TRADING FUTURES USING A CHANNEL RULE: A STUDY OF THE PREDICTIVE POWER OF TECHNICAL ANALYSIS WITH CURRENCY EXAMPLES

STEPHEN J. TAYLOR

INTRODUCTION

The channel rule of Irwin and Uhrig (1984) is a simple, technical, trading rule. A long futures position is replaced by a short position when the price is less than the minimum price during the preceding L days. Likewise, a short futures position is replaced by a long position when the price is more than the maximum price during the preceding L days. The integer L is the only parameter. This article uses theoretical methods to show that the channel rule can be surprisingly profitable when market prices almost follow random walks. The adjective *channel* here refers to two lines drawn parallel to the time-axis, one through the minimum closing price during the preceding L days and the other through the maximum closing price. Only closing prices are considered in this article. All intra-day prices are ignored when channels are defined and decisions taken.

■ *Stephen J. Taylor is a Professor of Finance at Lancaster University, England.*

The Journal of Futures Markets, Vol. 14, No. 2, 215–235 (1994)

CCC 0270-7314/94/020215-21

Lukac, Brorsen, and Irwin (1986, 1988) found that the channel rule performed best in their ex ante study of several technical trading systems, including the filter rule of Alexander (1961). They concluded that the channel rule should be preferred in academic studies of futures trading and market efficiency. Taylor (1992a) recently showed that the channel rule gives remarkably profitable results (after transactions costs) for ex ante currency futures trades, from December, 1981 to November, 1987. The channel rule outperformed decision rules based upon ARIMA predictions during this period. A theory of short-run disequilibrium, developed by Beja and Goldman (1980) and discussed by Lukac, Brorsen, and Irwin (1988), might explain the success of trend-following systems such as the channel rule.

The prediction rules called technical analysis by Neftci (1991) include rules like the channel rule. Neftci (1991, p. 550) remarked "To the best of my knowledge, there is no formal study of the predictive power of technical analysis." Presumably Neftci considered that the many published trading rule studies lacked a formal prediction methodology, although several studies apply a rigorous methodology for testing the efficient market hypothesis. See, for example, Sweeney (1986); Lukac, Brorsen, and Irwin (1988); Lukac and Brorsen (1990); and Taylor (1992a). The recent article by Brock, Lakonishok, and LeBaron (1992) shows that technical trading rules have predictive power for the Dow Jones Index.

The major contribution of this study is a theoretical explanation of why the channel rule can be successful compared with ARIMA decision rules when futures prices do not follow random walks. It is shown that when prices are generated by a particular ARIMA(1, 1, 1) model there are several values of L for which the channel rule correctly identifies the sign of conditional expected returns with probability well above 0.5. The probabilities attained by the channel rule are close to the best possible probabilities given by ARIMA predictions. The expected payoff from the channel rule as a function of L is shown to vary little in a fairly wide interval around the optimal value. The channel rule appears to do well in ex ante studies because it has a high probability of selecting a value for the single parameter L within a satisfactory interval even when very few contracts are available to guide the choice. Appropriate ARIMA rules, however, usually require the estimation of at least one AR and one MA parameter [Taylor (1988, 1992b)].

The next section provides definitions of the gross and net payoffs from the channel rule which are used throughout the article. The third section describes an ARIMA(1, 1, 1) model which can give a good description of the futures price process at some markets; while the fourth and fifth sections give the results of Monte Carlo studies of the channel rule, showing how expected payoffs vary with L and the number of daily futures prices available when choosing L. The sixth section summarizes new results for currency futures trades using the channel rule applied to daily closing prices.

PAYOFFS FROM CHANNEL RULE TRADES

Assumptions

The channel rule alternates between long and short futures positions and is never neutral. A method for summarizing the net payoff from a set of decisions for a single futures contract has been described in Taylor (1988, 1992a) and this method is used throughout this article. Several assumptions are made about a representative trader. First, the trader can deposit U.S. Treasury Bills as margin. Second, the futures price of goods traded equals a particular T-Bill price whenever a trade commences. Losses during a trade can only exceed the price of the margin security if a short position is taken and the futures price more than doubles. Third, transaction costs are taken to be a constant proportion of the futures price of goods traded.

The investment used to finance futures trades is assumed to be the capital required to purchase the T-Bill which is used as a margin security. All definitions of the investment capital required for futures trading are arbitrary to some degree. An alternative definition might define the investment to be the maximum drawdown but a rigorous ex ante definition of the maximum is problematic.

Measuring Trading Performance

First, suppose the channel rule generates N trades in some futures contract between times t_1 and t_2, with trade j begun at the price P_j and concluded at Q_j. The times t_1 and t_2 are defined ex ante. To distinguish between long and short positions, let $\pi_j = 1$ for the former and $\pi_j = -1$ for the latter.

Second, suppose there is a T-Bill which trades at both times t_1 and t_2. Let B_1 be the price of the T-Bill at time t_1 and let one futures contract be for F units of the relevant goods. It is assumed that when trade j commences, the quantity traded is $B_1/(FP_j)$ contracts. Thus B_1

is both the dollar price of the contracts when any trade starts and the price of the T-Bill when the first trade commences.

Third, suppose transaction costs for one round-trip in these contracts are a constant proportion c of B_1.

The *gross payoff* from the N trades is defined to be:

$$G = \sum_{j=1}^{N} \pi_j \frac{(Q_j - P_j)}{P_j} \tag{1}$$

The *net payoff* from the trades is defined to be

$$R = G - Nc \tag{2}$$

For example, if

1. t_1 is May 31st when a long position is commenced at 5000,
2. The long position is reversed to short at 5100 on July 10th, and
3. The short position is closed out at 5049 on November 30th which is t_2, then $N = 2$, $\pi_1 = 1$, $\pi_2 = -1$ and G is

$$(5100 - 5000)/5000 - (5049 - 5100)/5100 = 0.03 \text{ or } 3\%$$

When the proportional trading cost c is 0.2% the value of R is 2.6%.

The quantity R is simply the aggregate net trading profit divided by the futures price of goods at the commencement of trades. As the initial futures price of goods is the price of the T-Bill at time t_1, it follows that R equals the return on capital used to purchase a T-Bill which is then used as margin collateral for the futures trades minus the return when an identical T-Bill is purchased and not used as collateral. This interpretation allows R to be called an *excess return*. Further details are given in Taylor (1988, 1992a) where the relationship between return and risk is also discussed in a general portfolio context.

Currency Results to November, 1987

The average value of R is 3.93% over 6 months for ex ante Pound, Mark, Swiss Franc, and Yen channel rule trades using daily futures closing prices from December, 1981 to November, 1987, with trading costs c equal to 0.2% per trade [Taylor (1992a), Table 1]. All trades are executed at the close. No intra-day prices are considered. The assumed average trading cost is slightly more than $100 per contract, which is more than the total trading costs of a typical retail customer.

The channel parameter L is optimized over eight June and December contracts up to November, 1981 to obtain a value for trading the June, 1982 contract; then L is optimized again but over nine contracts to obtain a value to trade the December, 1982 contract, etc., with separate optimizations for each currency. Averaging R across currencies gives one average value for each of the 12 June and December delivery months from June, 1982 until December, 1987. From these 12 averages the figure above of 3.93% is computed. The null hypothesis that the average excess return is zero is rejected at the 1% significance level. Taylor (1992a) claims that risk cannot explain the significant average value of R. The high average value of R provides the motivation for the research reported in this article.

Decisions based upon ARIMA decision rules are less successful than those generated by the channel rule although the differences in performance between these rules are not statistically significant.

A PRICE MODEL FOR MONTE CARLO SIMULATIONS

To learn more about why the channel rule has the potential to give net payoffs whose expectations are positive, it is necessary to perform Monte Carlo simulations. The simulations require a time series model for daily futures prices which is consistent with the statistical evidence about prices. Daily changes in futures prices are known to be almost random and to display conditional heteroscedasticity. Consequently, a model is simulated that has very small positive correlations between price changes and also has changes in the conditional variances of price changes. As will be seen the small correlations are sufficient for the channel trading rule to obtain positive expected net returns.

The Model

Simulated futures prices f_t are obtained after first simulating logarithmic differences x_t defined by

$$x_t = \ln(f_t) - \ln(f_{t-1}) \tag{3}$$

The prices are given by $f_t = f_{t-1}\exp(x_t)$. The logarithmic differences x_t are called returns in this article, for convenience, even though capital is not necessary to trade futures. Futures prices would be a random walk if returns were uncorrelated at all lags; so $\text{cor}(x_t, x_{t+\tau}) = 0$ for all positive τ. Positive dependence between returns will create trend effects in prices and the possibility of trading profits—a set of positive returns will be more likely than a set of negative returns to be followed by positive returns. Evidence for positive autocorrelation among exchange rate returns can be found in Hodrick and Srivastava (1987); Mark (1988); and Taylor (1986, 1992b).

Positive dependence at all lags can be obtained by letting returns be the sum of two stationary and independent processes, the first an AR(1) component μ_t and the second a white noise component e_t, thus:

$$x_t = \mu_t + e_t \tag{4}$$

$$\mu_t = p\mu_{t-1} + \nu_t, \quad 1 > p > 0 \tag{5}$$

and
$$\operatorname{cor}(\nu_t, \nu_{t+\tau}) = \operatorname{cor}(e_t, e_{t+\tau}) = 0, \quad \tau > 0 \tag{6}$$

Then, for $\tau > 0$,

$$\operatorname{cor}(x_t, x_{t+\tau}) = Ap^{\tau}, \quad A = \operatorname{var}(\mu_t)/\operatorname{var}(x_t) \tag{7}$$

The term μ_t can be interpreted in several ways. A risk premium interpretation is one possibility but the predictable component might be attributable to inefficient pricing. As prices are at least very like random walks, only values for A, which are almost zero, can be considered.

Returns are conditionally heteroscedastic for exchange rates, stocks, and many other assets; therefore, $h_t = \operatorname{var}(x_t \mid I_{t-1})$ is a function of past returns $I_{t-1} = \{x_{t-1}, x_{t-2}, \ldots\}$. Recent reviews of empirical evidence and appropriate model specifications can be found in Bollerslev, Chou, and Kroner (1992) and Taylor (1994). The residual e_t must display conditional heteroscedasticity. An interesting issue is whether or not μ_t also displays conditional heteroscedasticity. There is certainly evidence that the conditional autocorrelations of returns decrease towards zero as the conditional variance of returns increases [Taylor (1986) and Kim (1989) for currencies and LeBaron (1992) and Sentana and Wadhwani (1990) for stocks]. These studies support models in which the conditional variance of μ_t is independent of past, present, and future residuals e_s.

A Gaussian, AR(1) process is simulated for μ_t. This process has mean 0, variance $A\sigma^2$ and autoregressive parameter p. The simulated process for e_t has unconditional variance $(1 - A)\sigma^2$ so that the unconditional variance of returns x_t is equal to σ^2. To simulate conditional heteroscedasticity effects, it is supposed that e_t is the product of a stochastic volatility variable V_t and an i.i.d. term ε_t; thus,

$$e_t = V_t\varepsilon_t, \tag{8}$$

and also

$$\ln(V_t) - \alpha = \phi[\ln(V_{t-1}) - \alpha] + \eta_t, \tag{9}$$

with

$$\varepsilon_t \sim \text{i.i.d. } N(0, 1 - A), \qquad \text{and} \qquad \eta_t \sim \text{i.i.d. } N(0, \beta^2(1 - \phi^2)) \tag{10}$$

Similar equations have been specified in Taylor (1986, 1994); Chesney and Scott (1989); Melino and Turnbull (1990); and Harvey, Ruiz, and Shephard (1994). The innovation processes $\{\nu_t\}$, $\{\varepsilon_t\}$ and $\{\eta_t\}$ are assumed to be stochastically independent.

Plausible Parameter Values

To perform the simulations it is necessary to select values for A, p, α, β, and ϕ. All the values chosen in this study are intended to be reasonable for the prices of exchange rate futures observed daily. Values for A and p should then satisfy two criteria. First, A and p should be consistent with observed sample autocorrelations for daily currency returns. Following Taylor (1986, 1992b), this suggests selecting $0 < A \leq 0.04$ and $0.9 \leq p < 1$. Second, the values must give a first-lag autocorrelation for monthly sums of daily returns which is consistent with the literature about returns from one-month forward contracts. Mark (1988) gives monthly, first-lag, autocorrelations which range from 0.08 to 0.19. Satisfactory choices for both criteria are given by $A = 0.02$ and $p = 0.95$. The maximum autocorrelation between returns is then a mere 0.019 but the sum of the autocorrelations over all positive lags is 0.38.

Choices for the remaining parameter, α, β, and ϕ, can be made by estimating the parameters of stochastic volatility models from time series of daily exchange rates. The "base case" simulations use values motivated by the DM futures results in an early draft of Taylor (1994). These values are $\alpha = -5.15$, $\beta = 0.422$, and $\phi = 0.973$.

It might at first be assumed that ARIMA decision rules would be better than the channel rule if prices are generated by the ARIMA(1, 1, 1) model implied by eqs. (3) to (10). The results of a small Monte Carlo study in Taylor (1992a) show that this is not so when the trading rule parameters are chosen by ex ante optimizations over a small number of simulated contracts. The average values of R are indistinguishable between the channel and ARIMA rules when 4 to 10 years of simulated daily prices are used to select decision and forecast parameters.

TRADING RESULTS AS A FUNCTION OF *L*

Assumptions

Daily prices f_t are simulated by using a reliable random number algorithm to generate values for the normally distributed residual terms in (5), (8), and (9). Every price should be interpreted as the price when the market closes. For each simulated contract, the time t runs from 1 to 189 (9 months, assuming 252 trading days in one year). Each contract is traded for six months, commencing when $t = 63$. The prices before time 63 are used only when they are needed for the calculation of minimum and maximum prices in an L-day period commencing before time 63. For a hypothetical June contract, one could associate $t = 1,\ 63$ and 189 with March 1st, May 31st, and November 30th.

The prices for one contract cannot be simulated without considering other contracts that would exist at the same time. Consequently, when 189 prices are generated for each of a series of contracts, each traded for 6 months, it is assumed that prices are equal across contracts during the overlapping 3-month periods, i.e., when $t < 64$, the price f_t for contract i equals the price f_{t+126} for contract $i - 1$; similarly, equal trend and volatility terms are assumed in (5) and (9). The assumption of perfect contemporaneous correlation between returns, across contracts, is satisfactory in a currency context because real market returns move closely together. They do this because futures returns are spot returns plus a relatively small term for changes in interest rates.

A long, short or neutral position is held by the channel rule at the end of each day. Neutral positions only occur for days 1 to 62, inclusive, and day 189. The choice of a long or short position on day 63 is made by selecting the position that would have been held on day 189 for the previous contract if further trading in that contract is allowed. From day 64 onwards, the previous contract is ignored in all the calculations. Of course the trading positions from days 63 to 188 depend on the channel length L.

Expected Payoffs

The price model given in the third section is simulated for 500,000 contracts using the price parameter values given in that section. The financial consequences of channel rule trading decisions are calculated for all channel lengths L from 2 to 60 days inclusive. Table I summarizes the average gross payoff $\overline{G}$, the average number of trades $\overline{N}$ and the average net payoff $\overline{R}$ when transactions costs are 0.2% per round-trip.

TABLE I

Channel Trading Results as the Channel Length L Varies

L days	$\overline{G}$ %	$\overline{N}$	$\overline{R}$ %	P %
2	2.44	35.68	−4.69	57.96
3	2.85	24.87	−2.12	59.35
4	3.13	19.08	−0.69	60.28
5	3.32	15.50	0.23	60.95
6	3.47	13.08	0.85	61.44
7	3.58	11.35	1.31	61.82
8	3.66	10.04	1.66	62.09
9	3.73	9.02	1.92	62.29
10	3.76	8.21	2.12	62.43
11	3.79	7.55	2.28	62.52
12	3.81	7.01	2.41	62.58
13	3.80	6.55	2.49	62.60
14	3.81	6.15	2.58	62.60
15	3.81	5.81	2.65	62.57
16	3.79	5.51	2.69	62.53
17	3.77	5.25	2.72	62.47
18	3.76	5.02	2.75	62.40
19	3.73	4.81	2.77	62.33
20	3.71	4.63	2.78	62.23
21	3.68	4.46	2.79	62.14
22	3.66	4.31	2.79	62.04
23	3.63	4.17	2.80	61.94
24	3.60	4.04	2.79	61.84
25	3.56	3.93	2.78	61.73
26	3.53	3.82	2.77	61.61
27	3.50	3.72	2.76	61.50
28	3.47	3.63	2.74	61.38
29	3.43	3.54	2.73	61.27
30	3.40	3.46	2.71	61.15
31	3.37	3.39	2.69	61.04
32	3.33	3.32	2.67	60.92
33	3.30	3.25	2.65	60.80
34	3.27	3.19	2.63	60.68
35	3.24	3.13	2.61	60.57
36	3.20	3.07	2.59	60.45
37	3.16	3.02	2.56	60.35
38	3.13	2.97	2.54	60.23
39	3.10	2.92	2.51	60.11
40	3.06	2.88	2.49	60.00

TABLE I (Continued)

L *days*	$\overline{G}$ %	$\overline{N}$	$\overline{R}$ %	P %
41	3.03	2.84	2.46	59.89
42	3.00	2.80	2.44	59.78
43	2.97	2.76	2.42	59.67
44	2.94	2.72	2.40	59.58
45	2.91	2.69	2.37	59.48
46	2.88	2.65	2.35	59.37
47	2.85	2.62	2.32	59.27
48	2.82	2.59	2.30	59.17
49	2.79	2.56	2.28	59.07
50	2.76	2.53	2.26	58.97
51	2.74	2.50	2.24	58.88
52	2.71	2.47	2.21	58.79
53	2.68	2.45	2.19	58.70
54	2.65	2.42	2.17	58.61
55	2.63	2.40	2.15	58.51
56	2.60	2.38	2.13	58.42
57	2.58	2.35	2.10	58.34
58	2.55	2.33	2.09	58.26
59	2.53	2.31	2.06	58.18
60	2.50	2.29	2.04	58.10

The numbers tabulated are the average gross payoff $\overline{G}$, the average number of trades $\overline{N}$, the average net payoff $\overline{R}$ when transactions costs c are 0.2%, and the probability P of taking the correct trading decision. These figures are calculated by simulating 500,000 contracts, each for 6 months. The correlation between daily futures returns on different days is always positive. The maximum correlation of 0.019 occurs when the time lag is one day.

The best choice of L is between 12 and 15 inclusive when there are no transaction costs. The expected value of G is then estimated to be 3.81%. This estimate implies that when T-Bills return 3% over 6 months, an average return of 6.81% can be obtained over 6 months on capital used to purchase T-Bills which are then used as margin collateral, always assuming that market prices are generated by the rules used for the simulations and that traders know the optimal range for L. The estimated standard deviation of G is 8.68% for the best L. The values of G have minimal autocorrelation (e.g., 0.008 for adjacent contracts) so the standard error of $\overline{G}$ is estimated to be 0.01%.

Figure 1 shows plots of $\overline{G}$ and $\overline{R}$ against L. It is striking how slowly $\overline{G}$ decreases as L is moved away from the optimal range. The expected value of G is within 90% of its maximum value whenever L is in the interval from 6 to 29.

The definition of the channel rule implies that the number of trades is a decreasing function of L. Very approximately the number of trades

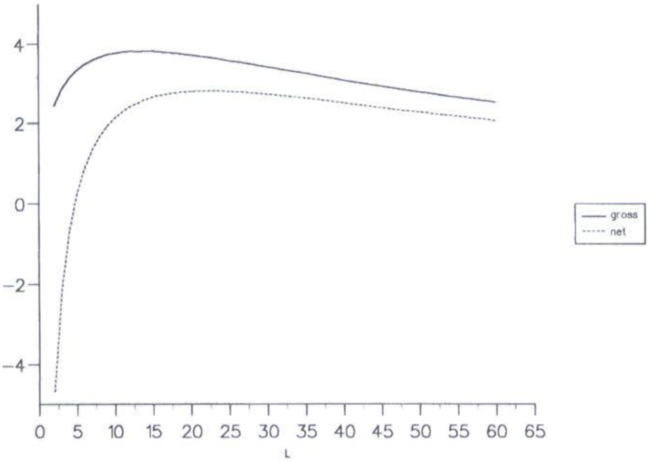

FIGURE 1
Gross and net average percentage payoffs. Vertical axis: average percentage payoff. Horizontal axis: channel length L (days).

is halved when L is doubled. There is about one trade per month, on average, when L maximizes the average gross payoff.

As transaction costs increase, so too does the optimal value of L. The optimal range is 21 to 24 when costs are 0.2% per trade. The estimated expected value of R is then 2.79% (standard error 0.01%) and 90% or more of the optimal expected net payoff can be obtained when L is in the interval from 14 to 38. At the optimum L, one-quarter of the gross payoff is lost to transaction costs and the number of trades over 6 months averages 4.2 and has standard deviation equal to 1.5. Figure 1 shows that $\overline{R}$ decreases sharply as L falls below 10 due to the rapidly increasing number of trades.

Small positive expected net payoffs can be obtained even for transaction costs as high as 1% per trade, providing L is at least 32.

Probabilities of Correct Trading Decisions

Trading rules aim to take long [short] positions when the futures price is expected to rise [fall]. The success of a trading rule applied to the simulated price process depends on how often long [short] positions are taken when the unobservable AR(1) component μ_t in returns x_t is positive [negative]. Trading rules use the returns history, $I_{t-1} = \{x_{t-1}, x_{t-2}, \ldots\}$, to implicitly provide forecasts of the sign of μ_t. The accuracy of a rule's sign forecasts can be measured by defining

$$d_t = 1 \quad \text{if long from close } t-1 \text{ to close } t$$

$$= -1 \quad \text{if short}\ldots \tag{11a}$$

$$s_t = \text{sign}(\mu_t) = 1 \quad \text{if } \mu_t > 0$$

$$= -1 \quad \text{if } \mu_t < 0 \tag{11b}$$

and then estimating

$$P = \text{Prob}(d_t = s_t) \tag{12}$$

Table I shows estimates of the probability P as a function of L. Not surprisingly the highest probabilities occur for the same L values that give the maximum expected gross payoff. The maximum probability of taking the correct trading decision is 62.6% and the probability is more than 60% for any L between 4 and 39 inclusive; here correct means being long [short] when μ_t is positive [negative]. There is an almost linear relationship between $\overline{G}$ and P. For all L considered, $100\overline{G}/(P - 0.5)$ is between 30 and 31.

Returns follow an ARMA(1, 1) process for the model simulated. An ARMA trading rule would use a forecast $\hat{x}_t$ of x_t, made at time $t - 1$ using I_{t-1}, as in Taylor (1988, 1992a, b). Let $d_t^* = \text{sign}(\hat{x}_t)$. It is of interest to compare

$$P^* = \text{Prob}(d_t^* = s_t) \tag{13}$$

with P. The usual optimal linear forecast, which ignores the ARCH effects in the simulated returns, achieves $P^* = 63.6\%$ when the correct AR and MA parameters are used to define the forecasts. This is similar to the theoretical probability (62.1%) which can be calculated for a Gaussian ARIMA process when the price volatility is constant ($\beta = 0$). Revising the forecasts to take account of conditional heteroscedasticity is possible in several ways. The best sign forecasts obtained have $P^* = 63.8\%$ compared with $P = 62.6\%$. It is concluded that the channel rule is only slightly inferior to ARMA forecasts when predicting the sign of the unobservable AR(1) component μ_t.

Sensitivity to Price Model Parameter Changes

All of the results summarized in Table I are for the chosen price process with the five parameters given by: $A = 0.02$, $p = 0.95$, $\alpha = -5.15$, $\beta = 0.422$, and $\phi = 0.973$. Table II summarizes further results when one or two parameters are changed. The price follows a random walk when $A = 0$ and, not surprisingly, $\overline{G}$ is almost zero for all L, while $\overline{R}$ is negative and is maximized by making L as large as possible.

TABLE II

Selected Channel Trading Results for Different Price Process Parameters

Variation	$\overline{G}$ %	*Best L*	$\overline{R}$ %	*Best L*	P %
None	3.81	14	2.80	23	62.6
Random walk	0.05	5	−0.48	60	50.0
$A = 0.005$	1.16	19	0.43	40	57.1
$A = 0.01$	2.13	14	1.28	26	59.6
$A = 0.03$	5.43	14	4.36	18	64.6
$p = 0.8$	2.10	4	0.25	19	56.7
$p = 0.9$	2.93	10	1.52	19	59.4
$p = 0.975$	4.84	24	4.17	31	66.2
$\phi = 0.95$	3.83	13	2.81	24	62.5
$\phi = 0.99$	3.92	14	2.96	24	62.8
$\beta = 0.3$, $\alpha = -5.062$	3.64	14	2.71	27	61.8
$\beta = 0.7$, $\alpha = -5.462$	4.64	12	3.50	18	65.5

Sensitivity of (a) the maximum average gross payoff $\overline{G}$ and the best L; (b) the maximum average net payoff $\overline{R}$ and the best L when transactions costs c are 0.2%; and (c) the probability P of taking the correct trading decision for the best L in (a), to changes from the simulation parameters used to produce Table I.
The first row of the table is the base case defined by:

$$A = 0.02,\ p = 0.95,\ \alpha = -5.15,\ \beta = 0.422, \text{ and } \phi = 0.973.$$

The averages are calculated from 60,000 contracts for the random walks, 20,000 contracts for each of the other variations, and 500,000 contracts for the base case.

The remaining variations are evaluated for 20,000 contracts, using the same sequence of random numbers to generate the prices for all these variations. The first set of variations only change the value of A, which is defined in (7). The optimal average gross payoff is very roughly a constant times A over the range $0.005 \leq A \leq 0.03$. Likewise, the optimal average net payoff has an approximate linear relationship with A, although the intercept is obviously negative and the best L decreases as A increases. The second set of variations only change the value of p, which is defined in (5). Now the optimal average gross payoff is approximately a constant times $\ln(1 - p)$. From the first and second sets of variations, a general result about the optimal expected gross payoff from the channel rule may be conjectured: when the returns process has autocorrelations ρ_τ satisfying $\rho_\tau > \rho_{\tau+1} > 0$ for all positive lags τ, then the optimal expectation is approximately some constant times the following function f of ρ_1 and $\rho_\Sigma = \sum_{\tau=1}^{\infty} \rho_\tau$:

$$f = \rho_1 \rho_\Sigma \left[\frac{\ln(\rho_\Sigma) - \ln(\rho_1)}{\rho_\Sigma - \rho_1} \right] \tag{14}$$

The multiplicative constant will depend primarily upon the unconditional standard deviation of the returns.

Table II also shows that the results are not sensitive to ϕ which measures how quickly volatility is expected to revert towards its median level. The final two rows of the table are for changes in both α and β which leave the unconditional standard deviation unchanged. The price volatility varies more through time as β increases and then the channel rule becomes slightly more profitable.

TRADING RESULTS WHEN L IS CHOSEN FROM PAST PRICES

Assumptions

The usual ex ante methodology for evaluating trading rules requires parameters to be chosen by optimization over some learning period. The average net payoff for contract $C + 1$ in a series of contracts is now estimated as a function of the amount of information, C contracts, used to choose L. The learning period ranges from $C = 1$ to $C = 30$ contracts inclusive, each traded for 6 months. The trading parameter L is chosen to be that value which maximizes the average net payoff during the learning period. Transaction costs are assumed to be 0.2% per trade. The choice of initial position on day 63 for the first contract in the learning period is arbitrary; a long position is always selected.

The price process of the third section is simulated with the parameters given in that section. It is therefore assumed that the process and the parameters do not change during the learning period and any subsequent ex ante evaluations.

Expected Net Payoffs

Prices for 30 consecutive contracts are simulated to give choices for L as C increases from 1 to 30. Three thousand replications are used to estimate the probability $p(l, C)$ of choosing $L = l$ when C contracts are used to make the choice. From Table I one knows the expected net payoff for any contract when $L = l$. Call this $E(R \mid l)$. The unconditional expected net payoff for contract $C + 1$ is then:

$$E(R \mid \text{learn from } C \text{ contracts}) = \sum_{l=2}^{60} p(C,l)E(R \mid l) \qquad (15)$$

Table III lists the estimates of the unconditional expected net payoff for contract $C + 1$ based upon learning from C contracts. The ultimate limit of these expectations, as $C \rightarrow \infty$, is 2.80% from Table I. It can be seen from Table III that what can be learned from past prices is learned remarkably quickly. Fifteen years of prices would give an expected net payoff of 2.49%; but 2.19% is expected if only 2 years of prices are used. More than two-thirds of the expected payoff from an infinite price history can be achieved from only one year of prices!

Other Criteria

Table III also gives the expected gross payoff and the probability of taking the correct trading decision (as discussed in the fourth section) when learning from C contracts. These numbers are almost constant. There is, of course, no reason why either the gross expectation or the probability should quickly improve when the objective is to maximize the net expectation.

The chance of taking the correct decision is 61.2% when seeking to use a few contracts to learn how to maximize net payoffs. This is close to the 61.9% attainable from an infinite learning set. It is conjectured that ARMA decision rules need more contracts to achieve a similar proportion of their predictive potential. If true, this could explain why channel and ARMA rules have comparable expected payoffs when learning from a moderate number of contracts with ARMA rules only superior for very long learning periods.

The Distribution of Choices for *L*

Some information about the distribution of the optimized choices for L is provided in Table III. The average value chosen for L varies between 25 and 27 when the learning period is between 5 and 30 contracts. As $C \rightarrow \infty$, the average converges to 23. The final two columns show that the inter-quartile ranges of the distributions are considerable. Learning from 4 years of prices would give half the choices for L either below 12 or above 38. Figure 2 shows how often the optimized channel

TABLE III
Channel Trading Results When L Depends on Past Prices

Learning Period C	*Expected Net Payoff* %	*Expected Gross Payoff* %	*Chance Correct Decision* %	*Choice for L*		
				Average	*Q1*	*Q3*
1	1.73	3.45	61.34	18.48	7	26
2	1.98	3.42	61.25	21.88	8	32
3	2.09	3.41	61.21	23.41	9	34
4	2.19	3.41	61.19	24.68	11	36
5	2.20	3.39	61.15	25.25	11	37
6	2.25	3.40	61.16	25.71	11	38
7	2.31	3.39	61.15	26.31	12	38
8	2.31	3.40	61.16	26.34	12	38
9	2.31	3.39	61.15	26.41	12	38
10	2.34	3.40	61.17	26.49	12	38
11	2.35	3.40	61.16	26.83	13	39
12	2.37	3.40	61.15	27.05	13	39
13	2.38	3.40	61.17	27.04	13	39
14	2.40	3.40	61.18	27.02	14	38
15	2.41	3.41	61.19	27.07	14	38
16	2.42	3.39	61.15	27.57	14	39
17	2.42	3.40	61.18	27.26	14	38
18	2.43	3.40	61.19	27.34	14	38
19	2.44	3.41	61.21	27.22	14	38
20	2.46	3.41	61.22	27.27	14	38
21	2.46	3.41	61.20	27.52	14	38
22	2.46	3.41	61.19	27.62	14	38
23	2.46	3.41	61.21	27.38	14	38
24	2.47	3.42	61.22	27.35	14	38
25	2.47	3.41	61.21	27.47	14	38
26	2.48	3.42	61.22	27.46	14	37
27	2.48	3.42	61.22	27.48	14	37
28	2.48	3.42	61.24	27.31	14	37
29	2.49	3.42	61.23	27.53	15	38
30	2.49	3.42	61.23	27.54	15	37
infinity	2.80	3.63	61.94	23	23	23

Expected net and gross payoffs when C contracts are used to choose the value of L for the next contract by maximizing the average net payoff for the C contracts. The expectations are estimated from 3000 replications of the learning process.

length equals the possible integers between 2 and 60 when $C = 2$ and $C = 20$. From Figure 1 it is clear that the channel rule will be most successful when small channel lengths L are avoided. Figure 2 shows that small values of L are often chosen when $C = 2$, but are infrequently chosen when $C = 20$.

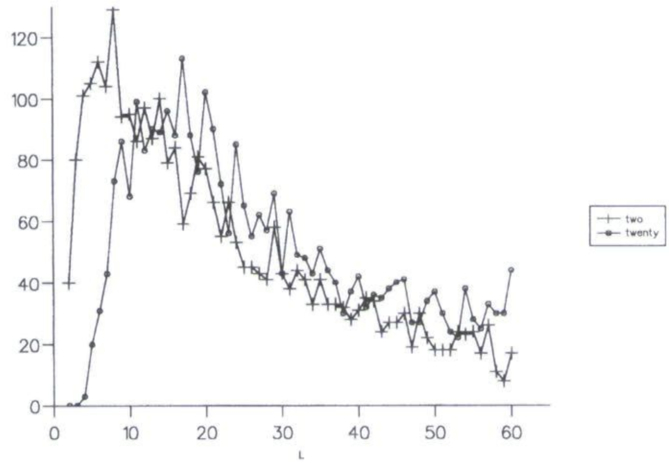

FIGURE 2
Distribution of optimized channel length (learning from 2 and 20 contracts). Vertical axis: frequency. Horizontal axis: channel length L (days).

FURTHER TRADING RESULTS FOR CURRENCY FUTURES

It is important to report the results of all trading rule evaluations, otherwise there is the possibility of biased conclusions with the direction of the bias depending on the inclination of the researcher either for or against the efficient market hypothesis. Taylor (1992a) gives channel results for IMM currency futures up to November, 1987. The channel results are here extended to November, 1990.

ARIMA trading rules would have given positive net returns for Sterling, Deutschemark, and Swiss Franc futures from December, 1978 to November, 1981 (contracts delivered in 1979–1981) according to Taylor (1986, p. 220). That methodology was different from the one applied here; although the results suggest an annual, average, net payoff of about 7%. Both technical and ARIMA trading rules would have been profitable from December, 1981 to November, 1987 (contracts delivered in 1982–1987) as discussed in the second section. Now consider whether positive average net payoffs persist into more recent years.

Channel trading results are calculated for March, June, September, and December contracts delivered in 1982–1990, with each contract traded for the 3 months preceding the delivery month. These trading periods are chosen to avoid any problems of relatively illiquid trading many months from delivery. Results are obtained for Sterling, Deutschemark, Yen, and Swiss Franc futures using daily closing prices. The choice of L for a contract is given by optimizing the average net payoff using all previous contracts for the currency, commencing with the March, 1980 contract. Transactions costs are assumed to be 0.2% of the futures price of the currencies traded. This is a conservative figure: most non-floor traders would pay less than 0.2% in costs. The initial position for each contract, long or short, is determined using the method described in the fourth section.

The average net payoff, over 3 months, is calculated for 3-year periods by averaging across 48 contracts, 12 for each of the 4 currencies. To calculate a standard error for a 3-year average, across currency averages are calculated for each of the 12 delivery months and these averages are assumed to be a random sample from some distribution. The conventional t-ratio used to test the null hypothesis that the average net payoff is at most zero is given by dividing the average net payoff by its estimated standard error. The results for the three, 3-year periods and all 9 years are as follows.

Delivery months	*Average net payoff %*	*Standard error %*	*t-ratio*
1982—1984	0.63	0.59	1.06
1985—1987	2.92	1.43	2.05
1988—1990	1.63	1.80	0.91
1982—1990	1.73	0.82	2.11

During the most recent period, 1988–1990, the average net payoff is positive, but statistically insignificant. Nevertheless, these results extend the sequence of profitable 3-year periods from three such periods (1979–1981, . . , 1985–1987) to four (1979–1981, . . , 1988–1990). It can be seen that net payoffs are more volatile during 1988–1990 than during the other two periods.

The channel results for all 9 years, 1982–1990, give an average net payoff equal to 6.9% per annum. A standard 95% confidence interval for this average is from 0.3% to 13.6%. A one-tail test is appropriate when seeking to decide whether the channel rule is worthless (the null hypothesis) or profitable on average (the alternative hypothesis). The null is then rejected at the 2.5% significance level. It is obvious that the best trading opportunities occur during the period 1985–1987 and this period alone can explain the significant t-ratio for all 9 years. Systematic equity risk cannot explain technical trading profits for currencies [LeBaron (1991); Taylor (1992a)].

The average net payoffs for real currency futures can be compared with the average figures obtained from the Monte Carlo simulations, but it must be remembered that the averages for real prices have high standard errors. The 9-year average of 6.9% per annum can be compared with 5.6% per annum for the best L in Table I. The Monte Carlo study might never have taken place if the results for 1982–1987 had not been so encouraging. A comparison of 1988–1990 with the Monte Carlo results from learning over 8 years may, therefore, be more objective. This gives 6.5% per annum for 1988–1990 compared with 4.8% expected from Table III. These figures are not far apart considering the standard errors involved.

CONCLUSIONS

Technical analysis often deserves unflattering criticism and its popularity has puzzled many researchers. Neftci (1991, p. 549) aptly commented that "The attention that technical analysis receives from financial markets is somewhat of a puzzle." and "...traders use technical analysis in day to day forecasting although it bears no direct relationship to Wiener–Kolmogorov prediction theory." The theoretical results presented in this article for the channel trading rule show that technical analysis can be constructive in certain circumstances and is not necessarily inferior to ARIMA prediction methodology. Indeed, the channel rule may be superior when a trading objective is evaluated because it may require less information to learn about a satisfactory value for its one parameter than an ARIMA rule needs to find satisfactory estimates of its AR and MA parameters. The channel rule does not provide predictions of future prices. It only predicts the sign of price changes, positive or negative; but this is sufficient to make net trading profits when the sign prediction is frequently correct. As shown in the fourth and fifth sections, more than 60% of the sign predictions can be correct for the channel rule even when the maximum correlation between returns on different days is less than 0.02.

Currency futures prices have not followed random walks [Taylor (1986, 1992b)] and there have been several studies which support the idea that technical rules can make net profits at currency markets. Some of these studies are Sweeney (1986); Bilson (1990); LeBaron (1991); and Taylor (1992a). The results of this study for an extended period for currency futures from 1988 to 1990 are not decisive, yet they give more credibility to the notion that the channel rule makes net profits. Net trading profits for currency futures transactions are now reported for four consecutive 3-year periods. These results may be explained by risk, by the activities of central banks, or by an inefficient market.

BIBLIOGRAPHY

Alexander, S. S. (1961): "Price Movements in Speculative Markets: Trends or Random Walks?," *Industrial Management Review,* 2:7–26.

Beja, A., and Goldman, M. B. (1980): "On the Dynamic Behavior of Prices in Disequilibrium," *Journal of Finance,* 34:235–247.

Bilson, J. F. O. (1990): "Technical Currency Trading," in *The Currency-Hedging Debate,* Thomas, L. R. (ed.), London: IFR Publishing Ltd., pp. 257–275.

Bollerslev, T., Chou, R. Y., and Kroner, K. F. (1992): "ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics,* 52:5–59.

Brock, W., Lakonishok, J., and LeBaron, B. (1992): "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance,* 47:1731–1764.

Chesney, M., and Scott, L. O. (1989): "Pricing European Currency Options: A Comparison of the Modified Block–Scholes Model and a Random Variance Model," *Journal of Financial and Quantitative Analysis,* 24:267–284.

Harvey, A. C., Ruiz, E., and Shephard, N. G. (1994): "Multivariate Stochastic Variance Models," *Review of Economic Studies,* forthcoming.

Hodrick, R. J., and Srivastava, S. (1987): "Foreign Currency Futures," *Journal of International Economics,* 22:1–24.

Irwin, S. H., and Uhrig, J. W. (1984): "Do Technical Analysts Have Holes in Their Shoes?," *Review of Research in Futures Markets,* 3:264–277.

Kim, C. M. (1989): "Volatility Effect on Time Series Behavior of Exchange Rate Changes," Working Paper, University of Chicago and Korea Institute for International Economic Policy.

LeBaron, B. (1991): "Technical Trading Rules and Regime Shifts in Foreign Exchange," Working Paper, Social Systems Research Institute, University of Wisconsin.

LeBaron, B. (1992): "Some Relations between Volatility and Serial Correlations in Stock Market Returns," *Journal of Business,* 65:199–219.

Lukac, L. P., Brorsen, B. W., and Irwin, S. H. (1986): "A Comparison of Twelve Technical Trading Systems with Market Efficiency Implications," Unpublished manuscript, Department of Agricultural Economics, Purdue University.

Lukac, L. P., Brorsen, B. W., and Irwin, S. H. (1988): "A Test of Futures Market Disequilibrium Using Twelve Different Trading Systems," *Applied Economics,* 20:623–639.

Lukac, L. P., and Brorsen, B. W. (1990): "A Comprehensive Test of Futures Market Disequilibrium," *Financial Review,* 25:593–622.

Mark, N. C. (1988): "Time-Varying Betas and Risk Premia in the Pricing of Forward Foreign Exchange Contracts," *Journal of Financial Economics*, 22:335–354.

Melino, A., and Turnbull, S. M. (1990): "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 45:239–265.

Neftci, S. N. (1991): "Naive Trading Rules in Financial Markets and Wiener–Kolmogorov Prediction Theory: A Study of Technical Analysis," *Journal of Business*, 64:549–571.

Sentana, E., and Wadhwani, S. (1990): "Feedback Traders and Stock Return Autocorrelations: Evidence from a Century of Daily Data," Working Paper, London School of Economics.

Sweeney, R. J. (1986): "Beating the Foreign Exchange Market," *Journal of Finance*, 41:163–182.

Taylor, S. J. (1986): *Modelling Financial Time Series*. Chichester: Wiley.

Taylor, S. J. (1988): "How Efficient are the Most Liquid Futures Contracts?—A Study of Treasury Bond Futures," *The Review of Futures Markets*, 7:574–592.

Taylor, S. J. (1992a): "Rewards Available to Currency Futures Speculators: Compensation for Risk or Evidence of Inefficient Pricing?," *Economic Record*, 68 (Supplement):105–116.

Taylor, S. J. (1992b): "Efficiency of the Yen Futures Market at the Chicago Mercantile Exchange," in *Rational Expectations and Efficiency in Futures Markets*, Goss, B. A., (ed.), London: Routledge, pp. 109–128.

Taylor, S. J. (1994): "Modelling Stochastic Volatility," *Mathematical Finance*, forthcoming.