Assessing Inefficiency in the Futures Markets

E. A. OLSZEWSKI*

INTRODUCTION

It is commonly believed that most individuals trading futures contracts lose money. Indeed, Babcock quotes statistics stating that the number of losers is between 75% and 90% (Babcock, 1989). An explanation is that there are a small number of knowledgeable traders who profit at the expense of the others. What type of analysis is employed by these few winners? Is it fundamental or technical analysis? Perhaps these successful traders have information about the markets, which is unavailable to the public, or know some secret about market behavior. In his book *Market wizards* Schwager has interviewed a number of highly successful traders (Schwager, 1989): As Schwager points out, the success of these traders is not due to a common approach such as fundamental or technical analysis, but rather to common trading attitudes and principles. One of the traders interviewed was Mr. Ed Seykota. In the interview Seykota stresses that the key to continued and prosperous trading is a technical system that includes sound money management techniques. Seykota then offers three primary components of his trading system in order of importance: (i) long-term trend, (ii) current chart pattern, (iii) a good price to buy or sell. He further adds that to reduce risk in a trade the selection of an entry point should be characterized by strong momentum in the direction of the trend.[1]

*For correspondence, Department of Physics, University of North Carolina at Wilmington, Wilmington, NC 28403-3297.

[1]The idea of a trend-following system with entry points based on market momentum has been exploited by a number of traders. A notable example is Mr. Kelly Angle, who is the principal of a successful trading advisory service.

■ *Edward Olszewski is a Professor in the Department of Physics at the University of North Carolina at Wilmington.*

The Journal of Futures Markets, Vol. 18, No. 6, 671–704 (1998)

CCC 0270-7314/98/060671-34

The primary purpose of this paper is to show that a technical trading system based on principles described by Mr. Seykota has the potential for being consistently profitable across a variety of markets. However, it is also shown that a judicious selection of which markets to trade can considerably improve the profitability of the trading system. The trading system is trend following, utilizing a mathematically rigorous momentum indicator to determine entry and exit points for trades. The rule for the selection of futures contracts to trade is based on a variant of the rescaled range analysis first described by Hurst (1951).

Drawing conclusions about market inefficiency based on a trading system is certainly not a new idea.[2] It is not uncommon to find tests of market inefficiency to be weak or lacking in rigor in some respect. This study is no exception. Attention will be drawn to such shortcomings as they occur in the subsequent analysis. Two potential weaknesses, which are especially relevant to any conclusions drawn about inefficiency, are whether the trading system can be implemented in practice and whether profits generated by the system "on paper" are realistic estimates of what would have been realized in actual trading. These two, in particular, are addressed in detail.

DATA AND METHODOLOGY

Data

The primary purpose is to show evidence for market inefficiency in the futures markets by presenting a trading system that generates a high rate of return. To help ensure that evidence supporting inefficiency is not an artifact of the market or the time frame chosen for investigation, ten different futures markets, each covering a relatively long period of time, have been selected. These futures markets are: cocoa, coffee, corn, cotton, crude oil, gold, the Japanese yen, live cattle, U.S. T-bonds, and the S&P 500 index. The data of each market represent 1767 trading days, beginning on 13 June 1986 and ending approximately on 30 June 1993. The ending date is approximate because the days on which various markets are closed may vary depending on the market. The data comprise daily high, low, and settlement prices of futures contracts for selected delivery months (see Table I).

For the ensuing analysis a time series covering the entire time period of each market is required. However, there is more than one futures con-

[2]In the discussion of inefficiency any effects due to risk premia have been neglected. The issue of whether a risk premium is priced into the value of a futures contract has not been unambiguously resolved. A discussion of risk premium in the context of futures contracts can be found in the following references: Chance, 1989; Bryant, 1995; Gencay, April 1995.

TABLE I

Futures Contracts and Delivery Months Used in Analysis

Contract	*Symbol*	*Delivery Months*
Cocoa	CC	H, K, N, U, Z
Coffee	KC	H, K, N, U, Z
Corn	CN	H, K, N, U, Z
Cotton	CT	H, K, N, V, Z
Crude Oil	CL	H, M, U, Z
Gold	GC	G, J, M, Q, V, Z
Japanese yen	JY	H, M, U, Z
Live cattle	LC	G, J, M, Q, V, Z
U.S. T-bonds	US	H, M, U, Z
S&P 500	SP	H, M, U, Z

Delivery Month Codes

Jan	*Feb*	*Mar*	*Apr*	*May*	*Jun*	*Jul*	*Aug*	*Sep*	*Oct*	*Nov*	*Dec*
F	G	H	J	K	M	N	Q	U	V	X	Z

tract trading at any given time and, furthermore, each of these contracts typically expires before the end of the time period. The procedure for constructing a single, continuous, bivariate time series, denoted (S_t, r_t), is now presented. First, for each futures contract of each market the following quantities are calculated,

$$s_t = \ln(P_t) - \ln(P_{t-1}), \tag{1}$$

and

$$r_t = \max \begin{cases} \ln(H_t) - \ln(L_{t-1}) \\ |\ln(H_t) - \ln(P_{t-1})| \\ |\ln(L_t) - \ln(P_{t-1})| \end{cases}, \tag{2}$$

where (H_t, L_t, P_t) are the high, low, and settlement prices on day t. A single bivariate time series is obtained by selecting consecutive pairs (s_t, r_t) from the contract, beginning with the earliest expiration date and continuing until the last trading day of the month prior to the expiration month of that contract. On that day pairs are then selected from the first forward contract. The procedure is continued in this manner until a single pair (s_t, r_t) is obtained for each trading day. This procedure provides a consistent way of selecting a single pair (s_t, r_t) out of all such pairs from the different futures contracts trading on a given day. Furthermore, the univariate series s_t approximates the daily percentage change in value of an actual trading position, which is continually rolled over from the near contract to the first forward contract. The date of rollover occurs on the last day of the month previous to the near contract's expiration month. The continuous bivariate time series (S_t, r_t) used in the analysis is then defined as follows:

$$S_t = \begin{cases} 100 & t = 0 \\ S_{t-1} \exp(s_t) & \text{otherwise}, \end{cases} \tag{3}$$

The series S_t is arbitrarily initialized so that $S_0 = 100$. The first component of the time series, S_t, is a synthetic price series reflecting percent change in the value of the futures contract. This series offers the advantage of being devoid of gaps in price, which typically occur when prices are rolled over from the near contract to the first forward contract. The second component, r_t, is a measure of daily market volatility and is similar to the true range described on page 21 of Wilder's book (Wilder, 1978). In Figure 1 we show an example of a continuous price series S_t so constructed.

Momentum

The trading strategy tested utilizes a momentum indicator for entering and exiting positions. The momentum indicator can be derived rigorously from theoretical considerations. Its derivation is now outlined here. Complete and concise details can be found in the appendix of a paper by Ingber (1984). The derivation begins with a set of Langevin equations, which are first-order, stochastic differential equations driven by Wiener noise.[3]

$$\dot{q}^{\alpha} = f^{\alpha}\,(t, q(t)) + \sigma_i^{\alpha}\,(t, q(t)\varepsilon^i\,(t)), \tag{4}$$

where $\{\alpha = 1, \ldots, \Lambda\}$ and $\{i = 1, \ldots, N\}$. Here $\varepsilon^i(t)$ are uncorrelated, Gaussian, white noise processes so that $\langle \varepsilon(t)^i \rangle_\varepsilon = 0$ and $\langle \varepsilon^i(t)\varepsilon^j(t')\rangle_\varepsilon = \delta^{ij}\delta(t - t')$. Derivatives are taken with respect to the time variable, t. The Langevin equations can be transformed to a Schrödinger equation whose solution is the conditional probability density of the $q^\alpha(t)$. The conditional probability density can be expressed as a path integral. From the expression for the path integral the Langrangian corresponding to the Langevin equations is obtained. The momentum, $p_\alpha(t)$ associated with each $q^\alpha(t)$ is given by

[3]Concerning conventions, unless otherwise stated the occurrence of repeated indices in a term of a mathematical expression indicates that a summation is performed over all values of the repeated index.

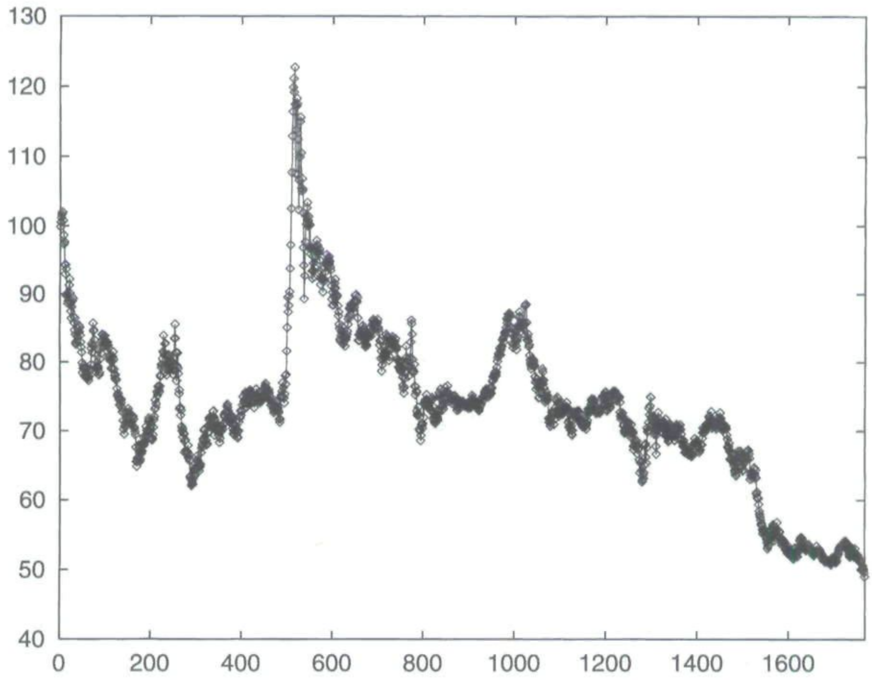

FIGURE 1
Continuation contract of daily closing prices for corn futures from 13 June 1986 until 11 June 1993. The initial value of the contract on 13 June has been arbitrarily set to 100. The contract contains 1767 data points.

$$p_\alpha(t) = \frac{\delta L(q_{\alpha,} \dot{q}_\alpha)}{\delta \dot{q}^\alpha} \tag{5}$$

For the univariate case, when

$$\dot{q} = f(t, q(t)) + \sigma(t, q(t))\varepsilon(t), \tag{6}$$

it can be shown that the Lagrangian is

$$L = \frac{1}{2}\frac{\Delta^2}{\sigma^2}. \tag{7}$$

The quantity Δ is

$$\Delta = \dot{q} - f(t, q(t)). \tag{8}$$

The momentum associated with $q(t)$ is obtained by differentiating Eq. 7

$$p(t) = \frac{\Delta}{\sigma^2}. \tag{9}$$

The interpretation of eq. 9 is straightforward. Momentum is a measure of how much the series $q(t)$ deviates from its conditional mean relative to its conditional variance.

The specific, univariate stochastic equation which is relevant to the subsequent analysis is

$$\dot{q} = \alpha(q(t) - m) - \alpha\sigma\varepsilon(t) \tag{10}$$

in the limit $\alpha \to \infty$. In its discrete form Eq. 10 is

$$q_t - \mu = \sigma\varepsilon_t, \tag{11}$$

where t is a non-negative integer. Here m, σ, μ are free parameters. The momentum corresponding to q_t is obtained from eqs. 9 and 11,

$$p_t = \frac{q_t - \mu}{\sigma^2}. \tag{12}$$

Long-Range Memory Processes

The hydrologist H. E. Hurst pioneered the study of long-range memory processes in his investigation of the storage capacity of reservoirs (Peters, 1994; Hurst, 1951). Extending Hurst's work Mandelbrot and others introduced a family of Gaussian random functions, designated fractional Brownian motions (fBm's), which exhibit long-term persistence similar to those studied by Hurst (Mandelbrot and Van Ness, 1960; Mandelbrot and Wallis, 1969a, b, c; Beran, 1994). A fractional Brownian motion, $B_H(t)$, of exponent H is defined as a moving average process in which past increments, $dB(s)$, are weighted by $(t - s)^{H-.5}$. Specifically,

$$B_H(t) = \frac{1}{\Gamma\left(H + \frac{1}{2}\right)} \int_0^t (t - s)^{H-.5}\, dB(s), \tag{13}$$

where $0 < H < 1$. If $H = 0.5$, Eq. 13 reduces to Brownian motion. FBm's possess a number of interesting properties. The spectral densities of fBm's are proportional to ω^{1-2H}, ω being the frequency. They exhibit the following scaling behavior in the variance $\mathrm{VAR}(B_H(t)) = t^{2H}\,\mathrm{VAR}(B_H(1))$. As detailed in Mandelbrot and Wallis (1969a) for $H > 0.5$ the future and past are positively correlated, with correlations approaching 1 as H approaches 1. These fBm's exhibit a long-range memory characterized by a persistance in positive correlation between events increasingly separated in time. Consequently, their autocorrelations decay much more slowly than short-range memory processes such as autoregressive or moving average. For $H = 0.5$ the future and past are uncorrelated. For $H < 0.5$ the future and past are negatively correlated, with correlations decreasing to -0.5. These fBm's are characterized by antipersistence, i.e., correlations are all negative. Such series are subject to more reversals than Brownian motion.

Detecting long-range memory dependence in time series, such as is found in fBm's, can be complicated by the presence of short-range memory dependence. Indeed, it has been pointed out by Lo that the classical *R/S* analysis used to identify long-range memory dependence can falsely indicate its presence in a time series when, in fact, only short-range memory dependence is present. Lo has derived a test statistic for detecting long-range memory dependence similar to the classical *R/S* statistic, R_n (Lo, 1991). It is less prone to the misspecification of classical *R/S* analysis, while also exhibiting power against rejecting long-range memory dependence when it is present. The test statistic takes account of short-range memory dependence by adjusting the variance used in R_n. Specifically, given n consecutive points from a discrete time series $\{x_t\}$ the test statistic Q_n is defined as[4]

$$Q_n \equiv \frac{r_n}{\sigma_n(q)}, \tag{14}$$

where

$$r_n = \max_{1\le k\le n} (x_1 + \ldots + x_k - k\bar{x}) - \min_{1\le k\le n} (x_1 + \ldots + x_k - k\bar{x}), \tag{15}$$

$$\bar{x} = \frac{1}{n}\sum_{k=1}^{n} x_k, \tag{16}$$

$$\sigma_n^2(q) = s_n^2 + \frac{2}{n}\sum_{j=1}^{q} \omega_j(q)\left\{\sum_{i=j+1}^{n} (x_i - \bar{x})(x_{i-j} - \bar{x})\right\}, \tag{17}$$

$$s_n = \sqrt{\frac{\sum_{k=1}^{n} (x_k - \bar{x})^2}{n}}, \tag{18}$$

and

[4]The x_t correspond to the s_t given in eq. 1.

TABLE II

Fractiles of the Limiting Distribution of the V Statistic under the Assumption of No Long-Range Memory

Prov($V < v$)	0.005	0.0250	0.050	0.100	0.200	0.300	0.400	0.500
v	0.721	0.809	0.861	0.927	1.018	1.090	1.157	1.223
Prob ($V < v$)	0.543	0.600	0.700	0.800	0.900	0.950	0.975	0.995
v	$\sqrt{\pi/2}$	1.294	1.374	1.473	1.620	1.747	1.862	2.098

$$\omega_j(q) = 1 - \frac{j}{q + 1}. \tag{19}$$

The value of the truncation lag, q, depends on the data being considered. Its value must be large enough to account for short-range memory dependence in the data, but not so large as to alter the finite sample distribution of Q_n radically. Andrews suggests the following data dependent rule for selecting q (Lo, 1991; Andrews, 1991):

$$q = [\kappa_n],\ \kappa_n = \left(\frac{3n}{2}\right)^{1/3} \left(\frac{2\hat{\rho}}{1 - \hat{\rho}^2}\right)^{2/3}, \tag{20}$$

where $[k_n]$ is the greatest integer less than or equal to k_n and $\hat{\rho}$ is the estimated first-order correlation of the data. Also, the weights $\omega_j\ (q)$ in Eq. 19 are replaced by

$$\omega_j(q) = 1 - \left|\frac{j}{k_n}\right|. \tag{21}$$

Lo defines the statistic $V_n(q)$, which is based on the test statistic Q_n,

$$V_n(q) \equiv \frac{Q_n}{\sqrt{n}}. \tag{22}$$

This statistic is identical to the statistic V_{n_j} used by Hurst and Peters, except that the variance has been adjusted to account for short-range memory effects (Peters, 1994). Lo derives the limiting distribution of V_n for which $V_n(q)$ is an estimator, under the assumption of no long-range memory dependence and a general set of conditions which include strong mixing (short-range memory) and conditional heteroskedasticity. In Table II the fractiles of the limiting distribution of the V statistic are presented. Small values of V correspond to antipersistence, and large values correspond to persistence. If $V_n(q)$ is less than 0.861, the hypothesis of antipersistence at the 5% confidence level would be accepted. Similarly, if $V_n(q)$ is greater than 1.747, the hypothesis of persistence at the 5% confidence level would be accepted.

TRADING SYSTEM

A trading system consistent with the following paradigm for market behavior is proposed. Market prices may be trending or nontrending, i.e., fluctuating about some equilibrium value. An event occurs that causes a trend change that persists for an extended period of time. The event is characterized by a price breakout in the direction of the trend, i.e., market momentum is strong. There are a variety of indicators traders use to measure market momentum, e.g., today's price less the price x days ago. A large positive/negative value suggests that there is strong force driving prices to increase or decrease. One problem with this type of indicator is that the value of momentum is not normalized with respect to a particular reference. In this case a simple solution may be to divide the difference by the price x days ago and express the result as a percentage. This would permit momentum comparisons to be made between different price levels or even different markets. There have been proposed a number of other normalized momentum indicators, such as Wilder's RSI, Lane's Stochastics, or Lambert's CCI. They are discussed in some detail on pages 282–283 of the book by Murphy (1986). The momentum construct given in Eq. 9 is also a normalized measure of momentum and is a principal component of the trading strategy used in the subsequent analysis. This construct measures the strength of momentum by subtracting the conditional mean from the price and dividing the result by the conditional standard variance. The construct is intuitatively appealing because momentum strength is based not only on how much prices deviate from what is expected but also on how much they deviate in comparison with current volatility. The drawback of this construct is that one needs estimates of the conditional mean, i.e., what is expected, and also the conditional variance, i.e., current volatility, to calculate its value.

The paradigm for market behavior is consistent with the assumption that market prices exhibit long-range memory dependence. Hurst, in his analysis of a biased random walk, has uncovered what has been referred to as the joker effect (This is discussed in some detail on pages 60 and 61 of Peters' (1994) book.) This is similar to the hypothesis that a significant event occurs that causes a change in trend of a market. The event corresponds to the joker. Indeed, in this same context Mandelbrot and Wallis (1969a) have pointed out a distinctive feature of fBm's with $H > 0.5$. They spend long periods of time either above or below their mean.

Mandelbrot likens this to a "dry spell" (a period of time when no rain falls) or a "wet spell" (a period of time when there is excessive precipitation). The statistic $V_n(q)$, is used for detecting long-range memory dependence in the markets under investigation.

Momentum-Based Trading System

To operationalize this paradigm the following model for market prices (eq. 11) is adopted

$$q_i = \mu + \sigma \varepsilon_i. \tag{23}$$

Here $q_i = \ln(S_i)$, where S_i is given by eq. 3, and $i = t, t - 1, t - 2, \ldots, t - (n_d - 1)$.[5] The quantity μ is estimated as

$$\mu = \frac{\sum_{i=t}^{t-(n_d-1)} q_i}{n_d}. \tag{24}$$

If it is assumed that during any given period t the series follows a random walk about its mean, then σ_i may be estimated using the result of Parkinson (1980),

$$\sigma_i = \left(\frac{\pi}{8}\right)^{1/2} E[r_i] = \left(\frac{\pi}{8}\right)^{1/2} r_i. \tag{25}$$

Here r_i is given in eq. 2, $E[\ \ ]$ is expected value, and $E[r_i]$ has been approximated by its best estimate r_i. σ is estimated by averaging over all σ_i

$$\sigma = \frac{\sum_{i=t}^{t-(n_d-1)} \sigma_i}{n_d}. \tag{26}$$

The momentum p_i is given by eq. 12. A proxy for the momentum that is dimensionless is used in the subsequent analysis. It is given by

$$\hat{p}_i = p_i \sigma = \frac{q_i - \mu}{\sigma}. \tag{27}$$

[5]It should be noted that the logarithm of price, q_i is used rather than price, S_i.

The momentum-based component of the trading system can now defined. There are five adjustable parameters: n_d, n_l, n_s, b_l, b_s subject to the constraint $n_s \leq n_l < n_d$. There are two average momentum variables,

$$p_t^l = \frac{\sum_{i=t}^{n_l-1} \hat{p}_i}{n_l} \tag{28}$$

and

$$p_t^s = \frac{\sum_{i=t}^{n_s-1} \hat{p}_i}{n_s}. \tag{29}$$

The purpose of p_t^l is to measure long-term momentum and p_t^s to measure short-term momentum. The underlying philosophy of the trading system is that whenever both short-term and long-term momenta become large in the positive or negative sense then a position is taken in the direction of the momentum. The momentum-based trading rules are as follows.[6] At the close of day t

1. If no position is held in the market, and
   a. if $p_t^l > b_t^l$ and $p_t^s > b_t^s$, buy one contract, or
   b. if $p_t^l < -b_t^l$ and $p_t^s < -b_t^s$, sell one contract;
   c. otherwise, do nothing.
2. If a long position is held in the market, and
   a. if either $p_t^l < -b_t^l$ or $p_t^s < -b_t^s$, close out the position, or
   b. if either $p_t^l < -b_t^l$ and $p_t^s < -b_t^s$, close out the position, and sell one contract;
   c. otherwise, do nothing.
3. If a short position is held in the market, and
   a. if either $p_t^l > b_t^l$ or $p_t^s > b_t^s$, close out the position, or
   b. if $p_t^l > b_t^l$ and $p_t^s > b_t^s$, close out the position and buy one contract;
   c. otherwise, do nothing.

The trading system is applied to a sample of data beginning on day t for a number of days n_{out}, and its profitability is evaluated. The five adjustable parameters of the trading system are determined by first optimizing the profits generated by the system in a sample of data of size

[6]These trading rules are similar, in some respects, to those in systems using so called volatility envelopes as discussed on pages 139–141 of Babcock's book (Babcock, 1989).

TABLE III

Parameters to Be Optimized and Restrictions on Their Search Values

Parameter	*Restrictions*
n_d (Eq. 23)	22, 33, 44, 55, 66
n_l (Eq. 28)	1, 2, 3, . . . 20
n_s (Eq. 29)	1, 2, 3, . . . 20
b_l (see trading rules)	0, 1 × 2.88, 2 × 2.88, 3 × 2.88, . . . 19 × 2.88
b_s (see trading rules)	0, 1 × 2.88, × 2 × 2.88, 3 × 2.88, . . . 19 × 2.88

n_{in} ending on day $t - 1$. The procedure is then reapplied with day t being replaced by day $t + n_{out}$. The procedure is followed until the trading system has been applied to the entire data set under consideration. At the end of each trading period all trades are closed out. This ensures that profits generated in different periods are independent of each other. One benefit of this methodology is that profitability of the trading system is assessed for a number of different, nonoverlapping samples, each with its own optimized parameters. This should help in avoiding the chance selection of a profitable set of adjustable parameters. Another benefit is that the adjustable parameters are allowed to evolve in time reflecting possible changes in market conditions.

Because of practical considerations, some of the parameters of the trading system have been assigned in an ad hoc manner. First, the size of each in-sample, n_{in}, used for optimization has been set to 250 days (approximately one year). The size of each out-of-sample used for evaluation has been set to 66 days (approximately one quarter). For each market this results in 22 independent samples for evaluating profitability. The 22 samples in each data set consist of points [316,381], [382,447], [448,513], . . . [1702,1767]. Points [1,66] are used for initializing average momenta, and points [66,315] are the first set of points used for optimizing the adjustable parameters which are applied in the first out-of-sample [316,381].

The calculations for performing optimizations in the ten markets are computationally intensive. To make the computations manageable the search over values of the adjustable parameters which optimize the profits of the trading system have been restricted.[7] The adjustable parameters and the restrictions placed on their search values are listed in Table III. If, in a given search, several different sets of parameters yield the same

[7] A sufficient number of points, the maximum of n_d (see Table III), from the beginning of the series are excluded from the analysis and used for obtaining initial values of average momenta.

TABLE IV

Futures Contracts and Delivery Months Used in Analysis

Contract	*Profit*	*Profit/Trade*	*Max. Drawdown*	*No. of Trades*
Cocoa	−4510	−92.04	−5580	49
Coffee	4736.25	87.71	−16492.5	54
Corn	10850	230.85	−2125	47
Cotton	36985	1120.76	−2185	33
Crude oil	−6960	−154.67	−20380	45
Gold	2120	40.77	−8650	52
Japanese yen	14975	348.26	−7637.5	43
Live cattle	−5401	−90.02	−8228	60
U.S. T-bonds	−7406.7	−142.44	−27656.4	52
S&P 500	−33550	−729.35	−75125	62
Total	11838.60	24.61		481

optimum profits, then the set of parameters to be used for out-of-sample testing is selected by applying each of the following criteria, successively, until only one set of parameters remains: select the set with the smallest value of (1) n_d, (2) n_l, (3) n_s, (4) b_l, (5) b_s.

Application of the trading system results in profits that are inconsistent across the ten markets. In Table IV the relevant trading results for each market are summarized.[8] It is noteworthy that the system does not generate large numbers of buy/sell signals per market, typically one or two trades on average per market per quarter. It is also noteworthy that in only three markets has the system shown respectable profits. The best results have been obtained in cotton, followed by the yen and corn. The worst results have been obtained in the S&P 500. With the exception of corn and cotton the maximum drawdown would probably be unacceptable to most traders. This is not surprising, because we have made no attempt to incorporate into the trading system a strategy for realistically managing risk. For all markets taken together the average profit per trade is only about \$25.00. This is probably not enough even to cover the cost of commissions and slippage, for which a fair value would be about \$50.00. This value for transaction costs does not include some of the trades that would, in practice, need to be rolled over to avoid trading in the expiration month.[9]

[8]Approximately six days of computing time have been required to perform the optimizations on the ten markets. Computations have been performed by a dedicated HP 735/125 workstation operating at approximately 55 mflops. More than two weeks have been required for similar computations using simulation data, described later.

[9]These results take account of limit moves which make it impossible to fill a position. When these

Modified *R/S* Test

The trading results obtained using the momentum-based trading system can be improved if trades are restricted according to the markets that show the most profit potential. Of use in this context is the *R/S* statistic $V_n(q)$ given in Eq. 22. Specifically, it is shown that some markets appear to be persistent, which may be attributed to their trending behavior, whereas others appear to be antipersistent, which may attributed their countertrending behavior. The statistic $V_n(q)$ has been calculated prior to each of the 22 out-of-sample trading periods using eqs. 20, 21, and 22. For each calculation the sample size has been taken to be $n = 1315$ (approximately five years). This has necessitated prepending an additional 1000 points to each data set. The decision to select this value of the sample size is based in part on Lo's study of the finite sample properties of $V_n(q)$ in which he has simulated data possessing various long-range or short-range dependencies (Lo, September 1991). In these studies Lo has used sample sizes ($n = 100, 250, 500, 750, 1000$). Lo has found that for $n = 1000$ asymptotic properties of the V statistic are reasonably well approximated.[10]

In Tables V and VI $V_{1315}(q)$ for all markets is presented.[11] According to Table II large values of V, ($V > 1.223$) correspond to persistence and small values, i.e. ($V < 1.223$), to antipersistence. One noteworthy feature of the V statistic, evidenced in Tables V and VI, is that, for the most part, its value changes slowly from period to period, rather than jumping about wildly. It is relatively stable. The implication is that a market that exhibits persistence is likely to continue in this manner, and similarly for antipersistence. Thus the V statistic may be useful in determining whether a trend-following or countertrend-following strategy is appropriate for trading a given market. Another feature gleaned from the tables is that markets that have been initially persistent, in general, have grown less persistent over time. According to the table of fractiles for the V statistic, some markets have, for some sample periods, even exhibited persistence or antipersistence, significant at the 5% level. Also apparent from the tables is that some markets have tended to exhibit persistence during the entire trading period, e.g., cotton, Japanese yen, and corn, whereas others have been consistently antipersistent, e.g., S&P 500 and cocoa. This pro-

---

have occurred, trades have been assumed filled on the open of the first day when trading occurs after the limit move. There are three such instances: two in crude oil and one in live cattle. The additional loss resulting from the limit moves in crude oil is about $3700 and in live cattle about $400.

[10]A more realistic estimate of transaction costs would be $75.00, because, on average, every trade actually corresponds to about 1.3 trades when rollovers are included.

[11]1ther than this motivation, the selection of $n = 1315$ has been an ad hoc decision.

TABLE V

$V_n(q)$ Statistic Prior to Each Out-of-Sample Trading Period

	Contract				
Period	*Cocoa*	*Corn*	*Coffee*	*Cotton*	*Crude Oil*
315	0.915	1.537	1.731	1.928	1.610
381	1.008	1.676	1.726	1.895	1.557
447	0.956	1.625	1.783	1.878	1.533
513	1.032	1.841	1.776	1.852	1.504
579	1.136	1.397	1.660	1.924	1.436
645	1.235	1.454	1.640	1.812	1.403
711	1.056	1.470	1.578	1.761	1.401
777	1.014	1.459	1.355	1.740	1.388
843	0.906	1.434	1.262	1.698	1.392
909	0.831	1.428	1.296	1.720	1.386
975	0.970	1.369	1.244	1.669	1.374
1041	0.998	1.362	1.254	1.628	1.351
1107	0.962	1.344	1.224	1.579	1.257
1173	0.933	1.366	1.190	1.539	1.517
1239	1.017	1.373	1.210	1.582	0.984
1305	0.904	1.355	1.111	1.605	1.026
1371	0.919	1.300	1.219	1.388	1.157
1437	0.922	1.309	1.151	1.474	1.190
1503	1.065	1.302	1.135	1.459	1.218
1569	0.941	1.440	1.219	1.653	1.283
1635	0.976	1.243	1.565	1.570	1.355
1701	0.979	1.201	1.083	1.465	1.292

vides an explanation why the momentum-based trading system has been profitable in the Cotton, Japanese yen, and Corn markets and unprofitable in S&P 500 and Cocoa markets. For these latter two markets, it appears that a countertrend following system would have been more appropriate.

Tables VII and VIII show trading results using the V statistic to filter trades. Specifically, for a given market, trades are initiated only when the value of V immediately prior to the trading period is greater than or equal to the value in the table. The values of V range from 0.7 (extreme anti-persistence to 1.70 (high persistence).[12] The profit per trade as a function of V improves for most of the markets. Although for values of V approximately 1.65 or greater, the profit per trade decreases, in general. In fact, for such values of V the trading system becomes extremely unprofitable in the case of coffee. It is doubtful that this failure can be explained

[12] In the case of crude oil there were insufficient additional data to calculate $V_{1315}(q)$. Instead, the the values of $V_n(q)$ corresponding to periods 315, 381, and 447 have been obtained using (n = 1119, 1185, 1251), respectively.

TABLE VI

$V_n\ (q)$ Statistic Prior to Each Out-of-Sample Trading Period

	Contract				
Period	*Gold*	*Japanese Yen*	*Live Cattle*	*U.S. T-Bonds*	*S&P 500*
315	1.578	1.806	1.633	1.486	1.189
381	1.642	2.024	1.708	1.612	0.956
447	1.514	1.984	1.802	1.672	0.961
513	1.518	1.957	1.618	1.717	0.957
579	1.657	2.018	1.750	1.655	0.952
645	1.532	1.914	1.774	1.653	0.954
711	1.542	2.013	1.630	1.647	0.956
777	1.623	2.164	1.631	1.438	0.966
843	1.510	2.003	1.670	1.446	0.973
909	1.292	2.021	1.608	1.332	0.959
975	1.287	1.997	1.602	1.403	0.954
1041	1.269	2.003	1.470	1.270	0.954
1107	1.182	2.124	1.277	1.177	0.923
1173	1.190	1.864	1.353	1.203	0.921
1239	1.138	1.579	1.325	1.024	0.920
1305	1.118	1.600	1.304	0.986	0.922
1371	0.970	1.508	1.249	1.014	0.927
1437	1.025	1.626	1.347	1.134	0.938
1503	0.926	1.528	1.052	1.199	0.935
1569	0.817	1.516	0.909	0.892	0.946
1635	0.806	1.518	0.936	0.918	0.896
1701	0.770	1.441	1.345	1.127	0.982

because of the small number of trades. This is contrary to what theory dictates. Furthermore, in the case of gold and live cattle the V statistic is useless. Specifically, profits per trade are less when the V statistic is used. Table IX shows the total profits per trade when the V statistic is used to filter trades. Profits per trade generally improve with increasing V, with the exception for high values of V as noted previously. When the V statistic is not used to filter trades, the profit per trade is approximately \$25. When trades are selected using a value of V such that $V \geq 1.25$, the value that separates persistence from antipersistence, profit per trade increases to an amount \$200 or better.

MODELING AND SIMULATION

In this section the profitability of the trading system utilizing the R/S statistic to filter trades is more rigorously tested. Because the test is computationally intensive, it is restricted to only a single market, corn. The reason for selecting corn is that it appears to have exhibited strong per-

TABLE VII

V Statistic and Trading Profits/Trade

	Contract				
V	*Cocoa*	*Coffee*	*Corn*	*Cotton*	*Crude Oil*
0.7	−92.041 (49)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
0.75	−92.041 (49)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
0.8	−92.041 (49)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
0.85	−164.22 (45)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
0.9	−164.22 (45)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
0.95	−87.727 (22)	87.708 (54)	230.85 (47)	1120.8 (33)	−154.67 (45)
1	415 (10)	87.708 (54)	230.85 (47)	1120.8 (33)	−156.98 (43)
1.05	−91.429 (7)	87.708 (54)	230.85 (47)	1120.8 (33)	174.29 (42)
1.1	236.67 (3)	72.332 (52)	230.85 (47)	1120.8 (33)	−174.29 (42)
1.15	865 (2)	130.27 (46)	230.85 (47)	1120.8 (33)	−174.29 (42)
1.2	865 (2)	169.48 (41)	230.85 (47)	1120.8 (33)	−257.69 (39)
1.25	0 (0)	113.07 (33)	281.6 (36)	1120.8 (33)	−250.79 (38)
1.3	0 (0)	400.18 (21)	281.6 (36)	1120.8 (33)	46.18 (34)
1.35	0 (0)	400.18 (21)	294.17 (30)	1120.8 (33)	46.18 (34)
1.4	0 (0)	425.92 (19)	367.05 (22)	1120.8 (33)	245 (24)
1.45	0 (0)	425.92 (19)	367.11 (19)	1120.8 (33)	609.33 (15)
1.5	0 (0)	425.92 (19)	635.42 (12)	1209.3 (29)	609.33 (15)
1.55	0 (0)	425.92 (19)	846.25 (10)	1263.1 (27)	202.86 (7)
1.6	0 (0)	−160.37 (17)	846.25 (10)	1060.3 (19)	770 (2)
1.65	0 (0)	−618.46 (13)	358.33 (6)	1650 (12)	0 (0)
1.7	0 (0)	−1183.1 (12)	459.38 (4)	1295.7 (7)	0 (0)

Profits per trade resulting when the *V* statistic obtained from data prior to the trading period is greater than or equal to the value of the *V* statistic in the table.

sistence, which is presumably why the trading system has been profitable. The procedure involves first constructing for the corn data a model that does not specifically take account of long-range memory dependence. Simulation data are derived from the model. Then the trading system is applied to the simulation data. This permits the following two hypotheses to be tested: (i) that the trading system is profitable when long-range memory dependence is absent from the data (the model for corn has been constructed to exclude long-range memory dependence explicitly.); (ii) that the profits generated when the trading system is applied to actual data are statistically greater than those when it is applied to simulation data. If the second hypothesis is true, then there is strong statistical evidence that judicious trade selection based on a market's propensity for long-range memory dependence would have significantly improved profits of the momentum-based trading system (at least in the case of corn). In order to test these hypotheses adequately the model must take account of a variety of systematic and nonstationary effects in the corn data, such

TABLE VIII

V Statistic and Trading Profits/Trade

	Contract				
V	*Gold*	*Japanese Yen*	*Live Cattle*	*U.S. T-Bonds*	*S&P 500*
0.7	40.7692 (52)	348.256 (43)	−90.02 (60)	−142.437 (52)	−729.348 (46)
0.75	40.7692 (52)	348.256 (43)	−90.02 (60)	−142.437 (52)	−729.348 (46)
0.8	−6.6 (50)	348.256 (43)	−90.02 (60)	−142.437 (52)	−729.348 (46)
0.85	−33.6364 (44)	348.256 (43)	−90.02 (60)	−142.437 (52)	−729.348 (46)
0.9	−33.6364 (44)	348.256 (43)	−90.02 (60)	−142.437 (52)	−525.595 (42)
0.95	18.8095 (42)	348.256 (43)	−81.89 (54)	−240.008 (50)	1858.33 (18)
1	14.1026 (39)	348.256 (43)	−81.89 (54)	−240.008 (50)	30225 (1)
1.05	10 (38)	348.256 (43)	−81.89 (54)	−148.947 (47)	30225 (1)
1.1	10 (38)	348.256 (43)	−103.53 (51)	−148.947 (47)	30225 (1)
1.15	−21.9444 (36)	348.256 (43)	−103.53 (51)	32.7659 (41)	30225 (1)
1.2	−25.2941 (34)	348.256 (43)	−103.53 (51)	34.0848 (33)	0 (0)
1.25	−25.2941 (34)	348.256 (43)	−180.4 (45)	176.403 (31)	0 (0)
1.3	22.5 (28)	348.256 (43)	−180.4 (45)	246.519 (27)	0 (0)
1.35	22.5 (28)	348.256 (43)	−143 (30)	246.519 (27)	0 (0)
1.4	22.5 (28)	348.256 (43)	−109.24 (29)	246.519 (27)	0 (0)
1.45	22.5 (28)	348.256 (43)	−109.24 (29)	579.536 (22)	0 (0)
1.5	22.5 (28)	348.256 (43)	−106.33 (27)	604.162 (21)	0 (0)
1.55	−413.333 (15)	383.929 (35)	−106.33 (27)	604.162 (21)	0 (0)
1.6	126 (5)	370.703 (32)	−106.33 (27)	604.162 (21)	0 (0)
1.65	−5 (4)	351.21 (31)	155.833 (12)	335.942 (12)	0 (0)
1.7	0 (0)	351.21 (31)	431.75 (4)	343.725 (4)	0 (0)

Profits per trade resulting when the *V* statistic obtained from data prior to the trading period is greater than or equal to the value of the *V* statistic in the table.

as the drought of 1988, seasonal variations in market volatility, changes in volatility over time, and gaps in the data, such as Friday to Monday. In addition, the model should take account of conditional heterogeneity, which is known to be present in long financial time series, and deviations from normality, which may be exhibited by the process underlying the data. The model, in principle, should include all effects other than long-range memory dependence. Thus, simulation data derived from such a model could then be used to test various hypotheses in which no long-range memory dependence is assumed to be present.

For the subsequent analysis the following bivariate time series is defined for the corn data,

$$\hat{y}_t' = (\hat{y}_t^1, \hat{y}_t^2) = (100s_t, \ln(r_t)), \tag{30}$$

where s_t and r_t are given by eqs. 1 and 2. The volatility, r_t has been logarithmically transformed to eliminate its extreme skewness. For conven-

TABLE IX

Total profits per Total Number of Trades

V	*Profit/Trade*	*Number of Trades*
0.70	24.61	481
0.75	24.61	481
0.80	19.60	479
0.85	11.42	469
0.90	36.20	465
0.95	187.44	408
1.00	212.03	374
1.05	215.22	366
1.10	219.29	357
1.15	259.39	342
1.20	174.69	323
1.25	180.17	293
1.30	266.81	267
1.35	299.82	246
1.40	347.22	225
1.45	415.28	208
1.50	439.94	194
1.55	429.40	161
1.60	373.91	133
1.65	340.87	90
1.70	172.57	62

The second column is the total profit for all markets divided by the total number of trades under the condition that trades are initiated only if the V statistic for a given market is greater than or equal to the value given in column 1.

ience in this Section $\hat{y}_t^1$ is referred to as the price differences and $\hat{y}_t^2$ as the volatility.

Data Adjustments

First, the data are analyzed for systematic or nonstationary effects that cannot be captured by the model. These effects are then removed using regression. The goal is to remove only those effects for which there is some reasonable basis for justification and not to engage in a witch hunt.

In Figures 2 and 3 the futures prices and volatility for corn as a function of time, t, are shown. Both graphs show a considerable degree of inhomogeneity in the data. The graph of price differences appears to exhibit a cyclical pattern in its volatility. This is confirmed in the graph of volatility as a cyclical pattern in level. In the plot of price differences there appear to be extreme values between $500 < t < 600$. This also occurs in the volatility plot. Also, it is appears that the variance of the data is compressed from left to right in the graph of price differences. This is confirmed in the volatility plot as its level declines gradually over

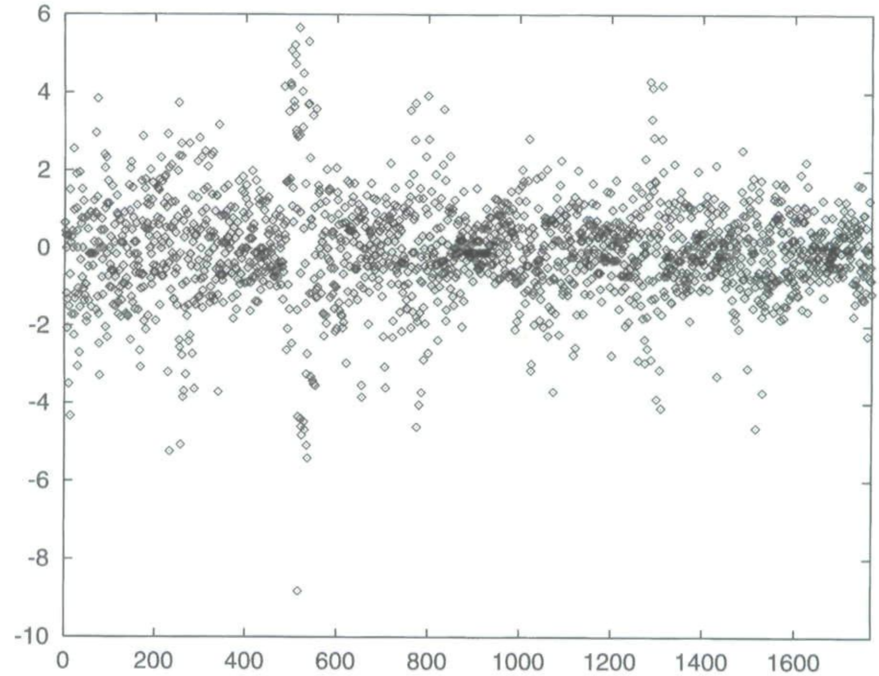

FIGURE 2
Unadjusted price differences for corn. The data have been obtained by differencing the logarithm of corn prices, S_t.

time. A short-term model may have difficulty capturing this long-term decline in volatility.

There are several reasons which may account for some of these inhomogeneities. The data contain gaps. A gap is the number of days between two successive trading days when the market is closed. Specifically, if trading has not occurred on the previous two calendar days immediately previous to the datum at t, that datum is assigned a gap of two. For example, a Monday on which trading occurs is assigned a gap of two if the market has been open on the previous Friday. Gaps range from zero to four. In Table X the gaps and their numbers of occurrences are listed. The data also exhibit cyclical variations in volatility. From a more detailed analysis it is found that, in general, volatility begins relatively low in January, peaks in midsummer, and then declines into December. In the summer of 1988 there was a drought. This corresponds to data points $500 < t < 600$ on the graphs.

The adjustment procedure involves regressing each univariate data series on variables characterizing each effect. This is referred to as a mean adjustment. Then the squares of the residuals from the first regressions

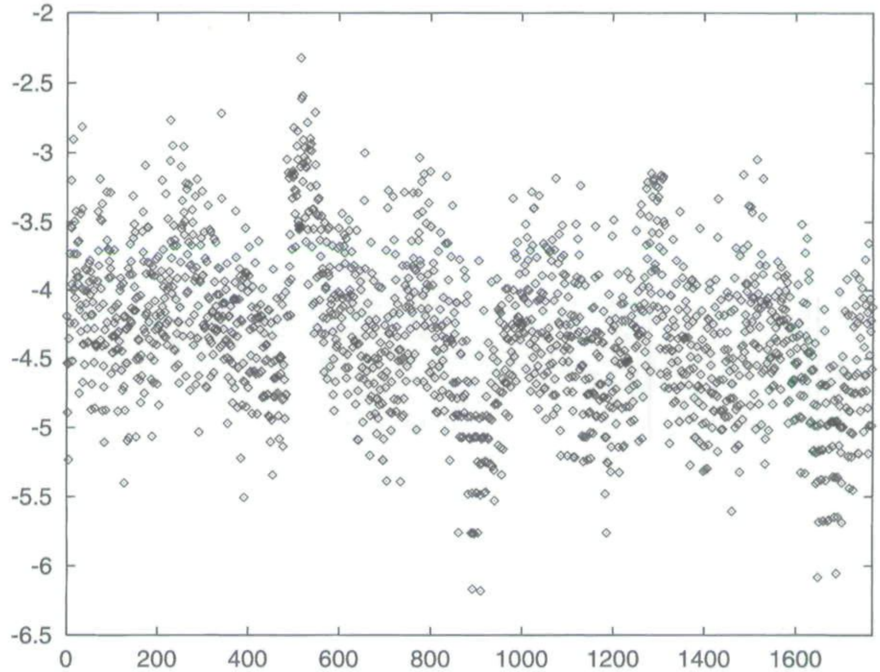

FIGURE 3

Unadjusted volatility for corn. The data consist of logarithmically transformed volatility, r_t.

are regressed on the effect variables. This is referred to as a variance adjustment. The regressors corresponding to the effect variables are then used to adjust the data. The details of performing the adjustment are given in the Appendix. After experimenting, the decision has been made to use the following scheme for adjusting the two data series.

The mean adjustment for the price difference, $\hat{y}_t^1$, involves a regression on two dummy variables which are associated with the drought in 1988

$$(d_1, d_2) = \begin{cases} (1,0) & \text{for June of 1988} \\ (0,1) & \text{for July of 1988}\ . \\ (0,0) & \text{otherwise} \end{cases} \tag{31}$$

The variance adjustment for $\hat{y}_t^1$ is based on several regressors:

1. A gap variable. The gap variable, g, assumes the values ($g = 0, 1, 2, 3, 4$), the size of the gap.
2. A dummy variable associated with the drought of 1988

$$d_3 = \begin{cases} 1 & \text{for June, July, and August of 1988} \\ 0 & \text{otherwise} \end{cases}. \tag{32}$$

TABLE X

Number of Occurrences of Gaps in the Data Series

Gap	*Number of Occurrences*
0	1386
1	16
2	324
3	40
4	1

A gap is a period of time between two successive trading days when the market is closed.

TABLE XI

Regression to Adjust the Mean of Price Differences, $\hat{y}_t^1$.

Variable	*Coefficient*	*Standard Error*	*T-Statistic*	*Significance*
Constant	−0.051285479	0.031017742	−1.65342	0.09842252
d_1	1.765824889	0.276404612	6.38855	0.00000000
d_2	−0.976388918	0.289729616	−3.37000	0.00076789

The regression $F(2, 1764) = 26.35$. The significance level of F is 0.00000000.

3. A trend variable. The variable T is the number of calendar days initialized so that 1 Jan 1986 corresponds to $T = 1$.
4. Two trigonometric variables to capture cyclical behavior in volatility,

$$\sin\left(\frac{2\pi T}{Y}\right), \quad \cos\left(\frac{2\pi T}{Y}\right), \tag{33}$$

where Y is either 365 or 366 depending on whether it is leap year.

The mean adjustment for the volatility, $\hat{y}_t^2$, uses several regressors:

1. A gap variable, g
2. A dummy variable associated with the drought of 1988

$$d = \begin{cases} 1 & \text{for June, July, and August of 1988} \\ 0 & \text{otherwise} \end{cases}. \tag{34}$$

3. A trend variables, T
4. Four trigonometric variables to capture cyclical behavior in volatility,

$$\sin\left(\frac{2\pi T}{Y}\right), \ \cos\left(\frac{2\pi T}{Y}\right), \ \cos\left(\frac{4\pi T}{Y}\right), \ \cos\left(\frac{6\pi T}{Y}\right), \ \sin\left(\frac{8\pi T}{Y}\right). \tag{35}$$

TABLE XII

Regression to Adjust the Variance of Price Differences, $\hat{y}_t^1$

Variable	*Coeff.*	*Std. Error*	*T-stat.*	*Signif.*
Constant	−0.667268212	0.110699832	−6.02773	0.00000000
g	0.287496214	0.055090203	5.21865	0.00000020
d_3	1.556150849	0.264778903	5.87717	0.00000000
T	−0.000412292	0.000065889	−6.25742	0.00000000
$\cos(2\pi T/Y)$	−0.645179418	0.070056571	−9.20941	0.00000000
$\sin(2\pi T/Y)$	−0.239796259	0.067964238	−3.52827	0.00042902

The regression $F(5, 1761) = 51.1634$. The significance level of F is 0.00000000

TABLE XIII

Regression to Adjust the Mean of Volatility, $\hat{y}_t^2$

Variable	*Coefficient*	*Standard Error*	*T-Statistic*	*Significance*
Constant	−4.095260106	0.025010135	−163.74402	0.00000000
g	0.076933586	0.012445640	6.18157	0.00000000
d_3	0.617926168	0.060573707	10.20123	0.00000000
T	−0.000190393	0.000014885	−12.79119	0.00000000
$\cos(2\pi T/Y)$	−0.270814549	0.015831929	−17.10559	0.00000000
$\sin(2\pi T/Y)$	−0.107501473	0.015350626	−7.00307	0.00000000
$\cos(4\pi T/Y)$	0.040111077	0.015526797	2.58335	0.00986483
$\cos(6\pi T/Y)$	−0.055733030	0.015380260	−3.62367	0.00029871
$\sin(8\pi T/Y)$	0.033787166	0.15217588	2.22027	0.02652749

The regression $F(8, 1758) = 111.7995$. The significance level of F is 0.00000000.

TABLE XIV

Regression to Adjust the Variance of Volatility, $\hat{y}_t^2$

Variable	*Coefficient*	*Standard Error*	*T-Statistic*	*Significance*
Constant	−2.937559836	0.059483384	−49.38454	0.00000000
g	0.167091460	0.060859407	2.74553	0.00610256

The regression $F(1, 1765) = 7.5379$. The significance level of F is 0.00610256.

The variance adjustment for $\hat{y}_t^2$ involves a regression on the gap variable, g.

In Tables XI, XII, XIII, and XIV the relevant statistics for the regression equations are presented. The F statistics and T-statistics are highly significant. In Figures 4 and 5 the adjusted data for price difference, y_t^1,

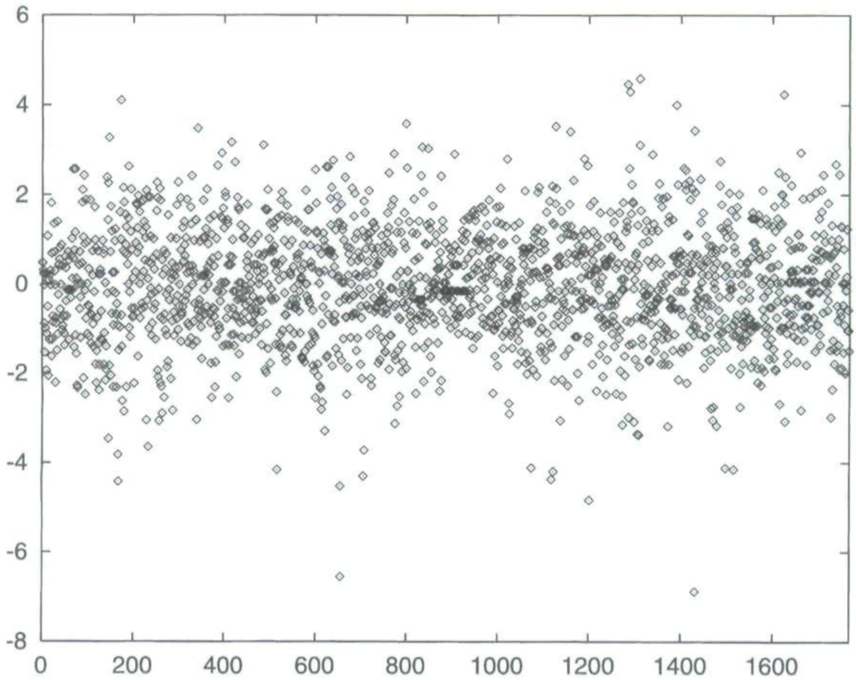

FIGURE 4
Adjusted price differences, y_t^1, for corn.

and volatility, y_t^2 are shown. It appears that the adjustment process has reduced the inhomogeneity in the data.

SNP Methodology

The adjusted Corn data are now modeled using the SNP methodology of Gallant and Tauchen (1995). The SNP methodology is a seminonparametric methodology for modeling an m-dimensional time series y_t, $(1 \leq t \leq n)$. The methodology is based on the fact that a probability density for m-dimensional innovations z_t can be approximated by an Hermite polynomial expansion. The lowest order polynomial corresponds to a multivariate normal distribution. As higher order polynomials are included in the expansion it can more closely approximate distributions that deviate from normal. The series y_t is assumed to be derivable from a VAR/ARCH-like process driven by the innovations z_t. Furthermore, the coefficients of the Hermite polynomials can depend on lagged values of y_t. There are a number of tuning parameters for controlling the nature of the distribution: L_u is the number of lags in the VAR component; L_r is the number

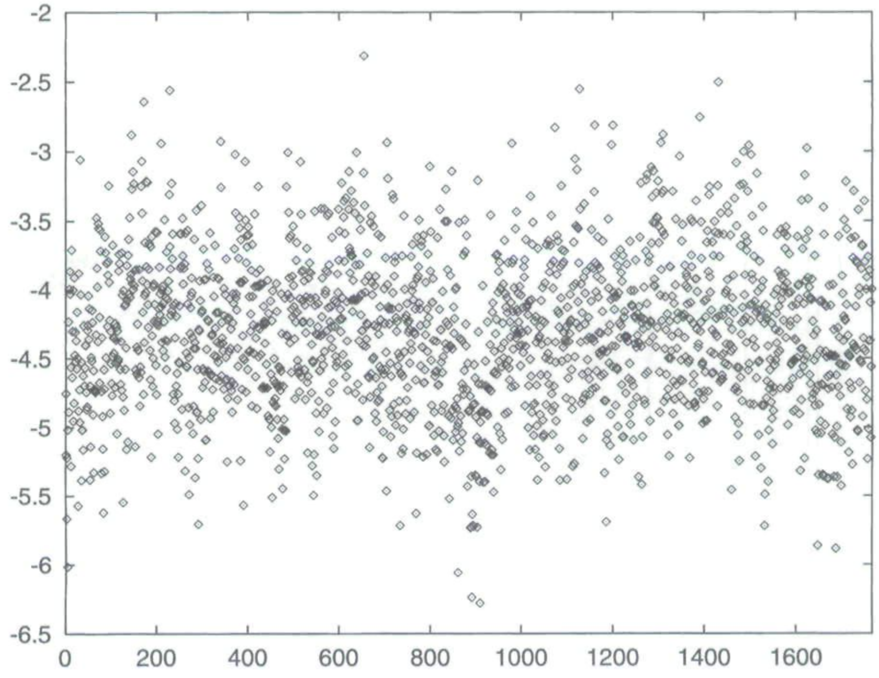

FIGURE 5
Adjusted volatility, y_t^2, for corn.

of lags in ARCH component; L_p is the number of lags of y_t in the coefficients of the polynomial; K_z is the highest degree in z_t of the Hermite polynomial; K_x is the highest degree in y_t of the coefficients of the Hermite polynomial; I_z and I_x are used for reducing the number of higher order interaction terms, i.e., products between different components of z_t and y_t, respectively, of the polynomial. The parameter K_z controls the shape of the homogeneous component of the distribution, and K_x the heterogeneous component. A suitable choice of the parameter L_r can greatly reduce the value of K_x in fitting a distribution that is conditionally heterogeneous. In Table XV the types of processes characterized by these parameters are summarized (Gallant and Tauchen, 1 January 1995).

The model building process is incremental in that one starts with a vector autoregressive model and gradually increases the number or value of parameters until a suitable fit to the to data is achieved. The goodness-of-fit is the Schwarz criterion,

$$BIC = s_n(\hat{\theta}) + (1/2)(p_\theta/n)\ \ln(n), \tag{36}$$

where s_n is the objective function for the model, $\hat{\theta}$ are the values of the estimated parameters defining the model, s_n is the objective function for

TABLE XV

Characterization of Processes According to Restrictions on the Tuning Parameters

Parameter Setting	*Characterization of y_t*
$L_u = 0, L_r = 0, L_p \geq 0, K_z = 0, K_x = 0$	iid Gaussian
$L_u > 0, L_r = 0, L_p \geq 0, K_z = 0, K_x = 0$	Gaussian VAR
$L_u > 0, L_r = 0, L_p \geq 0, K_z > 0, K_x = 0$	non-Gaussian VAR homongeneous innovations
$L_u \geq 0, L_r > 0, L_p \geq 0, K_z = 0, K_x = 0$	Gaussian ARCH
$L_u \geq 0, L_r > 0, L_p \geq 0, K_z > 0, K_x = 0$	non-Gaussian ARCH, homogeneous innovations
$L_u \geq 0, L_r > 0, L_p \geq 0, K_z > 0, K_x > 0$	general nonlinar process heterogeneous innovations

the fit to the data, p_θ is the number of parameters in the model. The Schwarz criterion is more conservative than other criteria like the Akaike information criterion in that it penalizes more for a model when a low value of the objective function is obtained at the expense of an excessive number of parameters.

The model for the adjusted corn data has been constructed using the FORTRAN code SNP 8.4 [8]. In Table XVI the results of constructing the model are summarized. The table gives model specifications and the value of the Schwarz criterion for each model. The optimum model according to the Schwarz criterion is a nonGaussian ARCH with homogeneous innovations, the values of the tuning parameters are $(L_u, L_r, L_p, K_z, K_x, I_z, I_x) = (2,1,1,8,0,0,0)$.

Finally, because of the conservative nature of the Schwarz criterion, the adequacy of the optimum model is checked for long-term and short-term misspecification as follows (Gallant and Tauchen, 1995). First, the one-day-ahead scaled residuals of the model are calculated

$$\hat{u}_t = [\mathrm{Var}_c(y_t)]^{-1/2} [y_t - \langle y_t \rangle_c], \tag{37}$$

where $\mathrm{Var}_c(\mathbf{y}_t)^{-1/2}$ is the Cholesky factor of the one-day-ahead conditional variance–covariance matrix, and $\langle \mathbf{y}_t \rangle_c$ is the one-day-ahead conditional mean, both being estimated from the model. Then, to test for short-term misspecification, the scaled residuals are regressed on the first, second, and third powers of all lagged components of $\mathbf{y}_t$ up to and including lag 10 (mean test). The squared components of the residuals are also regressed in a similar manner (variance test). If either multivariate regression accounts for a significant amount of variance according to the Wilks

TABLE XVI

A Representative Sample of Diagnostics for the Estimated Model of the Bivariate Series y_t.

L_u	L_r	L_p	K_z	I_z	K_x	I_x	p_θ	$s_n(\theta)$	BIC
2	1	1	0	0	0	0	15	2.78876	2.82050
2	1	1	8	0	0	0	59	2.35559	2.48042
2	11	1	8	0	0	0	79	2.33886	2.50600
2	12	1	8	0	0	0	81	2.33796	2.50933
3	8	1	8	0	0	0	77	2.33959	2.50250
3	9	1	8	0	0	0	79	2.33877	2.50592
4	7	1	8	0	0	0	79	2.33623	2.50338
4	8	1	8	0	0	0	81	2.33524	2.50662
5	5	1	8	0	0	0	79	2.33512	2.50227
6	3	1	8	0	0	0	79	2.33984	2.50698
7	1	1	8	0	0	0	79	2.34324	2.51038

test, then the values of the tuning parameters are systematically increased, and the parameters θ reestimated (Morrison, 1976). Testing for long-term misspecification proceeds in a similar manner. The scaled residuals are regressed on dummy variables assigned to each year. Squared, scaled residuals are regressed on dummy variables assigned to each year. Squared, scaled residuals are regressed on the annual dummies. In each case the Wilks test is again used to test for the significance of the regression. If either regression is significant, then the values of the tuning parameters are increased, and the model is re-estimated. This fine-tuning process is continued until no regression accounts for a significant amount of variance at the 5% level. In order to satisfy the short- and long-term misspecification criteria, we have had to augment the model. The final model is a 4818000 as shown in Tables XVII and XVIII. It has 81 parameters, resulting in a saturation ratio of approximately 44 ($2 \times 1767/81$) observations per parameter.

Simulations

Inefficiency in the corn data can be more rigorously tested by applying the trading system to simulated prices and volatility for Corn. First, data are simulated using the SNP model. The SNP code does this using the rejection method to generate random samples from the estimated probability distribution (Gallant and Tauchen, June 1995). Then, these simulations are 'unadjusted' using the regression equations obtained. The momentum-based trading system is applied to 30 such simulations of

TABLE XVII

Short-Term Diagnostics for the Bivariate Series y_t

							Mean		Variance	
L_u	L_r	L_p	K_z	I_z	K_x	I_x	Wilks Λ	Probability	Wilks Λ	Probability
Raw Data							0.53962	0.00000	0.41196	0.00000
Adjusted data							0.83385	0.00000	0.87858	0.00000
2	1	1	8	0	0	0	0.89817	0.00015	0.92002	0.07555
2	11	1	8	0	0	0	0.91663	0.03634	0.93659	0.67893
2	12	1	8	0	0	0	0.91649	0.03522	0.93613	0.65824
3	9	1	8	0	0	0	0.91657	0.03589	0.93816	0.74532
4	7	1	8	0	0	0	0.91630	0.03368	0.93739	0.71363
4	8	1	8	0	0	0	0.91897	0.06077	0.93916	0.78387
5	5	1	8	0	0	0	0.91630	0.03373	0.93601	0.65296
6	3	1	8	0	0	0	0.91243	0.01296	0.92994	0.37040

TABLE XVIII

Long-Term Diagnostics for the Bivariate Series y_t

							Mean		Variance	
L_u	L_r	L_p	K_z	I_z	K_x	I_x	Wilks Λ	Probability	Wilks Λ	Probability
Raw Data							0.86002	0.00000	0.81383	0.00000
Adjusted data							0.97334	0.00002	0.98664	0.05184
2	1	1	8	0	0	0	0.98369	0.01175	0.98806	0.10075
2	11	1	8	0	0	0	0.98847	0.12413	0.99261	0.52692
2	12	1	8	0	0	0	0.98859	0.13095	0.99256	0.52122
3	9	1	8	0	0	0	0.98854	0.12787	0.99168	0.40448
4	7	1	8	0	0	0	0.98847	0.12420	0.99079	0.30089
4	8	1	8	0	0	0	0.98889	0.14858	0.99104	0.32793
5	5	1	8	0	0	0	0.98907	0.15989	0.99069	0.29122
6	3	1	8	0	0	0	0.98749	0.08027	0.99019	0.24276

Corn prices and volatility, using the V statistic to filter trades. In Figure 6 a plot of profit per trade against the V statistic is shown for the simulations as well as the actual Corn data. For the simulated data the plot contains the profit per trade averaged over the 30 simulations and an error bar that is plus/minus one standard deviation of the mean. For the actual Corn data the plot contains the profit per trade obtained from Table VII. From the plot it appears that profits generated using actual data are significantly greater than for simulation data. This assertion can

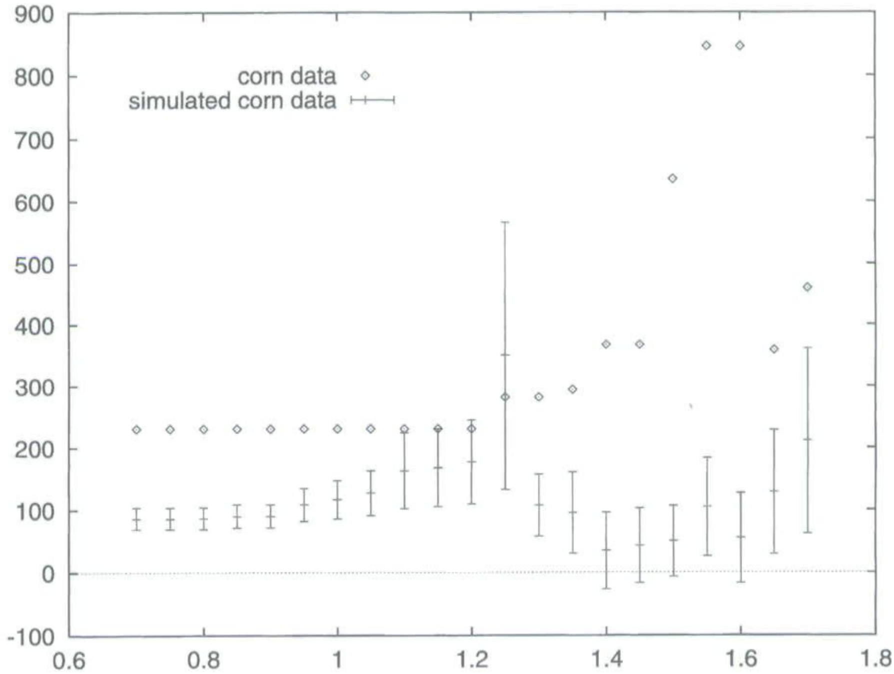

FIGURE 6
Profits per trade in dollars for corn.

be tested as follows. Given the model as the null hypothesis, the probability of the trading system generating profits using the actual data is calculated in the case when the filtering value is $V = 1.6$. The reason for selecting $V = 1.6$ is that profits derived using actual data and this filtering value appear to be the most different from profits generated using simulations. In only one set of simulation data out of 30 does the trading system generate profits greater than in the actual data. The value of the probability is approximately 4.5%, which is estimated by performing a linear interpolation. Thus, the null hypothesis can be rejected at the 5% level. Unfortunately, this is the only value of the filter for which the null hypothesis can be rejected at the 5% level. Although Figure 6 seems to suggest that the profit per trade of the actual data is not consistent with the model, statistically there is only weak evidence to this effect.

An alternative method is presented for testing whether actual trading profits are consistent with the model. Profit per trade for Corn is regressed against V. If using the V statistic to filter trades improves profits, then the regression coefficient, or correlation coefficient between V and profit per trade, should be significantly different from zero.[13] The result of the

[13]All V values of all markets are greater than 0.7.

regression is that the correlation between V and profit per trade is .610 with an $F(1, 9) = 5.31$, which is significant at the 4.7% level. This appears to be highly significant. Given this significant result and that the model seems to account for the behavior of the actual data (as suggested in the previous paragraph), then the correlation between V and profit per trade for simulation data should also be significantly different from zero. To test this hypothesis, profit per trade for each set of simulation data is regressed against V. The average correlation between V and profit per trade is 0.088, and the standard deviation is 0.120. This implies that the correlation is not significantly different from zero, i.e., using the V statistic to filter trades does not improve profit per trade for the simulations.[14] This apparent contradiction is resolved by noting that the correlation coefficient, 0.610, obtained from actual data, is significant only at the 31% confidence level when compared directly with correlation coefficients obtained from the simulation data. Consequently, one cannot reject at the 5% confidence level that the correlation coefficient derived from actual data is significantly different from zero, and that its value consistent with the model. The significant result obtained from regressing profit per trade against V for actual data may be the result of biasing due to correlation in the residuals for different values of V.

The results based on the simulations are disappointing in that they provide weak evidence that using the V statistic to filter trades improves profit per trade in the case of Corn. This is contrary to what one may conclude from visually inspecting Figure 6.[15]

PRACTICAL CONSIDERATIONS

Could the trading system described in this study generate profits comparable to those in Table IX in real time? The answer to this question is relevant because, if inefficiency has existed, this system should have been profitable, not only on paper but also in practice. The two primary factors determining if the trading system could have been applied profitably in real time are whether an individual would have sufficient time to calculate entry and exit points for trades and to enter the trades on the close and whether slippage would reduce profit per trade to values appreciably less than those presented in Table IX.

The procedure for applying the trading system in real time is essentially the same as in this analysis, with some modifications. First, system

[14]Because the values of the V statistic for Corn are all greater than 1.20 (see Table V), the regression is restricted such that $(1.20 \leq V \leq 1.70)$.
[15]This result seems apparent in Figure 6.

parameters need to be optimized for a year of data prior to the quarter in which trading system is be used. The time required for such optimizations is about six hours using a workstation such as a Hewlett Packard 9000 series 735/125. These optimizations can be performed, in principle, after the close on the last day prior to the beginning of that quarter. Once system parameters have been obtained, the time required to calculate entry and exit prices for the ten markets is less than one second. One technical problem is that entry and exit prices must be entered on the close; however, these prices can, in reality, only be obtained after the close so that they must be estimated as near to the close as possible while still allowing sufficient time for a trade to be entered on the close. The extent to which profits generated by the system in real time approximate theoretical profits depends on the accuracy of estimating the daily high and low prior to the close. It seems reasonable that, in most circumstances, good estimates can be obtained within the last fifteen minutes of trading, allowing ample time to perform necessary calculations and to enter trades on the close. Slippage should not be a factor which causes significant differences between actual profits and those presented in Table IX. The reason is that trades occur only at the close, rather than intraday, so that excessive slippage is less likely to occur because of extremely fast markets.

Consequently, it seems reasonable that the trading system could have been used in real time, with real-time profits being comparable to those in Table IX. However, the system would probably not be considered practical because it lacks a sound risk-management strategy. Specifically, no stops are even put in place intraday.

CONCLUSIONS

Inefficiency in ten futures markets cocoa, coffee, corn, cotton, crude oil, gold, Japanese yen, live cattle, U.S. T-bonds, and the S&P 500 index has been investigated during the period from June 1986 until June 1993. For this purpose a trend following trading system has been proposed that uses momentum to determine entry and exit points for trades. The trading system is found to be profitable in only a few markets. However, when the *V* statistic is used to filter trades, the profitability of the system is improved overall, as well as for some markets in particular. Furthermore, the *V* statistic seems to provide insight into why the momentum-based trading system is profitable in some but not other markets. Specifically, the trading system has generated the best profits for cotton, corn, and Japanese yen. This seems to be the result of long-range memory dependence in these markets. The system has generated large losses in the S&P

500. It has also generated losses in cocoa. The reason appears to be that these markets have exhibited a strong antipersistence, suggesting that a countertrend trading system may be more appropriate to use. In the case of live cattle or gold using the V statistic to filter trades has not improved profits. Perhaps adjusting the size of the sample when calculating the V statistic would improve profits in these markets. Another unexplained occurrence is that when trades are filtered using relatively large values of the V statistic ($V \approx 1.65$), profits begin to decline rather than continue to increase, contrary to what is suggested by theory. This is especially apparent in the case of coffee. Perhaps reducing the sample size used in calculating the V statistic may improve profits and help explain this anomaly.

To provide further evidence for inefficiency, a model for corn has been constructed and simulation data have been generated. Application of the trading system to the simulation data has yielded no statistical evidence that the trading system generates profits other than by chance. A test has been performed to determine whether the profits generated using real data are significantly greater than those generated using the simulation data. Only weak evidence has been found to this effect, implying that the model has accounted, to a large extent, for the behavior of the corn data.

Finally, it is concluded that, although evidence to the contrary exists in the case of corn, there is some evidence for inefficiency in the futures markets during the period from June 1986 until June 1993, and furthermore that a trend-following trading system based on market momentum and long-range memory dependence could have profited from this inefficiency.

APPENDIX

Adjustment for Systematic Effects

Let $\hat{y}_t$ be a univariate time series. Assume that the time series has been divided into groups which are suspected of exhibiting systematic differences. A procedure is now described for analyzing and adjusting for systematic differences in the means or variances of the groups (Gallant, Rossi, and Tauchen, 1990). First, dummy variables are assigned to each group. Then $\hat{y}_t$ is regressed on the dummy and any other continuous variables, denoted by the vector **g**,

$$\hat{y}_t = \mathbf{g} \cdot \beta + u_t. \tag{38}$$

When some regression coefficients, β^i, are found not to be significantly different from zero, the groups corresponding to these coefficients are combined with other groups, as appropriate. The regression is then rerun. This process is continued until the number of groups has been reduced, and all regression coefficients are significantly different from zero.

Next, a variance adjustment $g \cdot \gamma$ is obtained by regressing $\log(u_t^2)$ on g,

$$\log(u_t^2) = g \cdot \gamma + \varepsilon_t. \tag{39}$$

If some regression coefficients, γ^i, are found not to be significant, groups are combined and the regression is rerun as described previously. The series $\hat{y}_t$ is then adjusted so that

$$\tilde{y}_t = \frac{\hat{y}_t - g \cdot \beta}{\exp(g \cdot \gamma/2)}. \tag{40}$$

Finally, each $\hat{y}_t$ is transformed

$$y_t = a + b\tilde{y}_t, \tag{41}$$

where a and b are selected so that the mean and variance of y_t are the same as the unadjusted series $\hat{y}_t$.

BIBLIOGRAPHY

Andrews, D. (1991): "Heteroskedasticity and Autocorrelation Consistent Covariant Matrix Estimation," *Econometrics,* 59:817–858.

Babcock, B. Jr. (1989): *The Dow Jones Irwin Guide to Trading Systems.* Homewood, IL: Dow Jones-Irwin.

Beran, J. (1994): *Statistics for Long Memory Process.* New York: Chapman and Hall.

Bryant, R. C. (1995): "The 'Exchange Risk Premium,' Uncovered Interest Parity, and the Treatment of Exchange Rates in Multicountry Macroeconomic Models." An abstract of this paper and the paper itself are available at http://www.brook.edu/es/wp/bdp/bdp111/abstr111.htm.

Chance, D. M. (1989): *An Introduction to Options and Futures.* Orlando, FL: Dryden Press, pages 354–369.

Gallant, A. R., Rossi, P. E., and Tauchen, G. (1990): "Stock Prices and Volume," working paper, Department of Statistics, North Carolina State University.

Gallant, A. R., and Tauchen G. (1990): "Nonlinear Time Series Analysis: Estimation and Simulation," working paper, Department of Statistics, North Carolina State University.

Gallant, A. R., and Tauchen, G. (1995): "Snp: A Program for Nonparametric Time Series Analysis, Version 8.4, User's Guide." Available from ftp://www.ccon.duke.edu/home/get/papers/snpguide.ps.

Gencay, R. (1995): "Tests for the Risk Premium on Foreign Currency Futures Implied by the Intertemporal Asset Pricing Theory," *Applied Financial Economics,* 5(2):85–94.

Hurst, H. E. (1951): "The Long-Term Storage Capacity of Reservoirs," *Transactions of the American Society of Civil Engineers,* 116.

Ingber, L. (1984): "Statistical Mechanics of Nonlinear Nonequilibrium Financial Markets," *Journal of Mathematical Modelling,* 5(6):343–361. This paper is also available at http://www.ingber.com/markets84.statmech.ps.Z.

Lo, A. W. (1991): "Long-Term Memory in Stock Market Prices," *Econometrica,* 59(5):1279–1313.

Mandelbrot, B. B., and Van Ness, J. W. (1968): "Fractional Brownian Motions, Fractional Noises and Applications,": *SIAM Review,* 10(4):422.

Mandelbrot, B. B., and Wallis, J. R. (1969a): "Computer Experiments with Fractional Gaussian Noises. Part 1, Averages and Variances," *Water Resources Research,* 5(1) 229.

Mandelbrot, B. B., and Wallis, J. R. (1969b): "Computer Experiments with Fractional Gaussian Noises, Part 1, Mathematical Appendix," *Water Resources Research,* 5(1): 269.

Mandelbrot, B. B., and Wallis, J. R. (1969c): "Computer Experiments with Fractional Gaussian Noises. Part 2, Rescaled Ranges and Spectra," *Water Resources Research,* 5(1): 242.

Morrison, D. F. (1976): *Multivariate Statistical Methods,* (second edition). New York: McGraw-Hill, page 222.

Murphy, J. J. (1986): *Technical Analysis of the Futures Markets.* New York: New York Institute of Finance.

Parkinson, M. (1980): "The Extreme Value Method for Estimating the Variance of the Rate of Return," *Journal of Business,* 53(1):61–65.

Peters, E. E. (1994): "*Fractal Market Analysis.* New York: John Wiley & Sons, Inc.

Schwager, J. D. (1989): *Market Wizards: Interviews with Top Traders.* New York: New York Institute of Finance.

Wilder, J. W. Jr. (1978): *New Concepts in Technical Trading Systems.* Greensboro, NC: Trend Research.