and $Return_-$, is still significant. A prediction $\hat{A}_2^m(t+2) > 0.49$ is in average followed by a return of 0.317%, while a negative rank prediction $\hat{A}_2^m(t+2) < -0.49$ in average is followed by a return of 0.237%.

4.2 Simulated Trading

In this section the produced rank predictions are used as basis for a portfolio selection algorithm. The purpose is not to develop a "real" trading system, but rather to see how the predicted ranks can be used as a component in such a system, and how the transaction costs affect the performance. We use the 1-day predictions $\hat{A}_1^m(t+1)$ with a cutoff level $\gamma = 0.49$. The trading strategy is: on every trading day sell the entire portfolio and buy the stocks with predicted rank $\hat{A}_1^m(t+1) > 0.49$. Since we have 80 stocks available, this means that we buy one single stock with the highest predicted rank. A transaction cost of 0.15% (minimum 90 Swedish crowns ~ 10 USD) is assumed for every buy or sell order. The simulation is performed in the ASTA system, which is a general-purpose tool for development of trading and prediction algorithms. A general technical overview of the system can be found in [6] and examples of usage in [5] and [7]. The rank measure and also the prediction algorithm described in Section 3 is implemented in ASTA and therefore the test procedure is very straightforward. The ASTA command window for the discussed trading strategy is shown in Figure 2. The simulated trader executes the following trading rules:

Buy rule:	<code>prank1>0.49 & nstocks==0</code>	(6)
Sell rule:	<code>prank1<=0.49</code>	

The *prank1* function returns the prediction $\hat{A}_1^m(t+1)$ as defined in Section 3. A buy signal is issued for a stock if the predicted rank is > 0.49 and is not already in the portfolio (*nstocks* returns the number of stocks already in the portfolio). A sell signal is issued if the predicted rank is *not* > 0.49 . Simply telling the trader to sell the entire portfolio would cause unnecessary transaction costs for situations, in which a stock generates buy signals on consecutive days. The annual trading profit is presented in the table in Figure 2. A graph with the same trading results is shown in Figure 3. As can be seen, the performance is outstanding, to say the least. The trading strategy outperforms the benchmark (the Swedish Generalindex) consistently and significantly every year and the mean annual profit made by the trading is an overwhelming 123.6%. The mean annual index increase during the same period is 27.4%. By studying the annual performance for the trading in Figure 2, two important properties of the trading strategy can be observed:

1. The annual profit for the trading strategy is consistently higher than for Generalindex every year 1993-1997.
2. The number of trades every year is consistently high (roughly one buy and one sell per day), which increases the statistical credibility of the results.

A summary of the trading results and also some additional information is presented in Table 10.

Number of sell orders	812
Mean profit per trade	1.2%
Trades with profit>0	456 =56.2%
Trades with profit<0	212 =26.1%
Trades with profit=0	144=17.7%
Total Increase	5266.8 %
Total increase for index	221.9 %
Annual Profit	123.6%
Annual Profit for index	27.4%
Median Excess profit	115.8%
Number of ignored buy signals	148 (15.4%)

Table 10: Results for the simulated trading in Figure 2

Of particular interest is, apart from the high profit, the number of sell orders. The total number of trading days in the investigated time period is around 1270 while the number of sell orders is 812. This is due to the fact that the same stock might get selected for buy several days in a row. These unnecessary trades are, as previously mentioned, avoided by the Sell rule. We can also notice the distribution between profitable, non-profitable, and zero-profit trades. A mere 56.2% fraction profitable trades is obviously enough to generate a very high profit. This is of course due to the fact that the fraction of non-profitable trades is less than half of this number. Since we are doing 1-day trades, the fraction of stocks being sold with zero profit is as high as 17.7%.

ASTA

Stocks	From date	To date	Transaction cost (%)	Min transaction cost	Min buy (%) per trade	Max buy (%) per trade	Initial Cash
ALL	93	97	0.15	90	0	100	100000

Buy rule: < >

Sell rule: < >

Predict:

Predictor:

Market: 80 stocks. Dates: 890102-980409 (2387 days)

Dump Trades

☐ To window

☐ To file

☐ Diagram

Buy price:

☒ Today's

☐ Tomorrow's

Sell price:

☒ Today's

☐ Tomorrow's

Load

Generate

Save

Performance:

Annual profits:							
	93	94	95	96	97	Mean	Total
Strategy profit	98.6	122.2	89.4	167.9	139.7	123.6	5266.8
Index profit	52.1	4.6	18.3	38.2	23.8	27.4	221.9
Difference profit	46.5	117.6	71.1	129.7	115.8	96.2	5044.9
Number of trades	285	312	356	364	309	325	1626

Switch to graph window for performance plots

Multiple runs	Random	Parameter	Values	Save graph	Help
Run	<input type="text" value="1"/>	<input type="text" value="0"/>	<input type="text" value="Sweep"/>	<input type="text"/>	<input type="text"/>

Figure 2: The ASTA command window for the simulated trading with rank predictions. A stock is bought if the predicted rank is > 0.49 and is not already in the portfolio. A stock is sold if the predicted rank is *not* > 0.49 (to avoid buying and selling the same stock on the same day).

4.2.1 Where did we Cheat?

The results actually look too good. Are the results an example of market inefficiency and a refutation of the Efficient Market Hypothesis (EMH)? At first one might think so² but a closer analysis gives EMH a second chance. At each time step t the rank prediction $\hat{A}_1^m(t+1)$ is produced using the close price at time t for stock m . The model is computed in a window $[t_1 - k + 1, \dots, t_1]$ backwards (k equals 1000 in the presented example) where $t_1 < t$ denotes the time for the last performed sliding window modeling. This is normally not considered peeping into the future, but is nevertheless not realistic for a real trading situation. We cannot expect to know the close prices at time t and at the same time to be able to execute buy and sell orders with these prices (since the market is already closed)³. We can simulate a different scheme where this is taken into account, and the simulated trader has to use “tomorrow’s” close prices when executing the buy and sell orders. The result from such a strategy is shown in Figure 4. As can be seen, the superior performance has deteriorated into a stable loss. However, requiring

²Many readers would probably, and rightly, at first suspect a serious programming error. “Off-by-one-errors” are but too common in implementations of financial prediction algorithms.

³However, this fact is often ignored when stock prediction algorithms that work with daily predictions are evaluated.




Figure 3: Performance for the simulated trading with portfolio management based on 1-day rank predictions as shown in Figure 2. The predictions use close data up to the day of the rebalancing of the portfolio.

the trader to buy and sell using “tomorrow’s” close prices is not very realistic either. Using “tomorrow’s” open price, or the mean of the highest and lowest prices for tomorrow, would be other more realistic choices. Of course, in a real trading environment, the use of intraday data would make it easier to compute the predictions and do the actual trading during the same trading day. Neither of these alternatives has been further examined in this paper. However, these alternatives should be examined before anything conclusive can be said about the real usability of the presented trading strategy.

The same losing strategy as in Figure 4 with the transaction costs removed is simulated in Figure 5 and shows that the trading strategy once again is better than the bench mark. However, this time again with unrealistic trading conditions.

To get a feeling of the stability of the whole test procedure, a random predictor of ranks was also implemented. Results from simulated trading with these predictions are presented in Figure 6. The performance of daily predictions with the same random ranking is presented in Table 11. As we can see, the performance is what we should expect from a random predictor of ranks. No significant hit rates nor differences between positive and negative predictions can be seen and the profit made from the random trades is quickly eaten up by the broker’s fees.

5 Conclusions

We have successfully implemented a model for prediction of a new rank measure for a set of stocks. The shown result is clearly a refutation of the Random Walk Hypothesis (RWH). Statistics for the 1-day predictions of ranks show that we are able to predict the sign of the rank consistently over the investigated 5-year-period of daily predictions. Furthermore, the mean returns that accompany the ranks show a consistent difference for positive and negative predicted ranks which, besides refuting the RWH, indicates that the rank concept could be useful




Figure 4: Simulated trading with portfolio management based on 1-day rank predictions. The predictions use close data up to the day **before** the portfolio rebalancing. The huge profit shown in Figure 3 is clearly wiped out.

for portfolio selection in real trading. The trading simulation results confirm this and show a huge excess profit for a simple trading strategy based on rank predictions. However, a closer analysis of the differences between simulation and reality points out the often-overseen fact that 1-day predictions of stock close prices are not sufficient for implementing a real trading system. The huge profit for the simple system is turned into an equally huge loss when the executed trades use “tomorrow’s” close prices. However, using “tomorrow’s” close prices is actually also a simplification of a real trading situation and makes it too hard for the prediction algorithm. The huge profit in the original case must be regarded as extremely successful and further work with evaluation and modification of the trading algorithm looks very promising. Of course, the general idea of predicting ranks instead of returns can be implemented in many other ways than the one presented in this report. Replacing the linear models with neural networks and also adding inputs to the prediction model (3) are exciting topics for future research.

6 Acknowledgement

Many thanks to Zvi Gilula for invaluable discussions and many helpful comments. Part of this work was performed during an inspiring visit to the department of Statistics at the Hebrew University in Jerusalem, Israel. Also many thanks to Hans Georg Zimmermann for helpful discussions on the subject.

References

- [1] D. F. Bassi. Stock price predictions by recurrent multilayer neural network architectures. In A.-P. Refenes, Y. Abu-Mostafa, J. Moody, and A. Welgend, editors, *Neural Networks*




Figure 5: Same trading rules as in Figure 4 but the transaction cost has been set to zero. This clearly helps the trader to survive but of course, is not done in realistic trading conditions. The trading makes on average 47% annual profit while the index makes 20%.

- in *Financial Engineering, Proc. of the 3rd Int. Conf. on Neural Networks in the Capital Markets*, Progress in Neural Processing, pages 331–340, Singapore, 1996. World Scientific.
- [2] A. N. Burgess and A. N. Refenes. The use of error feedback terms in neural network modelling of financial time series. In C. Dunis, editor, *Forecasting Financial Markets*, Financial Economics and Quantitative Analysis, pages 261–274. John Wiley & Sons, Chichester, England, 1996.
 - [3] R. Gencay. Forecast performance of moving average rules with stock returns. In A.-P. Refenes, Y. Abu-Mostafa, J. Moody, and A. Welgend, editors, *Neural Networks in Financial Engineering, Proc. of the 3rd Int. Conf. on Neural Networks in the Capital Markets*, Progress in Neural Processing, pages 215–226, Singapore, 1996. World Scientific.
 - [4] R. Gencay and T. Stengos. The predictability of stock returns with local versus global nonparametric estimators. In A.-P. Refenes, Y. Abu-Mostafa, J. Moody, and A. Welgend, editors, *Neural Networks in Financial Engineering, Proc. of the 3rd Int. Conf. on Neural Networks in the Capital Markets*, Progress in Neural Processing, pages 259–269, Singapore, 1996. World Scientific.
 - [5] T. Hellström. ASTA - a Tool for Development of Stock Prediction Algorithms. *Theory of Stochastic Processes*, 5(21)(1-2):22–31, 1999.
 - [6] T. Hellström. ASTA - User's Reference Guide. Technical Report UMINF-00.xx ISSN-0348-0542, Department of Computing Science Umeå University, Umeå Sweden, 2000.
 - [7] T. Hellström and K. Holmström. Parameter Tuning in Trading Algorithms using ASTA. In Y. S. Abu-Mostafa, B. LeBaron, A. W. Lo, and A. S. Weigend, editors, *Computational Finance - Proceedings of the Sixth International Conference, Leonard N. Stern School of Business, January 1999*, Cambridge, MA, 1999. MIT Press.




Figure 6: Simulated trading with portfolio management based on 1-day random-rank predictions. The trader almost bankrupts after a few years due to the transaction costs.

Table 11: 1-day random predictions of 1-day ranks $|\hat{A}_1(t+1)| > 0.49$.

<i>Year :</i>	93	94	95	96	97	93-97
<i>Hirate</i> ₊	48.6	53.7	47.7	49.6	50.2	48.5
<i>Hirate</i> ₋	40.9	58.4	53.1	53.5	53.8	49.7
<i>Meanrank</i> ₊	0.007	-0.008	-0.011	0.008	-0.011	-0.014
<i>Meanrank</i> ₋	0.030	-0.039	0.001	-0.009	-0.025	-0.001
<i>Return</i> ₊	0.667	0.267	0.067	0.002	0.073	0.171
<i>Return</i> ₋	-0.006	-0.062	0.157	0.087	0.097	0.035
<i>Return</i> _{tot}	0.391	0.027	0.034	0.143	0.111	0.142
<i>#Pred</i> ₊	257	257	256	254	253	1281
<i>#Pred</i> ₋	257	257	256	254	253	1281
<i>#Pred</i>	20560	20560	20480	20320	20240	102480

- [8] H. Pi and T. S. Rönkvallsson. A neural network approach to futures trading. In A.-P. Refenes, Y. Abu-Mostafa, J. Moody, and A. Welgend, editors, *Neural Networks in Financial Engineering, Proc. of the 3rd Int. Conf. on Neural Networks in the Capital Markets*, Progress in Neural Processing, pages 17–25, Singapore, 1996. World Scientific.
- [9] C. Siriopoulos, R. N. Markellos, and K. Sirlantzis. Applications of artificial neural networks in emerging financial markets. In A.-P. Refenes, Y. Abu-Mostafa, J. Moody, and A. Welgend, editors, *Neural Networks in Financial Engineering, Proc. of the 3rd Int. Conf. on Neural Networks in the Capital Markets*, Progress in Neural Processing, pages 284–302, Singapore, 1996. World Scientific.
- [10] E. Steurer. Nonlinear modelling of the DEM/USD exchange rate. In A.-P. Refenes, editor, *Neural Networks in the Capital Markets*, chapter 13, pages 199–212. John Wiley & Sons, Chichester, England, 1995.
- [11] G. Tsibouris and M. Zeidenberg. Testing the efficient markets hypothesis with gradient descent algorithms. In A.-P. Refenes, editor, *Neural Networks in the Capital Markets*, chapter 8, pages 127–136. John Wiley & Sons, Chichester, England, 1995.
- [12] H. White. Economic prediction using neural networks: The case of IBM daily stock returns. In *IEEE International Conference on Neural Networks, San Diego*, pages 451–459, San Diego, 1988.