On Testing the Random-Walk Hypothesis:

A Model-Comparison Approach

Ali F. Darrat*
Department of Economics and Finance
Louisiana Tech University
Ruston, Louisiana 71272
darrat@cab.latech.edu

and

Maosen Zhong*
Department of Business Administration
University of Texas - Brownsville
Brownsville, TX 78520
mzhong@utb1.utb.edu

Abstract

The main intention of this paper is to investigate, with new daily data, whether prices in the two Chinese stock exchanges (Shanghai and Shenzhen) follow a random-walk process as required by market efficiency. We use two different approaches, the standard variance-ratio test of Lo and MacKinlay (1988) and a model-comparison test that compares the *ex post* forecasts from a NAÏVE model with those obtained from several alternative models (ARIMA, GARCH and Artificial Neural Network-ANN). To evaluate *ex post* forecasts, we utilize several procedures including RMSE, MAE, Theil's U, and encompassing tests. In contrast to the variance-ratio test, results from the model-comparison approach are quite decisive in rejecting the random-walk hypothesis in both Chinese stock markets. Moreover, our results provide strong support for the ANN as a potentially useful device for predicting stock prices in emerging markets.

*We wish to thank two anonymous referees for many helpful comments and suggestions. The usual disclaimer applies.

I. Introduction

By any standard, the Chinese economy has achieved a remarkable and, thus far, sustainable growth since the 1980s. To further enhance economic efficiency, China became the first Communist country in the early 1990s to establish a stock market and build a "Socialist Wall Street" to promote market economy.

China established two official stock exchanges: the Shanghai Exchange in December 1990, and the Shenzhen Exchange in July 1991. Both markets have enormously grown in terms of the number of companies listed, number of traders, and market capitalization.[1] The Chinese markets trade two main classes of stocks. Class A shares (denominated in local currency) are only available to Chinese nationals, and Class B shares (traded in U.S. or Hong Kong dollars) are only available to foreign investors. The A shares market is much larger than the B shares market in terms of the number of listed firms and market capitalization.[2] The two markets are effectively segmented markets and can be studied separately [Bailey (1994)].

The main purpose of this paper is to investigate, using new daily data from the inception of the markets, whether stock prices of the Shanghai and Shenzhen Exchanges follow a random-walk process and can thus be considered efficient. We use two different approaches; namely, the common variance-ratio test of Lo and MacKinlay (1988) and a model-comparison test that contrasts *ex post* forecasts from a random-walk (NAÏVE) model with those obtained from several alternative models.

---

[1]By the end of 1997, the number of listed companies in both markets grew from less than 20 to more than 800, the number of traders exceeded 32 million, and the market capitalization surpassed 1,770 billion renminbi (roughly $220 billion) or about one fourth of China's total GNP.

[2]Of the 800 firms listed in 1997, more than three fourths were registered with the A shares market, and the A shares represent more than 97% of total market capitalization.

Unlike the variance-ratio test, results from the model-comparison test are quite decisive in rejecting the random-walk hypothesis in both Chinese stock markets. Moreover, the results provide strong support for the Artificial Neural Network as a potentially useful device for predicting stock prices in emerging markets.

II. Data

The paper only focuses on Class A shares in the Shanghai (SHG) and Shenzhen (SHZ) Exchanges. We do not consider Class B shares since, as we mentioned earlier, the B shares market is relatively very small compared to the A shares market in terms of market capitalization and level of activity. Established research has shown that low-volume, thinly-traded markets are inappropriate for testing efficiency since they lack liquidity and do not provide smooth transfer of information. Moreover, price indices in small markets tend to exhibit inflated variances thereby complicating statistical inferences [Darrat (1990)].

We utilize new daily data of the A shares closing index prices of the Shanghai Exchange (SHG) from its inception on December 20, 1990 through October 19, 1998; and of the Shenzhen (SHZ) Exchange from its inception on April 4, 1991 through October 19, 1998.[3] Following Lo and MacKinlay (1988), we construct the corresponding weekly data from the daily figures in order to avoid well-known problems inherent in daily sampling (e.g., biases due to bid-ask spreads, non-trading, etc.). The weekly price series are based on the closing value for Wednesday of each week. If the Wednesday observation is missing, then the Tuesday's closing price (or Thursday's if Tuesday's

---

[3]The data are culled from various issues of *Chinese Securities Press* (in Chinese language), and gratefully supplied by Sima Tan and Bin Huang of the Guangzhou Securities Company.

is also missing) is used instead. When Tuesday's, Thursday's and Wednesday's prices are all missing, the return for that week is omitted. The weekly return is calculated as the logarithmic difference between two consecutive weekly prices, yielding continuously compounded returns. Excluding the missing weeks in both markets,[4] our sample yields 402 weekly observations for the Shanghai stock index, and 383 weekly observations for the Shenzhen stock index.[5] We follow Campbell *et al.* (1997) and express stock prices in natural logs in order to stabilize the variance of the series over time and incorporate their exponential growth behavior.

III. Variance-Ratio Test Results

To test the efficiency of the two Chinese stock markets, we first apply the standard variance-ratio test of Lo and MacKinlay (1988). If a given time series follows a random-walk process, the variance of its w-differences of overlapping stock prices is w times the variance of its first difference. Following Lo and MacKinlay (1988) as well as Campbell *et al*. (1997), we use overlapping (as opposed to non-overlapping) w-period returns in estimating the variances in order to obtain "a more efficient estimator and hence a more powerful test," Campbell *et al*. (1997, p. 52). An estimated variance ratio less than one implies negative serial correlation, while a variance ratio greater than one implies positive serial correlation. We also employ two other related statistics: the asymptotic normal

[4]Only six weeks are missing for the Shanghai market and only ten weeks are missing for the Shenzhen market.

[5]Distributional statistics of the weekly data on stock returns reveal that SHG exhibits a larger mean and standard deviation than those for SHZ (0.58% versus 0.37% for the means; and 8.05% versus 7.169% for the standard deviation). Further indication of higher volatility in the SHG market also comes from the coefficient of Kurtosis being almost double that of the SHZ market (15.81 versus 8.03). The skewness parameters are significantly positive (1.83 and 0.99 respectively for SHG and SHZ), implying fat-tailed distributions.

Z test statistic (assuming homoscedasticity), and the heteroscedasticity-consistent Z^c test statistic. Lo and MacKinlay (1988) demonstrate that both test statistics asymptotically follow standard normal distributions and they are thus amenable to conventional statistical inferences. Extensive Monte Carlo results reported in Lo and MacKinlay (1989) suggest that, under the heteroscedastic random-walk null, the Z^c test performs better than either the Box-Pierce test of serial correlation or the Dickey-Fuller test of unit roots.

Table 1 displays the values we obtained for the variance ratio, as well as for the Z and Z^c

----

Put Table 1 About Here

----

statistics using alternative weekly intervals (w=2, 4, 8, 16, and 32).[6] Across these intervals, the variance ratio tests for both SHG and SHZ markets indicate the presence of positive serial correlation in the weekly returns. For example, the variance ratio for the SHG (SHZ) market corresponding to w=2 is 1.10 (1.18). This implies a 10% (18%) first-order autocorrelation in the weekly returns and hence approximately 1% (3.2%) of next week's return variance can be predicted by the current week's return. The evidence from the Z and Z^c test generally concurs with the variance-ratio results.

It should, however, be noted that the test results appear unambiguous only for short intervals, but somewhat clouded for long horizons. Indeed, as lags lengthen beyond eight weeks, Z and Z^c statistics begin to lose significance. In view of that, and given recent debate over the power of such

[6]Following Lo and MacKinlay (1988) and Campbell *et al.* (1997), the variance ratio (VR) is computed as the ratio between the overlapping variance of its w-differences to the product of w times the variance of the first-difference, that is $VR=[var(P_{t+w} - P_t)]/[w \cdot var(P_{t+1} - P_t)]$. The random-walk hypothesis requires that VR=1 which can be tested asymptotically using the normal Z-distribution.

autocorrelation-based tests [see Poterba and Summers (1988) and Cuthbertson (1996)],[7] further evidence seems warranted to substantiate our claim. To do that, we examine *ex post* weekly forecasts from a random-walk model and compare them to forecasts obtained from alternative models. Results from the forecasting tests can shed further light on the efficiency of Chinese stock markets and also aid in judging the relative forecasting ability of several modeling devices.

IV. Further Tests

If Chinese stock prices follow a random-walk process, then a random-walk (NAÏVE) model should not be out-predicted by other models. A NAÏVE model maintains that the best forecast for next week's stock price is simply this week's price. Attempting to produce adequate forecasts of key financial variables, finance literature has witnessed an extensive use of two models; namely, the Auto-Regressive-Integrated-Moving-Average (ARIMA) and the Generalized-Auto-Regressive-Conditional-Heteroscedasticity (GARCH) models. Both ARIMA and GARCH models have achieved varying degrees of success as forecasting devices [see, for example, Domowiwtz and Hakkio (1985), Bollerslev (1986), Akgiray (1989), Alexander (1995), and Su and Fleisher (1998)]. Given their prominence in the literature, we employ the NAÏVE, ARIMA, and GARCH models to generate *ex post* weekly forecasts of Chinese stock prices. In addition, we also utilize another, potentially powerful, forecasting technique known as an Artificial Neural Network (ANN) model.

[7]Some researchers, e.g., Akgiray (1989), have also questioned the typical Gaussian assumption underlying return distributions. Observe also that autocorrelation in the SHG market becomes less pronounced and eventually negative (though insignificant) as interval increases. This may suggest that stock prices in that market exhibit both short-term momentum and long-term reversal behavior. Hence, investors in the SHG market appear to under-react to changes in price fundamentals for shorter periods, but over-react for longer horizons. Prices in the SHZ market, on the other hand, exhibit a momentum behavior throughout.

ANN is a non-parametric (non-linear) modeling technique in which the data series themselves identify the relationships among the variables. Derived from cognitive sciences, the ANN model processes information similar to the way human brains do. As a non-parametric model, ANN has the following important advantages over the more traditional parametric models. First, since ANN does not rely on restrictive parametric assumptions such as normality, stationarity, or sample-path continuity, it is robust to specification errors plaguing parametric models. Secondly, ANN is adaptive and, as such, responds to structural changes in the data-generating processes in ways that parametric models cannot. Finally, ANN is sufficiently flexible and can easily encompass a wide range of securities and fundamental asset price dynamics. Indeed, ANN has considerable flexibility to uncover hidden non-linear relationships among several classes of individual forecasts and realizations [Donaldson and Kamstra (1996)]. Such advantages have recently led to a growing interest in ANN models, and these models have proven successful in several situations like pricing initial public offerings [Jain and Nag (1993)], pricing derivatives [Hutchinson *et al.*(1994)], forecasting futures trading volume [Kaastra and Boyd (1995)], forecasting international equity prices [Cogger *et al.* (1997)], forecasting returns of large U.S. stocks and corporate bonds [Desai and Bharati (1998)], and also forecasting exchange rates [Hu *et al.* (1999)]. The Appendix provides further technical detail on the ANN methodology.[8]

Stock and Watson (1998) and Diebold and Kilian (1999) argue that a necessary prelude to obtaining reliable forecasts from parametric models like ARIMA and GARCH is a unit-root pretesting. We should, nevertheless, caution that the presence of a unit root (non-stationarity) in

[8] White (1992), Smith (1993), and Chauvin and Rumelhart (1995) contain lucid discussions of the ANN approach.

stock prices is only a necessary (but not sufficient) condition for a random-walk process. As Campbell *et al.* (1997) demonstrate, unit root tests only explore the permanent/temporary nature of shocks to the series and, as such, have no bearing on the random-walk hypothesis (or predictability).[9]

We apply the augmented Dickey-Fuller test (ADF) and the weighted symmetric test (WS) to the log-levels and first-differences of the two price series. While the ADF test is perhaps the most popular unit root test, the WS test is found more powerful by Pantula *et al.* (1994) for detecting unit roots. Table 2 reports the results from the unit root tests (with and without a deterministic time

----

Put Table 2 About Here

----

trend). The results there consistently indicate that the price series in both markets are non-stationary in log-levels, but achieve stationarity in first-differences. [That is, each is integrated of order one.]

These unit root results are used to specify the ARIMA model. We identify the autoregressive (AR) and the moving average (MA) terms using a number of criteria: the Akaike Information Criterion, the Schwert Baysian Information Criterion, and the absence of serial correlation in the errors. These criteria suggest an ARIMA (1,1,1) model. Such relatively short ARIMA terms may be adequate since simple t-tests reveal that price predictability primarily come from the first-order correlation component (results of these tests are available upon request). Additionally, Box and Jenkins (1976) argue that under-parameterized (parsimonious) ARIMA models produce better

[9]In this light, the use of unit root tests to examine the random-walk hypothesis appears doubtful. See Liu *et al*. (1997) and Long *et al.* (1999).

forecasts than over-parameterized variants. For the GARCH model, we specify the mean equation as an AR(2) process of price changes to satisfy the stationarity requirement and also to produce white-noise residuals.[10] The variance equation follows a GARCH (1,1) process with one autoregressive term for the conditional variance, and another moving-average term for the squared residuals.[11]

In generating the *ex post* forecast values from the ANN procedure, we design a three-layer backpropagation neural network. We use autoregressive weekly lags of stock prices as the input variables to forecast stock prices. To be compatible with the lag structures of ARIMA and GARCH models, the ANN input variables are specified as an AR(2) process. That is, we use two lags of the price series as inputs in the ANN.[12]

These four alternative procedures (NAÏVE, ARIMA, GARCH and ANN) are used to generate *ex post* dynamic one-week-ahead forecasts of the Chinese stock prices (both SHG and SHZ)[13]. Starting with the base sample of December 20, 1990 - July 15, 1998 for SHG; and of April 4, 1991 - July 15, 1998 for SHZ, we employ each of the models to forecast stock prices successively,

---

[10]This specification for the mean equation of the GARCH model is also compatible with that of the ARIMA model. At any rate, possible misspecification of the mean equation (e.g., omitting dividends or earnings) is not overly damaging in forecasting exercises since the conditional variance estimates are robust to incorrect specifications of the conditional mean [Nelson (1991)].

[11]Evidence reported in Akgiray (1989) for the U.S., and in Su and Fleisher (1998) for China also support the use of a GARCH (1,1) formulation to represent stock prices. Note further that our forecasting results from GARCH and ARIMA are strictly comparable since *both* models use similar (order one) auto-regressive and moving-average processes.

[12]We use different numbers of hidden nodes and autoregressive lags of the prices series, and the results are qualitatively similar to those reported in the text.

[13]To ensure stationary data, we use the logarithmic first-difference of prices in the ARIMA and GARCH models and then convert the obtained forecasts back to log-levels of prices which are then used in the comparison exercises. In the case of ANN, however, we start with prices in their log-levels since the model, being non-parametric, does not require stationary data.

one-week-ahead, for the 12 remaining weeks in both price series. The forecast values thus obtained are considered dynamic (conservative) since the base sample is enlarged successively by the forecast values generated from the preceding round rather than by actual price values.

We compare the *ex post* one-week-ahead forecasts of stock prices using three different evaluation statistics to ensure that our inferences regarding the relative efficiency of the forecasting models are not driven by the particular criterion used in these comparisons. The statistics are the root-mean-squared-error (RMSE), the mean-absolute-error (MAE), and Theil's Inequality Coefficient (U).[14] Besides these three individual *ex post* forecast statistics, we follow Curry *et al.* (1995) and also calculate sum statistics (of the entire forecast horizon) in order to evaluate the overall forecasting performance of the alternative models.

Although useful, these forecasting evaluation criteria cannot determine whether a given forecasting model is in fact "significantly" better than others.[15] Judging the statistical significance of

[14] As Green (2000, p. 310) explains, these evaluation statistics are defined as follows:

$$\mathbf{RMSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(\hat{P}_t - P_t)^2} \; ,$$

where $\hat{P}_t$ = forecast price values

P_t = actual price values

T = number of forecast horizons;

MAE = $\frac{1}{T}\sum |\hat{P}_t - P_t|$; and

$$\text{Theil's U} = \frac{\sqrt{\frac{1}{T}\sum_{t=1}^{T}(\hat{P}_t - P_t)^2}}{\sqrt{\frac{1}{T}\sum_{t=1}^{T}(P_t^2)}}$$

[15] We owe this and many other important insights to an anonymous referee.

rival forecasting models may require the use of encompassing tests [for details, see Donaldson and Kamstra (1996, 1997)].

The rationale behind these tests is that a model claiming to congruently represent the data-generating process must be able to account for the salient features of rival models. In more specific terms, a given model (k) can be considered superior to another model (j) if model k's forecasts significantly explain model j's forecasting errors, and further that model k incorporates relevant information neglected by model j. The encompassing test is implemented by testing the significance of the β and γ coefficients (using t-ratios) in the following two regression equations:

$$(\hat{P}_{jt} - P_t) = \beta_{jk} P_{kt} + \varepsilon_t \quad \textbf{and}, \tag{1}$$

$$(\hat{P}_{kt} - P_t) = \gamma_{kj} \hat{P}_{jt} + \eta_t \tag{2}$$

where $(\hat{P}_{jt} - P_t)$ and $(\hat{P}_{kt} - P_t)$ are the forecasting errors from model j and k, respectively; $\hat{P}_{jt}$ and $\hat{P}_{kt}$ are the forecasts of the two models; and ε and η are random errors. The null hypothesis is that neither model encompasses (outperforms) the other. If β is significantly different from zero but γ is not, then we reject the null hypothesis in favor of the alternative hypothesis that model k encompasses model j. Conversely, if γ is significant but β is not, this is evidence that model j encompasses model k. If both β and γ are not significant, or that both β and γ are significant, then we fail to reject the null hypothesis and conclude instead that neither model encompasses the other. If one of the models encompasses the NAÏVE model for predicting stock prices, it can be concluded that Chinese stock prices do not follow the random walk process.

V. Assessing the Forecasting Ability of Different Models

We employ the four alternative forecasting models to generate *ex post* one-week-ahead forecasts of stock prices in both SHG and SHZ markets. The holdout period contains 12 weeks in both markets. We use RMSE, MAE and Theil's U statistics to evaluate and compare these out-of-sample forecasts. Table 3 reports the results for the Shanghai stock prices, and Table 4 does the same for the Shenzhen market. These results provide credence to our preliminary inefficiency

----

Put Tables 3 and 4 About Here

----

finding derived earlier from the variance-ratio tests. As we can see from the tables, the NAÏVE model fares poorly relative to other models. Moreover, these results unambiguously support the ANN model as the dominant forecasting device in both Chinese stock markets. Looking at the sum values of the various statistics at the bottom of the two tables, the ANN consistently generates the best overall out-of-sample forecasts in both the SHG and the SHZ markets and according to all three evaluation criteria used (smallest sums of RMSE, MAE, and Theil's U). Even on a weekly basis, the ANN generally outperforms the NAÏVE, ARIMA, and GARCH models with much smaller forecasting errors throughout the forecasting horizons.[16]

The two central findings then are that the random-walk model does not receive support from the forecasting results, and that the ANN dominates other models in forecasting Chinese stock prices

[16]The only exception is for Shenzhen prices and then only at the early weeks of the forecasting horizon (up to 3 weeks). However, as the forecasting horizon lengthens beyond three weeks, the ANN begins to dominate all other forecasting procedures for predicting Shenzen stock prices. Observe also that the relative performance of the four forecasting procedures significantly diverge over weeks 4-6 of the holdout sample. We search for possible clues (events) that may explain this divergent behavior of the models but find none. Nevertheless, the relative superiority of the ANN over other models remains robust to using different holdout samples. These alternative results are available upon request.

in both Shanghai and Shenzhen markets. The evidence is underscored by the considerable percent improvement in the sum of RMSE from the ANN over the NAÏVE model, amounting to 57% improvement for the Shanghai market and 43% for the Shenzhen market. Such significant gains in forecasting Chinese stock prices attest to the departure of the Chinese markets from the random walk hypothesis and also further support the superiority of the ANN as a powerful forecasting device.

Results from the encompassing tests reported in Table 5 paint a similar picture. As is clear from the table (panel A for the SHG market and panel B for the SHZ market), forecasts from the

----

Put Table 5 About Here

ANN method significantly explain forecasting errors from each of the three alternative methods, (NAÏVE, ARIMA, and GARCH) in both markets. At the same time, the forecasting errors of the ANN model are not accounted for by any of the rival models. Hence, the ANN model significantly encompasses (outpredicts) other models and, as such, can be considered a dominant forecasting device. These results corroborate our earlier findings and provide another piece of evidence against the random-walk hypothesis in the context of the Chinese stock markets.

Two more implications present themselves. First, the ARIMA model tends to provide inferior forecasts than those from a NAÏVE model in both Chinese markets. Thus, the sole use of ARIMA model in testing market predictability would have, at least in our case, led to the incorrect inference that the Chinese markets follow a random walk.[17] Second, the relative success of the ANN model in forecasting Chinese stock prices suggest that these prices may be better captured by non-linear

[17]Studies that use ARIMA models to test market efficiency abound. Examples include Gau (1984), Mok (1993), and Hayri and Yilmaz (1997).

processes. Such an inference supports Hutchinson *et al* (1994) and Donaldson and Kamstra (1997) in their contention that non-linearities in financial data may be better approximated by the ANN structure and logistic transformation.

Moreover, the ANN process, being perceptual, can better account for cognitive errors generated by semi-rational investors. As Daniel *et al*. (1998) argue, modeling decision-making of semi-rational investors imposes numerous restrictions on trade distributions which are difficult to identify by traditional parametric time series models. The significant superior performance of the distribution-free ANN compared to the random-walk model and other conventional time-series models suggest that "hidden" relationships generated by potentially "irrational" investors are well accommodated for by the ANN algorithm.

VI. Some Possible Explanations for the Non-Random Walk Behavior in Chinese Markets

Our empirical results from variance-ratio and model-comparison tests consistently suggest that Chinese stock prices do not follow a random-walk process. Of course, finding evidence against the random-walk hypothesis is not unique with our study, and many other researchers reach similar conclusions [see, for example, Lo and MacKinlay (1988) for the U.S. equity market; Urrutia (1995) for several Latin American markets; Basci *et al.* (1996) and Antoniou *et al.* (1997) for Turkey; and Santis and Imrohorglu (1997) for many Latin American and Asian countries]. We also note that rejecting the random-walk hypothesis does not necessarily negate market efficiency. As Summers (1986) argues, contradicting the random-walk hypothesis in a given market may only mean that the obtained results are inconsistent with the particular martingale process of a random walk. In addition, some explanations for the rejection of the random-walk hypothesis in the Chinese context are not difficult to find. For example, Fama and French (1988) and Porterba and Summers (1988) suggest that stock prices may be described as the sum of both a random-walk component *and* a stationary (mean-reverting) component. Consequently, stock returns would tend to overreact to fundamental shocks and, as such, would exhibit negative autocorrelation (less-than-unity variance ratio). However, our results reported in Table 1 suggest exactly the opposite for both Chinese stock markets since stock returns display *positive* autocorrelation (larger-than-unity variance ratio). Therefore, Chinese stock prices do not seem to conform to the hypothesized mean-reverting behavior.

Non-synchronous stock trading could also provide another justification. Any new information that arrives after the last transaction of the day will only be reflected in the next day's or perhaps next week's closing price data. This "infrequent trading" could create some predictability in market returns since new information is not instantly embedded in traded stocks, thus allowing for exploitable lags.[18] However, with weekly returns data, the extent of infrequent trading necessary to produce a weekly autocorrelation of 18% (for SHZ) seems empirically unreasonable. According to Lo and MacKinlay's (1988) test of non-trading probability, such a high degree of weekly autocorrelation requires at least 50% of all Chinese stocks to be inactive. Therefore, while non-synchronous trading may be responsible for some of the observed autocorrelation, it appears insufficient to justify it all.

Another explanation for our inefficiency finding may lie in market imperfections that are common in emerging markets due to their ineffective legal structures and lack of transparency that prevent the smooth transfer of information. Interestingly, using Urrutia's (1995) posture, the presence of persistent autocorrelation in the Chinese stock markets may be the outcome of a growing economy rather than market inefficiency *per se*. Uncertainty in Chinese business and political

[18]Note that this non-synchronous trading should be distinguished from the missing data problem we alluded to earlier. The former causes low trading volume, while the latter only means that there is no data for particular days, although trading may be active.

environment could also contribute to the inefficiency results [Su and Fleischer (1998)]. As Campbell *et al.* (1997) argue, a certain degree of market predictability may be necessary to reward investors for bearing certain dynamic risks associated with business and political instability.

Finally, recent research in behavioral finance also suggests other possible reasons for the apparent predictability of stock prices. For example, the behavioral model proposed by Daniel *et al.* (1998) reconciles short-term momentum with long-term reversal behavioral of stock prices. Daniel *et al*. argue that investors tend to be quite confident of the precision of private information and thus overreact to private information signals. Hence, price movements in reaction to the arrival of private information are on average partially reversed in the long-run. Furthermore, investors' confidence rises when public information is in agreement with private information, though does not seem to fall commensurately when contradiction occurs. This "biased self-attribution" behavior could explain the presence of positive short-lag autocorrelation (momentum) in stock prices [see Odean (1999) and Gervais and Odean (2000) and for further discussion of these overconfidence effects]. In emerging markets, such as the Chinese market where securities are less liquid and information is asymmetric, the effects of overconfidence and biased self-attribution can be especially pronounced.[19] In this light, our finding of inefficiency in the Chinese stock market does not appear too surprising.

VII. Conclusions

[19]Other relevant behavioral-finance stories include the psychological model of representativeness heuristic and conservatism of Barberis *et al.* (1998), the prospect theory of Barberis *et al*. (1999), and the theory of information traps and misaligned beliefs proposed by Noth *et al.* (2000)].

Results from variance ratio tests applied on new daily stock price data of China's two official stock exchanges (Shanghai and Shenzhen) appear at odds with the random-walk hypothesis. Irrespective of a constant or a changing variance, Chinese stock prices display a pronounced tendency for positive autocorrelation, raising the potential for predictability.

Besides the standard variance-ratio test, we argue that another interesting approach for testing the random-walk hypothesis is to compare the *ex post* forecasts from the NAÏVE model with those generated from rival models. The random-walk hypothesis would be negated if the NAÏVE model fails to outpredict alternative models. We follow this model-comparison approach and generate *ex post* (one-week-ahead) forecasts of Chinese stock prices from four different forecasting models; namely, NAÏVE, ARIMA, GARCH, and also ANN (artificial neural network). We compare the *ex post* forecasting ability of these models on the basis of alternative evaluation criteria (RMSE, MAE, and Theil's U). In addition, we perform encompassing tests that are particularly useful for assessing statistical superiority among rival forecasting models. The results unambiguously reject the random-walk hypothesis in both Chinese stock markets. We also find consistent evidence supportive of the ANN approach over other models as a useful device for forecasting Chinese stock prices.

Our inference against the random-walk hypothesis should not be totally surprising especially in the case of an emerging market, and many studies have also reported similar results. We have examined several plausible reasons as to why Chinese stock prices may be predictable and highlight the semi-rational behavior of Chinese investors as a likely candidate. We should finally caution that improved forecast performance with the ANN algorithm does not necessarily imply profitable trading rules unless such rules can also provide risk-adjusted excess returns after controlling for transaction costs.

Appendix
Artificial Neural Network

An artificial neural network consists of one input layer, one output layer, and a number of hidden layers in between. Typically, one hidden layer is sufficient to produce acceptable results [Hornik, *et al.* (1992)]. As Figure 1 shows, each of the layers comprises a number of processing

----

Put Figure 1 About Here

----

units known as neurons (or nodes). Each neuron in the input layer stores information provided by the user (e.g., past stock prices). The hidden layer also contains neurons, but information stored in these neurons continuously change as the network trains. The output layer consists of a single neuron since, in our case, we are forecasting only one variable – Chinese stock prices.

The neurons in all three layers are interconnected such that the input layer is only connected forward to the hidden layer, and the output layer is only connected backward to the hidden layer. All connections are assigned certain weights to characterize the strength of the connection. In addition to the processing units, there is a bias dummy neuron connected to hidden and output layers that allows for a more rapid convergence of the training process.

In our application of ANN, we use a particularly popular algorithm called the "backpropagation" to forecast Chinese stock prices. The backpropagation algorithm proceeds as follows. First, inputs (past stock prices) are passed forward to the hidden layer and multiplied by their respective weights to compute a weighted sum. Next, the weighted sum is modified by a transfer function (usually a logistic -- or sigmoid -- function) and then sent to the output layer. Third, the output layer neuron re-calculates the weighted sum and applies the transfer function to produce the output value of this forward pass. Finally, an error signal, which is computed as the difference between the output value of the forward pass and the target value, is "backpropagated" to the hidden layer and then to the input layer. Every weight that connects the hidden and output layers is adjusted proportionally to each neuron's contribution to the forecast error with the objective to minimize the mean squared-error.

The above training process continues interactively until an acceptable (minimum) mean squared-error target chosen by the user is achieved. When setting the target for the mean squared-error, the user should compare the forecast accuracy with the necessary time for convergence. Below is a brief account of the backpropagation algorithm of ANN training process.

Taking a one-hidden layer network as an example, the backpropagation algorithm of ANN works as follows. Available information (in our case, past stock prices) is stored in the input layer as a signal i_m, where m = 1, 2, ..., k, where k is the number of neurons in the input layer (i.e., k is the number of lags in stock prices). The first neuron in the hidden layer forms a sum of the connection weights times the input signals over the connections with all neurons including the bias neuron in the input layer, that is

$$SUM_{hn} = \sum_{m=1}^{k} (w_{im-hn} i_m) + w_{bias-hn} \qquad (n = 1, 2, ... j) \qquad \text{(A1)}$$

where, SUM_{hn} is the weighted sum formed in the n^{th} neuron in the hidden layer;

w_{im-hn} is the weight assigned to the connection between the m^{th} neuron in the input layer and the n^{th} neuron in the hidden layer,

i_m is the signal from the m^{th} neuron in the input layer, and

$w_{bias-hn}$ represents the contributory effect of the bias neuron to the n^{th} neuron in the hidden layer.

The sum is then transformed to a value (OUT_{hn}) between 0 and 1, using the following logistic (sigmoid) function:

$$OUT_{hn} = 1/(1 + \exp(-SUM_{hn})) \qquad (n = 1, 2, \ldots j) \qquad \text{(A2)}$$

The output values from the hidden layer (OUT_{hn}) again become the input values for the output layer, and the following weighted sum is calculated:

$$SUM_O = \sum_{n=1}^{j} (w_{hn-O} OUT_{hn}) + w_{bias-O} \qquad \text{(A3)}$$

The output value of the entire epoch (OUT_o) is calculated as:

$$OUT_O = 1/(1 + \exp(-SUM_O)) \qquad \text{(A4)}$$

Next, the mean-squared error between the actual series and the output series obtained from this epoch is computed. If it exceeds the specified minimum error objective, the ANN adjusts the connection weights for this training epoch accordingly. The ANN then begins another epoch until the mean-squared error reaches the specified minimum error objective.

After each epoch, the connection weights are adjusted as follows. Starting with the weights that connect the output and hidden layers, the weight adjustments are propagated backward using the formula:

$$\delta_O = OUT_O(1 - OUT_O)(TARGET - OUT_O) \qquad \text{(A5)}$$

where δ_O is the delta value (error signal) of the neuron in the output layer, and TARGET is the target output value based on the pre-specified mean-squared error. Based on (A5), the change of the weight connecting the n^{th} hidden neuron to the output layer is also calculated:

$$\Delta w_{hn-O}(q) = \eta \delta_O OUT_{hn} + \alpha[\Delta w_{hn-O}(q-1)] \qquad (n = 1, 2, \ldots j) \qquad \text{(A6)}$$

where, $\Delta w_{hn\text{-}O}(q)$ is the change in weight for the connection between the n^{th} hidden neuron and the output layer neuron in the q^{th} epoch,

η is the user-specified learning rate that controls changes in the weights,

δ_O is the delta value for the neuron in the output layer, and

α is the user-specified momentum rate to the adaptive training process, that helps prevent temporary changes in direction from adversely affecting the learning process (i.e., mitigating wild fluctuations in different directions).

The new weight assigned to this connection is computed as:

$$w_{hn-O}(q) = w_{hn-O}(q-1) + \Delta w_{hn-O}(q) \quad (n = 1, 2, ... j) \qquad \text{(A7)}$$

To adjust the weights in the hidden layer, another formula is used to calculate the delta value (δ_{hn}) since no targets are established for this layer. The formula is:

$$\delta_{hn} = OUT_{hn}(1 - OUT_{hn})(\sum_{n=1}^{j} \delta_O w_{hn-O}) \quad (n = 1, 2, ... j) \qquad \text{(A8)}$$

The weight connecting the input layer to the hidden layer is adjusted using an equation similar to (A6). That is:

$$\Delta w_{im-hn}(q) = \eta \delta_{hn} i_m + \alpha[\Delta w_{im-hn}(q-1)] \qquad (n = 1, 2, ... j, m=1, 2, ...k) \qquad \text{(A9)}$$

This procedure continues until a specified minimum mean-squared error is reached.

References

Antoniou, A., N. Ergul, P. Holmes, and R. Priestly, 1997, Technical Analysis, Trading Volume and Market Efficiency: Evidence from an Emerging Market, *Applied Financial Economics*, **7**, 361-365.

Akgiray, V., 1989, Conditional Heteroscedasticity in Time Series of Stock Returns: Evidence and Forecasts, *Journal of Business*, **62**, 55-80.

Alexander, J., Jr., 1995, Refining the Degree of Earnings Surprise: A Comparison of Statistical and Analysts Forecasts, *Financial Review*, **30**, 469-506.

Bailey, W. 1994, Risk and Return on China's New Stock Markets: Some Preliminary Evidence, *Pacific-Basin Finance Journal*, **2**, 243-260.

Barberis, N., M. Huang, and T. Santos, 1999, Prospect Theory and Asset Prices, *NBER Working Paper 7220*, Cambridge, Mass.

Barberis, N., A. Shleifer, and R. Vishny, 1998, A Model of Investor Sentiment, *Journal of Financial Economics*, **49**, 307-343.

Basci, E. S., S. Ozyildirim, and K. Aydogan, 1996, A Note on Price-Volume Dynamics in an Emerging Market, *Journal of Banking and Finance*, **20**, 389-400.

Bollerslev, T. 1986, Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, **31**, 307-327.

Box, G.E.P. and G. M. Jenkins, 1976, *Time Series Analysis: Forecasting, and Control*, San Francisco, CA: Holden Day.

Campbell, J. Y., A. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, New Jersey.

Chauvin, Y. And D. E. Rumelhart, 1995, *Backpropagation: Theory, Architectures, and Applications*, Erlbaum Publisher, Hillsdale, NJ.

Cogger, K. O., P. D. Koch, and D. M. Lander, 1997, A Neural Network Approach to Forecasting Volatile International Equity Markets, *Advances in Financial Economics*, **3**, 117-157.

Curry, D, Divakar, S., Mathur, S., and Whiteman, C., 1995, BVAR As a Category Management Tool: An Illustration and Comparison with Alternative Techniques, *Journal of Forecasting*, **14**, 181-199.

Cuthbertson, K., 1996, *Quantitative Financial Economics*, John Wiley & Sons, New York.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam, 1998, Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance*, **53**, 1839-1885.

Darrat, A. F. , 1990, Stock Prices, Money, and Fiscal Deficits, *Journal of Financial and Quantitative Analysis*, **25**, 387-398.

Desai, V. S. and R. Bharati, 1998, The Efficacy of Neural Networks in Predicting Returns on Stock and Bond Indices, *Decision Sciences*, **29**, 405-424.

Diebold, F. X. and L. Kilian, 1999, Unit Root Tests Are Useful for Selecting Forecasting Models, *NBER Working Paper 6928*, Cambridge, MA.

Domowitz, I. and C. Hakkio, 1985, Conditional Variance and the Risk Premium in the Foreign Exchange Market, *Journal of International Economics*, **19**, 47-66.

Donaldson and Komstra, 1996, A New Dividend Forecasting Procedure that Rejects Bubbles in Asset Prices, *Reveiw of Financial Studies*, **9**, 333-383.

________, 1997, An Artificial Neural Network-GARCH model for International Stock Return Volatility, *Journal of Empirical Finance, **4**, 17-46.*

Fama, E. F. and K. R. French, 1988, Permanent and Temporary Components of Stock Prices, *Journal of Political Economy*, **96**, 251-276.

Gau, G., 1984, Weak Form Tests of the Efficiency of Real Estate Investment Markets, *Financial Review*, **19**, 301-320.

Gervais, S. and T. Odean, 2000, Learning to Be Overconfident, Review of Financial Studies (forthcoming).

Greene, W. H., 2000, *Econometric Analysis,* Fourth Edition. Upper Saddle River, Vew Jersey: Prentice-Hall, Inc.

Hayri, A. And K. Yilmaz, 1997, Privatization and Stock Market Efficiency: The British Experience, *Scottish Journal of Political Economy*, **44**, 113-133.

Hornik, K., M. Stinchcombe, and H. White, 1989, Multilayer Feedforward Networks Are Universal Approximators, *Neural Networks,* **2**, 359-366.

Hu, M. Y., G. Zhang, C. Jiang and B. E. Patuwo, A Cross Validation Analysis of Neural Network Out-of-Sample Performance in Exchange Rate Forecasting, *Decision Sciences*, **30**, 197-216.

Hutchinson, J., A. Lo, and T. Poggio, 1994, A Nonparametric Approach to the Pricing and Hedging of Derivative Securities Via Learning Networks, *Journal of Finance*, **49**, 851-889.

Jain, B. A. and B. N. Nag, 1993, Artificial Neural Network Models for Pricing Initial Public Offerings, *Decision Sciences*, **26**, 283-302.

Kaastra, I. and M. S. Boyd, 1995, Forecasting Futures Trading Volume Using Neural Networks, *Journal of Futures Markets*, **15**, 953-970.

Liu, X., H. Song, and P. Romilly, 1997, Are Chinese Stock Markets Efficient? A Cointegration and Causality Analysis, *Applied Economics Letters*, **4**, 511-515.

Lo, A.W. and A.C. MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies*, **1**, 41-66.

__________, 1989, The Size and Power of the Variance Ratio Test in Finite Samples: A Monte Carlo Investigation, *Journal of Econometrics*, **40**, 203-238.

Long, D. M., J. D. Payne, and C. Feng, 1999, "Information Transmission in the Shanghai Equity Market," *Journal of Financial Research*, **22**, 29-46.

Mok, H. M. K., 1993, Causality of Interest Rate, Exchange Rate and Stock Prices at Stock Market Open and Close in Hong Kong, *Asia Pacific Journal of Management*, **10**, 123-143.

Nelson, D. B., 1991, Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica*, **59**, 347-370.

Noth, M., C. F. Camerer, C. R. Plott, and M. Weber, 2000, Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs, American Financial Association Annual Conference Proceeding, Boston, MA, January 2000.

Odean, T., 1999, Do Investors Trade Too Much? *American Economic Review*, **89**, 1279-1298.

Pantula, S., G. Gonzales-Farias, and W.A. Fuller, 1994, A Comparison of Unit-Root Test Criteria, *Journal of Business and Economic Statistics*, **12**, 449-459.

Poterba, J.M., and L.H. Summers, 1988, Mean Reversion in Stock Prices: Evidence and Implications, *Journal of Financial Economics*, **22**, 27-59.

Santis, G. and S. Imrohoroglu, 1997, Stock Returns and Volatility in Emerging Financial Markets, *Journal of International Money and Finance*, **16**, 561-579.

Smith, M., 1993, *Neural Networks for Statistical Modeling*, Van Nostrand Reinhold Publisher, New York, NY.

Stock, J. H. and M.W. Watson, 1998, A Comparison of Linear and Nonlinear Univariate Models For Forecasting Macroeconomic Time Series, *NBER working paper 6607*, Cambridge, Mass.

Su, D. and B.M. Fleisher, 1998, Risk, Return and Regulation in Chinese Stock Markets, *Journal of Economics and Business*, **50**, 239-256.

Summers, L. H., 1986, Does the Stock Market Rationally Reflect Fundamental Values? *Journal of Finance*, **41**, *591-601*.

Urrutia, J. L. 1995, Tests of Random-Walk and Market Efficiency for Latin American Emerging Markets, *Journal of Financial Research*, **18**, 299-309.

White, H., 1992, *Artificial Neural networks: Approximation and Learning Theory*, Blackwell Publishers, New York.

Table 1. Variance Ratio Test Results for China's Weekly Stock Returns

	w=2	w=4	w=8	w=16	w=32
A. Shanghai Stock Market (SHGE)					
Variance Ratios	1.10	1.20	1.30	1.25	0.87
Z statistics	1.92*	2.11**	2.05**	1.13	-0.41
Z^c statistics	1.11	1.79*	2.18**	1.34	-0.47
B. Shenzhen Stock Market (SHZ)	w=2	w=4	w=8	w=16	w=32
Variance Ratios	1.18	1.29	1.49	1.52	1.41
Z statistics	3.36**	2.94**	3.18**	2.28**	1.24
Z^c statistics	2.70**	3.28**	3.78**	2.82**	1.77

Notes: w is the number of weekly intervals aggregated to compute the variance ratios; Z statistics are the asymptotic normal test statistics under homoscedasticity; Z^c statistics are the asymptotic normal test statistics under heteroscedasticity.
An * indicates rejection of the null hypothesis of no autocorrelation at the 10% significance level, while ** indicates rejection at the 5% level.

Table 2. Unit Root Test Results for China's Weekly Stock Prices (P) and Stock Returns (ΔP)

	Augmented Dickey-Fuller Test		Weighted Symmetric Test	
A. Shanghai Stock Market (SHG)				
	With Trend	Without Trend	With Trend	Without Trend
P_t	-2.58	-2.46	-2.00	-0.65
ΔP_t	-6.83**	-6.77**	-6.93**	-6.86**
B. Shenzhen Stock Market (SHZ)				
	Augmented Dickey-Fuller Test		Weighted Symmetric Test	
	With Trend	Without Trend	With Trend	Without Trend
P_t	-1.79	-1.50	-2.06	-1.34
ΔP_t	-9.83**	- 9.84**	-9.89**	-9.89**

Notes: See notes to Table 1. P_t=log of the stock price index; ΔP_t = first differences of P_t . The numbers of lags in the testing equations are selected by Akaike Information Criterion (provided the residuals are also white-noise).

Table 3. Out-of-Sample, One-Week-Ahead, Forecasting Performance of Alternative Models
(The variable being forecast is the logs of stock prices)

The Shanghai Stock Market (SHG)

	NAÏVE			**ARIMA**			**GARCH**			**ANN**		
Week-Ahead	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U
1	0.0286	0.0286	0.0040	0.0302	0.0302	0.0042	0.0536	0.0536	0.0075	0.0045	0.0045	0.0006
2	0.0238	0.0232	0.0033	0.0268	0.0266	0.0037	0.0396	0.0349	0.0055	0.0194	0.0158	0.0027
3	0.0229	0.0225	0.0032	0.0284	0.0282	0.0040	0.0370	0.0337	0.0052	0.0284	0.0241	0.0040
4	0.0608	0.0456	0.0085	0.0698	0.0538	0.0098	0.0689	0.0557	0.0096	0.0312	0.0277	0.0044
5	0.0936	0.0706	0.0131	0.1061	0.0814	0.0149	0.0722	0.0614	0.0101	0.0160	0.0385	0.0065
6	0.1030	0.0822	0.0145	0.1187	0.0959	0.0167	0.0659	0.0514	0.0093	0.0452	0.0389	0.0064
7	0.1131	0.0935	0.0159	0.1323	0.1100	0.0186	0.0638	0.0511	0.0090	0.0463	0.0408	0.0065
8	0.1133	0.0961	0.0160	0.1352	0.1155	0.0190	0.0600	0.0468	0.0084	0.0433	0.0360	0.0061
9	0.1089	0.0926	0.0153	0.1327	0.1149	0.0187	0.0570	0.0440	0.0080	0.0454	0.0386	0.0064
10	0.1066	0.0916	0.0150	0.1329	0.1169	0.0187	0.0561	0.0443	0.0079	0.0456	0.0395	0.0064
11	0.1039	0.0898	0.0146	0.1326	0.1181	0.0187	0.0537	0.0417	0.0076	0.0476	0.0417	0.0067
12	0.1032	0.0903	0.0145	0.1351	0.1216	0.0190	0.0534	0.0424	0.0075	0.0474	0.0420	0.0067
(12)	**0.9818**	**0.8267**	**0.1380**	**1.1808**	**1.0131**	**0.1659**	**0.6812**	**0.5610**	**0.0956**	**0.4203**	**0.3881**	**0.0634**

Notes: NAÏVE is the random-walk model, ARIMA is the autoregressive integrated moving average model; GARCH is the generalized autoregressive conditional heteroscedasticity model, and the ANN is the artificial neural network. The out-of-sample, one-week-ahead, forecasts are generated dynamically whereby the sample-base period of December 20, 1990-July 15, 1998 is successively expanded by the forecast values from the previous round until the forecast of the 12^{th} week is generated.

Table 4. Out-of-Sample, One-Week-Ahead, Forecasting Performance of Alternative Models (the variable being forecast is the logs of stock prices)

The Shenzhen Stock Market (SHZ)												
	NAÏVE			**ARIMA**			**GARCH**			**ANN**		
Week-Ahead	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U	RMSE	MAE	Theil's U
1	0.0116	0.0116	0.0014	0.0074	0.0074	0.0009	0.0668	0.0668	0.0081	0.0213	0.0213	0.0026
2	0.0085	0.0073	0.0010	0.0054	0.0045	0.0007	0.0882	0.0861	0.0107	0.0400	0.0368	0.0049
3	0.0102	0.0091	0.0012	0.0112	0.0090	0.0014	0.0845	0.0829	0.0103	0.0442	0.0417	0.0054
4	0.0508	0.0319	0.0062	0.0555	0.0340	0.0068	0.0732	0.0636	0.0089	0.0403	0.0377	0.0049
5	0.0814	0.0557	0.0099	0.0886	0.0601	0.0108	0.0682	0.0594	0.0083	0.0473	0.0438	0.0058
6	0.0897	0.0669	0.0109	0.0990	0.0734	0.0121	0.0800	0.0700	0.0098	0.0454	0.0423	0.0057
7	0.1027	0.0802	0.0126	0.1144	0.0888	0.0140	0.0741	0.0602	0.0091	0.0491	0.0458	0.0060
8	0.1064	0.0864	0.0131	0.1202	0.0970	0.0148	0.0720	0.0596	0.0088	0.0474	0.0442	0.0058
9	0.1053	0.0874	0.0129	0.1207	0.1001	0.0148	0.0779	0.0657	0.0096	0.0447	0.0397	0.0055
10	0.1068	0.0906	0.0131	0.1243	0.1054	0.0153	0.0773	0.0663	0.0095	0.0428	0.0375	0.0053
11	0.1092	0.0943	0.0134	0.1289	0.1111	0.0158	0.0783	0.0683	0.0096	0.0416	0.0367	0.0051
12	0.1145	0.0999	0.0141	0.1365	0.1187	0.0168	0.0758	0.0659	0.0093	0.0431	0.0383	0.0053
(12)	**0.8972**	**0.7214**	**0.1101**	**1.0121**	**0.8095**	**0.1242**	**0.9163**	**0.8148**	**0.1120**	**0.5072**	**0.4658**	**0.0623**

Notes: See notes to Table 3. The out-of-sample, one-week-ahead, forecasts are generated dynamically whereby the sample-base period of April 4, 1991 - July 15, 1998 is successively expanded by the forecast values from the precious round-until the forecast of the 12th week is generated.

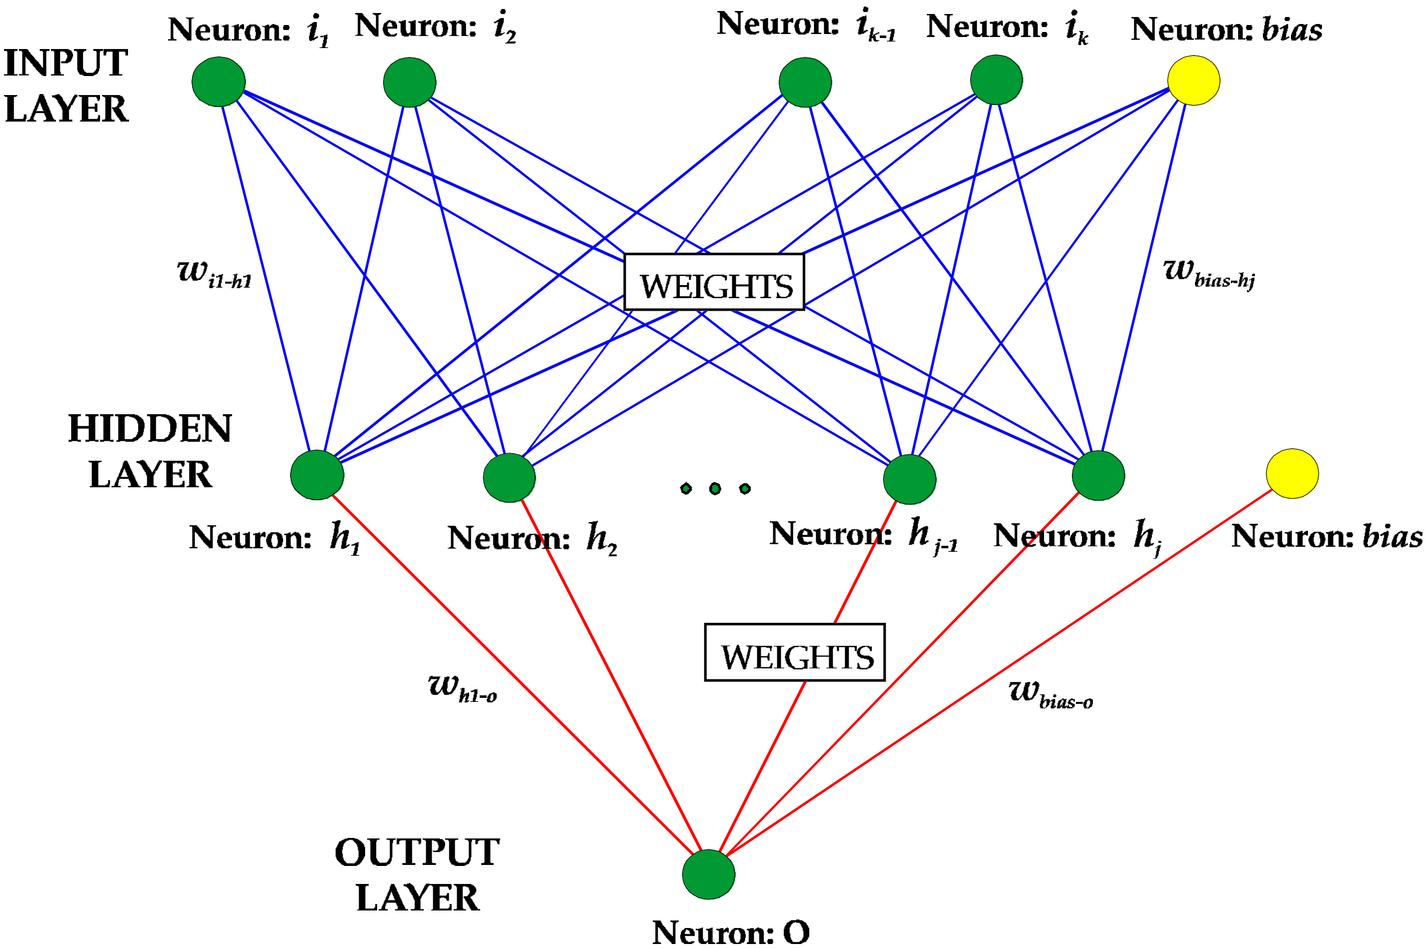

Figure 1. Three-Layer Artificial Neural Network

Table 5. Encompassing Tests of Out-of-Sample Forecasting Performance of Alternative Models

Panel A: The Shanghai Stock Market (SHG)				
Dependent Variable: Forecasting Errors from	Independent Variable: Forecasts from			
	NAÏVE	ARIMA	GARCH	ANN
NAÏVE	--	6.00**	5.90**	5.93**
ARIMA	6.84**	--	6.72**	6.76**
GARCH	3.00**	3.00**	--	3.01**
ANN	0.41	0.40	0.40	--

Panel B: The Shenzhen Stock Market (SHZ)				
Dependent Variable: Forecasting Errors from	Independent Variable: Forecasts from			
	NAÏVE	ARIMA	GARCH	ANN
NAÏVE	--	5.83**	5.71**	5.76**
ARIMA	5.78**	--	5.67**	5.73**
GARCH	5.74**	5.74**	--	5.74**
ANN	0.64	0.65	0.62	--

Notes: Variables being forecast are the logs of stock prices. The test statistics are heteroscedastic-consistent t-ratios. An ** indicates statistical significance at the 5% level. See notes to Tables 3 & 4 for further forecasting details.