The Economic Journal, **109** (*October*), 636–661. © Royal Economic Society 1999. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 350 Main Street, Malden, MA 02148, USA.

METHODICAL MADNESS: TECHNICAL ANALYSIS AND THE IRRATIONALITY OF EXCHANGE-RATE FORECASTS*

P. H. Kevin Chang and Carol L. Osler

Substantial empirical research documents that exchange-rate forecasts are not formed rationally. This paper identifies a common technical trading signal, the head-and-shoulders pattern, as a potential source of departures from rationality in exchange-rate forecasts. Forecasts based on this pattern are evaluated for daily dollar exchange rates over 1973 to 1994, using two criteria for rationality: profitability and efficiency. Resulting profits, replicable in real-time, are tested for statistical significance using a bootstrap technique. We find that the rule is profitable, but not efficient, since it is dominated by simpler trading rules.

Since Goodman (1979) and Levich (1980) found that the forward rate outperformed most foreign exchange forecasting services over 1973–8, substantial evidence has accumulated that currency forecasts are not 'rational' because they incorporate predictable errors (see, for example, Frankel and Froot, 1987). Despite the plethora of such evidence, little is known about the sources of predictable exchange-rate forecast errors. This paper documents a specific behaviour common among foreign exchange market participants that generates forecast errors. In particular, we show that a certain, widespread approach to forecasting the direction of short-term exchange-rate movements relies exclusively upon irrelevant information. The information in question is the observation of a nonlinear pattern in recent exchange-rate movements called a 'head-and-shoulders', which occurs when the second of three consecutive peaks is higher than the first and third. After identifying such a pattern, forecasters anticipate a reversal of any previous sustained trend.

A broad definition of rationality is provided by McCallum (1980), for whom a rational agent 'behave[s] purposefully in collecting and using information.' Similarly, Friedman (1979) states that rational 'agents use efficiently whatever information is available.' There is general agreement that this requires the absence of *ex ante* predictable forecast errors. Applying the concept of rationality to any practical context, however, requires one to be more specific about the forecaster's goal. A common goal is to generate point forecasts as close as possible to the realised outcome, a goal which we will refer to as 'numerical accuracy', following Levich (1980). A forecast whose goal is numerical accuracy must be unbiased and efficient in order to be rational. A biased forecast could be costlessly improved by noting and eliminating the bias. An inefficient forecast could be costlessly improved by using a different information set.

Direct support for the hypothesis that exchange-rate forecasts do not conform to the requirements of rationality has come from numerous empirical

* The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

papers that test their accuracy according to the criteria of unbiasedness and efficiency. In tests for bias, which typically involve a regression of actual exchange-rate changes against forecasts by survey respondents, the null hypothesis of no bias is consistently rejected across currency, sample period, and firm compiling the survey (Dominguez, 1980; Blake *et al.* 1986; Cavaglia *et al.* 1994). More subtle tests of bias, conducted by Frankel and Froot (1987), show that forecasts place incorrect weights on potential indicators of future exchange rates. Ito (1990) shows that people's expectations are biased in the direction that benefits them. Tests for rationality based on the efficiency criterion require forecast errors to be orthogonal to all available information, such as interest rates. As discussed in Takagi (1991), relatively powerful efficiency tests consistently reject the null with respect to exchange-rate forecasts.

Further, indirect support for the hypothesis that exchange-rate forecasts are not rational has come from the so-called 'forward premium anomaly', or the repeated rejection of the joint hypothesis of uncovered interest parity and rational expectations. Since survey data indicate that uncovered interest parity itself may be consistent with the data, at least up to a constant risk premium (Froot and Frankel, 1989), one infers that the failure of the joint hypothesis may well be due to the failure of rational expectations. Further support for this idea is provided by Lewis (1996), who notes that the forward premium anomaly is only partially explained by alternative interpretations of rationality involving learning or peso problems.

Despite such evidence against exchange-rate forecast rationality, little is known about the specific market behaviours that give rise to predictable forecast errors. Technical analysis, an approach to forecasting prices in which the sole information sources are historical patterns of prices and trading volume, seems a natural place to look. It is almost universally used by practitioners in formulating short-term exchange-rate expectations: over 90% of participants in the London and Hong Kong markets rely on technical strategies (Allen and Taylor, 1990; Lui and Mole, 1996). Nonetheless, many economists consider technical analysis as implausible or far-fetched: Malkiel (1990), for example, considers it analogous to 'alchemy'.

Foreign exchange market practitioners use many technical trading techniques. Some of these techniques rely on calculated constructs like moving averages and others rely on visual chart patterns like the head-and-shoulders. According to technical traders, a head-and-shoulders pattern following a sustained up-trend indicates an imminent down-trend in prices, and they recommend selling the currency immediately after such a pattern is complete. The pattern is considered by practitioners to be one of the most reliable of all chart patterns.

Since the goal of technical trading is profitability, the best technical forecast is not necessarily the one which is most numerically accurate. Instead, the best technical forecast is the most profitable one, and profitability must replace unbiasedness as the first criterion by which to judge the price forecasts implicit in any technical trading rule. (This difference is elaborated in Leitch and Tanner (1991)).

The efficiency criterion remains a requirement for rationality regardless of a forecast's objective-any forecast that fails to use all available information, or which uses irrelevant information, could be costlessly improved and cannot be considered rational. In the case of technical trading, inefficiency can be demonstrated if the trading strategy is dominated completely by another strategy.

To test the rationality of forecasts based on head-and-shoulders patterns, we use daily spot rates for six currencies vs. the dollar: the yen, mark, Canadian dollar, Swiss franc, French franc, and U.K. pound. Our data cover 19 March 1973 to 13 June 1994, a twenty-one year period with over 5,500 daily observations.[1] Currency markets seem especially appropriate for testing technical signals because of their very high liquidity, low bid-ask spreads, and round-the-clock decentralised trading. Further, technical analysts claim that 'the principles that underlie analysis of currencies from a technical aspect are basically the same as those used in any other financial market or for individual stocks' (Pring, 1985, p. 466).

Our results indicate that the head-and-shoulders pattern is not profitable for four out of the six currencies analysed, and thus fail the first condition for rationality for those currencies. For the mark and the yen, however, the head-and-shoulders pattern produces statistically significant profits, so that the first condition for rationality is satisfied. However, for these two currencies, trading based on head-and-shoulders patterns is dominated by simpler trading strategies, so the pattern is an inefficient forecasting tool. Continued wide reliance on input from head-and-shoulders patterns thus appears to constitute a source of predictable errors in currency forecasts.

Note that this paper documents that foreign exchange forecasts do not conform to the requirements of rationality *without relying on survey data*, unlike most earlier evidence suggesting that foreign exchange forecasts are not generally formed rationally (Dominguez, 1980; Blake and Brasse, 1986; Cavaglia *et al.* 1994; Frankel and Froot, 1987; Ito 1990; Takagi, 1991). Since the reliability of survey data is open to question, many international economists remain unconvinced of the relevance of imperfect rationality for the functioning of currency markets. By showing that a widely practiced technical trading strategy incorporates imperfectly rational exchange-rate forecasts, the results of the present paper support the view that deviations from perfect rationality may affect exchange-rate dynamics.

We know of only three studies that evaluate the visual, nonlinear patterns that are the focus of this paper, and none of these examine them in the context of currency markets. The paucity of research on this subject, which contrasts sharply with the ubiquity of their market use, may be attributable to the highly nonlinear and complex nature of these rules. The three existing

[1] Japanese data were sampled as the 'Tokyo market closing middle rate'. Mark data were taken at 1:00 p.m. Frankfurt time. The pound and Swiss franc rates represent 'current market rates at 2:15 p.m. Swiss time'. French franc rates are 'indicate rates at 2:15 p.m'. Canadian dollar rates represent 'London market middle rates at around noon Swiss Time'. We use the average of daily bid and ask quotes. (Source: BIS Data Bank, Courtesy of the Federal Reserve Bank of New York.)

studies come to different conclusions: Osler (1998) tests the head-and-shoulders pattern in U.S. equity markets, and finds that they have no predictive power. The same conclusion is drawn in Levy (1971), which tests the predictive power of 32 'five-point chart patterns', including the head-and-shoulders, in U.S. stock prices. This study may not be reliable, however, because it does not follow the rules for head-and-shoulders trading delineated by practitioners. Brock *et al.* (1992) finds that breakouts from observed trading ranges are meaningful predictors of short-term returns in the Dow Jones Index during 1897–1986, a result that corroborates technicians' claims regarding 'support' and 'resistance' levels. In short, research on visual trading patterns is both scarce and inconclusive and, as noted by Neftci (1991), these strategies represent 'a broad class of prediction rules with unknown statistical properties'.

Section 1, which follows, describes our methodology for evaluating the profitability of trading based on the head-and-shoulders pattern. We construct an objective algorithm for identifying and trading on head-and-shoulders patterns, and use a bootstrap technique to ascertain the statistical significance of the associated profits. Section 2 describes our data and shows that head-and-shoulders trading is not profitable for the U.K. pound, the Swiss franc, the French franc, and the Canadian dollar, but it is indeed profitable for the mark and the yen. Thus exchange-rate forecasts based on head-and-shoulders patterns fail the first of the two necessary conditions for rationality listed above for four currencies, but satisfy that first condition for the two most widely traded currencies besides the U.S. dollar itself. Section 3 shows that the head-and-shoulders trading rule is dominated by other, simpler trading rules, implying that the information the pattern provides for exchange-rate forecasts is irrelevant. Thus the head-and-shoulders pattern is not an efficient tool for forecasting the mark and the yen, and forecasts based on the pattern fail the second necessary condition for rationality. Section 4 summarises the results.

1. Testing Profitability: Methodology

This section examines whether forecasts based on head-and-shoulders patterns conform to the first necessary condition for rationality: profitability. After creating a replicable pattern-recognition strategy, we establish rules for entering and exiting positions. Finally, we statistically evaluate the profits derived from these strategies.

1.1. *Identifying a Head-and-Shoulders Pattern*

Eight technical analysis manuals were consulted to ensure that the head-and-shoulders identification algorithm used in this study conforms closely to the descriptions of practicing technical analysts. The manuals differ only trivially in their description of a head-and-shoulders pattern and in their recommended strategy for trading once the pattern has been identified. As illustrated in Fig. 1, the pattern comprises a sequence of three peaks with the highest in the middle.[2] The left and right peaks are the 'shoulders', the centre peak is the 'head', and the straight line connecting the troughs on either side of the head is called 'the neckline'. The pattern is 'complete' or 'confirmed' when the price path croses the neckline after forming the right shoulder. Head-and-shoulders can occur both at peaks, where they are called 'tops', and at troughs, where they are called 'bottoms'. All indicators of head-and-shoulders tops apply equally, with 'troughs' replacing 'peaks' and *vice versa,* to head-and-shoulders bottoms. Head-and-shoulders is a reversal pattern, that is, a head-and-shoulders top indicates that an earlier upward trend is about to be reversed and *vice versa.*

To identify peaks and troughs in the data, we first trace out a zigzag pattern—a smoothed version of the original data consisting of peaks and troughs connected with diagonal trend lines. We define a peak as a local maximum at least $x\%$ higher than the preceding trough, and a trough as a local minimum at least $x\%$ lower than the preceding peak, where x is referred to as 'the cutoff'. The number of peaks and troughs in a given data sample is inversely related to the cutoff. Varying the cutoff generates sets of head-and-shoulders patterns of different sizes.

To identify head-and-shoulders patterns, we scan the data 10 times, with a different cutoff each time. The 10 different cutoffs are chosen with reference to each currency's volatility. Specifically, we set one standard deviation of daily exchange-rate changes as a lower bound for cutoff, and consider as well the following multiples of that standard deviation: 1.25, 1.50, 1.75, 2.00, 2.50, 3.00,

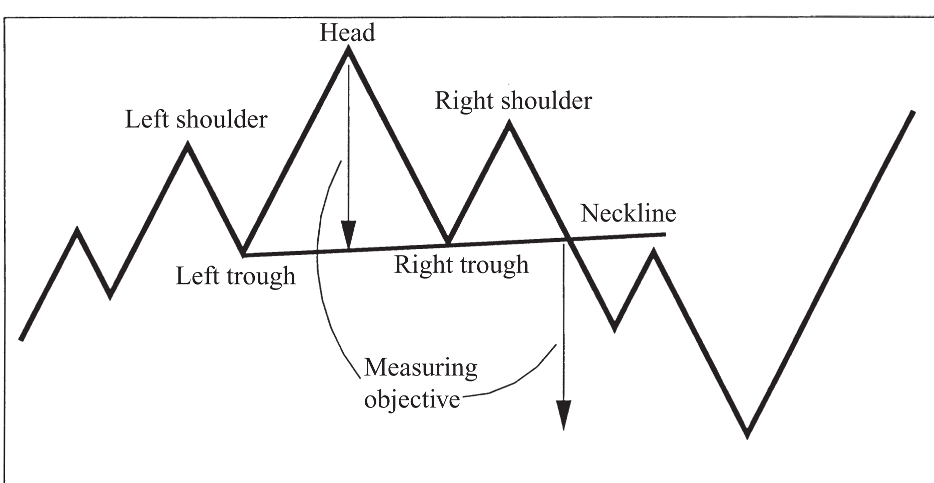

Fig. 1. *Hypothetical Head-and-Shoulders Pattern*

[2] Those manuals are Arnold and Rahfeldt (1986), Edwards and Magee (1966), Hardy (1978), Kaufman (1978), Murphy (1986), Pring (1985), Shabacker (1930), and Sklarew (1980).

3.50, 4.00, and 4.50.[3] Each time we scan the data at a new cutoff we eliminate duplicate patterns.[4]

To qualify as a head-and-shoulders, a given set of consecutive peaks must satisfy certain criteria identified only qualitatively in technical manuals. For example, the manuals only suggest that the pattern should not be very steep, without specifying what would be too steep. Details on how we formalise those criteria are provided in Appendix A. As discussed later, our results are not sensitive to the precise formalisation used.

Belief in the ability of head-and-shoulders patterns to forecast the direction of financial price movements originated with pre-1930 equity data (Shabacker, 1930). Since we use post-1972 exchange rate data (a later time period and a different set of financial prices), our test data are independent from those that originally suggested the pattern's significance, and our tests can be considered as 'out of sample'.

In considering a nonlinear pattern as a specific source of exchange-rate forecast errors, we contribute to a growing body of research on nonlinear dependence in financial prices. Early tests for the presence of nonlinearities, testing the null hypothesis of i.i.d. behaviour, indicate that nonlinearities are indeed present in stock markets (Hsieh, 1991) and in floating exchange rates (Hsieh, 1989). The form of these nonlinearities remains unclear. One potential source of nonlinearity is 'chaos', although the few available tests fail to confirm its existence in exchange rate data. Modelling financial prices as a GARCH process seems to capture some nonlinearities; more specifically, it is helpful for predicting volatilities (Hsieh, 1989).

The head-and-shoulders trading rule is one of a large class of 'nonlinear prediction rules' potentially deriving from nonlinear versions of structural models. These include the monetary model (Meese and Rose, 1991), target-zone models (Krugman, 1991), chaos models (Gilmore, 1991),[5,6] econometrically-based models with structural content such as ARCH-in-Mean (Diebold and Pauly, 1988), or other models such as the Self-Exciting Threshold

[3] Increasing the cutoff beyond 4.50 times the standard deviation of daily exchange rate changes resulted in very few additional positions taken.

[4] If a head-and-shoulders pattern using one cutoff suggested an entering position within two days of a previously identified entry date, we do not include the new position in our analysis. Our results are unaffected by whether we identify such positions with successively higher cutoffs or successively lower cutoffs.

[5] For more information on chaos, see Brock *et al.* (1991).

[6] Aczel and Josephy (1991) calculate correlation dimensions and associated standard errors for four European currencies and the Singapore dollar, all against the U.S. dollar. (The 'correlation dimensions' are intended to capture the fact that phase diagrams of chaotic series occupy a lower share of the phase space than do the phase diagrams of purely random processes.) Though they find that the correlation dimensions of the European currencies are all statistically indistinguishable from each other and statistically different from the Singapore dollar, their evidence does not rule out the possibility that nonlinearities in exchange rates are due to other, non-chaotic generating processes. Further, correlation dimensions are considered unreliable in small samples, and for this study the authors have only 115 observations (Hsieh, 1991). Gilmore (1991) considers an alternative approach to testing for chaos in exchange rates, the method of close returns, which he asserts is more robust in small samples. On a sample of 935 daily observations for four exchange rates vs the U.S. dollar, Gilmore finds no evidence to indicate the presence of chaos.

Autoregressive model (Krager and Kugler, 1993). Many of these models have been shown to fit the data with some acceptable level of explanatory power within sample, and some appear to be helpful in forecasting conditional exchange rate variances.

1.2. *Measuring Profits*

Once a head-and-shoulders pattern has been identified, we compute the profits of a hypothetical market participant acting on this information. Profits are calculated as the cumulative percent change in the exchange rate between entry and exit, adjusting the sign to reflect whether the simulated speculator was short or long. Note that these returns correspond to unlevered positions and that the trading strategy is carefully designed so that it would be replicable in real time.

Entry: The technical manuals explicitly state that one should not enter a position until the price line 'penetrates' the neckline.[7] We take as our entry price the recorded price on the day of the neckline's penetration. Since peaks and troughs are identified only after they occur, the pattern could cross the neckline before the price has declined by the cutoff percent required to identify the right shoulder as a peak. When this occurs, we enter only when cutoff has been reached to ensure that positions do not benefit from future information.[8]

Since technical analysis manuals are ambiguous about when to exit, we use two different approaches, one endogenous, the other exogenous.

An Endogenous Exit Rule: We exit when it would appear to real-world observers that the price has conclusively stopped moving in the predicted direction. For a head-and-shoulders top, we formalise this criterion by exiting on the day when a new trough is identified, i.e., when the price has risen above a local minimum by the cutoff percentage. (For a head-and-shoulders bottom the criteria are analogous.) The holding period depends on how prices evolve after the day of entry. There are two important exceptions to this endogenous exit rule:

Exception 1 – Stoploss: To limit losses, we follow standard trading practice and close all positions when the exchange rate moves sufficiently far in the 'wrong' direction. We set this stoploss limit at 1%, although the actual loss could exceed this because we exit at the day's recorded price. Since this

[7] In practice we find that 25-40% of all head-and-shoulders patterns are confirmed by penetration of the neckline, depending on the currency.

[8] Neftci (1991) asserts that acting on the head-and-shoulders pattern necessarily requires future information. This assertion could be correct only in the context of his particular representation of the head-and-shoulders pattern, which does not correspond to representations in the technical analysis manuals we studied. The assertion also relies on the requirement that the right and left trough be very nearly identical, which is not required in the manuals and which we do not require in our algorithm.

stoploss parameter is set somewhat arbitrarily, we examine the implications of the alternative value of $\frac{1}{2}\%$ in one of our sensitivity analyses.

Exception 2 – Bounce: The technical analysis manuals consistently stress that, after entry, prices may temporarily move back toward the neckline and that to exit at this point would be premature because the trend reversal would not yet be complete. Therefore, we allow for a possible 'bounce' or interruption in the reversal pattern. Specifically, after a head-and-shoulders top, if the trend turns upward before falling by at least 25% of the vertical distance from the head to the neckline (a distance referred to in the manuals as the 'measuring objective' or 'price objective'), then our simulated investors hold their positions unchanged until one of two conditions is met: (1) the price crosses back over the neckline by at least 1%, triggering the stoploss discussed above; or (2) a second trough (of any size) is reached in the zigzag.[9] (These rules hold with reverse sign for head-and-shoulders bottoms.) Though we attempted to ascertain a 'best' number for the bounce parameter via conversations with practicing technical analysts and a close reading of the manuals, these sources provided only vague guidance, and the 25% bounce parameter is somewhat arbitrary. For this reason, the implications of raising the parameter to 50% are examined in one of the sensitivity analyses.

An Exogenous Exit Rule: Alternatively, we close any position after an exogenously specified number of days (1, 3, 5, 10, 20, 30, and 60). For clarity we focus on *marginal* returns accruing to the position between one exit point and the next (e.g. after ten days, the returns accruing in days 11-20).

1.3. *Evaluating the Statistical Significance of Profits*

After running the head-and-shoulders identification and profit-taking programmes on the actual data, we evaluate the statistical significance of the resulting profits. Since the distribution from which such profits are drawn is not known, we derive an appropriate statistical distribution via a bootstrap methodology, following Levich and Thomas (1993), Brock *et al.* (1992), and Allen and Karjalainen (1993). This non-parametric approach to statistical testing provides confidence intervals consistent with the null hypothesis that head-and-shoulders trading is not profitable.[10]

Our application of the bootstrap method, described further in Efron (1979, 1982) and Efron and Tibshirani (1993), involves comparing actual profits with

[9] In practice we found that for most currencies the 'bounce' possibility was invoked about 15% of the time.

[10] An alternative approach to calculating confidence intervals would be to combine an assumed distribution of exchange rate changes and our profit-taking programme. There is no consensus in the profession, however, regarding the closest approximation to the distribution of exchange rate changes (Westerfield, 1977; Boothe and Glassman, 1987, Hsieh, 1988). For this reason we choose instead a strategy for finding confidence intervals that seems the most agnostic: the random walk.

the distribution of profits in 10,000 artificial exchange rate series, each the same length as the original series for that currency. These new series are constructed so that their mean, variance, skewness, and kurtosis are representative of the population that underlies the actual exchange rate series. Yet, there is a critical difference: the artificial series are constructed so that any head-and-shoulders pattern is meaningless, whereas in the original exchange rate series this may or may not be true. We run the head-and-shoulders identification and profit-taking programmes described above on each of the artificial series. The distribution of profits from these 10,000 simulated series provides us with confidence intervals against which to measure the corresponding profit values from *actual* exchange rates. If technical analysts are correct, then head-and-shoulders patterns predict future trends, and statistical tests should highlight a difference between the simulated and original data series.

In constructing the artificial exchange rate series, we begin with a null hypothesis for exchange-rate behaviour. Our baseline results use the random walk model, suggested by Mussa (1979), and Meese and Rogoff (1983). In particular, this means we begin each series at the level corresponding to the actual exchange rate on the first day of our sample period, and for each simulated series we choose each day's (percent) exchange-rate change by drawing with replacement from the original series of exchange-rate changes.

To confirm empirically that the random walk model describes our data with reasonable accuracy, we compute autocorrelation coefficients on daily returns for ten lags. These individual autocorrelation coefficients, reported in Table 1*a*, are all quite small. Though a few of them are statistically significant (using heteroskedasticity-corrected standard errors, as in Hsieh (1988)), this is not surprising in a sample of 60 statistics, and need not imply rejection of the null hypothesis of zero autocorrelation. We evaluate the joint significance of these autocorrelation coefficients by currency, using the Box-Pierce *Q*-statistic (also corrected for heteroskedasticity). These measures, shown in Table 1*b*, indicate no significant autocorrelation except in the case of the yen. Table 1*b* also reports results from augmented Dickey-Fuller tests on daily exchange rate levels. These tests are unable to reject a unit root for any of the six currencies we study.[11]

Because exchange rate innovations may not be homoskedastic, even if they are serially uncorrelated, we also perform a modified GARCH estimation to allow for conditional heteroskedasticity, following Hsieh (1988, 1989) (see Appendix B). We model the exchange rate level as an AR(1) process, and model the exchange rate's conditional standard deviation as a linear function of its own first lag and of the lagged absolute residual from the level regression. The relevant coefficients, shown in Table 1*c*, provide strong support for the serial correlation of conditional volatility in our data. The coefficients on lagged standard deviation and lagged absolute residuals are statistically differ-

[11] In the presence of heteroskedasticity, the Dickey-Fuller test is biased towards an over-rejection of the unit root, as discussed in Kim and Schmidt (1993). Since, despite the bias, we do *not* reject a unit root, it seems reasonable to conclude that the data are indeed characterised by a unit root.

Table 1

Statistical Properties of Daily Exchange Rate Changes

Lag number	JPY	DEM	CAD	CHF	FRF	GBP
	a. Autocorrelation coefficients for daily exchange rate returns					
1	0.0201	−0.0172	−0.0404*	−0.0164	−0.0339	0.0370*
2	0.0329	−0.0022	−0.0249	−0.0122	0.0031	−0.0102
3	0.0110	0.0213	0.0136	0.0254	0.0194	0.0116
4	0.0148	0.0148	0.0030	0.0015	0.0158	0.0139
5	0.0214	0.0124	0.0172	0.0051	0.0083	0.0183
6	−0.0132	0.0068	0.0243	−0.0002	−0.0142	0.0117
7	0.0081	−0.0102	0.0259	0.0023	−0.0164	−0.0372*
8	0.0038	0.0138	−0.0108	−0.0103*	0.0036	−0.0009
9	0.0310	0.0249	−0.0025	0.0317	0.0305	0.0261
10	0.0557*	0.0299	0.0161	0.0139	0.0270	0.0282
	b. Summary measures of autocorrelation					
Adjusted Box-Pierce *Q*-statistic (on 10 lags)†	24.05*	11.03	13.53	8.28	13.61	17.37
Augmented Dickey-Fuller Test‡	−0.40	−1.73	−1.04	−2.30	−0.96	−2.02
	c. Coefficients from GARCH estimation of daily exchange rates (standard errors in parentheses)					
Coefficient on Lagged Variance	0.737 (0.017)	0.861 (0.006)	0.910 (0.004)	0.719 (0.011)	0.842 (0.006)	0.812 (0.006)
Coefficient on Lagged Residual Squared	0.108 (0.008)	0.136 (0.007)	0.084 (0.004)	0.206 (0.012)	0.149 (0.007)	0.163 (0.007)

* Denotes significance in the 5% level using heteroskedasticity-corrected standard errors (Hsieh, 1988).

† This is distributed as chi-squared with *k* degrees of freedom, where *k* is the number of lags. The critical value at the 0.05 significance level is 18.31.

‡ In this one-sided t-test, the critical value at the 0.05 significance level is −2.86.

ent from zero for all six currencies. Our final results differed only trivially across the random walk and GARCH cases, so we focus the rest of our discussion on results using the null hypothesis of a random walk.

Technical analysts do not agree on the appropriate price data to use for charting. Some prefer daily bar charts (which use high, low, and closing prices), while others 'believe that the closing price is the most critical price of the trading day' (Murphy, 1986, p. 36). Daily closing prices seem especially appropriate for currencies, since the daily highs and lows are not recorded in foreign exchange spot markets. Since technical analysis manuals maintain that their techniques apply equally to all financial price series, the absence of daily highs and lows should not compromise the accuracy of our analysis.

2. Testing Profitability: Results

This section presents the results of the statistical tests described in Section 1. Before describing profits, we present some related information on the frequency and length of positions. For the mark, head-and-shoulders signals during the sample period resulted in 32 speculative positions, held on average 17 business days. For the yen, head-and-shoulders signals resulted in 20 speculative positions, held on average 22 business days. For the other currencies, the number of positions and the average holding periods do not differ markedly from those of the mark and yen. Overall, the algorithm would lead a hypothetical trader to have open positions (either short or long) about 10% of the time.

2.1. *Results of Profits Test*

Size of Profits from Endogenous Exit Rule: For the pound and the Canadian dollar, average profits per position are actually negative. Average profits for the Swiss franc are positive but small. Average profits for the French franc are sizable, averaging more than 0.50%. Profits for the mark averaged 0.78% per position and reached a maximum of 9.0% for an individual position; those for the yen averaged 1.52% per position and reached a maximum of 12.8%.

Statistical Significance of Profits from Endogenous Exit Rule: Average percent returns for the U.K. pound, the Canadian dollar, the Swiss franc, and the French franc do not differ statistically from those found in the simulated data. The significance of these profits is shown in Table 2*a*, which reports profits from each currency from the endogenous exit rule strategy and corresponding p-values. Thus, for these four currencies the head-and-shoulders trading rule does not satisfy our first criterion for rationality, profitability. Actual profits for the mark and the yen are significantly higher than profits calculated in the simulated series, thus supporting

Table 2*a*

Profitability of Positions Following Head-and-Shoulders Patterns, 1973–1994

(Endogenous exit rule)*

	JPY	DEM	CAD	CHF	FRF	GBP
Average % profit						
value	1.5158	0.7784	−0.0377	0.0959	0.5653	−0.0699
p-value	0.0035	0.0438	0.5263	0.3638	0.1109	0.5022
Positions						
Number	20	32	29	27	33	27
p-value	0.90	0.48	0.53	0.73	0.20	0.63
Average holding period (business days)	22.2	16.6	18.7	11.5	16.2	20.0

* Head-and-shoulders identification and profit-taking algorithms applied to actual exchange rate data, assuming exit takes place when trough (or peak) has been identified following entry. *p*-values indicate the marginal significance of head-and-shoulders returns under the null hypothesis that the exchange rate follows a random walk. *p*-values are calculated by comparing profits from the actual data with those from data simulated by sampling with replacement from original exchange rate changes.

Table 2*b*

Profitability of Positions Following Head-and-Shoulders Patterns, 1973–1994

(Endogenous exit rule)*

Daily Returns:	JPY	DEM	CAD	CHF	FRF	GBP
1-day return						
Profits (%)	0.0247	0.0403	−0.1016	−0.0119	−0.1117	−0.0995
p-Value	0.3465	0.3171	0.0228	0.4690	0.7704	0.7479
3-day marg.ret.						
Profits	0.4793	−0.0299	0.0161	0.0973	−0.2581	0.0908
p-Value	0.0023	0.5303	0.3959	0.2895	0.9122	0.2814
5-day marg.ret.						
Profits	0.0939	−0.1673	−0.0667	−0.4059	−0.0619	0.2214
p-Value	0.2850	0.8324	0.8148	0.9756	0.6352	0.1019
10-day marg.ret.						
Profits	−0.2481	0.2350	−0.0366	−0.1647	0.3247	0.0781
p-Value	0.8268	0.1991	0.6151	0.6968	0.1329	0.3789
20-day marg.ret.						
Profits	0.9179	0.1103	−0.0771	−0.2308	0.2748	−0.0280
p-Value	0.0222	0.3848	0.6610	0.6869	0.2574	0.5297
30-day marg.ret.						
Profits	−0.4967	0.7919	−0.0673	0.3223	0.6324	0.3735
p-Value	0.8964	0.0274	0.6412	0.2350	0.0646	0.8315
60-day marg.ret.						
Profits	0.5688	−0.5323	0.0055	−0.1751	−0.7319	0.3000
p-Value	0.1852	0.7799	0.4803	0.5776	0.8448	0.3294
10-day total ret.						
Profits	0.3525	−0.1509	0.0145	−0.4851	−0.1069	0.2907
p-Value	0.1558	0.7812	0.4329	0.8375	0.5684	0.2087
30-day total ret.						
Profits	0.1690	0.9802	−0.1299	−0.3936	0.8003	−0.1107
p-Value	0.3558	0.0767	0.6407	0.6699	0.1328	0.5423

* Head-and-Shoulders identification and profit-taking algorithms applied to actual exchange rate data, assuming exit takes place a fixed number of days following entry. *p*-values indicate the marginal significance of head-and-shoulders returns under the null hypothesis that the exchange rate follows a random walk. *p*-values are calculated by comparing profits from the actual data with those from data simulated by sampling with replacement from original exchange rate changes.

technical analysts' claims about the ability of head-and-shoulders patterns to predict trends for these two currencies.[12]

Profits from Exogenous Exit Rule: Profits measured according to the exogenous exit rules, shown in Table 2*b*, confirm the poor predictive power of the head-and-shoulders signals for the Canadian dollar, Swiss franc, French franc, and the U.K. pound. For the mark and the yen, most exogenously specified exit intervals also produce insignificant profits. The contrast

[12] The contrast between the mark and the French franc, pound, and Swiss franc may seem surprising, since these currencies tended to move in line against the dollar during much of our sample period. During this time period, the correlation of daily changes were: 0.583 French franc-mark, 0.307 pound-mark, and 0.468 Swiss franc-mark.

between these results and those based on the endogenous exit rule support technical analysts' claims regarding the timing of exit.

It may be helpful to put the profitability of head-and-shoulders trading in the mark and yen into perspective, by comparing it with some natural benchmarks.

Annualised Profits: The average 0.8% profits in the mark over 17 days correspond to a compound annual return of roughly 13%. This exceeds both the 2.5% average annual buy-and-hold return to holding marks and the corresponding average S&P 500 annual return of 6.8%. The 1.5% average return on yen positions corresponds to a compound annual return of roughly 19%. This outperforms U.S. equity yields and the 4.4% average annual buy-and-hold return to holding yen.

Profits From Multi-Currency Strategies: A hypothetical speculator trading in all six currencies would have earned total profits (the sum of percentage profits from individual positions) over the period of 69.9%, significantly higher than profits generated in our simulated data (p-value of 0.0155). Because each random draw in our simulations corresponds to the same historical day's change across currencies, we avoid overstating the significance of this result. Such a bias could have arisen from positive historical correlation across dollar exchange rates. If simulated profits from each currency had been constructed from uncorrelated exchange rate series, then profits from each currency would likewise have been uncorrelated. Simulated aggregate profit (over all six currencies) would therefore have been relatively concentrated around the mean, resulting in more extreme p-values for actual profit.

These results are strikingly robust to modifications of the methodology described above. Details of these many modifications explored are reported in Appendix C, where the reader can see that they can be categorised according to whether changes were made in (1) the parameter configuration defining a head-and-shoulders pattern, (2) the sample period, (3) the allowed size of head-and-shoulders patterns, or (4) the assumed underlying behaviour of exchange rates. The results from these alternatives conform closely to those reported above.[13]

2.2. *Discussion*

It is natural to investigate possible sources of the profitability of trading on the head-and-shoulders pattern for the mark and the yen. The explanations in technical analysis manuals themselves are limited and vague. This subsection first explores whether trading profits represent compensation to investors for international interest differentials, risk, or transactions costs. Since these

[13] To save space, the results of these sensitivity analyses are summarised verbally. Full details are available from the authors upon request.

factors do not seem sufficient to account for the profits, we then turn to some other possible explanations.

Interest Differentials: We calculate daily differentials in overnight Eurocurrency rates between the foreign currency in question and the U.S. dollar, and apply them to the holding periods associated with our head-and-shoulders positions.[14] Adjusting for interest differentials *raises* average profits for the mark and the yen. For the mark, profits rise from 0.78% to 0.85% per position. For the yen the increase in profits is quite small, and gets lost in the rounding.

Risk: Although profitable for the mark and the yen, the head-and-shoulders strategy results in extremely variable returns on a position-by-position basis. While there is no consensus on appropriate risk adjustment of excess returns to foreign exchange speculation, one approach is to calculate the Sharpe ratio of annualised excess returns to their annualised standard deviation. These are shown below, along with the comparable statistic for the S&P 500:

	H&S Mark	H&S Yen	S&P 500
Sharpe Ratio	1.00	1.47	0.32

By this measure, risk-adjusted rewards to speculating on the head-and-shoulders pattern in yen and marks dominates those from speculating on the S&P 500, although this approach does not provide confidence intervals indicating statistical significance.

To investigate whether the observed profits can be explained as a reward for bearing systematic risk, we calculated the beta of excess returns from our head-and-shoulders positions against excess U.S. stock market returns (S&P 500) and foreign stock market returns (DAX for mark positions, Nikkei for yen positions). None of the point estimates of beta were statistically significantly different from zero, and the largest was 0.03. In sum, it seems unlikely that these profits represent compensation for bearing systematic risk.

Transaction Costs: To investigate whether these profits would vanish if we subtracted transactions costs, i.e. bid-ask spreads, we take as our benchmark a mid-sized spread of 0.0010 marks or 0.1 yen per dollar traded. This reduces returns about 0.05% for each round-trip transaction in either currency.[15] After adjustment for both interest differentials and transactions costs, mark returns would have averaged 0.80% per position, and yen

[14] In performing this adjustment, we accounted for the cost of funds over weekends, and for the direction of our position.

[15] For both currencies the average percent bid-ask cost was calculated by taking a given point spread and dividing it by the average exchange rate over the sample period. The 'given' point spreads, provided by foreign exchange salespersons at Natwest Markets, New York, actually represented a very generous (large) estimate of the typical spreads for 'good corporate customers'.

returns would have averaged 1.48%, still remarkably high given the brevity of our holding periods.

Interest differentials, risk, and transactions costs apparently fail to explain the significance of head-and-shoulders profits for the mark and the yen based on the head-and-shoulders pattern. We turn next to three other possible, though less readily tested, interpretations.

Uncertainty and Imperfect Information: Brown and Jennings (1990) suggest that, when markets are characterised by uncertainty and imperfect information, technical analysis may be useful because the current price is not a sufficient statistic for the market's full information set. While their model incorporates only two periods and is thus not directly applicable to the head-and-shoulders pattern, its overall insight might generalise to the many-period setting of that pattern.

Self-Fulfilling Predictions: The presence of self-fulfilling expectations could help explain why we find significant profits only for the mark and yen: since technical traders who follow head-and-shoulders trading rules sell as soon as they perceive downtrends and *vice versa*, they could potentially create persistence in what might otherwise be purely transitory price movements.[16] In practice, if such speculative activity is concentrated on the mark and the yen, the self-fulfilling tendency would be most apparent in those currencies. Related evidence in Osler (1998) indicates that there is some self-fulfilling influence of head-and-shoulders trading in U.S. stocks, but that the associated price movement may be insufficient for profitable speculation.

Central Bank Intervention: The introduction of a major player with non-profit-maximising objectives, such as a central bank, can introduce departures from random walk behaviour. The hypothesis that central bank intervention leads to predictability in exchange rates has received support from a number of sources. Silber (1994) examines profits from moving-average trading rules in futures contracts for currencies, short-term interest rates, the S&P 500, and commodities. He finds abnormal profits in the markets characterised by government intervention such as currencies and interest rates, but not in the other markets. Stronger support comes from LeBaron (1996), who finds that significant profits using moving average trading rules decline and lose statistical significance when one removes days in which the Federal Reserve intervened in that currency. Similarly, Szakmary and Mathur (1997) find that moving-average rule profits are largely explained by intervention intended to lean against the wind of unwanted exchange-rate trends. For intervention to explain the difference in our results across currencies, intervention would presumably have to be

[16] This speculative activity provides an example of 'positive feedback trading', of the sort examined by De Long *et al.* (1990) and Frankel and Froot (1990).

particularly concentrated in the mark and yen. Since most central banks do not report their intervention, regardless of currency, this is an unverifiable condition.

3. Testing Efficiency

The analysis so far has demonstrated that trading based on head-and-shoulders patterns, as recommended by technical analysis manuals, is profitable for the mark and the yen. Such profitability is a necessary but not sufficient condition for forecasts based on head-and-shoulders patterns to be rational. For forecasts based on trading signals, another necessary condition for rationality is *efficiency*. This section shows that forecasts based on head-and-shoulders patterns are inefficient because simpler trading strategies dominate the trading strategy based on the head-and-shoulders pattern. In particular, these simpler trading strategies generate higher profit and have lower risk. Moreover, adding information from head-and-shoulders signals to any of the simpler strategies significantly increases risk without significantly increasing profitability. Hence, trading based on the head-and-shoulders pattern is not efficient and cannot be rational.

In this section, we also demonstrate the unreliability of technical analysts' claim that the head-and-shoulders patterns can predict the magnitude of price movements. This provides further evidence against the reliability this trading strategy's proponents.

3.1. *Comparing Head-and-Shoulders with Simple Trading Rules*

We compare the head-and-shoulders rule with two other types of trading rules common in the marketplace, 'oscillators' and momentum rules. When using an oscillator, one takes a long position whenever a short-horizon moving average is higher than a long-horizon moving average, and *vice versa*. When using a momentum rule, one takes a long position whenever the current exchange rate exceeds the exchange rate lagged z periods, where z is specified exogenously. In contrast with the head-and-shoulders rule, which often advocates a zero position, these rules always involve either a long or short position. The profitability of these rule when applied to dollar exchange rates has been widely documented (Logue *et al.* 1978; Dooley and Shafer, 1984; Sweeney, 1986; Levich and Thomas, 1993; Menkhoff, 1995).

These rules can be considered 'simpler rules' in that one's long or short position is based on a simple objective condition: usually the comparison of one measure (e.g. the spot rate or a moving average) against another (e.g. the lagged spot rate or another moving average). The single inequality condition involved here can be easily expressed algebraically, and has relatively little ambiguity. In contrast, trading based on head-and-shoulders patterns is based on a number of conditions fraught with ambiguities, as noted earlier.

We first report the profitability of the simple trading rules when applied to our data sample. For the oscillators we take positions upon the crossing of the

5-day moving average by the current exchange rate, and also consider the following four pairs: 1 vs 20 days, 5 vs 20 days, 5 vs 50 days, and 20 vs 50 days. For the momentum rules we use 5, 20, and 50-day lags. As indicated in Table 3*a*, profits from these rules are extremely statistically significant: in fact, they are significant at the 1% level in 45 out of 48 cases. Simple technical trading strategies thus seem to have significant predictive power not just the mark and yen, but for all six of the currencies considered.

Table 3

Profitability of Simple Technical Trading Rules

(*a*) Marginal significance levels resulting from 10,000 simulations

Rule	JPY	DEM	CAD	CHF	FRF	GBP
H&S	0.004	0.044	0.526	0.364	0.111	0.502
Oscillator						
1 vs 5 day MAs	0.000	0.015	0.123	0.033	0.032	0.000
1 vs 20 day MAs	0.000	0.000	0.005	0.001	0.000	0.000
5 vs 20 day MAs	0.000	0.000	0.004	0.001	0.000	0.000
5 vs 50 day MAs	0.000	0.000	0.292	0.000	0.002	0.000
20 vs 50 day MAs	0.001	0.000	0.343	0.005	0.001	0.000
Momentum						
5-day lag	0.000	0.000	0.019	0.000	0.000	0.000
20-day lag	0.000	0.000	0.010	0.003	0.000	0.003
50-day lag	0.001	0.002	0.425	0.000	0.014	0.001

(b) Total profits (%)

(t-statistics for the difference between head-and-shoulders profits and profits for each alternative rule are presented in parentheses)

H&S	33.3	22.8	−1.52	2.2	7.8	−2.1
Oscillator						
1 vs 5 day MAs	218.5 (−3.99)	116.7 (−1.73)	23.9 (−1.30)	111.9 (−1.82)	97.9 (−1.69)	184.4 (−3.73)
1 vs 20 day MAs	246.1 (−4.59)	239.1 (−4.00)	52.0 (−2.75)	177.2 (−2.90)	228.2 (−4.15)	236.6 (−4.78)
5 vs 20 day MAs	222.6 (−4.08)	222.8 (−3.70)	48.0 (−2.54)	184.1 (−3.02)	215.1 (−3.90)	166.4 (−3.37)
5 vs 20 day MAs	213.7 (−3.89)	181.2 (−2.93)	13.7 (−0.78)	183.2 (−3.01)	132.6 (−2.35)	178.5 (−3.62)
20 vs 50 day MAs	139.9 (−2.30)	176.8 (−2.85)	12.5 (−0.72)	157.7 (−2.58)	148.0 (−2.64)	160.0 (−3.25)
Momentum						
5-day lag	248.0 (−4.63)	218.4 (−3.61)	40.1 (−2.13)	225.7 (−3.71)	198.1 (−3.58)	250.9 (−5.07)
20-day lag	229.1 (−4.22)	226.8 (−3.77)	47.3 (−2.51)	162.4 (−2.66)	193.6 (−3.50)	139.1 (−2.83)
50-day lag	151.4 (−2.54)	133.8 (−2.05)	8.0 (−0.49)	166.4 (−2.73)	111.1 (−1.94)	152.7 (−3.10)

Oscillator rules (long whenever short-horizon moving average exceeds long-horizon moving average, and *vice versa*) and momentum strategies (long whenever current exchange rate exceeds exchange rate at *z*-day lag) applied to actual exchange-rate data. *p*-values indicate the marginal significance of returns under the null hypothesis that the exchange rate follows a random walk. *p*-values are calculated by comparing profits from the actual data with those from data simulated by sampling with replacement from original exchange-rate changes.

Profits from the simpler technical trading strategies not only have greater statistical significance than profits from the head-and-shoulders strategy, they also are substantially larger, as shown in Table 3*b*. Every one of the simpler rules produce positive profits, for each currency, in contrast to the head-and-shoulders rule, which produced negative profits for the Canadian dollar and the U.K. pound. In all but 6 of the 48 cases, the difference in total profits is statistically significant.

Because the simpler technical trading rules involve more active speculation, one might think that they have a natural advantage in producing profits. It is important to stress, however, that a more active level of market participation does not necessarily confer any advantage. If the exchange rates follow a martingale process, for example, as many observers have concluded is true for the dollar during the floating rate period, more active speculation would have no *a priori* advantage over less active speculation. More generally, if the simpler strategies were not profitable overall, their extra market activity would not increase profitability.

On a risk-adjusted basis, the simpler rules still outperform the head-and-shoulders-based trading rules. The standard deviation of daily profits is lower for the head-and-shoulders strategy than for the simpler rules, as shown in Table 4*a* (this difference in standard deviation is to a large degree due to the infrequency of head-and-shoulders trading). However, the difference in risk does not fully offset the difference in profitability across the strategies, as shown by the Sharpe ratios, which incorporate both excess returns *and* risk (see Table 4*b*). For the yen, the Sharpe ratios for all eight simpler rules, which range from 0.68 to 1.21, are higher than the 0.55 Sharpe ratio of the head-and-shoulders rule. For the mark, the other profitable currency for head-and-shoulders trading, only the 1 vs. 20 day moving-average rule generated a lower Sharpe ratio (0.10) than the 0.32 obtained for the head-and-shoulders rule. The other seven of the simpler rules generated Sharpe ratios ranging from 0.49 to 0.95. In short, when returns and risk are jointly considered in this way, the simpler trading rules easily dominate the head-and-shoulders rules, even for the mark and the yen.

Although they generate less profit than the simpler rules on both an absolute and a risk-adjusted basis, using information from the head-and-shoulders rule could nonetheless prove rational if it *incrementally* added value. For example, a trader could conceivably increase his/her profit by combining information from the head-and-shoulders rule with information from the more profitable simpler rules. If such a combined trading strategy increased profit without increasing risk, the head-and-shoulders rule could still be rational.

We thus evaluate the incremental contribution of the head-and-shoulders strategy when combined with each of simpler rules, focusing only on the mark and yen, where head-and-shoulders trading generated a significant profit. The combined strategy works as follows: if the two strategies agree, open a position twice as large as one might otherwise have taken; if they disagree, take no position. Profits from each of these eight combinations are reported in Table

Table 4

Riskiness of Head-and-Shoulders and Simpler Trading Strategies

(*a*) Standard deviation of daily profits, %.

Rule	JPY	DEM	CAD	CHF	FRF	GBP
H&S	0.177	0.21	0.087	0.169	0.177	0.140
Oscillator						
1 vs 5 day MAs	0.598	0.697	0.273	0.792	0.692	0.657
1 vs 20 day MAs	0.579	0.696	0.273	0.792	0.691	0.656
5 vs 20 day MAs	0.598	0.696	0.273	0.792	0.691	0.657
5 vs 50 day MAs	0.598	0.695	0.273	0.791	0.692	0.656
20 vs 50 day MAs	0.598	0.695	0.273	0.791	0.691	0.656
Momentum						
5-day lag	0.597	0.696	0.273	0.792	0.692	0.656
20-day lag	0.598	0.696	0.273	0.792	0.692	0.657
50-day lag	0.598	0.696	0.273	0.791	0.692	0.656

(*b*) Sharpe ratios for daily profits.

H&S	0.55	0.32	−0.06	0.04	0.13	−0.04
Oscillator						
1 vs 5 day MAs	1.07	0.49	0.31	0.41	0.41	0.82
1 vs 20 day MAs	1.20	0.10	0.66	0.65	0.96	1.05
5 vs 20 day MAs	1.09	0.94	0.62	0.68	0.91	0.74
5 vs 50 day MAs	1.05	0.76	0.18	0.68	0.56	0.80
20 vs 50 day MAs	0.68	0.74	0.16	0.58	0.63	0.71
Momentum						
5-day lag	1.21	0.92	0.51	0.83	0.84	1.12
20-day lag	1.12	0.95	0.61	0.60	0.82	0.62
50-day lag	0.74	0.56	0.10	0.62	0.47	0.68

5*a*. For the mark and yen, each combination strategy produces a slightly higher profit than the simple strategy alone, but the differences are not significant (as revealed by difference-in-means tests). In other words, the head-and-shoulders rules do not incrementally increase the profits of the simpler trading rules to a statistically significant degree.

In contrast, the head-and-shoulders rules do significantly increase the riskiness of the simpler trading rules' profits (Table 5*b*). We compare the variance of daily profits of the simpler rules with the variance of daily profits of the simpler rules combined with the head-and-shoulders rules. The ratio of the combined rules' variance to the simpler rules' variance will be distributed as $F_{m,n}$, with m and n corresponding to the number of days each rule is used. These F-statistics, which range from 1.16 to 1.36, have p-values of virtually zero (differing from zero in only the tenth decimal). On this basis we reject the hypothesis of equal variances in favour of the hypothesis that the combined strategies have higher variances than the simpler strategies.

In sum, the head-and-shoulders strategy is not efficient because it is dominated by simpler trading rules. On a stand-alone basis, the head-and-shoulders strategy is inferior to simpler rules because the latter provide positive, significant profits for all six of our currencies, rather than only for the mark and the

Table 5(*a*)

Total Profits (%) to Combination Strategies

t-statistics for the difference between simple strategy and combined strategy are presented in parentheses

Rule	JPY	DEM
Oscillator		
1 vs 5 day MAs	251.8	139.5
	(0.505)	(0.299)
1 vs 20 day MAs	279.4	261.9
	(0.499)	(0.295)
5 vs 20 day MAs	255.9	245.6
	(0.499)	(0.295)
5 vs 5 day MAs	247.0	204.0
	(0.504)	(0.297)
20 vs 50 day MAs	173.2	199.6
	(0.507)	(0.298)
Momentum		
5-day lag	281.3	241.2
	(0.504)	(0.298)
20-day lag	262.4	249.6
	(0.500)	(0.296)
50-day lag	184.7	156.6
	(0.506)	(0.298)

Oscillator rules (long whenever short-horizon moving average exceeds long-horizon moving average, and *vice versa*) and momentum strategies (long whenever current exchange rate exceeds exchange rate at z-day lag) applied to actual exchange rate data. p-values indicate the marginal significance of returns under the null hypothesis that the exchange rate follows a random walk. p-values are calculated by comparing profits from the actual data with those from data simulated by sampling with replacement from original exchange rate changes.

yen. Adjusted for risk, the head-and-shoulders strategy also underperforms the simpler rules, which have consistently lower Sharpe ratios. On an incremental basis, the head-and-shoulders rule does not significantly raise profitability, but does significantly increase risk, as measured by variation in daily profit. Since it is clearly dominated by the simpler rules on a stand-alone basis, and adds risk without significantly adding profit when used in combination with the simpler rules, the use of the head-and-shoulders trading rule does not appear to be rational.

3.2. *Testing Technical Analysts' Other Claims*

Technical analysts claim that head-and-shoulders patterns predict not only the direction of exchange-rate movements but also their magnitude. In particular, the vertical distance from the head to the neckline (referred to as a 'price

Table 5(*b*)
Riskiness of Combination Trading Strategies

Variances for simple rules and combinations of simple rules and head-and-shoulders rule. Also, F-statistics and associated p-values for the hypothesis that the variance of the combined strategy exceeds the variance of the simpler strategy.

Simple rule	JPY					DEM				
Oscillator	Simple	Comb.	F-stat.	df	*p*-value*	Simple	Comb.	F-stat.	df	*p*-value*
1 vs 5 day MAs	0.598	0.653	1.192	(5536, 5536)	0.000	0.697	0.750	1.158	(5536, 5536)	0.000
1 vs 20 day MAs	0.597	0.669	1.256	(5521, 5521)	0.000	0.696	0.769	1.221	(5521, 5521)	0.000
5 vs 20 day MAs	0.598	0.670	1.255	(5521, 5521)	0.000	0.696	0.769	1.221	(5521, 5521)	0.000
5 vs 50 day MAs	0.598	0.657	1.207	(5491, 5491)	0.000	0.695	0.760	1.196	(5491, 5491)	0.000
20 vs 50 day MAs	0.598	0.648	1.174	(5491, 5491)	0.000	0.695	0.758	1.190	(5491, 5491)	0.000
Momentum										
5-day lag	0.597	0.658	1.215	(5535, 5535)	0.000	0.696	0.754	1.174	(5535, 5535)	0.000
20-day lag	0.598	0.665	1.237	(5520, 5520)	0.000	0.696	0.765	1.208	(5520, 5520)	0.000
50-day lag	0.598	0.651	1.185	(5490, 5490)	0.000	0.696	0.759	1.189	(5490, 5490)	0.000

* *p*-values are computed in SAS using PROBF function.

objective') indicates the minimum distance the price should decline after it has penetrated the neckline (See Fig. 1).[17] (An identical condition is associated with head-and-shoulders bottoms, but for simplicity this section refers to all head-and-shoulders as tops.) To test this, we compare the 'price objective' to the post-head-and-shoulders downward vertical movement, measuring that movement as the maximum fall to the next trough after entry. In every case the measure is chosen to give the benefit of the doubt to technicians' claims. For example, the maximum fall will be greater than actual profits earned under the endogenous exit rule, which requires one to wait until a trough has been identified by a subsequent upward price movement. In the case of a 'bounce', the measure is based on the larger of the difference between (1) the price at entry and at the first trough and (2) the difference between the price at entry and at the second trough. Finally, when the head-and-shoulders signal fails utterly and a trader would exit at a stoploss, the maximum fall is assigned a value of zero. The ratio of the maximum fall to the price objective is referred to as 'mirror'.

We find no empirical support of technical analysts' claims about this price objective. As shown in Table 6, the maximum post-entry fall is usually smaller than the proclaimed price objective, or equivalently, mirror usually falls below unity. In fact, for all currencies but the yen, even the *average* value of mirror is well below one. The continued use of this price standard in exchange-rate forecasting, despite its lack of predictive power, does not seem to be consistent with rationality.

Table 6
Distribution of Values for 'Mirror'

	JPY		DEM		CAD		CHF		FRF		GBP	
Average Mirror	1.21		0.87		0.74		0.51		0.69		0.55	
Range:	No.	(%)	No.	(%)	No.	(%)	No.	(%)	No.	(%)	No.	(%)
0.0	5	(23)	11	(34)	8	(28)	11	(41)	14	(42)	11	(41)
0.0–0.5	5	(23)	5	(16)	7	(24)	7	(26)	4	(12)	8	(30)
0.5–1.0	3	(14)	9	(28)	4	(14)	4	(15)	3	(9)	2	(7)
1.0 or higher	7	(32)	8	(25)	10	(34)	5	(19)	12	(36)	6	(22)
Total Positions	20		32		29		27		33		27	

'Mirror' represents the ratio of (a) the difference between the entry price and the final trough before exit to (b) the vertical distance from the head-and-shoulder's peak to its neckline. Technical analysts claim this should be at least unity.

[17] Our conversations with practising technical analysts indicate that these claims are supported by practitioners, as well as by the written manuals.

4. Concluding Comments

This paper assesses the rationality of a commonly used trading signal, the head-and-shoulders pattern. The pattern comprises three peaks, the highest of which is in the middle. According to a large group of financial market participants called 'technical analysts', such patterns precede trend reversals and can be used profitably as a trading signal. We find that trading on the pattern in the manner recommended by technical analysts is profitable for the mark and the yen, but not for four other currencies (all vis-à-vis the v.s. dollar). We also find, however, that head-and-shoulders trading is dominated by simple technical trading rules, that are readily available. That is, once one incorporates information from the simpler rules, additional information provided by head-and-shoulders patterns is of no value in predicting the mark or the yen. Thus, the continued reliance on the head-and-shoulders pattern appears to represent a source of predictable exchange-rate forecast errors.

Credit Suisse First Boston, New York

Federal Reserve Bank of New York

Date of receipt of first submission: February 1997
Date of receipt of final typescript: January 1999

Appendix A. Criteria for Identification of Head-and-Shoulders

These criteria listed below apply to head-and-shoulders tops. Analogous criteria apply to head-and-shoulders bottoms.

1. The height of the head must exceed that of the left and right shoulders.
2. Since the pattern should *presage a trend reversal*, it must occur following an uptrend. Thus, (i) the left shoulder must exceed the previous peak and (ii) the first trough within the head-and-shoulders must exceed the previous trough.
3. To avoid extreme *horizontal asymmetries*, the time between the left shoulder and the head must be no greater than 2.5 times the time between the head and the right shoulder; likewise the time between the head and the right shoulder should be no greater than 2.5 times the time between the left shoulder and the head.
4. To avoid extreme *vertical asymmetries*, the pattern must be only moderately sloped. The right shoulder must exceed, and the right trough must not exceed, the midpoint between the left shoulder and the left trough. Likewise, the left shoulder must exceed, and the left trough must not exceed, the midpoint between the right shoulder and right trough.
5. Since the head-and-shoulders pattern in principle indicates an *imminent* trend reversal, the time required for the price to cross the neckline must be no longer than the time interval between right and left shoulders.

Appendix B

To create GARCH-based exchange rate series we began by estimating the parameters for the following exchange rate process:

$$s_t = \alpha_0 + \alpha_1 s_{t-1} + \epsilon_t,$$

where

$$\epsilon_t \sim \mathrm{N}(0, \sigma_t^2)$$

and

$$\eta \sim \mathrm{N}(0, \sigma_\eta^2)$$

From this we were able to reconstruct the series of volatility residuals, η_t, and a series of standardised exchange rate residuals, $e_t = \epsilon_t/\sigma_t^2$. Then, drawing randomly with replacement from these two residual series and combining these draws according to the process specified above, we construct our 10,000 simulated exchange rate (and volatility) series.

Appendix C: Sensitivity of Our Results to Alternative Parameterisation

To determine whether our findings depend on specific parameter values, we recalculate our results under various alternative parameter configurations. First, we vary the horizontal symmetry requirement (number 3 in Appendix A). (All the modifications here will be illustrated with respect to head-and-shoulders tops. Analogous modifications were made for head-and-shoulders bottoms.) We use 2.0 or 3.0, instead of the base value of 2.5, for the maximum ratio of 'number of days between the left (right) shoulder and the head' and 'number of days between the head and right (left) shoulder'. Second, we relax the vertical symmetry requirement regarding the height of the left and right shoulders relative to other local peaks and troughs (number 4 in Appendix A). Specifically, we require that the right (left) shoulder exceed the left (right) trough, rather than exceed the midpoint between the left (right) shoulder and left (right) trough. Third, we reduce by half stoploss, the parameter determining the time of exit when exchange rates fail to reverse as predicted by the head-and-shoulders signal. Fourth, we double the 'bounce' parameter, from 0.25 to 0.5. None of these parameter adjustments has any noticeable effect on the significance of our profit results. We also split the sample approximately in half, the first period covering 19 March 1973 to 31 December 1983, the second period covering 1 January 1984 to 13 June 1994. This results in no consistent increase or decrease in significance of profits.

To determine whether smaller or larger cutoffs tend to lead to higher profits, we also perform these tests using only the five smallest cutoffs and the five largest cutoffs. Though the number of positions declines as cutoff increases, the overall profitability results are largely unchanged by this breakdown.

To investigate whether the random walk assumption about exchange rate behaviour could have affected our results, we also generate simulated exchange rate series in which the data follow a GARCH (1,1) process,[18] a choice motivated by the extensive literature documenting the serial correlation of exchange rate volatility.[19] Results (not reported) from our GARCH-based simulations, the methodology for which is described in Appendix B, yield virtually the same conclusions as those from simulations based on the random walk.

[18] The standard deviation was used in place of the more familiar variance because randomly drawn shocks to variance often resulted in negative variance.

[19] For example, see Domowitz and Hakkio (1985), Bollerslev (1987), and Diebold and Pauly (1988).

References

Aczel, A. D. and Josephy, N. H. (1991). 'The chaotic behavior of foreign exchange rates', *American Economist*, vol. 35, pp. 16–24.

Allen, F. and Karjalainen, R. (1993). 'Using genetic algorithms to find technical trading rules', Rodney L. White Working Paper 20-93, Wharton School.

Allen, H. and Taylor, Mark P. (1990) 'Charts, noise, and fundamentals in the foreign exchange market', ECONOMIC JOURNAL vol. 100, pp. 49–59.

Arnold, C. and Rahfeldt, D. (1986). 'Timing the market: how to profit in bull and bear markets with technical analysis' Chicago: Probus Publishing.

Blake, D., Beenstock, M. and Brasse, V. (1986). 'The performance of UK exchange-rate forecasters', ECONOMIC JOURNAL, vol. 96, pp. 986–99.

Bollerslev, T. (1987). 'A conditionally heteroskedastic time series model for speculative prices and rates of return', *Review of Economics and Statistics*, vol. 69, pp. 542–7.

Boothe, P. and Glassman, D. (1987). 'The statistical distribution of exchange rates', *Journal of International Economics*, vol. 22, pp. 297–319.

Brock, W., Lakonishok, J. and LeBaron, B. (1992). 'Simple technical trading rules and the stochastic properties of stock returns', *Journal of Finance*, vol. 48, pp. 1731–64.

Brown, D. P. and Jennings, R. H. (1990). 'On technical analysis', *Review of Financial Studies*, vol. 2, pp. 527–51.

Cavaglia, S., Verschoor, W. and Wolff, C. (1994). 'On the biasedness of forward foreign exchange rates: irrationality or risk premia?, *Journal of Business*, vol. 67, pp. 321–43.

De Long, J. B., Shleifer, A., Summers, L. and Waldmann, R. (1990). 'Positive feedback investment strategies and destabilizing rational speculation', *Journal of Finance*, vol. 45, pp. 379–95.

Diebold, F. and Pauly, P. (1988). 'Endogenous risk in a portfolio-balance rational-expectations model of the deutschemark-dollar rate', *European Economic Review*, vol. 32, pp. 27–53.

Dominguez, K. (1980). 'Are foreign exchange forecasts rational? New evidence from survey data', *Economics Letters*, vol. 21, pp. 277–81.

Domowitz, I. and Hakkio, C. S. (1985). 'Conditional variance and the risk premium in the foreign exchange market', *Journal of International Economics*, vol. 19, pp. 47–66.

Dooley, M. P. and Shafer, J. (1984). 'Analysis of short-run exchange rate behavior: March 1973 to November 1981', in (Bigman and Taya, eds.), *Floating Exchange Rates and the State of World Trade and Payments.* Cambridge MA: Ballinger Publishing Company, pp. 43–70.

Edwards, R. D. and Magee, J. (1966). *Technical Analysis of Stock Trends*, Fifth edition, Boston: John Magee Inc.

Efron, B. (1979). 'Bootstrap methods: another look at the jackknife', *Annals of Statistics*, vol. 7, pp. 1–26.

Efron, B. (1982). 'The jackknife, the bootstrap, and other resampling plans', Society for Industrial and Applied Mathematics, Philadelphia.

Efron, B. and Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*, New York: Chapman and Hall.

Frankel, J. and Froot, K. (1987). 'Using survey data to test propositions regarding exchange rate expectations', *American Economic Review*, vol. 77, pp. 133–53.

Frankel, J. and Froot, K. (1990). 'Explaining the demand for dollars: international rates of return and the expectations of chartists and fundamentalists', in (R. Chambers and P. Paarlberg, eds.) *Agriculature, Macroeconomics and the Exchange Rate*, Boulder, Colorado: Westfield Press.

Friedman, B. (1979). 'Optimal expectations and the extreme information assumptions of ''rational expectations'' macromodels', *Journal of Monetary Economics*, vol. 5, pp. 23–41.

Froot, K. A. and Frankel, J. A. (1989). 'Forward discount bias: is it an exchange risk premium?', *Quarterly Journal of Economics*, vol. 104, pp. 139–61.

Gilmore, C. (1991). 'A new test for chaos with application to exchange rate data, mimeo, Drexel University.

Goodman, S. H. (1979). 'Foreign exchange-rate forecasting techniques: implications for business and policy', *Journal of Finance*, vol. 34, pp. 415–27.

Hardy, C. C. (1978). *The Investor's Guide to Technical Analysis*, New York: McGraw Hill.

Hsieh, D. (1988). 'The statistical properties of daily foreign exchange rates: 1974-1983', *Journal of International Economics*, vol. 24, pp. 129–45.

Hsieh, D. (1989). 'Testing for nonlinear dependence in daily foreign exchange rates', *Journal of Business*, vol. 62, pp. 339–68.

Hsieh, D. (1991). 'Chaos and nonlinear dynamics: application to financial markets', *Journal of Finance*, vol. 46, pp. 1839–77.

Ito, T. (1990). 'Foreign exchange rate expectations: micro survey data', *American Economic Review*, vol. 80, pp. 434–49.

Kaufman, P. (1978). *Commodity Trading Systems and Methods*, New York: Ronald Press.

Kim, K. and Schmidt, P. (1993). 'Unit root tests with conditional heteroskedasticity, *Journal of*

Econometrics, vol 59, pp. 287–300.
Krager, H. and Kugler, P. (1993). 'Nonlinearities in foreign exchange markets: a different perspective', *Journal of International Money and Finance*, vol. 12, pp. 195–208.
Krugman, P. (1991). 'Target zones and exchange rate dynamics', *Quarterly Journal of Economics*, vol. 106, pp. 669–82.
LeBaron, B. (1991). 'Technical trading rules and regime shifts in foreign exchange', mimeo, University of Wisconsin.
LeBaron, B. (1996). 'Technical trading rule profitability and foreign exchange intervention', mimeo, University of Wisconsin.
Leitch, G. and Tanner, J. E. (1991). 'Economic forecast evaluation: profits versus the conventional error measures', *American Economic Review*, vol. 81, pp. 580–90.
Levich, R., (1980). 'Analyzing the accuracy of foreign exchange advisory services: theory and evidence, in (R. M. S. Levich and C. Wihlborg, eds.) *Exchange Risk Exposure*, Lexington: D. C. Heath.
Levich, R. and Thomas, L. (1993). 'The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach', *Journal of International Money and Finance*, vol. 12, pp. 451–74.
Levy, R. A. (1971). 'The predictive significance of five-point chart patterns', *Journal of Business*, vol. 44, pp. 316–23.
Lewis, K. K. (1996). 'Puzzles in international financial markets', in (G. Grossman and K. Rogoff, eds), *Handbook of International Economics*, Amsterdam: North Holland.
Logue, D., Sweeney, R. and Willett, T. (1978). 'The speculative behavior of foreign exchange rates during the current float', *Journal of Business Research*, vol. 6, pp. 159–74.
Lui, Y.-H. and Mole, D. (1996). 'The use of fundamental and technical analyses by foreign exchange dealers: Hong Kong evidence', mimeo, City University of Hong Kong.
Malkiel, B. G. (1990). *A Random Walk Down Wall Street*, New York: W. W. Norton and Company.
McCallum, B. T. (1980). 'Rational expectations and macroeconomic stabilization policy', *Journal of Money, Credit, and Banking*, vol. 12 (2), pp. 716–46.
Meese, R. and Rogoff, K. (1983). 'Empirical exchange rate models of the seventies: do they fit out of sample?' *Journal of International Economics*, vol. 14, pp. 3–24.
Meese, R. and Rose, A. (1991). An empirical assessment of nonlinearities in models of exchange rate determination', *Review of Economic Studies*, vol. 58, pp. 603–19.
Menkhoff, L. and Schlumberger, M. (1995). 'Persisent profitability of technical analysis on foreign exchange markets?', *Banca Nazionale del Lavaro Quarterly Review*, no. 193, pp. 189–216.
Murphy, J. J. (1986). *Technical Analysis of the Futures Market: A Comprehensive Guide to Trading Methods and Applications*, New York: Prentice Hall.
Mussa, M. (1979). 'Empirical regularities in the behavior of exchange rates and theories of the foreign exchange market', in (K. Brunner and A. H. Meltzer, eds.) Carnegie-Rochester Series on Public Policy vol. 11, Amsterdam: North-Holland.
Neftci, S. (1991). 'Naive trading rules in financial markets and Wiener-Kolmogorov prediction theory: a study of technical analysis', *Journal of Business*, vol. 64, pp. 549–71.
Osler, C. L. (1998). 'Identifying noise traders: the head-and-shoulders pattern in U.S. equities', Federal Reserve Bank of New York, mimeo.
Pring, M. (1985). *Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points*, Third Edition, New York: McGraw-Hill.
Shabacker, R. W. (1930). *Stock Market Theory and Practice* New York: B. C. Forbes Publishing Company.
Silber, W. L. (1994). 'Technical trading: when it works and when it doesn't,' *Journal of Derivatives* vol. ? (Spring), pp. 34–44.
Sklarew, A. (1980). 'Techniques of a professional commodity chart analyst,' Commodity Research Bureau, New York.
Sweeney, R. J. (1986). 'Beating the foreign exchange market,' *Journal of Finance*, vol. 41, pp. 163–82.
Szakmary, A. C. and Mathur, I. (1997). 'Central bank intervention and trading rule profits in foreign exchange markets,' *Journal of International Money and Finance*, vol. 16, pp. 513–35.
Takagi, S. (1991). 'Exchange rate expectations,' *IMF Staff Papers*, vol. 38, pp. 156–83.
Westerfield, J. M. (1977). 'An examination of foreign exchange risk under fixed and floating rate regimes,' *Journal of International Economics*, vol. 7, pp. 181–200.