

Chasing Trends down Wall Street

Hannu A. Kahra*

Hannu J. Schadewitz†

and

Dallas R. Blevins‡

This draft: February 8, 2005

Abstract

Theoretically superior investing performance should not be possible in the efficient asset markets. In other words, there should not exist a trading rule that is systematically beating the market. Using a wider concept, an efficient market is a market in which predictability of asset returns, after adjusting for time-varying risk-premia and transaction costs, can still exist but only 'locally in time.' But once predictable patterns are discovered by a wide group of investors, they will rapidly disappear through these investors' transactions. We propose a price-trend model as an alternative to the random walk hypothesis. We find that a trading rule based on the model and applied to the S&P 500 index from 1940 to 2005 beats the buy-and-hold strategy during the 20-year period from the early 60s to the late 70s or early 80s. Thereafter it is possible that the forecastable pattern has become widely known or the market has experienced structural changes that the model was unable to take into account. The poor performance of the trading rule during the first 20 years may be attributed to the estimation method applied.

EFM Codes: 310, 350, 320.

*Assistant Professor, Turku School of Economics and Business Administration, Turku, Finland and Monte Paschi Asset Management, Milan, Italy. E-mail address: kahra@mpsg.it.

†Professor, Turku School of Economics and Business Administration, Turku, Finland. E-mail address: hannu.schadewitz@tukkk.fi.

‡CCA, CFE, CIA, CMA, C.P.M., Professor Emeritus, University of Montevallo, Pelham, AL, USA. E-mail address: DallasBlevins@charter.net.

1 Introduction

Theoretically superior investing performance should not be possible in the efficient asset markets. In other words, there should not exist a trading rule that is systematically beating the market. Paul Samuelson [9] is probably the pioneer, who analytically argues that randomness is achieved through the active participation of many investors seeking greater wealth. The idea is summarized in the title of the article: "Proof that Properly Anticipated Prices Fluctuate Randomly." In an informationally efficient market price changes must be unforecastable if they are properly anticipated. In other words, they are unforecastable if they fully incorporate the expectations and information of all market participants. Later, the key idea was modified as "prices fully reflect all available information."

The Random Walk Hypothesis (RWH) and its close relative the Efficient Market Hypothesis (EMH) have become icons in modern financial economic. Before Samuelson, connected to the birth and development of probability theory, the RWH has remarkable intellectual forbears such as Bachelier, Einstein, Lévy, Kolmogorov and Wiener. The early papers in this area are contained in the collection of Cootner [3]. A study by Granger and Morgenstern [4] provides a detailed development and empirical examination of the random walk model and its various refinements. For practitioners, Malkiel's [8] "A Random Walk Down Wall Street" Malkiel shows why a broad portfolio of stocks selected at random will match the performance of one carefully chosen by financial experts.

The RWH and the EMH continue to fire imagination of academics and investment professionals alike. Lo and MacKinlay [7] is a classic text in the theory of finance, where using sophisticated econometric and statistical techniques, the authors show that the market is not completely random at all. Lo and MacKinlay note that their study to reject the RWH was not the first one, but earlier studies were largely ignored by the academic community and they were unknown to the authors until after their own papers were published.¹ Lo and MacKinlay note that mainstream economists are trained to study the data through filtered lenses of classical market efficiency. Many financial economist would still agree with Jensen's [6] belief that "there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Market Hypothesis."

Timmermann and Granger [12] define a market as being *efficient locally in time* with respect to information set Ω_t and the forecasting model $m_{it}(z_t, \hat{\theta})$ drawn from a set of available models M_t , if

$$E \left[f_t \left\{ R_{t+1}^*, m_{it}(z_t; \hat{\theta}_t), \mathbf{c}_t \right\} \right] = 0, \quad (1)$$

¹Lo and MacKinlay [7] refer a few earlier studies in their book "A Non-Random Walk Down Wall Street."

where f_t embody the set of possible transactions at time t , R_{t+1}^* is the risk-adjusted return, $\hat{\theta}$ is the vector of parameters estimated using data up to time t , $\mathbf{z}_t \in \Omega$, and \mathbf{c}_t is the vector of transaction cost parameters. $E[\cdot|\Omega_t]$ is the mathematical expectation operator, or population expectation, conditional on the information set Ω_t . Timmerman and Granger intend 'model' to be interpreted in the broadest sense to incorporate both the functional form, prediction variables, estimation method and choice of sample period. The latter can be, e.g. an expanding window, rolling window or based on exponential discounting.

Timmermann and Granger allow that some models have predictive power before their discovery. This does not, however, violate the EMH, since such models would not be elements in the relevant set of available models, M_t . This means that the definition of the EMH does not rule out profits from new forecasting techniques. These techniques may have a 'honeymoon' period before they become widely known and they cease to generate profits. Thus, individual forecasting models are likely to go through stages of success, declining value, and finally disappear.

Three forms of market efficiency are commonly defined in terms of variables contained in the information set Ω_t . First, when Ω_t only comprises past and current asset prices, the EMH in its weak form is being tested. Second, when Ω_t is expanded to include all publicly available information gives rise to the EMH in its semi-strong form. Third, if all public and private information is included in Ω_t , market efficiency in the strong form is being tested. A private model is like private information. Most of the tests of EMH are designed to rule out private information that is hard to measure and perhaps also more expensive to acquire.

Timmermann and Granger argue (p. 16) that forecasting experiments testing the EMH have to specify at least five factors:

1. the set of forecasting models available at any given point in time, including estimation methods;
2. the search technology used to select the best, or a combination of best, forecasting model(s);
3. the available 'real time' information set, including public versus private information and ideally the cost of acquiring such information;
4. an economic model for the risk premium reflecting economic agents' trade-off between current and future payoffs;
5. the size of transaction costs and the available trading technologies and any restrictions on holdings of the asset in question.

In this framework, Timmermann and Granger define that an efficient market is thus a market in which predictability of asset returns, after ad-

justing for time-varying risk-premia and transaction costs, can still exist but only 'locally in time'. This means that once predictable patterns are discovered by a wide group of investors, they will rapidly disappear through these investors' transactions.

Timmermann and Granger argue that consideration needs to be turned to quickly changing models that can detect and utilize any instances of temporary forecastability that might arise and quickly disappear as learning opportunities arise and close down.

Any model is, however, only an approximation to the rules which convert relevant information and numerous beliefs and actions into markets. Taylor [11] (p. 14) argues that plausible models should satisfy five criteria. First, models should be consistent with past prices. Second, the model should be feasible. In other words, hypotheses implied by a model ought to be amenable to rigorous testing. Third, model should be as simple as possible: we prefer parsimonious models. Fourth, a model should provide forecasts of future returns and prices, which are statistically optimal assuming the model is correct. Fifth, it is beneficial if a model can be used to aid rational decision making.

One traditional way to test the RWH is to examine the autocorrelation properties of price changes. A more general perspective has been to view the price process as a particular model within the class of ARIMA models popularized in the 70s. Martingale processes lead naturally on to non-linear stochastic processes that are capable of modeling higher order conditional moments, such as the ARCH model introduced in the 80s. At least these popular models have been in the set of available models for a few decades.

When forecasters constantly search for predictable patterns, their behavior affects market prices when they attempt to exploit trading opportunities in a large scale. These patterns are unlikely to persist for long periods when discovered by a large number of investors. This gives rise to nonstationarities in the time series of asset returns.

Bossaerts [2] (p. 43) distills the following two components to the traditional concept of EMH. Firstly, market beliefs are correct: *ex ante* expectations coincide with true expectations, and *ex ante* covariances correspond to true covariances. Secondly, return distributions are time-invariant, i.e., stationary. The latter implies that the law of large numbers holds such that the sample moments estimate the population moments, if they exist. Tests to verify asset-pricing theory require even stronger assumptions that the data exhibit the right 'mixing' conditions: the memory cannot be too long, so that central limit theorems obtain.

In this paper we propose a price-trend model that is assumed to offer the best chance to find improvements upon random walk forecasts. Trends will only occur if some information is reflected in several consecutive returns. It is assumed that each item of information is reflected either quickly or slowly into prices.

The existence of trends can be explained by the theories of behavioral finance. Trends will occur if information is used imperfectly. This can occur if enough people are irrational or rational but unable to interpret all information quickly and correctly. If trends truly exist, trends and trend reversals may occur at all time scales. We examine if it is possible to find predictable structures in daily data.

The paper is organized as follows. Section 2 describes the price-trend hypothesis. In Section 3 the model is estimated. The model is applied to the daily returns of the S&P 500 index. In Section 4 a simple trading rule is applied to test the hypothesis. Section 5 concludes.

2 The price-trend hypothesis

The price-trend hypothesis is an alternative to the RWH. The professional analysts commonly operate in the belief that there exist certain trend generating facts that will guide a speculator to profit if only the facts can be interpreted correctly. The two main schools of professional analysts; the 'fundamentalists,' represented by financial economists, and the 'technicians,' represented by 'financial alchemists;' agree on the basic assumption but differ in the methods used to gain knowledge.

Both schools of analysts assume the existence of trends that represent the gradual recognition by the agents of emergent factual situations. If trends exist, they must reflect a lagged response of the market prices to the underlying factors governing the price process. Thus, the existence of trends imply that prices do not adjust fully and instantaneously when new information becomes available. Instead, some new information is incorporated only gradually into prices.

In recent years, research in behavioral finance has shed some important light on the causes and implications of the presence of trends. Firstly, trends will occur if information is used imperfectly, for example if enough investors are irrational. A feed back loop is a typical case. In the most popular version of the theory, feedback takes place because past price changes generate expectations of further price changes. This explanation relies on adaptive expectations.

Andreassen and Krauss [1] report psychological experiments, in which subjects are shown real stock prices from the past and asked to forecast subsequent changes. It was found that subjects track the past averages, when the stock prices are stable. However, as prices began to show consistent trends, the subjects began to switch to a trend-chasing strategy. 'Technical analysis' is even more compelling evidence of the presence of trend-chasing. These techniques try to spot trends and trend reversals by using technical indicators associated with past price movements.

Another version of the theory states that feedback takes place because of

increased (decreased) investor confidence in response to past price increases (decreases).

Secondly, the investors may well be rational but unable to interpret all information quickly and correctly.

Thirdly, there are limits of arbitrage dedicated to exploiting noise traders' misperceptions. Arbitrageurs are likely to be risk averse and have reasonable short horizons. As a result their willingness to take positions against noise traders is limited. The power of arbitrage is limited by fundamental risk. Another source of risk borne by short-horizon investors engaged in arbitrage against noise traders is the risk that noise traders' beliefs will not revert to their mean for a long time and might in the meantime become even more extreme. An arbitrageur faces the risk that he has to liquidate before the price change occurs. Fear of this loss tends to limit the original arbitrage position.

Slow interpretation of information will cause consecutive returns to be partially determined by the same information. This causes positive autocorrelation such that several returns are all influenced in the same way towards a conditional mean. The impact of the current information, however, diminishes as time goes on. Hence, the autocorrelations should decrease as the lag length increase.

The alternative hypothesis requires a theoretical autocorrelation function that is consistent with the estimated sample autocorrelations. Taylor [10] defines the autocorrelation function of the price-trend hypothesis as

$$H_1 : \rho_\tau = Ap^\tau, \quad A, p, \tau > 0. \quad (2)$$

Here H_1 refers to the alternative hypothesis to the random walk hypothesis H_0 . Parameter A measures the proportion of information not reflected by prices within one day, while parameter p measures the speed at which imperfectly reflected information is incorporated into prices. The limiting cases are $A \rightarrow 0$ or $p \rightarrow 0$, when information is used perfectly. The mean trend duration is defined by $m = 1/(1 - p)$.

We follow Taylor [11] and assume that returns X_t have stationary mean μ such that

$$X_t = \mu + (\mu_t - \mu) + e_t. \quad (3)$$

In the return process e_t is the response to quickly reflected information and $\mu_t - \mu$ is the response to slowly reflected information. By definition e_t is a zero mean uncorrelated irregular component, the error or disturbance. The μ_t have mean μ and they are autocorrelated such that $\{\mu_t\}$ is an AR(1) process

$$\mu_t - \mu = p(\mu_{t-1} - \mu) + \zeta_t \quad (4)$$

having autocorrelations p^τ . The error term ζ_t represents the effect of all the slowly reflected news first available on day t .

The return process $\{X_t\}$ is a state-price process such that (3) is the measurement equation with an unobserved state variable μ_t that has a transition equation given by (4). We may interpret that the smart money forms optimal forecasts of the future price on the basis of the transition equation (4). Assuming $\{e_t\}$ is stationary, the returns have autocorrelations given by (2).

The state μ_s and the e_t are uncorrelated. Hence, noise traders may respond to e_t that is uncorrelated with the unobserved state μ_s , but not necessarily independent, for all s and t . Using the standard formula for uncorrelated processes $V(X_t) = V(\mu_t) + V(e_t)$ and since $\{e_t\}$ is stationary, we obtain for A in (2): $A = V(\mu_t)/V(X_t)$. The latter is a measure of the proportion of slowly reflected information. In the trend hypothesis it is assumed that slowly reflected information has influence to several returns in the same way.

Financial returns often exhibit time-varying conditional variance. We define the standardized return by

$$U_t = (X_t - \mu)/\sqrt{V_t}, \quad (5)$$

that have mean zero and variance equal to one.

To obtain optimal linear forecasts for returns, we have to estimate the trend parameters A and p . An ARMA(1,1) process is a parsimonious model that has autocorrelations Ap^τ . Assuming $\{\xi_t\}$ is white noise and $\{Y_t\}$ can be defined by

$$Y_t - pY_t = \xi_t - q\xi_{t-1}, \quad (6)$$

then the corresponding autocorrelations are Ap^τ such that

$$A = (p - q)(1 - pq)/[p(1 - 2pq + q^2)]. \quad (7)$$

Thus, in our case the return process $\{X_t\}$ is described by the ARMA(1,1) model

$$X_t - \mu = p(X_{t-1} - \mu) + \xi_t - q\xi_{t-1}. \quad (8)$$

The optimal linear forecast is given by

$$F_{t,1} = \mu + (p - q) \sum_{i=0}^{\infty} q^i (X_{t-i} - \mu) \quad (9)$$

$$= \mu + (p - q)(X_t - \mu) + q(F_{t-1,1} - \mu). \quad (10)$$

The variance of the forecast error is

$$V(\xi_{t+1}) = MSE(F_{t,1}) = (1 - p^2)\sigma^2/(1 - 2pq + q^2) \quad (11)$$

and thus the variance of the forecast is

$$V(F_{t,1}) = (p - q)^2\sigma^2/(1 - 2pq + q^2) \quad (12)$$

$$= Ap(p - q)\sigma^2/(1 - pq). \quad (13)$$

The random walk forecast for X_t is μ and it has MSE equal to σ^2 . The proportional reduction in MSE obtained using the optimal linear forecast is given by

$$W = \frac{\sigma^2 - MSE(F_{t,1})}{\sigma^2} = \frac{V(F_{t,1})}{\sigma^2} = \frac{Ap(p-q)}{1-pq}. \quad (14)$$

Once the parameters are estimated an appropriate forecast of x_{t+1} is defined by

$$f_{t,1} = \left(\sqrt{\hat{v}_{t+1}/\sqrt{\hat{v}_t}} \right) [(\hat{p} - \hat{q})x_t + \hat{q}f_{t-1,1}], \quad (15)$$

where \hat{v}_{t+1} is the forecast of the conditional variance.

Using the properties of the variance of the forecast $F_{t,1}$

$$V(F_{t,1}) = Cov(F_{t,1}, X_{t+1}) = Cov(F_{t,1}, \mu_{t+1}), \quad (16)$$

we obtain the correlation between the forecast $F_{t,1}$ and the unobservable trend component μ_{t+1} that is denoted by λ :

$$\lambda = \sqrt{V(F_{t,1})/V(\mu_{t+1})} \quad (17)$$

$$= \sqrt{[p(p-q)/(1-pq)]}. \quad (18)$$

The optimal linear forecasts are unbiased. Using standard statistical calculus, if the forecast $f_{t,1}$ is k standard deviations from its mean μ , the expected return over the next period is $\mu + k\sigma_F$, i.e. the expected change in the price logarithm over h periods is

$$h\mu + k\sigma_F(1-p^h)/(1-p), \quad (19)$$

where

$$\sigma_F = \lambda\sigma_\mu = \lambda\sigma_X\sqrt{A}. \quad (20)$$

Our main interest is to forecast the sign of the price change correctly. Assuming that μ_{t+1} and $F_{t,1}$ has a bivariate normal distribution

$$\mu_{t+1}|f_{t,1} \sim N(f_{t,1}, (1-\lambda^2)\sigma_\mu^2) \quad (21)$$

we can estimate the probability of correctly forecasting the direction: up or down.

Assuming $\mu = 0$ and that the forecast $f_{t,1}$ is a positive number $k\sigma_F$, the estimated probability that the trend is positive can be obtained by computing

$$P \left\{ N \left[k\lambda\sigma_\mu, (1-\lambda^2)\sigma_\mu^2 \right] > 0 \right\} = \Omega \left(k\lambda/\sqrt{1-\lambda^2} \right), \quad (22)$$

where $\Omega(\cdot)$ is the cumulative distribution of the standard normal distribution.

3 Data and Estimation

We apply the model to the daily observations of Standard and Poor's S&P 500 index. The data set is from December 30, 1927 to December 31, 2005. The period ending December 29, 1939 is used for the first estimation period, having 3024 observations, approximately a 12 year's history. Thereafter an observation is added one by one and the model is re-estimated. In other words, we apply an expanding window. Thus, the information set grows by one observation on each trading day.

In the case of a linear returns process the parameters can be estimated by maximum likelihood. This is not, however, possible for non-linear processes. The parameters A and p can be estimated by matching theoretical autocorrelations Ap^τ and estimated autocorrelations r_τ :

$$\min_{\{A,p\}} F(A,p) = n \sum_{\tau=1}^K (r_\tau - Ap^\tau)^2, \quad (23)$$

where n is the number of returns and K is the number of estimated autocorrelations. Given \hat{A} and \hat{p} , \hat{q} is the solution of

$$q^2 - q[1 + (1 - 2A)p^2]/[(1 - A)p] + 1 = 0. \quad (24)$$

The model is fitted to rescaled returns, obtained by dividing the log return by its estimated conditional standard deviation. It is obtained by fitting a GARCH(1,1) model to the return series before estimating the parameters of the return process. The lag length is set $K = 50$. The object function (23) is optimized for each value of $m = 1/(1 - p)$, $m = 1, 2, \dots, 40$; and the estimates (\hat{A}, \hat{p}) are selected that correspond to the minimum value of the object function (23). Thereafter \hat{q} is obtained by solving (24). Figure 1 displays the estimation results.

[Figure 1 about here]

Both \hat{p} and \hat{q} exhibit similar behavior. Both parameter estimates are close to 0.9 till the mid 80s, in fact till the year 1984. Thereafter, the parameter estimates experience a sudden jump down, followed by a jump up and back down, again. The behavior of \hat{A} is their mirror image. The average trend duration decreases gradually from 24 days in the early 40s into 11 days in the beginning of 1984, jumping into only two days in the middle of 1984, the minimum accepted value for the trend duration.

The hit ratio displays the relative frequency of forecasting next day's direction, up or down, correctly. The time t hit ratio refers to the parameter estimates \hat{A}, \hat{p} and \hat{q} estimated at time t , using history from the end of 1939 up to t , which are used to calculate the *ex post* performance since the beginning of the data. Interestingly, the hit ratio experiences a sudden jump

upwards that coincides with the jumps in the parameter estimates. After the jump, the hit ratio starts to decrease gradually. Before the jump, the hit ratio varies between 57 and 60 percent, but when the estimated average trend duration drops to two days, the corresponding hit ratio jumps to 72 percent.

We find two plausible explanations to these observations. Firstly, the model may well be correct, but the estimation method can be unstable and it may provide biased parameter estimates. Since standard econometric theory assumes asymptotic efficiency, it is surprising that one single observation can cause a sudden jump in the parameter estimates. In this case a better estimation method should be applied.

Secondly, it is very likely that the underlying return process experiences structural changes as time goes by. This may occur through learning. The agents may well be 'rational', but unlike under the rational expectations hypothesis, they do not know the true parameters, e.g. of the equations for dividends or returns, and have to estimate them from limited data. If this is the case, agents optimally update their estimates of the true population parameters as more data arrives. The agents may apply a rational valuation formula and produce weakly rational forecasts based on the limited information set. Thus, the agents' price estimates are obtained under learning. In this case the model should be amended to take into account the structural changes in the parameters.

Questions of model stability relate to the properties of a process with a specific generating mechanism after it has been running for a long period. Granger and Teräsvirta [5] examine various forms of process instability. For example the underlying probability distribution $F_t(x)$ can be explosive in mean or variance, or in higher order moments; or it can be periodic, such that for example

$$F_{t+p}(x) = F_t(x) \neq F_{t+j}(x), \quad j = 1, \dots, p - 1.$$

In this example the mean and/or some of the moments will be periodic, and the process is said to contain a limit cycle. Other plausible explanations are state-dependent or regime-switching behavior and various forms of long-memory dependence.

In order to have a second look at the data, we run an AR(1) model with an intercept on the data using a moving 252-day window. Figure 2 displays the behavior of the parameter estimates. The solid lines correspond to the original return data and the dotted line corresponds to the standardized returns.

[Figure 2 about here]

The estimates $\hat{\phi}$ and $\hat{\mu}$ correspond to the AR(1) and the constant term, respectively. Both parameter estimates exhibit unstable behavior suggesting

that the risk premium may well be time-varying. Thus, either the quantity of risk or the price of risk, or both, are not constant.

4 A Test based on a trading rule

The price-trend model can be used to construct a fictitious trading rule. The behavior of the hit ratio suggests that a profitable trading rule can be found. A market is said to use information efficiently if there is no way to use the information to increase expected wealth by frequent trading. The efficient market hypothesis is said to be true if the risk adjusted return, net of all costs, from the best trading rule is not more than the comparable figure when assets are traded infrequently, as stated by Jensen [6]. We do not, however, consider the costs and the risks, but instead, use the trading rule to examine, if the series exhibits a predictable pattern.

Using the procedure described in the previous Section, on each day, starting on December 29, 1939; we estimate the model and provide a forecast (15) for the next day's return. The trading decision depends on a standardized forecast calculated by assuming the non-linear trend model is correct, computed as

$$k = f_{t,1}/\hat{\sigma}_F \quad (25)$$

$$\hat{\sigma}_F = \sqrt{\hat{v}_t} [Ap(p-q)/(1-pq)]^{1/2}. \quad (26)$$

The value of k (25) is inserted into (22) and a value for the probability of an up direction is obtained. If the value exceeds 0.5 a long position is opened. If the value is less than 0.5, a short position is opened. It is assumed that the trades can be transacted at published closing prices. The trades are closed at the following day's closing prices.²

Figure 3 displays the time series of cumulative S&P 500 log returns, cumulative log returns of the trading rule and their difference. The cumulative log return, following the trading rule, is 546.8 percent; while the buy-and-hold strategy yields a 458.2 percent cumulative log return. However, when looking at their difference, it can be noticed that the trading rule underperforms the market during the 20-year period from 1940 to the early 60s. Thereafter the trading rule beats the buy-and-hold strategy during the next 20-year period, ending in the late 70s or early 80s. Thereafter the trading rule loses the edge and it does not work any better than the benchmark random walk model. During the 1982-1999 bull market the trading rule even underperforms the buy-and-hold strategy.

[Figure 3 about here]

²The trade can be transformed into a binary option, betting the sign of the market move. For example London based spread betting companies quote S&P 500 binary bets. The bid-ask spread is their transaction cost.

Figure 3 carries the same message as Figures 1 and 2. Either the market experiences structural changes in the late 70s/early 80s so that it can not be explained by a constant parameter model and/or the agents have learned the process. The poor performance of the trading rule during the first 20 years can be attributed to the poor performance of the estimation algorithm. When the parameter estimates exhibit sudden jumps in the early 80s, the *ex post* hit ratio experiences a large jump upwards. The estimated values of p , q were far too high and the value of A was far too low to fit the data properly.

Interestingly, Lo and MacKinlay [7] note that in a recent update of their original variance ratio test for random walk for weekly US stock market indexes, they discover that the most current data (1986-1996) conforms more closely to the random walk than their original 1962-1985 sample period. They also report that several investment houses, most notably Morgan Stanley and D.E. Shaw, have been engaged in high-frequency equity trading strategies based on 'financial engineering.' These strategies were specifically designed to take advantage of the patterns the authors uncovered in 1988.

Statistical arbitrage strategies have fared reasonably well until recently, when it is regarded as a very competitive and thin-margin business conducted by hedge funds. The conclusion obtained by Lo and MacKinlay and our findings provide a plausible explanation for the trend towards randomness in recent data.

Samuelson [9] argues that randomness is achieved through the active participation of many investors seeking greater wealth. It is also achieved through and active search for new models. The discovery of a new forecasting model that enjoys its 'honeymoon' is not, however, inconsistent with the practical version of the EMH. Market opportunities need not mean market inefficiencies.

5 Conclusions

A price-trend model is a natural alternative hypothesis for random walk. We find that a trading rule based on the price-trend model beats the buy-and-hold strategy during the 20-year period from 1940 to the early 70s or late 80s. Thereafter it is possible that the predictive pattern has been discovered by a wide group of investors, such that it has disappeared through investors' transactions. The alternative is that the pattern still exists, but it does not stay constant. Thus, forecasting techniques may exist even if the EMH is correct. The process may well experience structural changes that should be modeled by a varying parameter model that takes into account regime shifts and state dependencies. It is quite obvious that standard constant-parameter models are not up to the task since simple specifications,

such as the ARMA models, assume stationarity. There is also strong evidence that the estimation method used is unstable. Advances in simulation based estimation, especially in the Monte Carlo Markov Chain (MCMC) techniques may provide a better solution to the estimation problem. It can be applied in estimating the parameters of the autocorrelation function or directly the parameters of the underlying, more realistic state-space model. While stable forecasting patterns are unlikely to persist for long periods of time and they will be self-destructing when discovered by a large number of investors, it may pay off to search for these patterns. This search needs to be turned to quickly changing models that can detect and utilize any instances of temporary forecastability. However, this arising property is expected to be of short duration and quickly disappear as learning opportunities arise and close down.

References

- [1] P. B. Andreassen and S. J. Krauss, *Judgemental extrapolation and the salience of change*, Journal of Forecasting (1990), no. 9, 347–372.
- [2] P. Bossaert, *The paradox of asset pricing*, Princeton University Press, Princeton, 2002.
- [3] P. Cootner (ed.), *The random character of stock market prices*, MIT Press, Cambridge, 1964.
- [4] C. W. J. Granger and O. Morgenstern, *Predictability of stock market prices*, Lexington, Heath, 1970.
- [5] C. W. J. Granger and T. Tervirta, *Modelling nonlinear economic relationships*, Oxford University Press, New York, 1997.
- [6] M. Jensen, *Some anomalous evidence regarding market efficiency*, Journal of Financial Economics (1978), no. 6, 95–102.
- [7] A. W. Lo and A. C. MacKinlay, *A non-random walk down wall street*, Princeton University Press, Princeton, 1999.
- [8] B. G. Malkiel, *A random walk down wall street*, W.W. Norton & Company, New York, 1999.
- [9] P. Samuelson, *Proof that properly anticipated prices fluctuate randomly*, Industrial Management Review (1965), no. 6, 41–49.
- [10] S. Taylor, *Conjectured models for trends in financial prices - tests and forecasts*, Journal of Financial and Quantitative Analysis (1980), no. 17, 37–61.

- [11] ———, *Modelling financial time series*, John Wiley & Sons, Chichester, 1986.
- [12] A. Timmermann and C. W. J. Granger, *Efficient market hypothesis and forecasting*, International Journal of Forecasting (2004), no. 20, 15–27.

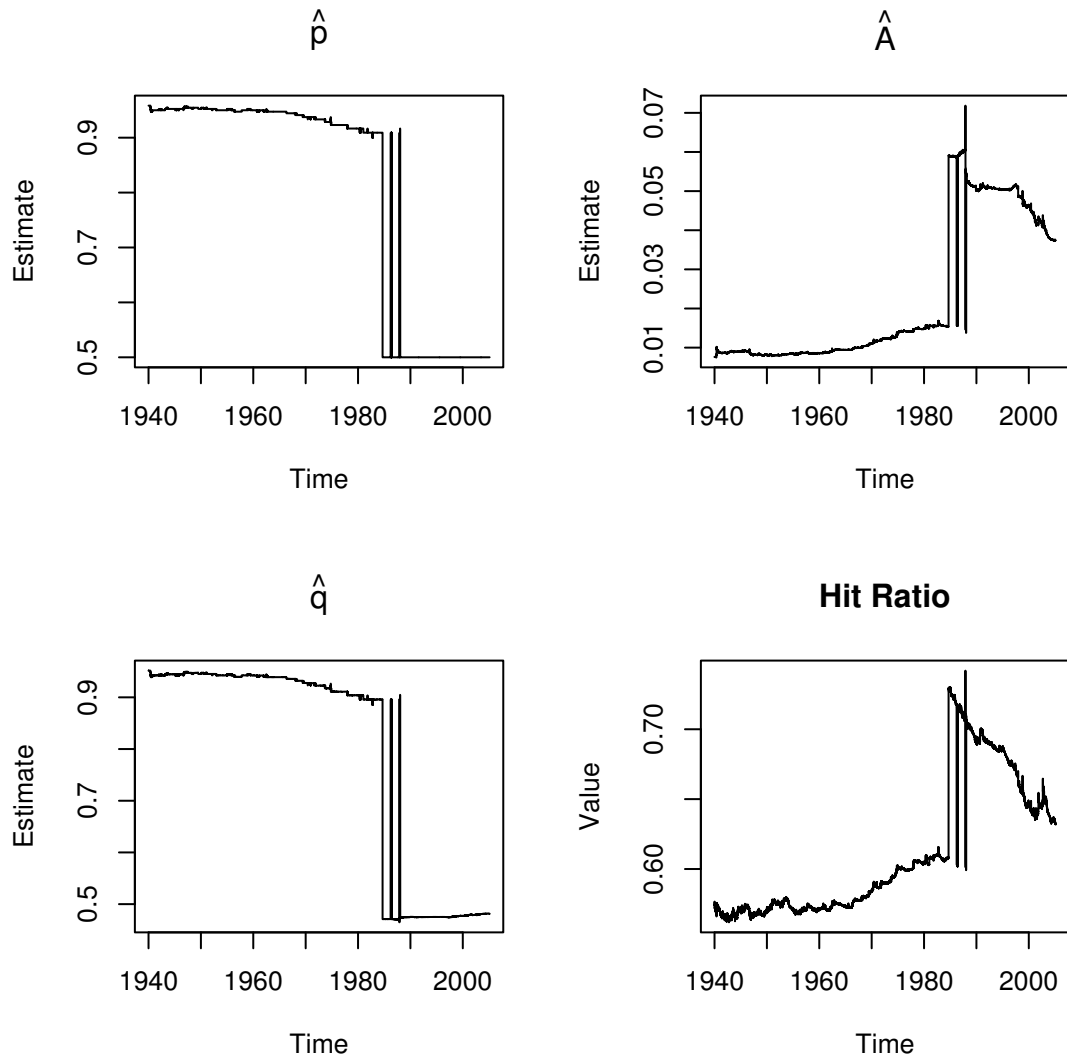


Figure 1: Parameter estimates and the hit ratio

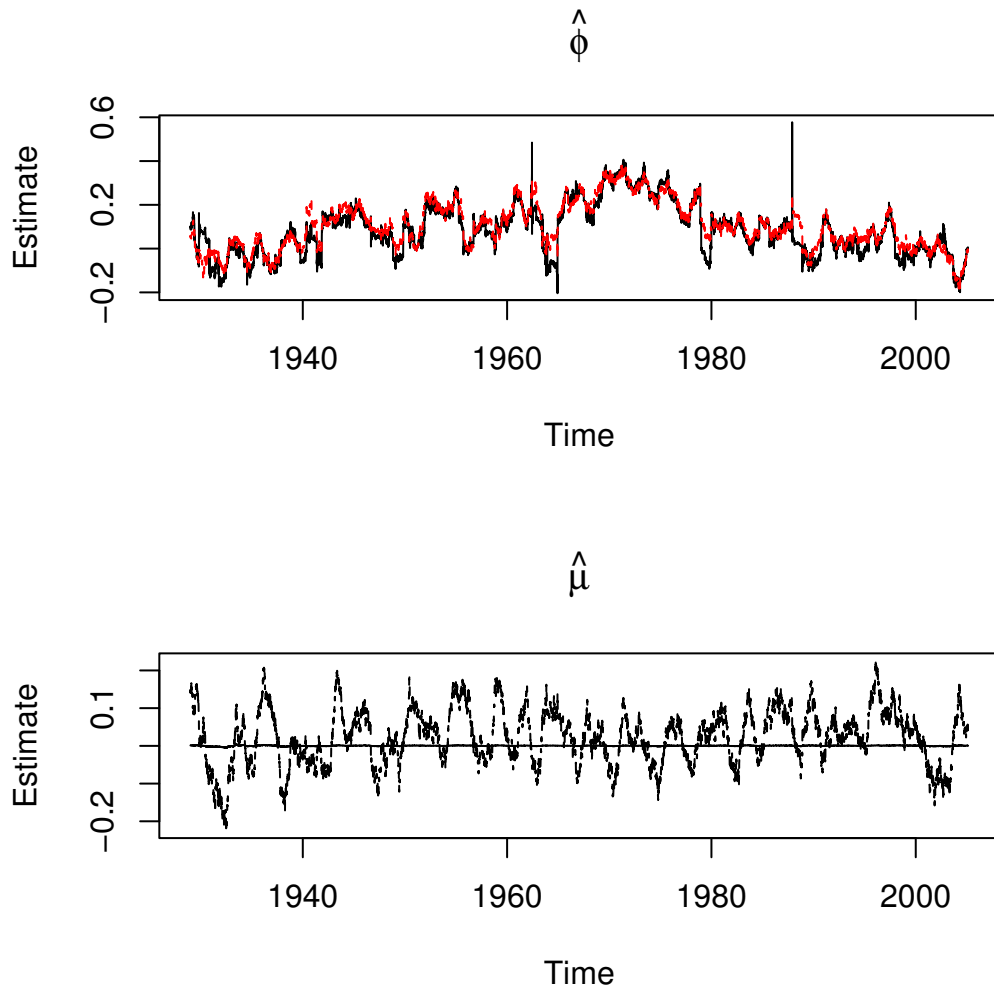


Figure 2: Parameter estimates in a moving 252-day window

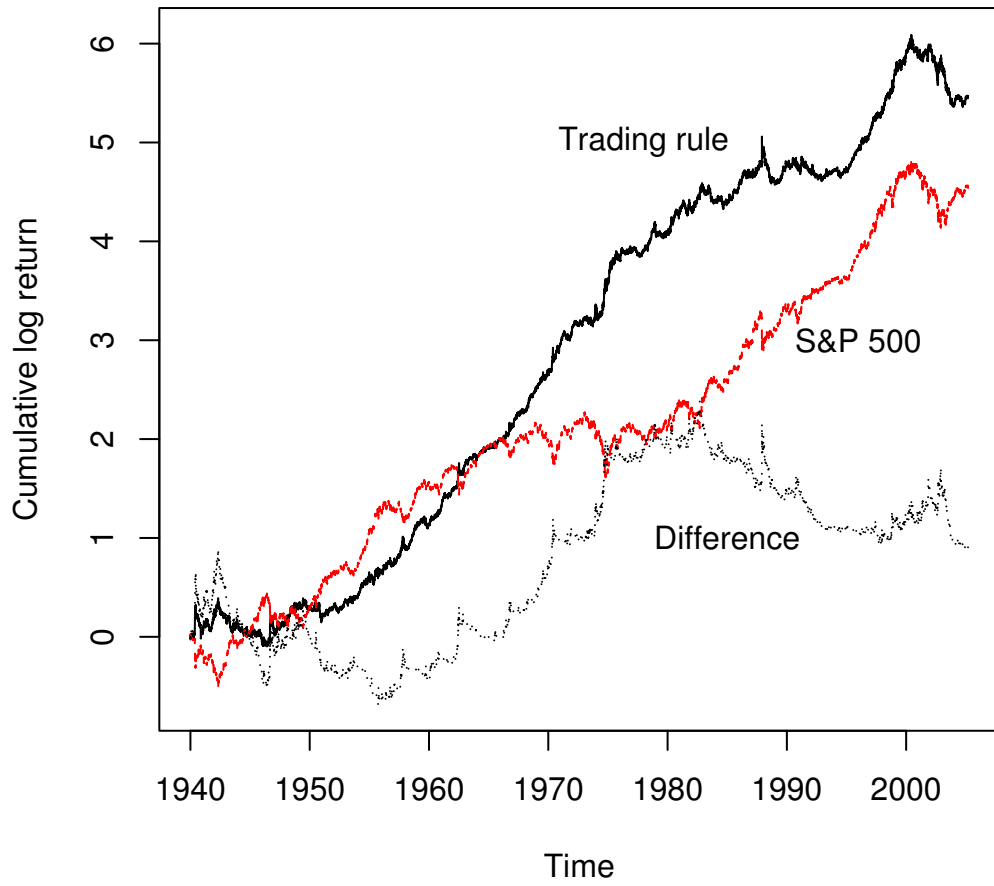


Figure 3: Trading rule performance