Journal of Forecasting
J. Forecast. **19**, 39–52 (2000)

Forecasting for the Generation of Trading Signals in Financial Markets

KIN LAM[1]* AND KING CHUNG LAM[2]
[1]*Hong Kong Baptist University, Hong Kong*
[2]*The University of Hong Kong, Hong Kong*

ABSTRACT
In this paper we show that optimal trading results can be achieved if we can forecast a key summary statistic of future prices. Consider the following optimization problem. Let the return r_i (over time $i = 1, 2, \ldots, n$) for the ith day be given and the investor has to make investment decision d_i on the ith day with $d_i = 1$ representing a 'long' position and $d_i = 0$ a 'neutral' position. The investment return is given by $r = \sum_{i=1}^{n} r_i d_i - c \sum_{i=1}^{n+1} |d_i - d_{i-1}|$, where c is the transaction cost. The mathematical programming problem of choosing $d_1, \ldots, d_n$ to maximize r under a given transaction cost c is shown to have an analytic solution, which is a function of a key summary statistic called the largest change before reversal. The largest change before reversal is recommended to be used as an output in a neural network for the generation of trading signals. When neural network forecasting is applied to a dataset of Hang Seng Index Futures Contract traded in Hong Kong, it is shown that forecasting the largest change before reversal outperforms the k-step-ahead forecast in achieving higher trading profits. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS finance; forecasting; trading rules; largest change before reversal; neural network

'Buy low, sell high' is an investment dictum which is easier said than done. Trading decisions have to be made without the knowledge of future prices and therefore it is impossible to determine whether the current price is a low or a high. For a trading rule to be practicable, it has to satisfy the Markov property as defined in Neftci (1991), which is equivalent to saying that trading decisions can only utilize past but not future price information. For the profitability of some technical trading rules using past information only, see Taylor (1994) and Corrado and Lee (1992). However, even if we are allowed to make use of future information, it may still be a non-trivial problem to devise an optimal strategy, as we will explain below.

Consider an asset whose prices fluctuate from day to day and the closing price on the t*th* day ($t = 0, 1, 2, \ldots, n$) is q_t. Let $p_t = \ln q_t$ be the log-price and $r_t = p_t - p_{t-1}$ be the continuously compounded return on day t. An investment decision is a vector $\mathbf{d} = (d_1, d_2, \ldots, d_n)$, where

* Correspondence to: Kin Lam, Department of Finance and Decision Sciences, School of Business, Hong Kong Baptist University, 224 Waterloo Road, Kowloon, Hong Kong.

CCC 0277–6693/2000/010039–14$17.50

Received July 1997
Accepted December 1998

$d_i = 0$ or 1, meaning that the investor maintains a 'neutral' or a 'long' position respectively on day *i*. Here, we assume that the strategy of short-selling the asset is not available to the investor. This restriction can be lifted by allowing $d_i = -1$, but we will carry out the discussion here using the simplified version of restricting d_i to be 0 or 1. Assuming that the investor starts with a certain sum of money and is fully invested when a 'long' position is taken, the investment return corresponding to the decision **d** is given by $r = \sum_{i=1}^{n} r_i d_i$. Assuming that $r_1, \ldots, r_n$ are known, r can be maximized by choosing $d_i = 1$, if $r_i > 0$ and $d_i = 0$, if $r_i \leqslant 0$.

The optimization problem becomes non-trivial when transaction cost is taken into account. Whenever the investment decision d_{i+1} is such that $d_i \neq d_{i+1}$, a transaction cost of 100c% is incurred. Assuming that the investor has a neutral position in the beginning and in the end, i.e. $d_0 = d_{n+1} = 0$, the investment return is given by

$$r = \sum_{i=1}^{n} r_i d_i - c \sum_{i=1}^{n+1} |d_i - d_{i-1}|$$

The mathematical programming problem (*) of maximizing r with given $r_1, \ldots, r_n$ and $d_0 = d_{n+1} = 0$ becomes non-trivial. The solution of (*) will be discussed in later sections of this paper. In the next section, we give a justification of why (*) is of practical interest and why it has important implications of forecasting for the generation of trading signals in a financial market.

FORECASTING FOR THE GENERATION OF TRADING SIGNALS

At first sight, the solution of (*) is of no practical interest because the optimal trading decision at time t depends on the values of $r_{t+1}, r_{t+2}, \ldots$ and is hence not Markov in the sense of Neftci (1991). However, the solution of (*) can give great insights as to which quantity is to be forecasted for the generation of trading signals. When the market satisfies the weak-form market efficiency in the sense of Fama (1971), the efforts of generating trading signals would be futile. However, according to Brock, Lakonishok and LeBaron (1992), the efficient market hypothesis has come under serious siege in recent years. Serial correlation in returns over various investment horizons are reported for individual stocks as well as for various portfolio of stocks; see, for example, Fama and French (1986), Poterba and Summers (1988), Jegadeesh (1990), and Cutler, Poterba and Summers (1990). The correlations are statistically significant and sometimes are of magnitudes that could not be explained by non-synchronous trading or by the existence of market friction such as bid/ask spread. This provides evidence for the predictability of equity returns from past returns. Utilizing the Dow Jones Index from 1897 to 1986, Brock, Lakonishok and LeBaron (1992) showed that returns obtained from some popular trading strategies are not consistent with the random walk model and the GARCH families of models. Beja and Goldman (1980) argued that prices do not always adjust quickly to information shocks and markets can be in short-run disequilibrium. Grossman and Stiglitz (1980) attributed the slow adjustment process to the cost of acquiring and evaluating information and the need to adjust to new information. The profitability of technical trading systems on financials was reported by Dale and Workman (1981), and Taylor and Tari (1989). Lukac and Brorsen (1990) carried out a comprehensive test of futures market disequilibrium. Twenty-three trading systems were tested in thirty futures markets for eleven years. All but two trading systems had significant gross returns. For currencies futures, Taylor (1994) reported that currency futures could have been traded profitably from 1982 to 1990 using the channel rule. The average payoff, net of transaction costs, was both statistically and economically significant. Recently, Bessembinder and Chan (1995) pointed out that some Asian markets may not be as informationally efficient as their US or European counterparts. They found that some simple technical analysis can be quite successful in the emerging markets of Malaysia, Thailand and Taiwan. The 'break-even' transaction cost for some technical trading rules was found to be as high as 1.57%.

Recently, there has been considerable research in using neural networks to generate trading signals in order to profit from a possibly inefficient market. Such applications can be found in Refenes *et al.* (1993), Moody and Utans (1992), Mehta (1994) and in the collected work of Trippi and Turbon (1992). Historical data forming what is called a training set are needed for the training of a neural network. In the training set, the whole price series is given, satisfying the assumption in (*) that $r_1, \ldots, r_n$ are known. Typically, the neural network is trained to forecast the k-step-ahead price and trading signals will be generated with the help of the forecasted prices. In the training set, the k-step-ahead price, p_{t+k}, is known and the network is trained to give a forecast as $\hat{p}_{t+k}$. After training, trading signals can be generated by the following rule:

$$\begin{cases} \text{long} & \text{if } \hat{p}_{t+k} > p_t \\ \text{short} & \text{if } \hat{p}_{t+k} < p_t \end{cases} \tag{1}$$

This approach has been widely used in the literature; see, for example, Mehta (1994), Bjorn (1994) and Hsu *et al.* (1993).

The use of a neural network to forecast with the purpose of generating trading signals raises two general but important questions: (1) Which quantity should be forecasted? and (2) How do we make use of the forecasted quantity to construct trading signals? In the k-step-ahead forecast, it is not clear how k should be chosen and we are not sure whether strategy (1) is the best way to generate trading signals. Instead of making a single-step-ahead forecast, one can make multi-period forecasts, which help in tracing the path of the prices in detail. However, we cannot be sure how many periods should be used and how these forecasts can be utilized to construct trading signals. In Kimoto and Asakawa (1990), the authors chose to forecast a summary statistic of future prices, which reflect the future dynamics of the price changes. The quantity to be forecasted is the weighted sum of the future weekly returns of TOPIX on the Tokyo Stock Exchange. As the weighted sum captures more future information, this may enhance the performance of the trading signals generated. However, there is no theoretical reason why one summary statistic should be preferred over another. By deriving the solution to the optimization problem (*), it is discovered that its analytic solution depends on a summary statistic called the 'largest change before reversal'. This summary statistic will be described in detail in the next section. Since the optimal solution is a function of this summary statistic, it is clear how it can be used to generate trading signals.

On the other hand, it is possible to generate trading signals directly, bypassing the forecasting step. One can train the network to generate trading signals directly from the inputs. In other words, the output of the network is not a quantity to be forecasted, but is a trading signal. Some applications using this type of output are Chauvin (1994), Binks and Allinson (1991), Bergerson and Wunsch (1991), Margarita (1991) and Zaremba (1990). What we need here is a set of target trading signals for the training of the neural network. In most cases, these target trading signals are expert generated, as in Bergerson and Wunsch (1991), or from price patterns recognized by technical chartists, or from a patterns library, as in Binks and Allinson (1991). However, if it pays to train the computer to learn the technical analysts' timing for when to buy and sell, it may be better to train the computer to learn the optimal buy/sell times which should be known as far as the training set is concerned. Thus, an analytic solution for (*) is useful and can be used as output in the training set so as to train the network to trade in the future.

LARGEST CHANGE BEFORE REVERSAL

In this section we assume that prices $p_0, p_1, \ldots, p_n$ are known. At the end of day t, we define a statistic depending on $p_{t+1}, \ldots, p_n$ which summarizes the price movements after day t. This summary statistic is called the largest change before reversal and is denoted by C_t. To give a rigorous definition of C_t, we first define T_t, the time for first reversal, as follows. Let p_t be the closing price at the end of day t. When $p_{t+1} > p_t$, T_t is defined as the first day after $(t+1)$ in which the price reverses to a level below p_t, i.e. $p_{Tt} \leqslant p_t$. Similarly, when $p_{t+1} < p_t$, T_t is defined as the first day after $(t+1)$ in which the price reverses to a level above p_t, i.e. $p_{Tt} \geqslant p_t$. If the reversal never happens, define $T_t = n$. Mathematically,

$$T_t = \begin{cases} \min\{T : (p_{t+1})(p_t - p_t) \leqslant 0 \quad t+1 < T \leqslant n\} \wedge n & \text{if } p_{t+1} \neq p_t \\ t+1 & \text{if } p_{t+1} = p_t \end{cases}$$

Here the minimum of an empty set is taken to be ∞ and $a \wedge b$ denotes min (a,b). We now define C_t as follows:

$$C_t = \begin{cases} \max\limits_{t+1 \leqslant j \leqslant T_t} (p_j - p_t) & \text{if } p_{t+1} - p_t > 0 \\ 0 & \text{if } p_{t+1} - p_t = 0 \\ \min\limits_{t+1 \leqslant j \leqslant T_t} (p_j - p_t) & \text{if } p_{t+1} - p_t < 0 \end{cases}$$

C_t is called the largest change before reversal. For a pictorial illustration of C_t, see Figure 1. If the reversal is from above p_t to below p_t, C_t is positive, meaning that the price first goes up and rises to a maximum equal to $p_t + C_t$ before it drops below p_t again. If the reversal is from below p_t to above p_t, C_t is negative, meaning that the price first drops to a minimum of $p_t + C_t$ before it rises to a level above p_t.

The statistic C_t summarizes a lot of information about the price evolution after time t. Its sign indicates whether the price is rising or not in the immediate future. Its magnitude indicates whether it is worthwhile to buy or sell the asset. If the forecast is such that the price is rising but not to a level enough to recover transaction costs, an investor should not switch from a neutral position to a long position. Thus, it is natural to compare C_t with the transaction cost c. In the next section, we will establish the theoretical importance of C_t by showing that it is closely related to the solution of (*).

OPTIMIZING TRADING PROFIT WITH FULL KNOWLEDGE OF PRICE CHANGES

In this section we state the theorem which gives an analytical solution for the optimization problem (*). The solution is given more generally in that the given value of d_0 can either be 0 or 1.

Theorem Let $r = \sum_{i=1}^{n} r_i d_i - c \sum_{i=1}^{n+1} |d_i - d_{i-1}|$ in which $r_0, r_1, \ldots, r_n$, d_0 (=0 or 1) and d_{n+1} (=0) are given. Consider the problem (*) of maximizing r as a function of $d_1, \ldots, d_n$ where d_i is constrained to take the value 0 or 1 for $i = 1, 2, \ldots, n$. The optimal solution of (*) can be obtained as follows. Let C_t be the largest change before reversal at day t. Define $\tilde{d}_t$ sequentially as follows:

$$\begin{aligned}
&\text{(i) } \tilde{d}_0 = d_0 \\
&\text{(ii) If } \tilde{d}_{t-1} = 0, \text{ define } \quad \tilde{d}_t = \begin{cases} 1 & C_{t-1} > 2c \\ 0 & C_{t-1} \leqslant 2c \end{cases} \\
&\text{(iii) If } \tilde{d}_{t-1} = 1, \text{ define } \quad \tilde{d}_t = \begin{cases} 0 & C_{t-1} < -2c \\ 1 & C_{t-1} \geqslant -2c \end{cases}
\end{aligned} \tag{2}$$

By slightly modifying $\tilde{\mathbf{d}}$, we can obtain the optimal solution $\mathbf{d}^*$ of (*) as follows:

(1) If $\tilde{d}_n = 0$, define $\mathbf{d}^* = \tilde{\mathbf{d}}$.

(2) If $\tilde{d}_n = 1$, let $T = \min\{0 \leqslant i \leqslant n,\ \tilde{d}_i = \tilde{d}_{i+1} = \cdots = \tilde{d}_n = 1\}$. Suppose $p_T, p_{T+1}, \ldots, p_n$ attains a maximum at $p_{T'}$. Define, for $t = 0, 1, \ldots, n+1$,

$$d_t^* = \begin{cases} \tilde{d}_t & t \leqslant T' \\ 0 & t > T' \end{cases}$$

The above theorem states that the optimal solution $\mathbf{d}^*$ equals $\tilde{\mathbf{d}}$ with a possible adjustment at the ending. Since $\tilde{d}_t$ depends on C_t, the optimal solution can be regarded as a function of C_t. The way $\tilde{d}_t$ is determined by C_t in model (2) is consistent with the observation in the previous section that C_t should be compared with the transaction cost. The idea is that if price changes are not large enough to cover a round trip (a buy and a sell) of transaction cost (=$2c$), one should maintain the original position instead of switching from neutral to long or from long to neutral. The final decision needs to be adjusted by modifying $\tilde{d}_t$ to obtain d_t^* for t close to n because of the finite investment horizon n which imposes a restriction that the investment exercise ends on day n. When n is large, the optimal solution is basically $\tilde{\mathbf{d}}$.

Before we prove the theorem, we first establish some properties associated with the decision $\mathbf{d}^* = (d_1^*, d_2^*, \ldots, d_n^*)$ defined in the theorem. We consider here the case $d_0 = 0$ only. $\mathbf{d}^*$ divides the time interval $[0, n+1]$ into k sub-intervals $I_1 = [T_0, T_1]$, $I_2 = [T_1 + 1, T_2]$, $I_3 = [T_2 + 1, T_3], \ldots, I_k = [T_{k-1} + 1, T_k]$ where $T_0 = 0$, $T_k = n + 1$, $d_t^* \equiv 0$ for $t \in I_1$, $d_t^* \equiv 1$ for $t \in I_2$, $d_t^* \equiv 0$ for $t \in I_3, \ldots$, and $d_t^* \equiv 0$ for $t \in I_k$. We first state and prove the following lemma.

Lemma

(1) If $d_j^* \equiv 1$ for $j \in I_i$, then for any $[s, t] \subseteq I_i$, we have

$$\text{(a) } \sum_{j=s}^{t} r_j \geqslant -2c (\Leftrightarrow p_t - p_{s-1} \geqslant -2c)$$

$$\text{(b) } \sum_{j=T_{i-1}+1}^{s} r_j \geqslant 0 (\Leftrightarrow p_s - p_{T_{i-1}} \geqslant 0)$$

$$\text{(c) } \sum_{j=1}^{T_i} r_j \geqslant 0 (\Leftrightarrow p_{T_i} - p_{t-1} \geqslant 0)$$

(2) If $d_j^* \equiv 0$ for $j \in I_i$, then for any $[s, t] \subseteq I_i$, we have

$$\text{(a)} \ \sum_{j=s}^{t} r_j \leqslant 2c (\Leftrightarrow p_t - p_{s-1} \leqslant -2c)$$

$$\text{(b)} \ \sum_{j=T_{i-1}+1}^{s} r_j \leqslant 0 (\Leftrightarrow p_s - p_{T_{i-1}} \leqslant 0) \text{ for } i \neq 1$$

$$\text{(c)} \ \sum_{j=t}^{T_i} r_j \geqslant 0 (\Leftrightarrow p_{T_i} - p_{t-1} \leqslant 0) \text{ for } i \neq k$$

Proof of Lemma Instead of presenting the full proof, we give the proof for (1)(a) only. Suppose (1)(a) is false and we have $p_t - p_{s-1} < -2c$. It follows that among $C_{S-1}, C_S, \ldots, C_{t-1}$, one of them, say C_w, has to be less than $-2c$. This implies that d_{w+1}^* is equal to 0, contradicting the assumption that $d_j^* \equiv 1$ on $[s, t] \subseteq I_i$. The proofs of the other cases are similar.

Proof of Theorem We will prove the theorem by mathematical induction on n. Consider first the case $n = 1$. Obviously, $C_0 = p_1 - p_0$ and

$$d_1^* = \begin{cases} 1 & p_1 - p_0 > 2c \\ 0 & p_1 - p_0 \leqslant 2c \end{cases}$$

It can be seen easily that **d*** is an optimal solution when $n = 1$.

Suppose the theorem has been established for n and we want to establish it for $n + 1$. We assume that the theorem is not valid for $n + 1$ and let **d** be an optimal solution for $n + 1$ with investment return better than **d*** and **d** $\neq$ **d***. We will establish contradiction under two separate cases. Note that $d_{n+2} = d_{n+2}^* = 0$.

Case 1 $d_i = d_i^*$ for some $i \in [1, n + 1]$.
Let $s(\geqslant 1)$ be the smallest index satisfying $d_s = d_s^*$. If $s = 1$, contradiction arises because of the induction assumption. If $s > 1$, let s' be the largest index such that $d_1^* = d_2^* = \cdots = d_{s'}^*$. Let $s^* = \min(s', s)$ and define **d**′ as follows:

$$d_i' = \begin{cases} 1 - d_i^* & i \leqslant s^* \\ d_i^* & i > s^* \end{cases}$$

Claim d' has an investment return not less than that of **d** and hence not less than that of **d***.

We will first see how contradiction follows from the claim. This is because $d_1' = d_1^*$ and by an induction assumption, **d**′ cannot have an investment return better than that of **d***.

To establish the claim, we consider one case in detail and omit the other cases, the proofs of which are similar. The case we consider is $s > s' + 1$ and $d_1' = d_2' = \cdots = d_{s'}^* = 0$. For this case, $d_{s'+1}^* = 1$, $d_{S'+1} = 0$, $d_1 = d_2 = \cdots = d_{S'} = 1$. By definition of **d**′, $d_1' = d_2' = \cdots = d_{S'+1}' = 1$. By part (2)(c) of the lemma, it can be shown that the investment return of **d**′ is not less than that of **d**.

Case 2 $d_t \neq d_t^*$ for all $t \in [1, n+1]$.
Let s' be the largest index satisfying $d_1^* = d_2^* = \cdots = d_{s'}^*$. If $s' < n+2$, we define $\mathbf{d}'$ by

$$d_t = \begin{cases} 1 - d_t & t \leqslant s' \\ d_t & t > s' \end{cases}$$

As in the proof of the claim above, we can show that $\mathbf{d}'$ has an investment return better than $\mathbf{d}$. This leads to a contradiction because $d_1' = d_1^*$. If $s' = n+2$, we have $d_0^* = d_1^* = d_2^* = \cdots = d_{n+2}^* = 0$ and $d_1 = d_2 = \cdots = d_{n+1} = 1$, $d_1' = d_2' = \cdots = d_{n+1}' = 0$. By part (2)(a) of the lemma, $\mathbf{d}'$ has an investment return not less than that of $\mathbf{d}$. This leads to a contradiction because $d_1' = d_1^*$. The proof of the theorem is now complete by mathematical induction.

LARGEST CHANGE BEFORE REVERSAL AS OUTPUT IN A NEURAL NETWORK

As indicated in previous sections, we can make use of C_t, the largest change before reversal, as an output in a neural network. The advantage is that once C_t is known, the optimal trading rule is a function of C_t, as indicated by (2) in the theorem. We can compare its performance with the neural network which uses a fixed horizon future return as output. The trading rule associated with the neural network is the obvious one that the investor will go long if the fixed horizon return is forecasted to be positive and go short otherwise (see equation (1)). The trading rule based on a forecast of C_t is illustrated in Figure 1.

If the investor is in a 'neutral' position on day t and if near the close of day t, the forecasted value $\hat{C}_t$ of C_t is larger than $2c$, the investor will switch positions at the close of day t to attain a 'long' position on day $t+1$.

To compare the trading performance of a neural network with output C_t versus those with output p_{t+k}, we carry out the following empirical study. We apply neural network forecasting to the daily closing prices of the Hang Seng Index Futures Contract in Hong Kong. The index futures market in Hong Kong is a very active market. The dataset used is the daily closing prices of the Hang Seng Index Futures from 1 January 1987 to 31 December 1994. The first three quarters of the dataset is used for the training of the neural network and the last quarter is used for evaluation of the network's performance.

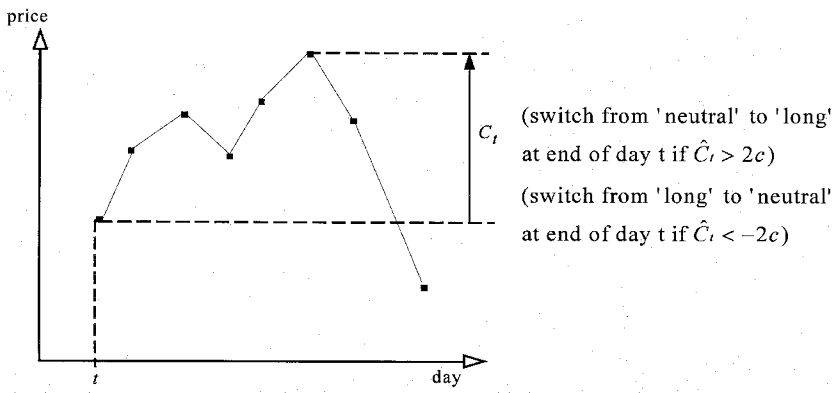

Figure 1.

The neural network used is a multi-layer perceptron network with one hidden layer having five units, a hyperbolic tangent transfer function, a sum of square error function and regulation training. Here, we only use a very simple network, the multi-layer perceptron network. The reason is that in Chen *et al.* (1995) this simple network is shown to be a universal approximator and can approximate any arbitrary function. We use the hyperbolic tangent transfer function because it can incorporate both positive and negative outputs. Also, it is found that a network with the hyperbolic tangent transfer function can be trained faster than one with a sigmoidal transfer function. The reason for using regulation training is that it is found to be very effective on the short and noisy dataset by Weigend *et al.* (1992). Our dataset is short and noisy so the regulation training is very suitable.

For inputs to the neural network, we use the historical one-, two-, three-, ten-, twenty- and thirty-day returns, i.e. $\ln p_t - \ln p_{t-i}$ with $i = 1, 2, 3, 10, 20$, and 30. We will use these inputs to forecast (1) C_t, the largest change before reversal and (2) $\ln p_{t+k} - \ln p_t$, the future return with a fixed forecasting horizon of $k = 5, 10, 20$, and 30 days.

To compare the trading performance of (1) and (2), we rely on two performance measures, the profit in index points and the investment return. When the networks are trained to generate buy/sell signals for the last quarter of the data period, two trading performance measures will be computed as follows. (1) We assume the investor to buy one contract on the first recommended buy signal, cover up on the next 'sell' signal, and buy another contract on the next 'buy' signal, etc. Here, no reinvestment of profit is assumed. Let the m buy/sell signals occur at days with closing indices $x_1, x_2, \ldots, x_m$, the profit in index points are, after taking transaction costs into consideration, calculated by

$$\begin{cases} (x_2 - x_1) + (x_4 - x_3) + \cdots + (x_m - x_{m-1}) - c(x_1 + \cdots + x_m) & \text{if } m \text{ is even} \\ (x_2 - x_1) + (x_4 - x_3) + \cdots + (x_{m-1} - x_{m-2}) - c(x_1 + \cdots + x_{m-1}) & \text{if } m \text{ is odd} \end{cases}$$

(2) To calculate the second performance measure called the investment return, note that the margin system in a futures market may affect the way how investment return can be computed. Here we follow Lukac, Brorsen and Irwin (1988) and assume no margining, i.e. the invested capital is the full amount for a contract and is not just the initial margin. Thus, in the first occurrence of a 'buy' signal, the investor goes long in one futures contract by putting up an amount V which is the worth of the contract. Assume that the 'buy' signal occurs when the index is x_1. If a 'sell' signal then occurs at index value x_2, the futures contract is sold. The capital which remains, assuming a round-trip transaction cost of 100(2c)%, is $V(\ln x_n - \ln x_1 - 2c)$. This amount of money is then fully invested in buying $(\ln x_2 - \ln x_1 - 2c)$ futures contract on the next 'buy' signal, without relying on the margining system. The accounting continues in this manner until the investor is left with an amount equal to V_f at the end. The investment return is then calculated by $\ln V_f - \ln V$. In this paper, a round-trip transaction cost of 0.2% for a futures contract is assumed.

In Table I, both performance measures are reported in the training period as well as in the testing period (see rows 1 and 3 in Table I). Since the training of the network involves a random start in the beginning, both performance measures will depend on the random start. Ten random starts are alternative and the average performance measures as well as their standard deviations (over 10 trials) are reported (see rows 2 and 4 in Table I).

It can easily be seen from Table I that for most cases using C_t as output gives rise to a more profitable trading strategy than that using fixed-period return as output. The only exception is

Table I. Comparison of trading performance with historical return as inputs

	Largest change before reversal as output		Future 5-day return as output		Future 10-day return as output		Future 20-day return as output		Future 30-day return as output	
	Training	Testing	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Average profit in index points	11264	2614	5233	1717	4809	787	4237	720	4723	340
SD of profit in index points	387	732	171	182	1039	658	233	613	1730	180
Average investment return (%)	333.3	32.0	193.9	34.9	126.5	13.3	170.1	21.2	169.5	17.4
SD of investment return (%)	12.3	7.8	7.6	1.9	21.7	6.5	8.0	4.9	46.4	30.4

Table II. *T*-statistics for a two-sample test

	C_t versus 5-day	C_t versus 10-day	C_t versus 20-day	C_t versus 30-day
Comparison of average profit	3.761*	5.870*	6.273*	9.540*
Comparison of average return	−1.142	5.824*	3.708*	1.471

that C_t loses out to future five-day-return when average return is used as a performance measure (32.0% versus 34.9%).

We can test the hypotheses

$$\begin{cases} H_0: & \text{average performance of } C_t \leqslant \text{average performance of } k\text{-day forecast} \\ H_1: & \text{average performance of } C_t > \text{average performance of } k\text{-day forecast} \end{cases} \quad (3)$$

The *t*-statistics for comparing the two independent samples are provided in Table II. Those values which are significant at the 0.01 levels are marked with *. It can be seen that the average performance of C_t is comparable to that of a five-day forecast and is significantly better than ten-, twenty-, and thirty-day forecasts.

To see whether the results are robust with respect to the choice of inputs, another set of inputs is tried. Let m_t be the 20 days moving average price at day t. Since technical traders often based their trading decisions on the comparison between m_t and p_t, we use as inputs to the neural network the following quantities:

$$\begin{array}{lll} \ln m_{t-2} - \ln p_{t-2} & \ln m_t - \ln p_t & \ln m_{t-1} \ln p_{t-1} \\ \ln m_{t-3} - \ln p_{t-3} & \ln m_{t-10} - \ln p_{t-10} & \ln m_{t-20} - \ln p_{t-20} \end{array}$$

Comparison of trading performance of (1) and (2) under this new set of inputs is reported in Table III.

The same observation holds that trading performance is higher for a trading strategy based on C_t than that based on a k-day return. In fact profits more than double when a trading strategy is

Table III. Comparison of trading performance with moving averages as inputs

	Largest change before reversal as output		Future 5-day return as output		Future 10-day return as output		Future 20-day return as output		Future 30-day return as output	
	Training	Testing	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Average profit in index points	10957	3978	10375	508	2487	1342	4036	590	6240	−479
SD of profit in index points	913	1788	340	820	372	1320	233	613	576	317
Average investment return (%)	344.4	45.8	308.7	16.3	68.6	18.7	119	11.2	177.7	−0.9
SD of investment return (%)	14.3	15.9	12.8	8.5	11.0	11.6	8.5	5.8	13.7	3.2

Table IV. Comparison of trading performance with moving averages as inputs

	C_t versus 5-day	C_t versus 10-day	C_t versus 20-day	C_t versus 30-day
Comparison of average profit	5.578*	3.751*	5.668*	7.762*
Comparison of average return	5.174*	4.354*	6.465*	9.105*

based on C_t. The intuitive explanation to this fact is that when the trading strategy is based on C_t, transaction cost plays a key role in market timing. On the other hand, the other approach devises a trading strategy with a fixed horizon in mind. However, in the course of investment decisions, the actual investment period may deviate from the horizon preassigned in the forecasting exercise. For example, the average holding period of a long/short position under the ten-day fixed horizon approach is 33.3 days for inputs $\ln p_t - \ln p_{t-n}$ and 29.4 days for inputs $\ln m_t - \ln p_{t-n}$, all quite different from the preassigned horizon of 10 days. This may be one of the reasons why the fixed horizon approach fails to perform well.

In testing hypothesis (3), the *t*-statistics for comparing the two independent samples are provided in Table IV. It can be seen that the performance of C_t is significantly better than those of five, ten, twenty and thirty days.

LARGEST CHANGE BEFORE REVERSAL VERSUS OTHER SUMMARY STATISTICS

Instead of forecasting the future return over a fixed horizon, we can follow Kimoto and Asakawa (1990) and choose to forecast the average future return. In other words, instead of using $\ln p_{t+k} - \ln p_t$ as output, we use

$$\frac{1}{k}\sum_{i=1}^{k}(\ln p_{t+i} - \ln p_t)$$

Table V. Comparison of trading performance with moving averages as outputs

	Largest change before reversal as output		Average return over 10 days		Average return over 20 days		Average return over 30 days	
	Training	Testing	Training	Testing	Training	Testing	Training	Testing
Average profit in index points	10957	3978	7581	1223	7713	376	4119	1643
SD of profit in index points	913	1788	1259	317	203	188	270	1248
Average investment return (%)	344.4	45.8	202.2	21.6	210.7	24.9	118.1	26.5
SD of investment return (%)	14.3	15.9	29.9	3.5	15	3.8	9.5	12.2

Table VI. Comparison of trading performance with moving averages as inputs

	C_t versus 10-day	C_t versus 20-day	C_t versus 30-day
Comparison of average profit	4.798*	6.336*	3.386*
Comparison of average return	4.700*	4.043*	3.045*

as an output of the neural network. Its trading performance can be compared with that of C_t and the results are reported in Table V. Here, we assume that the inputs used are the historical moving averages similar to those used to produce Table III.

The t-statistics for comparing the performance of the 'largest change before reversal' and the 'average return over k days' are presented in Table VI. C_t again performs significantly better than other forecasts.

DISCUSSION

The above empirical results show that the trading strategy based on neural forecast of the largest change before reversal, or the neural forecast of a fixed horizon future return can all derive positive profits on average. However, this does not mean that we can reject the weak-form efficient market hypothesis. The reason is that the naive strategy can also be very profitable. Using the naive buy-and-hold strategy over the whole testing period, one can earn a profit of 2820 index points and a return of 41.5%. In Table III, although the average profit (return) derived by the neural forecast of largest change before reversal with moving average inputs is 3978 index points (45.8%), the large standard deviation involved makes it difficult to confirm with high statistical significance that the proposed trading strategy actually beats the buy-and-hold strategy. Neural forecasting is usually not very stable as the training process may depend on the choice of a random start. Thus the degree of success may fluctuate from one training to another. Although it yields a positive profit on average, its profit on an individual trial can be negative. This also raises concern that the high average profit is derived at the expense of exposing the investor to higher risk. Because of the unstable performance of neural forecasting, it is not conclusive that the market is inefficient.

Although we may not be able to reject the weak-form market efficiency, the empirical findings in this paper show that the theoretical result derived in the theorem can be used to devise a successful trading rule which compares favourably with those based on fixed–horizon forecasts. We should note here that although we forecast the largest change before reversal using a neural network, the same approach can be used together with other forecasting techniques. To forecast the 'largest change before reversal' has the advantage that it advises the investor to trade only when the profit from doing so exceeds the cost, as pointed out by a referee of this paper. This concept is of course an obvious one. However, it has not been practised by technical analysts in formulating their trading rule. If a financial analyst only forecasts the price ten days from now, there is no way to find out whether the potential profit exceeds the cost in days other than the given time horizon. The empirical work carried out in this paper demonstrates that this intuitive idea is implementable and it can strengthen the applicability of any potential technique in financial forecasting.

SUMMARY AND CONCLUSION

In this paper we show that optimal trading results can be achieved if we can forecast a key summary statistic of future prices. The key summary statistic C_t, called the largest change before reversal, is defined. Assuming future prices are completely known, the problem of constructing an optimal investment decision in the presence of a transaction cost c is formulated as a mathematical programming problem. It is shown that the optimal solution for this programming problem is essentially a function of C_t. Thus the statistic C_t acts in some sense as a sufficient statistic for an optimal trading strategy with transaction cost. Hence it is natural to focus attention on the forecast of C_t if the purpose of forecasting is to obtain a good trading rule. Empirical results also support the use of C_t as output in a neural network when it is applied to the financial market for the generation of trading signals.

REFERENCES

Beja, A. and Goldman, M. B. 'On the dynamic behavior of prices in disequilibrium', *Journal of Finance*, **34** (1980), 235–47.

Bergerson, K. and Wunsch, D. C. II. 'A commodity trading model based on a neural network–expert system hybrid', *IJCNN91*, **1** (1991), 289–93.

Bessembinder, H. and Chan, K. 'The profitability of technical trading rules in the Asian stock markets', *Pacific Basic Finance Journal*, **3** (1995), 257–84.

Binks, D. L. and Allinson, N. M. 'Financial data recognition and prediction using neural networks', *Artificial Neural Networks*, **1** (1991), 1709–12.

Bjorn, V. 'Optimal multiresolution decomposition of financial time series', in Refenes, A. N. (ed.), *Neural Networks in the Capital Markets*, Chichester: Wiley, 1994.

Brock, W., Lakonishok, J. and LeBaron, B. 'Simple technical tracing rules and the stochastic properties of stock returns', *Journal of Finance*, **5** (1992), 1731–65.

Chauvin, Y. 'Trading Decision Learning: From Theory to Personal Traders', in Refenes, A. N. (ed.), *Neural Networks in the Capital Markets*, Chichester, U.S.A: John Wiley, 1994.

Chen, T. P., Chen, H. and Liu, R. W. 'Approximation capability in C(R) by multilayer feedforward networks and related problems', in Refenes, A. N. (ed.), *IEEE Transactions on Neural Networks*, **6** Chichester: Wiley, 1995, 25–30.

Corrado, C. J. and Lee, S. H. 'Filter rule tests of economic significance of serial dependencies in daily stock return', *Journal of Financial Research*, **15** (1992), 369–87.
Cutler, D. M., Poterba, J. M. and Summers, L. H. 'Speculative dynamics', *Review of Economic Studies*, **58** (1991), 529–46.
Dale, C. and Workman, R. 'Patterns of price movement in Treasury bill futures', *Journal of Economics and Business*, **33** (1981), 81–7.
Fama, E. F. 'Efficient capital markets: a review of theory and empirical work', *Journal of Finance*, **25** (1970), 383–417.
Fama, E. F. and French, K. R. 'Permanent and temporary components of stock prices', *Journal of Political Economy*, **98** (1986), 246–74.
Grossman, S. J. and Stiglitz, J. E. 'On the impossibility and informationally efficient markets', *American Economic Review*, **70** (1980), 393–408.
Hsu, W., Hsu, L. S. and Tenorio, M. F. 'Parameter significance estimation and financial prediction', *Neural Computing and Application*, **1** (1993), 280–6.
Jegadeesh, N. 'Evidence of predictable behavior of securities returns', *Journal of Finance*, **45** (1990), 881–98.
Kimoto, T. and Asakawa, K. 'Stock market prediction system with modular neural networks', *Proc IJCNN*, San Diego, CA: **1**, 1990, 1–6.
Lukac, L. P. and Brorsen, B. W. 'A comprehensive test of futures market disequilibrium', *Journal of Finance*, **25** (1990), 593–622.
Lukac, L. P., Brorsen, B. W. and Irwin, S. H. 'A comparison of twelve technical trading systems with market efficiency implications', *Station Bulletin* (1988), No. 495, Department of Agricultural Economics, Purdue University.
Margarita, S. 'Neural network, genetic algorithms and stock trading', *Artificial Neural Networks*, **1** (1991), 1763–6.
Mehta, M. 'Trading in foreign exchange markets using Kalman filter and neural network', in Refenes, A. N. (ed.), *Neural Networks in the Capital Markets*, Chichester: Wiley, 1994.
Moody, J. and Utans, J. 'Principled architecture selection for neural networks: application to corporate bond rating prediction', *Advances in Neural Information Processing Systems*, **4** (1992), 683–90.
Neftci, N. 'Naive trading rules in financial markets and Wiener–Kolmogorov prediction theory: a study of "technical analysis" ', *Journal of Business*, **64** (1991), 549–71.
Poterba, J. L. and Summers, L. 'Mean reversion in stock prices: evidence and implications', *Journal of Financial Economics*, **22** (1988), 27–59.
Refenes, A. N., Azema-Barac, M., Chan, L. and Karoussos, S. A. 'Currency exchange rate prediction and neural network design strategies', *Neural Computing and Applications*, **1** (1993), 46–58.
Taylor, S. J. 'Trading futures using a channel rule: a study of the predictive power of technical analysis with currency examples', *The Journal of Futures Market*, **14** No. 2 (1994), 215–35.
Taylor, S. J. and Tari, A. 'Further evidence against the efficiency of futures markets', in Guimaraes, R. M. C., Kingsman, B. G. and Taylor, S. J. (eds), *A Reappraisal of the Efficiency of Financial Markets*, New York: Springer-Verlag, 1989.
Trippi, R. and Turban, E. *Neural Network Applications in Investment and Financial Services*, New York: Probus Publishing, 1992.
Weigend, A. S., Huberman, B. A. and Rumelhart, D. E. 'Predicting sunspots and exchange rates with connectionist networks', *Nonlinear Modeling and Forecasting, SFI Studies in the Sciences of Complexity*, **7** (1992), 395–432.
Zapranis, A. D. and Refenes, A. N. 'Neural networks in tactical asset allocation: towards a methodology for hypothesis testing and confidence intervals', *Proceedings of Neural Networks in the Capital Markets*, 1994.
Zaremba, T. 'Case Study III: Technology in search of a buck', *Neural Network PC Tools*, (1990), 251–83.

Authors' biographies:
K. Lam obtained a PhD in Mathematics from the University of Wisconsin in 1972 and worked as a lecturer, senior lecturer and reader in the Department of Statistics in the University of Hong Kong. He took up the chair professorship in the Department of Finance and Decisions Sciences of the Hong Kong Baptist University in January 1995 and has headed the department since then.

Charles K. C. Lam is studying for an MPhil degree and was a studentship holder in the Department of Statistics in the University of Hong Kong from September 1994 to August 1996. He joined the Dao Heng Bank in Hong Kong where he has worked as an investment advisor since September 1996.

Authors' addresses:
K. Lam, Department of Finance and Decision Sciences, School of Business, Hong Kong Baptist University, 224 Waterloo Road, Kowloon, Hong Kong.

Charles K. C. Lam, Department of Statistics, The University of Hong Kong.

