Christopher J. Neely is an economist at the Federal Reserve Bank of St. Louis. Kent A. Koch provided research assistance.

Technical Analysis in the Foreign Exchange Market: A Layman's Guide

Christopher J. Neely

Technical analysis suggests that a long-term rally frequently is interrupted by a short-lived decline. Such a dip, according to this view, reinforces the original uptrend. Should the dollar fall below 1.5750 marks, dealers said, technical signals would point to a correction that could pull the dollar back as far as 1.55 marks before it rebounded.

Gregory L. White
Wall Street Journal
November 12, 1992

Technical analysis, which dates back a century to the writings of *Wall Street Journal* editor Charles Dow, is the use of past price behavior to guide trading decisions in asset markets. For example, a trading rule might suggest buying a currency if its price has risen more than 1 percent from its value five days earlier. Such rules are widely used in stock, commodity, and (since the early 1970s) foreign exchange markets. More than 90 percent of surveyed foreign exchange dealers in London report using some form of technical analysis to inform their trading decisions (Taylor and Allen, 1992). In fact, at short horizons—less than a week—technical analysis predominates over fundamental analysis, the use of other economic variables like interest rates, and prices in influencing trading decisions.

Investors and economists are interested in technical analysis for different reasons.

Investors are concerned with "beating the market," earning the best return on their money. Economists study technical analysis in foreign exchange markets because its success casts doubt on the efficient markets hypothesis, which holds that publicly available information, like past prices, should not help traders earn unusually high returns. Instead, the success of technical analysis suggests that exchange rates are not always determined by economic fundamentals like prices and interest rates, but rather are driven away from their fundamental values for long periods by traders' irrational expectations of future exchange rate changes. These swings away from fundamental values may discourage international trade and investment by making the relative price of U.S. and foreign goods and investments very volatile. For example, when BMW decides where to build an automobile factory, it may choose poorly if fluctuating exchange rates make it difficult or impossible to predict costs of production in the United States relative to those in Germany.

Despite the widespread use of technical analysis in foreign exchange (and other) markets, economists have traditionally been very skeptical of its value. Technical analysis has been dismissed by some as astrology. In turn, technical traders have frequently misunderstood what economists have to say about asset price behavior. What can the two learn from each other? This article provides an accessible treatment of recent research on technical analysis in the foreign exchange market.

A PRIMER ON TECHNICAL ANALYSIS IN FOREIGN EXCHANGE MARKETS

Technical analysis is a short-horizon trading method; positions last a few hours or days. Technical traders will not hold

Figure 1

Peaks, Troughs, Trends, Resistance and Support Levels Illustrated for the $/DM

$ per DM
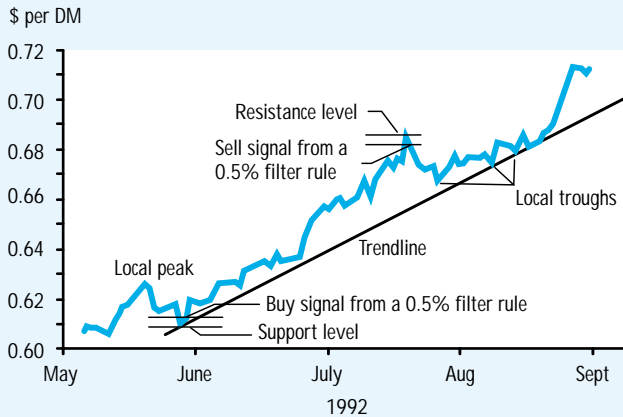
NOTES: Not all buy and sell signals from the filter rule are identified.

positions for months or years, waiting for exchange rates to return to where fundamentals are pushing them. In contrast, fundamental investors study the economic determinants of exchange rates as a basis for positions that typically last much longer, for months or years. Some traders, however, use technical analysis in conjunction with fundamental analysis, doubling their positions when technical and fundamental indicators agree on the direction of exchange rate movements.

Three principles guide the behavior of technical analysts.[1] The first is that market action (prices and transactions volume) "discounts" everything. In other words, all relevant information about an asset is incorporated into its price history, so there is no need to forecast the fundamental determinants of an asset's value. In fact, Murphy (1986) claims that asset price changes often precede observed changes in fundamentals. The second principle is that asset prices move in trends. Predictable trends are essential to the success of technical analysis because they enable traders to profit by buying (selling) assets when the price is rising (falling), or as technicians counsel, "the trend is your friend." Practitioners appeal to Newton's law of motion to explain the

existence of trends: Trends in motion tend to remain in motion unless acted upon by another force. The third principle of technical analysis is that history repeats itself. Asset traders will tend to react the same way when confronted by the same conditions. Technical analysts do not claim their methods are magical; rather, they take advantage of market psychology.

Following from these principles, the methods of technical analysis attempt to identify trends and reversals of trends. These methods are explicitly extrapolative; that is, they infer future price changes from those of the recent past. Formal methods of detecting trends are necessary because prices move up and down around the primary (or longer-run) trend. An example of this movement is shown in Figure 1, where the dollar/deutsche mark ($/DM) exchange rate fluctuates around an apparent uptrend.[2]

To distinguish trends from shorter-run fluctuations, technicians employ two types of analysis: *charting* and *mechanical rules.* Charting, the older of the two methods, involves graphing the history of prices over some period—determined by the practitioner—to predict future patterns in the data from the existence of past patterns. Its advocates admit that this subjective system requires the analyst to use judgement and skill in finding and interpreting patterns. The second type of method, mechanical rules, imposes consistency and discipline on the technician by requiring him to use rules based on mathematical functions of present and past exchange rates.

Charting

To identify trends through the use of charts, practitioners must first find *peaks* and *troughs* in the price series. A peak is the highest value of the exchange rate within a specified period of time (a local maximum), while a trough is the lowest value the price has taken on within the same period (a local minimum). A series of peaks and troughs establishes downtrends and uptrends, respectively. For example, as

[1] These principles and a much more comprehensive treatment of technical analysis are provided by Murphy (1986) and Pring (1991). Rosenberg and Shatz (1995) advocate the use of technical analysis with more economic explanation.

[2] Figure 1 shows only closing prices. In this, it differs from most charts employed by technical traders, which might show the opening, closing, and daily trading range.

shown in Figure 1, an analyst may establish an uptrend visually by connecting two local troughs in the data. A trendline is drawn below an apparent up trend or above an apparent downtrend. As more troughs touch the trendline without violating it, the technician may place more confidence in the validity of the trendline. The angle of the trendline indicates the speed of the trend, with steeper lines indicating faster appreciation (or depreciation) of the foreign currency.

After a trendline has been established, the technician trades with the trend, buying the foreign currency if an uptrend is signaled and selling the foreign currency if a downtrend seems likely. When a market participant buys a foreign currency in the hope that it will go up in price, that participant is said to be *long* in the currency. The opposite strategy, called *shorting* or *selling short*, enables the participant to make money if the foreign currency falls in price. A short seller borrows foreign currency today and sells it, hoping the price will fall so that it can be bought back more cheaply in the future.

Spotting the reversal of a trend is just as important as detecting trends. Peaks and troughs are important in identifying reversals too. Local peaks are called *resistance levels*, and local troughs are called *support levels* (see Figure 1). If the price fails to break a resistance level (a local peak) during an uptrend, that may be an early indication that the trend may soon reverse. If the exchange rate significantly penetrates the trendline, that is considered a more serious signal of a possible reversal.

Technicians identify several patterns that are said to foretell a shift from a trend in one direction to a trend in the opposite direction. An example of the best-known type of reversal formation, called "head and shoulders," is shown in Figure 2. The head and shoulders reversal following an uptrend is characterized by three local peaks with the middle peak being the largest of the three. The line between the troughs of the shoulders is known as the "neckline." When the exchange rate penetrates the neckline of a head and

Figure 2

### The Head and Shoulders Reversal Pattern Illustrated for the $/DM

shoulders, the technician confirms a reversal of the previous uptrend and begins to sell the foreign currency. There are several other similar reversal patterns, including the V (single peak), the double top (two similar peaks) and the triple top (three similar peaks). The reversal patterns of a downtrend are essentially the mirrors of the reversal patterns for the uptrend.

Mechanical Rules

Charting is very dependent on the interpretation of the technician who is drawing the charts and interpreting the patterns. Subjectivity can permit emotions like fear or greed to affect the trading strategy. The class of mechanical trading rules avoids this subjectivity and so is more consistent and disciplined, but, according to some technicians, it sacrifices some information that a skilled chartist might discern from the data. Mechanical trading rules are even more explicitly extrapolative than charting; they look for trends and follow those trends. A well-known type of mechanical trading rule is the "filter rule," or "trading range break" rule which counsels buying (selling) a currency when it rises (falls) x percent above (below) its previous local minimum (maximum). The size of the filter, x, which is

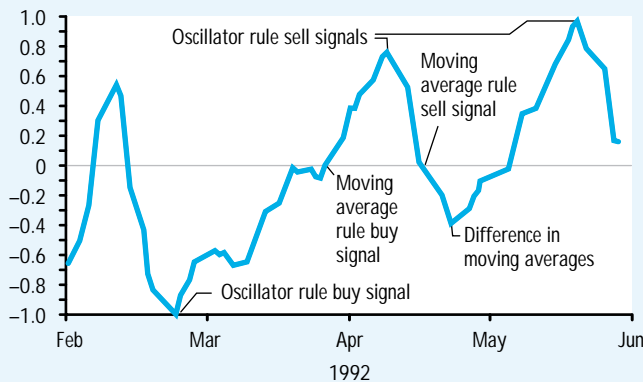## Figure 3

5- and 20-Day Moving Averages

$ per DM
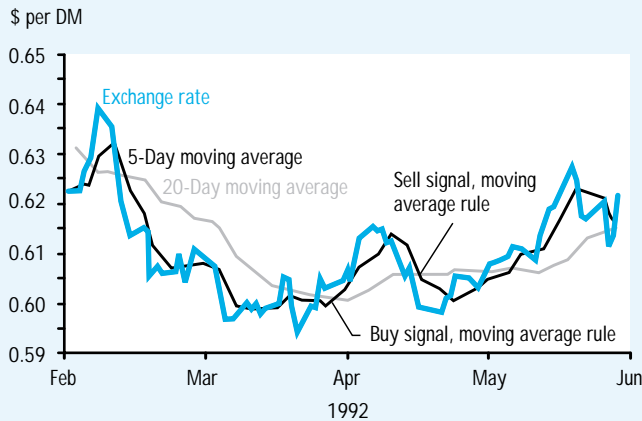
NOTES: These moving averages smooth the exchange rate and can be used to generate buy and sell signals in the foreign exchange market.

Figure 4

The Oscillator Index

Normalized difference in moving averages

NOTES: The 5-day moving average minus the 20-day moving average can also be used to generate buy and sell signals.

[3] For example, the five-day moving average of an exchange rate series is given by:

$$M(5)_t = \frac{1}{5} \sum_{i=0}^{4} S_{t-i}$$

where S_t denotes the closing price of the spot exchange rate at day t.

exchange rate over a given number of previous trading days. The length of the moving average "window"—the number of days in the moving average—governs whether the moving average reflects long- or short-run trends.[3] Any moving average will be smoother than the original exchange-rate series, and long moving averages will be smoother than short moving averages. Figure 3 illustrates the behavior of a 5-day and a 20-day moving average of the exchange rate in relation to the exchange rate itself. A typical moving average trading rule prescribes a buy (sell) signal when a short moving average crosses a longer moving average from below (above)—that is, when the exchange rate is rising (falling) relatively fast. Of course, the lengths of the moving averages must be chosen by the technician. The length of the short moving average rule is sometimes chosen to equal one, the exchange rate itself.

A final type of mechanical trading rule is the class of "oscillators," which are said to be useful in *non-trending* markets, when the exchange rate is not trending up or down strongly. A simple type of oscillator index, an example of which is shown in Figure 4, is given by the difference between two moving averages: the 5-day moving average minus the 20-day moving average. Oscillator rules suggest buying (selling) the foreign currency when the oscillator index takes an extremely low (high) value. Note that the oscillator index, as a difference between moving averages, also generates buy/sell signals from a moving average rule when the index crosses zero. That is, when the short moving average becomes larger than the long moving average, the moving average rule will generate a buy signal. By definition, this will happen when the oscillator index goes from negative to positive. Therefore, an oscillator chart is also useful for generating moving average rule signals.

Other Kinds of Technical Analysis

Technical analysis is more complex and contains many more techniques than those described in this article. For

chosen by the technician from past experience, is generally between 0.5 percent and 3 percent. Figure 1 illustrates some of the buy and sell signals generated by a filter rule with filter size of 0.5 percent.

A second variety of mechanical trading rule is the "moving average" class. Like trendlines and filter rules, moving averages bypass the short-run zigs and zags of the exchange rate to permit the technician to examine trends in the series. A moving average is the average closing price of the

example, many technical analysts assign a special role to round numbers in support or resistance levels. When the exchange rate significantly crosses the level of 100 yen to the dollar, that is seen as an indication that further movement in the same direction is likely.[4] Other prominent types of technical analysis use exotic mathematical concepts such as Elliot wave theory and/or Fibonacci numbers.[5] Finally, traders sometimes use technical analysis of one market's price history to take positions in another market, a practice called inter-market technical analysis.

EFFICIENT MARKETS AND TECHNICAL ANALYSIS

Technical analysts believe that their methods will permit them to beat the market. Economists have traditionally been skeptical of the value of technical analysis, affirming the theory of efficient markets that holds that no strategy should allow investors and traders to make unusual returns except by taking excessive risk.[6]

Investing in the Foreign Exchange Market

To understand the efficient markets hypothesis in the context of foreign exchange trading, consider the options open to an American bank (or firm) that temporarily has excess funds to be invested overnight. The bank could lend that money in the overnight bank money market, known as the federal funds market. The simple net return on each dollar invested this way would be the overnight interest rate on dollar deposits. The bank has other investment options, though. It could instead convert its money to a foreign currency (e.g., the deutsche mark), lend its money in the overnight German money market (at the German interest rate) and then convert it back to dollars tomorrow. This return is the sum of the German overnight interest rate and the change in the value of the DM. Which investment should the bank choose? If the bank were not concerned about risk, it would choose

the investment with the higher expected return. While the U.S. and German interest rates are known, the bank must base its decision on its forecast of the rate of appreciation of the DM. If market participants expect the return to investing in the German money market to be higher than that of investing in the U.S. money market, they will all try to invest in the German market, and none will invest in the U.S. money market. Such a situation would tend to drive down the German return and raise the U.S. return until the two were equalized. The *excess return* on a German investment over an investment in the U.S. money market (R_t^{DM}), at date t, from the point of view of a U.S. investor is defined as

$$(1) \qquad R_t^{DM} \equiv i_t^{DM} + \Delta S_t - i_t^{\$},$$

where i_t^{DM} is the German overnight interest rate, ΔS_t is the percentage rate of appreciation of the DM against the dollar overnight, and $i_t^{\$}$ is the U.S. overnight interest rate.[7] If market participants cared only about the expected return on their investments, and if their expectations about the change in the exchange rate were not systematically wrong, the expected excess return on foreign exchange should equal zero, every day.

The assumption that market participants care only about the expected return is too strong, of course. Surely, participants also care about the *risk* of their investment.[8] Risk can come from either the risk of default on the loan or the risk of sharp changes in the exchange rate, or both. If investing in the German market is significantly riskier than investing in the U.S. market, investors must be compensated with a higher expected return in the German market, or they will not invest there. In that case, the expected excess return would be positive and equal to a risk premium. The expected risk-adjusted excess return would be equal to zero. That is,

$$(2) \qquad E[R_t^{DM}] - RP_t = 0,$$

where $E[*]$ is a function that takes the expected value of the term inside the

[4] "The 100 yen level for the dollar is still a very big psychological barrier and it will take a few tests before it breaks. But once you break 100 yen, it's not going to remain there for long. You'll probably see it trade between 102 and 106 for a while," said Jorge Rodriguez, director of North American Sales at Credit Suisse, as reported by Creswell (1995).

[5] Murphy (1986) discusses Elliot wave theory, Fibonacci numbers, and many other technical concepts.

[6] Samuelson (1965) did seminal theoretical work on the modern theory of efficient markets.

[7] The excess return may also be considered the return to someone borrowing in dollars and investing those dollars in German investments.

[8] Market participants may be concerned about the *liquidity* of their position as well as the expected return and risk. Liquidity is the ease with which assets can be converted into cash.

brackets [*] and RP_t is the risk premium associated with the higher risk of lending in the German market.

Efficient Markets

The idea that the expected risk-adjusted excess return on foreign exchange is zero implies a sensible statement of the efficient markets hypothesis in the foreign exchange context: Exchange rates reflect information to the point where the potential excess returns do not exceed the transactions costs of acting (trading) on that information.[9] In other words, you can't profit in asset markets (like the foreign exchange market) by trading on publicly available information.

This description of the efficient markets hypothesis appears to be a restatement of the first principle of technical analysis: Market action (price and transactions volume) discounts all information about the asset's value. There is, however, a subtle but important distinction between the efficient markets hypothesis and technical analysis: The efficient markets hypothesis posits that the current exchange rate adjusts to all information to prevent traders from reaping excess returns, while technical analysis holds that current and past price movements contain just the information needed to allow profitable trading.

What does this version of the efficient markets hypothesis imply for technical analysis? Under the efficient markets hypothesis, only current interest rates and risk factors help predict exchange rate changes, so past exchange rates are of no help in forecasting excess foreign exchange returns—i.e., if the hypothesis holds, technical analysis will not work. Malkiel's summary of the attitude of many economists toward technical analysis in the stock market is based on similar reasoning:

> The past history of stock prices cannot be used to predict the future in any meaningful way. Technical strategies are usually amusing, often comforting, but of no real value. (Malkiel, 1990, p. 154.)

How do prices move in the hypothetical efficient market? In an efficient market, profit seekers trade in a way that causes prices to move instantly in response to new information, because any information that makes an asset appear likely to become more valuable in the future causes an immediate price rise today. If prices *do* move instantly in response to all new information, past information, like prices, does not help anyone make money. If there were a way to make money with little risk from past prices, speculators would employ it until they bid away the money to be made. For example, if the price of an asset rose 10 percent every Wednesday, speculators would buy strongly on Tuesday, driving prices past the point where anyone would think they could rise much further, and so a fall would be likely. This situation could not lead to a predictable pattern of rises on Tuesday, though, because speculators would buy on Monday. Any pattern in prices would be quickly bid away by market participants seeking profits. Indeed, there is considerable evidence that markets often *do* work this way. Moorthy (1995) finds that foreign exchange rates react very quickly and efficiently to news of changes in U.S. employment figures, for example.

Because the efficient markets hypothesis is frequently misinterpreted, it is important to clarify what the idea does *not* mean. It does not mean that asset prices are unrelated to economic fundamentals.[10] Asset prices may be based on fundamentals like the purchasing power of the U.S. dollar or German mark. Similarly, the hypothesis does not mean that an asset price fluctuates randomly around its intrinsic (fundamental) value. If this were the case, a trader could make money by buying the asset when the price was relatively low and selling it when it was relatively high. Rather, "efficient markets" means that at any point in time, asset prices represent the market's best guess, based on all currently available information, as to the fundamental value of the asset. Future price *changes*, adjusted for risk, will be close to unpredictable.

But if any pattern in prices is quickly bid away, how does one explain the

[9] There are a number of versions of the efficient markets hypothesis. This version is close to that put forward by Jensen (1978).

[10] For an example of an incorrect interpretation of the efficient markets hypothesis, see Murphy (1986, p. 20-21) who offers, "The theory is based on the *efficient markets hypothesis*, which holds that prices fluctuate randomly about their intrinsic value. . . . it's just unrealistic to believe that *all* price movement is random."

apparent trends seen in charts of asset prices like those in Figure 1? Believers in efficient markets point out that completely random price changes—like those generated by flipping a coin—will produce price series that seem to have trends (Malkiel, 1990, or Paulos, 1995). Under efficient markets, however, traders cannot exploit those trends to make money, since the trends occur by chance and are as likely to reverse as to continue at any point. (For example, some families have—purely by chance—strings of either boys or girls, yet a family that already has four girls and is expecting a fifth child still has only a 50 percent chance of having another girl.)

EVALUATING TECHNICAL ANALYSIS

The efficient markets hypothesis requires that past prices cannot be used to predict exchange rate changes. If the hypothesis is true, technical analysis should not enable a trader to earn profits without accepting unusual risk. This section examines how two common types of trading rules are formulated and how the returns generated by these rules are measured. Problems inherent in testing the rules, measuring risk, and drawing conclusions about the degree of market efficiency are discussed.[11]

Finding a Trading Rule

A basic problem in evaluating technical trading strategies is that rules requiring judgement and skill are impossible to quantify and therefore unsuitable for testing. A fair test requires fixed, objective, commonly used trading rules to evaluate. An "objective" rule does not rely on individual skill or judgement to determine buy or sell decisions. The rule should be commonly used to reduce the problem of drawing false conclusions from "data mining"— a practice in which many different rules are tested until, purely by chance, some are found to be profitable on the data set. Negative test results are ignored, while positive results are

published and taken to indicate that trading rule strategies can yield profits. For example, there is a vast literature on pricing anomalies in the equity markets, summarized by Ball (1995) and Fortune (1991), but Roll (1994) has found that these aberrations are difficult to exploit in practice; he suggests that they may be partially the result of data mining.

Trading Rules

With these considerations, two kinds of trading rules have been commonly tested: *filter rules* and *moving average rules.* As a preceding section of this article explained, filter rules give a buy signal when the exchange rate rises x percent over the previous recent minimum. The analyst must make two choices to construct a filter rule: First, how much does the exchange rate have to rise, or what is the size of the filter? Second, how far back should the rule go in finding a recent minimum? The filter rules studied here will use filters from 0.5 percent to 3 percent and go back five business days to find the extrema.[12] A moving average rule gives a buy signal when a short moving average is greater than the long moving average; otherwise it gives a sell signal. This rule requires the researcher to choose the lengths of the moving averages. The moving average rules to be tested will use short moving averages of 1 day and 5 days and long moving averages of 10 days and 50 days. Both the filter rules and the moving average rules are extrapolative, in that they indicate that the trader should buy when the exchange rate has been rising and sell when it has been falling.

Profits

The trading rules switch between *long* and *short* positions in the foreign currency. Recall that a long position is a purchase of foreign currency—a bet that it will go up—while a short position is the reverse, selling borrowed foreign currency now in the hope that its value will fall. Denoting the percentage change in the exchange rate

[11] A number of previous studies have documented evidence of profitable technical trading rules in the foreign exchange market: Sweeney (1986); Levich and Thomas (1993); Neely, Weller, and Dittmar (1997).

[12] As with most aspects of technical analysis, the choice of filter size and window lengths has been determined by practitioners through a process of trial and error.

Table 1

Technical Trading Rule Results for the $/DM
Moving Average Rule Results

Short MA	Long MA	Annual Return	Monthly Standard Deviation	Number of Trades	Sharpe Ratio	Estimated CAPM Beta	Standard Error of Est. Beta
1	10	6.016	2.979	928	0.583	− 0.022	0.091
1	50	7.546	3.155	268	0.690	− 0.135	0.085
5	10	6.718	3.064	576	0.633	− 0.144	0.084
5	50	6.671	3.236	146	0.595	− 0.134	0.080

Filter Rule Results

Filter	Annual Return	Monthly Standard Deviation	Number of Trades	Sharpe Ratio	Estimated CAPM Beta	Standard Error of Est. Beta
0.005	5.739	3.057	1070	0.542	− 0.071	0.089
0.010	6.438	2.951	584	0.630	− 0.092	0.093
0.015	3.323	3.255	382	0.295	− 0.037	0.085
0.020	1.934	3.348	234	0.167	− 0.128	0.087
0.025	0.839	3.236	142	0.075	− 0.118	0.082
0.030	− 1.541	3.578	92	− 0.124	− 0.086	0.077

NOTES: The first two columns of the top panel characterize the length of the short and long moving averages used in the moving-average trading rule. The third column is the annualized asset return to the rule, while the fourth column is the monthly standard deviation of the return. The fifth column is the number of trades over the 23-year sample. The sixth column is the Sharpe ratio, and the last two columns provide the CAPM beta with the S&P 500 and the standard error of that estimate. The lower panel has a similar structure, except that the first column characterizes the size of the filter used in the rule. All extrema for filter rules were measured over the previous five business days.

[13] The estimate of transactions costs used here is consistent with recent figures. Levich and Thomas (1993) consider a round-trip cost of 0.05 percent realistic, as do Osler and Chang (1995).

[14] The exchange rate data were obtained from DRI and were collected at 4:00 p.m. local time in London from Natwest Markets and S&P Comstock. Daily overnight interest rates are collected by BIS at 9:00 a.m. London time. Interest rates for Japan were unavailable before 3/1/82, so the interest rates before this date were set to 0 for the $/¥ case.

($ per unit of foreign currency) from date t to $t+1$ by ΔS_t, and the domestic (foreign) overnight interest rate by $i_t^\$$ (i_t^{DM}), then the overnight return from a long position is approximately given by Equation 1:

$$(1) \qquad R_t^{DM} \equiv i_t^{DM} + \Delta S_t - i_t^\$.$$

The return to a short position is the negative of the return to a long position. The return to a trading rule over a period of time is approximately the sum of daily returns, minus transactions costs for each trade. Transactions costs are set at 5 basis points (0.05 percent) for each *round trip* in the currency. A round trip is a move from a long position to a short position and back or vice-versa.[13]

Evidence from Ten Simple Technical Trading Rules

Six filter rules and four moving average rules were tested on data consisting of the average of daily U.S. dollar bid and ask quotes for the DM, yen, pound sterling, and Swiss franc.[14] All exchange rate data begin on 3/1/74 and end on 4/10/97. These four series are called $/DM, $/¥, $/£, and $/SF. Because the results for the four exchange rates were similar, full results from only the $/DM will be reported in the tables.

Table 1 shows the annualized percentage return, monthly standard deviation (a measure of the volatility of returns), number of trades per year, and two measures of risk, the *Sharpe ratio* and the *CAPM beta*, for each of the 10 trading strategies for the $/DM. The Sharpe ratio and CAPM betas are discussed in some detail in the shaded insert. The mean annual return to the 10 rules was 4.4 percent, and 38 of the 40 trading rules were profitable (had positive excess return) over the whole sample. These results cast doubt on the efficient markets hypothesis, which holds that no trading strategy should be able to consistently earn positive excess

returns. The number of trades over the 23-year sample varied substantially over the 10 rules, ranging from 4 trades per year to almost 50 trades per year. The moving average rules were somewhat more profitable than the filter rules.

There is little evidence that these excess returns are compensation for bearing excessive risk. The first measure of risk, the Sharpe ratio, is the mean annual return divided by the mean annual standard deviation. The moving average rules had higher Sharpe ratios (0.6 vs. 0.25) than the filter rules. Six of the 10 Sharpe ratios are better than the 0.3 obtained by a buy-and-hold strategy in the S&P 500 over approximately the same period. This result indicates that the average return to the rules is very good compared to the risk involved in following the rules.
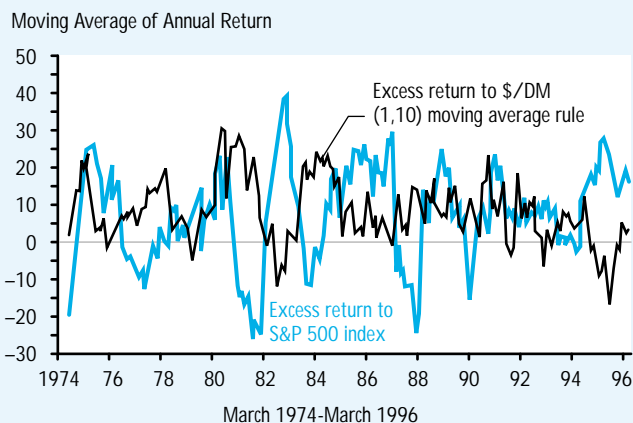
The second measure of risk, the CAPM betas, reflects the correlation between the monthly trading rule returns and the monthly returns to a broad portfolio of risky assets (the S&P 500). Significantly positive betas indicate that the rule is bearing undiversifiable risk. These CAPM betas estimated from the 10 rules generally indicate negative correlation with the S&P 500 monthly returns. None of them is significantly positive, statistically or economically. In other words, there is no systematic risk in these rules that could explain the positive excess returns.

For Whom is Technical Trading Appropriate?

The discussion of risk and returns suggests that technical analysis may be very useful for banks and large financial firms that can borrow and lend freely at the overnight interbank interest rate and buy and sell in the wholesale market for foreign exchange, where transactions sizes are in the millions of dollars. Technical trading is much less useful for individuals, who would face much higher transactions costs and must consider the opportunity cost of the time necessary to become an expert on foreign exchange speculating and to keep

Figure 5

One-Year Moving Average, Forward-Looking Excess Returns to the (1,10) Moving Average Trading Rule

Moving Average of Annual Return

Excess return to $/DM (1,10) moving average rule

Excess return to S&P 500 index

March 1974-March 1996

up with the market on a daily basis. How large would transactions costs have to be to eliminate the excess return to the technical rules? If we assume a 6 percent annual excess return to the rule and 230 trades (10 trades a year), round-trip transactions costs would have to be greater than 0.6 percent to produce zero excess returns.

In addition to higher transactions costs, individual investors following technical rules also must accept the risk that such a strategy entails. Figure 5 illustrates the risk by depicting, at monthly intervals, the one-year-ahead excess return from 1974 through 1996 for the (1,10) moving average rule on the $/DM and, for comparison, the total excess return on buying and holding the S&P 500 index, a popular measure of returns to a stock portfolio. The figure shows that the excess returns to both portfolios vary considerably at the annual horizon, often turning negative. While the technical trading rule excess return is less variable than the S&P excess return, it can still lead to significant losses for some subperiods. Two ways to measure losses over subperiods are the maximal single-period loss (maximum drawdown) and maximum loss in a calendar year. Over the period from March 1974 through March 1997, the maximum

that an investor could have lost by using the moving average trading rule was –28.2 percent; this loss, which would have occurred between March 7, 1995, and August 2, 1995 (a period of 149 days), translates into an annual rate of –69.2 percent. In other words, an investor using this rule would have lost almost 30 percent of his capital over this five-month period. Similarly, the maximum loss for this technical trading rule in a complete calendar year was –9.8 percent in 1995, but –17.8 percent for the S&P 500 in 1981.[15]

Perhaps the biggest obstacle to exploiting technical rules is that while the returns to stocks depend ultimately on the profitability of the firms in which the stock is held, the *source* of returns to technical analysis is not well understood; therefore, the investor does not know if the returns will persist into the future or even if they continue to exist at the present. Indeed, Figure 5 shows that the post-1992 return to the (1,10) moving average rule for the $/DM has been negative.

Do These Results Measure the Degree of Market Efficiency?

There are a number of problems associated with inferring the degree of market efficiency from the apparent profitability of these trading rules. The first problem is the data. To test the profitability of a trading rule, the researcher needs actual prices and interest rates from a series of simultaneous market transactions. Unfortunately, simultaneous quotes for daily exchange rates and interest rates are not generally available for a long time span. For example, these exchange-rate data were collected late in the afternoon, while the interest rates were collected in the morning. Although most economists judge this problem to be very minor, some argue that the trading rule decisions could not have been executed at the exchange rates and interest rates used.

The second problem is that without a good model of how to price risk, positive excess returns resulting from the use of trading rules cannot be used to measure the degree of inefficiency. Risk is notoriously difficult to measure. In fact, a major area of study for macro and financial economists for the last 10 years has been to explain why the return on stocks is so much higher than that on bonds, a phenomenon called the *equity premium puzzle*. Of course, at least part of the answer is that stocks are much riskier than bonds, but there is no generally accepted model of risk that will explain the *size* of the return difference.[16] Defenders of the efficient markets hypothesis maintain that the discovery of an apparently successful trading strategy may not indicate market inefficiency but, rather, that risk is not measured properly.

Another problem is that of "data mining": If enough rules are tested, some—purely by chance—will produce excess returns on the data. These rules may not have been obvious to traders at the beginning of the sample. In fact, the rules tested here are certainly subject to a data-mining bias, since many of them had been shown to be profitable on these exchange rates over at least some of the subsample. Closely related to the data-mining problem is the tendency to publish research that overturns the conventional wisdom on efficient markets, rather than research that shows technical analysis to be ineffective. One solution to the data-mining problem is suggested by Neely, Weller, and Dittmar (1997), who apply *genetic programming* techniques to the foreign-exchange market. Genetic programming is a method by which a computer searches through the space of possible technical trading rules to find a group of good rules (i.e., rules that generate positive excess return). These good rules are then tested on out-of-sample data to see if they continue to generate positive excess returns.

RETHINKING THE EFFICIENT MARKETS HYPOTHESIS

Early research in finance on the efficient markets hypothesis was very supportive; little evidence was found of profitable trading rules after transactions costs were accounted for (Fama, 1970).

[15] The returns for complete calendar years were available from 1975 through 1995.

[16] Kocherlakota (1996) and Siegel and Thaler (1997) discuss the equity premium puzzle extensively.

The success of technical trading rules shown in the previous section is typical of a number of later studies showing that the simple efficient markets hypothesis fails in important ways to describe how the foreign exchange market actually functions. While these results did not surprise market practitioners, they have helped persuade economists to examine features of the market like sequential trading, asymmetric information, and the role of risk that might explain the profitability of technical analysis.

The Paradox of Efficient Markets

Grossman and Stiglitz (1980) identified a major theoretical problem with the hypothesis termed the *paradox of efficient markets*, which they developed in the context of equity markets. As applied to the foreign exchange market, the argument starts by noting that exchange rate returns are determined by fundamentals like national price levels, interest rates, and public debt levels, and that information about these variables is costly for traders to gather and analyze. The traders must be able to make some excess returns by trading on this analysis, or they will not do it. But if markets were perfectly efficient, the traders would not be able to make excess returns on any available information. Therefore, markets cannot be perfectly efficient in the sense of exchange rates' always being *exactly* where fundamentals suggest they should be. Of course, one resolution to this paradox is to recognize that market analysts can recover the costs of some fundamental research by profiting from having marginally better information than the rest of the market on where the exchange rate should be. In this case, the exchange rate remains close enough to its fundamental value to prevent less informed people from profiting from the difference. Partly for these reasons, Campbell, Lo, and MacKinlay (1997) suggest that the debate about perfect efficiency is pointless and that it is more sensible to evaluate the degree of inefficiency than to test for absolute efficiency.

Empirical Reasons to Suspect Failure of Efficient Markets

The miserable empirical performance of standard exchange rate models is another reason to suspect the failure of the efficient markets hypothesis. In an important paper, Meese and Rogoff (1983) persuasively showed that no existing exchange rate model could forecast exchange rate changes better than a "no-change" guess at forecast horizons of up to one year. This was true even when the exchange rate models were given true values of future fundamentals like output and money. Although Mark (1995) and others have demonstrated some forecasting ability for these models at forecasting horizons greater than three years, no one has been able to convincingly overturn the Meese and Rogoff (1983) result despite 14 years of research. The efficient markets hypothesis is frequently misinterpreted as implying that exchange rate changes should be unpredictable; that is, exchange rates should follow a *random walk*. This is incorrect. Equation 2 shows that interest rate differentials should have forecasting power for exchange rate changes, leaving *excess returns* unpredictable. There is, however, convincing evidence that interest rates are *not* good forecasters of exchange rate changes.[17] According to Frankel (1996), this failure of exchange rate forecasting leaves two possibilities:

- Fundamentals are not observed well enough to allow forecasting of exchange rates.

- Exchange rates are detached from fundamentals by (possibly irrational) swings in expectations about future values of the exchange rate. These fluctuations in exchange rates are known as *bubbles*.[18]

Which of these possibilities is more likely? One clue is given by the relationship between exchange rates and fundamentals when expectations about the value of the exchange rate are very stable, as they are under a *fixed exchange rate*

[17] Engel (1995) reviews the failure of this theory, called *uncovered interest parity*.

[18] Swings in expectations that are subsequently justified by changes in the exchange rate are known as *rational bubbles*. Swings that are not consistent with the future path of exchange rates are *irrational bubbles*.

Figure 6

Monthly Percentage Changes in the $/DM Real Exchange Rate
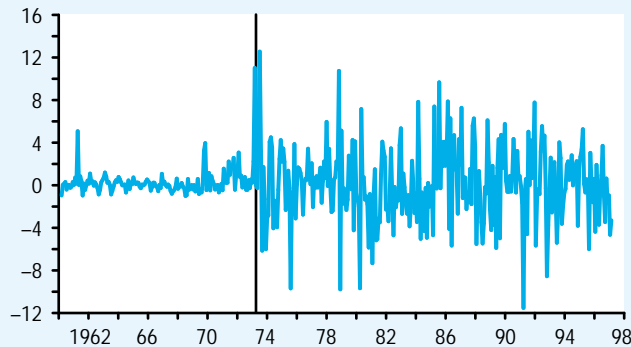
NOTES: These changes become much more volatile after the end of the Bretton Woods system of fixed exchange rates in March 1973. The vertical line denotes this break date in the series. Data cover January 1960–February 1997.

regime. A fixed exchange rate regime is a situation in which a government is committed to maintaining the value of its currency by manipulating monetary policy and trading foreign exchange reserves. Fixed exchange rate regimes are contrasted to *floating* regimes, in which the government has no such obligation. For example, most countries in the European Union had a type of fixed exchange rate regime, known as a target zone, from 1979 through the early 1990s. Fixed exchange rates anchor investor sentiment about the future value of a currency because of the government's commitment to stabilize its value. If fundamentals, like goods prices, or expectations based on fundamentals, rather than irrationally changing expectations, drive the exchange rate, the relationship between fundamentals and exchange rates should be the same under a fixed exchange rate regime as it is under a floating regime. This is not the case. Countries that move from floating exchange rates to fixed exchange rates experience a dramatic change in the relationship between prices and exchange rates. Specifically, real exchange rates (exchange rates adjusted for inflation in both countries) are much more volatile under floating exchange rate regimes, where expectations are not tied down by promises of government intervention. Figure 6 illustrates a typical case:

When Germany and the United States ceased to fix their currencies in March 1973, the variability in the real $/DM exchange rate increased dramatically. This result suggests that, contrary to the efficient markets hypothesis, swings in investor expectations may detach exchange rates from fundamental values in the short run.

Why Do Bubbles Arise?

If traders might profit by anticipating swings in investor expectations, then the efficient markets hypothesis needs significant adjustment. The structure of the foreign exchange market has several features that might help drive these swings in expectations that produce bubbles. Most foreign exchange transactions are conducted by large commercial banks in financial centers like London, New York, Tokyo, and Singapore. These large banks "make a market" in a currency by offering to buy or sell large quantities (generally more than $1 million) of currencies for a specific price in another currency (e.g., the dollar) on request. The exchange rates at which they are willing to buy or sell dollars are known as the bid and ask prices, respectively. The market is highly competitive, and transactions occur 24 hours a day over the telephone and automated trading systems. The first feature of this market that might influence technical trading is that specific transactions quantities and prices are not public information; the market is *non-transparent*. But the bid and ask exchange rates are easy to track, as banks freely quote them to any participant. Second, the trades take place *sequentially*—i.e., there is time to learn from previous trades. Third, the participants in this market differ from one another in the information they have and their willingness to tolerate risk.[19] In other words, the participants are heterogeneous.

How might these features combine to produce bubbles? To the extent that some participants are better informed about certain fundamentals than other agents (for instance, they will know more about their own and their customers' demand for foreign exchange), the trading behavior of

[19] It has long been assumed that there is little or no private information in foreign exchange markets, but this view has been forcefully challenged with respect to intraday trading by Ito, Lyons, and Melvin (1997).

the informed participants will reveal some of their private information to the uninformed agents. For example, if the informed agents know of fundamental forces that are likely to make the exchange rate rise in the future, they are likely to buy the foreign currency and thereby bid up the publicly observed bid and ask prices. The uninformed agents might infer from the rise that the rate will continue to rise and, as a result, they might buy more foreign exchange, pushing the rate up themselves in a self-fulfilling prophecy.[20] This inference from past price behavior is extrapolative technical analysis: It assumes that the exchange rate will continue moving as it has in the recent past. The uninformed traders may continue to buy foreign exchange past the point where it is supported by fundamentals. Although this story is most plausible for very high-frequency (intraday) trading, it might also generate longer-term swings in the exchange rate.

There are other explanations for extrapolative trading that jettison the assumption of rational behavior in favor of the study of how people really make decisions. This field, called *behavioral finance*, has concentrated on examples of seeming irrationality in decision making. Two findings of this field are that (1) experimental participants seem unusually optimistic about their chances for success in games and (2) the behavior and opinions of members of a group tend to reinforce common ideas or beliefs.[21] For example, members of a jury may become more confident about their individual verdicts if the other members of the group agree.

Either explanation for extrapolative trading implies that bubbles may be produced by slow dissemination of private information into the market, coupled with extrapolative trading rules. There is some evidence to support this explanation. Eichenbaum and Evans (1995) found that foreign exchange markets reacted gradually to money supply shocks, over a period of many months, instead of instantly incorporating the new information. Surveys revealed that foreign exchange market participants' expectations are extrapolative at horizons up to six months. That is, if the exchange rate has risen recently, market participants expect it to continue to rise in the near future (Frankel and Froot, 1987). Also, the success of extrapolative traders tends to feed on itself. Frankel and Froot (1990) argue that extrapolative traders' success during the early part of the large dollar appreciation of 1981-1985 convinced many other traders to follow extrapolative rules, driving the dollar up even further.

Central Bank Intervention

The other popular explanation for the apparent profitability of technical trading rules is that technical traders are able to profit consistently from central bank intervention. Some central banks frequently intervene (buy and sell currency) in the foreign exchange market to move the exchange rate to help influence other variables like employment or inflation.[22] Because these actions are designed to control macroeconomic variables rather than to make money, central banks may be willing to take a loss on their trading. Trading rule profits may represent a transfer from central banks to technical traders. Lebaron (1996) found that most trading rule profits were generated on the day before a U.S. intervention. Neely and Weller (1997) find that "intelligent" trading rules tend to trade against the Fed; that is, they tend to buy dollars when they find out the Fed is selling dollars. This tantalizing story does not fit all the facts, however. For example, Leahy (1995) finds that U.S. foreign exchange operations make positive profits, on average.[23] This finding is inconsistent with the idea that central banks are giving money away to technical traders.

Why Are the Profits Not Arbitraged Away?

Whether the trends or inefficiencies in exchange rates are created by swings in expectations or by central bank intervention, efficient market advocates would ask why any predictable returns in exchange rates

[20] Treynor and Ferguson (1985), Brown and Jennings (1989), Banerjee (1992), and Kirman (1993) construct models of behavior in which information is inferred from the actions of others. One easily understood example is the problem of consumers who must choose between two restaurants. One seemingly sensible strategy for choosing would be to go to the more crowded restaurant on the theory that it is likely to be crowded because it has better food. This phenomenon depends on *asymmetric information*.

[21] Shiller (1988) and Shleifer and Summers (1990) discuss behavioral finance in more detail. Ohanian (1996) considers the reasons for the collapse of bubbles.

[22] In the United States, the Federal Reserve and the U.S. Treasury generally collaborate on foreign exchange intervention decisions, and operations are conducted by the Federal Reserve Bank of New York on behalf of both.

[23] See Szakmary and Mathur (1996) for more on central bank intervention and trading rule profits.

HOW TO MEASURE RISK?

The simplest widely used measure of risk is the *Sharpe ratio* or the ratio of the average annual excess return to a measure of excess return volatility called the standard deviation. Higher Sharpe ratios are desirable because they indicate either higher average excess returns or less volatility. A commonly used benchmark of a good Sharpe ratio is that of the S&P 500, which Osler and Chang (1995) estimated to be about 0.32 from March 1973 to March 1994.

A major drawback to Sharpe ratios is that they ignore an important idea in finance: An investment is risky only to the extent that its return is correlated with the return to a broad measure of the investments available. To see this, consider the risk associated with holding a portfolio of assets whose returns are each individually volatile but completely independent of each other. Each year, the assets in the portfolio that do unusually well will tend to offset those that do unusually poorly. The portfolio as a whole will be much less risky than any of the individual assets. The more assets in the portfolio, the less risky it will be. In fact, if enough of these independent assets are grouped together into a portfolio, the return on this portfolio becomes certain. This means that investors do not need to be compensated for holding risky assets that are not correlated with all the other assets they can buy (the market portfolio), because the risk of each uncorrelated asset can be reduced to zero if the portfolio contains a large enough variety of these assets. On the other hand, assets for which returns are positively correlated with those of the other assets on the market need a higher expected return to convince investors to hold them.

This idea motivates the second measure of riskiness, the *CAPM beta*: the coefficient from the linear regression of an asset's (or trading rule's) excess return on the excess return of a proxy for the market portfolio, the return to a broad equity index like the S&P 500. An estimated beta equal to zero means that the trading rule is bearing no systematic risk, while significantly positive betas indicate that a trading strategy is bearing some risk, and a beta equal to one means that the trading rule moves closely with the market, so that following it requires the investor to accept significant risk.

---

should not be arbitraged away. One answer to this question is that speculators have short horizons and are deterred from speculating against the trends by the risk that such a strategy would incur. There are several reasons for this: First, traders typically operate on margin, borrowing some of the money with which they trade. With a limited line of credit, the borrowing costs would add up if traders were not able to turn a quick profit. Second, a trader's performance is typically evaluated on relatively short horizons (less than a year). Third, there may be institutional or legal restrictions that prevent some types of enterprises from taking on "excessive" exchange risk. And finally, traders do not

know the equilibrium value of the exchange rate with any certainty, so they cannot distinguish bubbles from movements in fundamentals. Investors who bet on long-run reversion to fundamental values in exchange rates may be wiped out by short-run deviations away from those values.[24]

Explaining the success of technical trading rules with bubbles begs one more question: Why do destabilizing extrapolative traders not lose their money? Friedman (1953) showed that destabilizing speculation is doomed to lose money and so drive the speculators out of the market. Friedman argued that speculation can only destabilize asset prices if the speculators consistently buy when the asset price is above its equi-

[24] Both Shleifer and Summers (1990) and Shleifer and Vishny (1997) discuss the importance of risk in speculating against bubbles.

[25] Essentially the same argument is presented more simply in Shleifer and Summers (1990).

librium value (driving the price up further) and sell when the asset price is below its equilibrium value; as the destabilizing speculators lose their money, he maintained, they will have less effect on the market. The corollary to this argument is that all successful speculation is stabilizing. Delong, Schleifer, Summers, and Waldman (1989) constructed a "noise trader" model that questioned this logic, however.[25] They showed that irrational ("noise") traders could create so much risk in asset markets that the returns to those assets would have to be unusually high for rational traders to trade in them at all. In other words, the irrational traders make unusually high returns (on average) by foolishly pursuing risky strategies. Some go out of business, but, on average, this group increases its market position.

CONCLUSION

Technical analysis is the most widely used trading strategy in the foreign exchange market. Traders stake large positions on their interpretations of patterns in the data. Economists have traditionally rejected the claims of technical analysts because of the appealing logic of the efficient markets hypothesis. More recently, however, the discovery of profitable technical trading rules and other evidence against efficient markets have led to a rethinking about the importance of institutional features that might justify extrapolative technical analysis such as private information, sequential trading, and central bank intervention, as well as the role of risk.

The weight of the evidence now suggests that excess returns have been available to technical foreign exchange traders over long periods. Risk is hard to define and measure, however, and this difficulty has obscured the degree of inefficiency in the foreign exchange market. There is no guarantee, of course, that technical rules will continue to generate excess returns in the future; the excess returns may be bid away by market participants. Indeed, this may already be occurring. Continued research on high-frequency transactions data or experimental work on expectations formation may provide a better understanding of market behavior.

REFERENCES

Ball, Ray. "The Theory of Stock Market Efficiency: Accomplishments and Limitations," *Journal of Applied Corporate Finance* (Spring 1995), pp. 4-17.

Banerjee, Abhijit V. "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* (August 1992), pp. 797-817.

Brown, David P., and Robert H. Jennings. "On Technical Analysis," *The Review of Financial Studies* (1989), pp. 527-51.

Campbell, John Y., Andrew W. Lo, and Archie Craig MacKinlay. *The Econometrics of Financial Markets*, Princeton University Press, 1997.

Creswell, Juli. "Currency Market Expects Rate Cut By Bank of Japan," *Wall Street Journal*, September 5, 1995, p. C16.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. "Noise Trader Risk in Financial Markets," *Journal of Political Economy* (August 1990), pp. 703-38.

Eichenbaum, Martin, and Charles L. Evans. "Some Empirical Evidence on the Effects of Shocks to Monetary Policy on Exchange Rates," *Quarterly Journal of Economics* (November 1995), pp. 975-1009.

Engel, Charles. "Why is the Forward Exchange Rate Forecast Biased? A Survey of Recent Evidence," Federal Reserve Bank of Kansas City Working Paper 95-06, September 1995.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* (May 1970), pp. 383-417.

Fortune, Peter. "Stock Market Efficiency: An Autopsy?," *New England Economic Review*, Federal Reserve Bank of Boston (March/April 1991), pp. 18-40.

Frankel, Jeffrey. "How Well Do Foreign Exchange Markets Function: Might a Tobin Tax Help?," National Bureau of Economic Research Working Paper 5422, January 1996.

_____ and Kenneth A. Froot. "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *The American Economic Review* (March 1987), pp. 133-53.

_____ and _____. "Chartists, Fundamentalists and the Demand for Dollars," *Private Behavior and Government Policy in Interdependent Economies*, Anthony S. Courakis and Mark Taylor, eds., Clarendon Press, 1990.

Friedman, Milton. "The Case for Flexible Exchange Rates," *Essays in Positive Economics*, University of Chicago Press, 1953.

Grossman, Sanford J., and Joseph E. Stiglitz. "On the Impossibility of Informationally Efficient Markets," *The American Economic Review* (June 1980), pp. 393-408.

Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin. "Is There Private Information in the FX Market? The Tokyo Experiment," National Bureau of Economic Research Working Paper 5936, February 1997.

Jensen, Michael C. "Some Anomalous Evidence Regarding Market Efficiency," *Journal of Financial Economics* (June/September 1978), pp. 95-101.

Kirman, Alan. "Ants, Rationality and Recruitment," *Quarterly Journal of Economics* (February 1993), pp. 137-56.

Kocherlakota, Narayana R. "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature* (March 1996), pp. 42-71.

Leahy, Michael P. "The profitability of US intervention in the foreign exchange markets," *Journal of International Money and Finance* (December 1995), pp. 823-44.

Lebaron, Blake. "Technical Trading Rule Profitability and Central Bank Intervention," National Bureau of Economic Research Working Paper 5505, March 1996.

Levich, Richard M., and Lee R. Thomas. "The Significance of Technical Trading-Rule Profits in the Foreign Exchange Market: A Bootstrap Approach," *Journal of International Money and Finance* (October 1993), pp. 451-74.

Malkiel, Burton G. *A Random Walk Down Wall Street*, Fifth Edition, W. W. Norton & Company, 1990.

Mark, Nelson C. "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *The American Economic Review* (March 1995), pp. 201-18.

Meese, Richard A., and Kenneth Rogoff. "Empirical Exchange Rate Models of the Seventies: Do They Fit out of Sample?," *Journal of International Economics* (February 1983), pp. 3-24.

Moorthy, Vivek. "Efficiency Aspects of Exchange Rate Response to News: Evidence from U.S. Employment Data," *Journal of International Financial Markets, Institutions and Money* (1995), pp. 1-18.

Murphy, John J. *Technical Analysis of the Futures Markets*, New York Institute of Finance, Prentice-Hall, New York, 1986.

Neely, Chris, and Paul Weller. "Technical Analysis and Central Bank Intervention," Federal Reserve Bank of St. Louis Working Paper 97-002A, January 1997.

_____, _____, and Robert Dittmar. "Is Technical Analysis Profitable in the Foreign Exchange Market? A Genetic Programming Approach," Forthcoming in *Journal of Financial and Quantitative Analysis* (December 1997).

Ohanian, Lee E. "When the Bubble Bursts: Psychology or Fundamentals?," *Business Review*, Federal Reserve Bank of Philadelphia (January/February 1996), pp. 3-13.

Osler, Carol L., and P. H. Kevin Chang. "Head and Shoulders: Not Just a Flaky Pattern," Federal Reserve Bank of New York Staff Paper 4, August 1995.

Paulos, John Allen. *A Mathematician Reads the Newspaper*, Basic Books, 1995.

Pring, Martin J. *Technical Analysis Explained*, Third Edition, McGraw-Hill, 1991.

Roll, Richard. "What Every CFO Should Know About Scientific Progress in Financial Economics: What is Known, and What Remains to be Resolved," *Financial Management* (Summer 1994), pp. 69-75.

Rosenberg, Michael R., and Eric A. Shatz. *The Merrill Lynch Guide to Understanding and Using Technical Analysis*, Merrill Lynch and Co., Global Securities Research & Economics Group, 1995.

Samuelson, Paul A. "Proof that Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review* (1965), pp. 41-49.

Shiller, Robert J. "Fashions, Fads and Bubbles in Financial Markets," *Knights, Raiders and Targets: The Impact of the Hostile Takeover*, J. Coffee, S. Ackerman, and L. Lowenstein, eds., Oxford University Press, 1988. Reprinted in *Market Volatility*, by Robert J. Shiller, MIT Press, 1989.

Shleifer, Andrei, and Lawrence Summers. "The Noise Trader Approach to Finance," *Journal of Economic Perspectives* (Spring 1990), pp. 19-33.

_____ and Robert W. Vishny. "The Limits of Arbitrage," *Journal of Finance* (March 1997), pp. 35-55.

Siegel, Jeremy J., and Richard H. Thaler. "Anomalies: The Equity Premium Puzzle," *Journal of Economic Perspectives* (Winter 1997), pp. 191-200.

Sweeney, Richard J. "Beating the foreign exchange market," *Journal of Finance* (March 1986), pp. 163-82.

Szakmary, Andrew C., and Ike Mathur. "Central Bank Intervention and Trading Rule Profits in Foreign Exchange Markets," *Journal of International Money and Finance* (August 1997), pp. 513-35.

Taylor, Mark P., and Helen Allen. "The use of technical analysis in the foreign exchange market," *Journal of International Money and Finance* (June 1992), pp. 304-14.

Treynor, Jack L., and Robert Ferguson. "In Defense of Technical Analysis," *Journal of Finance* (July 1985), pp. 757-73.