

Simple Technical Trading Rules and the Stochastic Properties of Stock Returns

WILLIAM BROCK, JOSEF LAKONISHOK, and
BLAKE LeBARON*

ABSTRACT

This paper tests two of the simplest and most popular trading rules—moving average and trading range break—by utilizing the Dow Jones Index from 1897 to 1986. Standard statistical analysis is extended through the use of bootstrap techniques. Overall, our results provide strong support for the technical strategies. The returns obtained from these strategies are not consistent with four popular null models: the random walk, the AR(1), the GARCH-M, and the Exponential GARCH. Buy signals consistently generate higher returns than sell signals, and further, the returns following buy signals are less volatile than returns following sell signals, and further, the returns following buy signals are less volatile than returns following sell signals. Moreover, returns following sell signals are negative, which is not easily explained by any of the currently existing equilibrium models.

THE TERM “TECHNICAL ANALYSIS” is a general heading for a myriad of trading techniques. Technical analysts attempt to forecast prices by the study of past prices and a few other related summary statistics about security trading. They believe that shifts in supply and demand can be detected in charts of market action. Technical analysis is considered by many to be the original form of investment analysis, dating back to the 1800s. It came into widespread use before the period of extensive and fully disclosed financial information, which in turn enabled the practice of fundamental analysis to develop. In the United States, the use of trading rules to detect patterns in stock prices is probably as old as the stock market itself. The oldest technique is attributed to Charles Dow and is traced to the late 1800s. Many of the techniques used today have been utilized for over 60 years. These techniques for discovering hidden relations in stock returns can range from extremely simple to quite elaborate.

*Brock and LeBaron are from the University of Wisconsin and Lakonishok is from the University of Illinois. We are grateful to Tim Bollerslev, K. C. Chan, Louis Chan, Eugene Fama, Bruce Lehmann, Mark Ready, Jay Ritter, William Schwert, Theo Vermaelen, and the editor, René Stulz, and anonymous referees. The paper was presented at the AFA Meetings at New Orleans, The Amsterdam Institute of Finance, the NBER Summer Institute, the University of Limburg, and The Wharton School. We are thankful to Hank Pruden, a leading technical analyst, for guidance on some of the technical analysis literature. Brock was partially supported by the National Science Foundation (SES 87-20671), the Vilas Trust, and the University of Wisconsin Graduate School. LeBaron was partially supported by the National Science Foundation (SES 91-09671), and the University of Wisconsin Graduate School.

The attitude of academics towards technical analysis, until recently, is well described by Malkiel (1981):

Obviously, I am biased against the chartist. This is not only a personal predilection, but a professional one as well. Technical analysis is anathema to the academic world. We love to pick on it. Our bullying tactics are prompted by two considerations: (1) the method is patently false; and (2) it's easy to pick on. And while it may seem a bit unfair to pick on such a sorry target, just remember: it is your money we are trying to save.

Nonetheless, technical analysis has been enjoying a renaissance on Wall Street. All major brokerage firms publish technical commentary on the market and individual securities, and many of the newsletters published by various "experts" are based on technical analysis.

In recent years the efficient market hypothesis has come under serious siege. Various papers suggested that stock returns are not fully explained by common risk measures.¹ A line of research directly related to this work provides evidence of predictability of equity returns from past returns.² In general, the results of these studies are in sharp contrast with most earlier studies that supported the random walk hypothesis and concluded that the predictable variation in equity returns was economically and statistically very small. Two competing explanations for the presence of predictable variation in stock returns have been suggested: (1) market inefficiency in which prices take swings from their fundamental values, and (2) markets are efficient and the predictable variation can be explained by time-varying equilibrium returns.³ There is no evidence so far that unambiguously distinguishes these two competing hypotheses.

¹ A significant relationship between expected return and fundamental variables such as price-earnings ratio, market-to-book ratio, and size was documented. Another group of papers has uncovered systematic patterns in stock returns related to various calendar periods such as the weekend effect, the turn-of-the-month effect, the holiday effect, and the January effect.

² Chopra, Lakonishok, and Ritter (1992), De Bondt and Thaler (1985), Fama and French (1986), and Poterba and Summers (1988) find negative serial correlation in returns of individual stocks and various portfolios over three- to ten-year intervals. Rosenberg, Reid, and Lanstein (1985) provide evidence for the presence of predictable return reversals on a monthly basis at the level of individual securities. Jegadeesh (1990) finds negative serial correlation for lags up to two months and positive serial correlation for longer lags. Lo and MacKinlay (1990a) report positive serial correlation in weekly returns for indexes and portfolios and a somewhat negative serial correlation for individual stocks. Lehmann (1990) and French and Roll (1986) report negative serial correlation at the level of individual securities for weekly and daily returns. Cutler, Poterba, and Summers (1990) present results from many different asset markets generally supporting the hypothesis that returns are positively correlated at the horizon of several months and negatively correlated at the 3-to-5 year horizon.

³ Other explanations for predictability of returns over short horizons often mentioned are based on market microstructure stories. According to these explanations, reversals in recorded returns can be accounted for by movements from the bid to the ask. However, our strategies are not based on reversals, hence the microstructure explanation is implausible.

Although many earlier studies concluded that technical analysis is useless, the recent studies on predictability of equity returns from past returns suggest that this conclusion might have been premature.⁴ In this paper we explore two of the simplest and most popular technical rules: moving average-oscillator and trading-range break (resistance and support levels). In the first method, buy and sell signals are generated by two moving averages, a long period, and a short period. In the second method signals are generated as stock prices hit new highs and lows. These rules will be evaluated by their ability to forecast future price changes. For statistical inferences, standard tests will be augmented with the bootstrap methodology inspired by Efron (1979), Freedman and Peters (1984a, 1984b), and Efron and Tibshirani (1986). Following this methodology, returns from an artificial Dow series are generated and the trading rules are applied to the series. Comparisons are then made between returns from these simulated series and the actual Dow Jones series.

Neither the bootstrap methodology nor the use of technical analysis to evaluate model specifications are in particular new to the finance literature. The contribution of this paper lies in the combination of these two techniques. This procedure allows testing a wide range of null models. When models are rejected by such a statistical test, information is provided on how to modify the model to achieve a better description of the series. In addition, the trading rules used in this paper may have power against certain alternatives that are difficult to detect using standard statistical tests.

Few, if any, empirical tests in financial economics are free of the data-instigated pre-test biases discussed in Leamer (1978).⁵ The more scrutiny a collection of data receives, the more likely "interesting" spurious patterns will be observed. Stock prices are probably the most studied financial series and, therefore, most susceptible to data snooping. In addition, Merton (1987) suggests that individuals have a tendency to come up with "exciting" spurious results (anomalies):

All this fits well with what the cognitive psychologists tell us is our natural individual predilection to focus, often disproportionately so, on the unusual... This focus, both individually and institutionally, together with little control over the number of tests performed, creates a fertile environment for both unintended selection bias and for attaching greater significance to otherwise unbiased estimates than is justified.

Therefore, the possibility that various spurious patterns were uncovered by technical analysis cannot be dismissed. Although a complete remedy for data-snooping biases does not exist, we mitigate this problem: (1) by report-

⁴ Some of the earlier work on technical analysis includes papers by Alexander (1961, 1964), Fama and Blume (1966), Levy (1967a, 1967b), Jensen (1967), and Jensen and Bennington (1970).
⁵ recent paper in this area is by Sweeney (1988).

⁵ Data-snooping issues are also discussed in Lakonishok and Smidt (1988) and Lo and MacKinlay (1990b).

ing results from all our trading strategies, (2) by utilizing a very long data series, the Dow Jones index from 1897 to 1986, and (3) emphasizing the robustness of results across various nonoverlapping subperiods for statistical inference.⁶

Our study reveals that technical analysis helps to predict stock price changes. The patterns uncovered by technical rules cannot be explained by first order autocorrelation and by changing expected returns caused by changes in volatility. To put it differently, the trading profits are not consistent with a random walk, an AR(1), a GARCH-M model, or an Exponential GARCH. The results generally show that returns during buy periods are larger than returns during sell periods. Moreover, returns during buy periods are less volatile than returns during sell periods. For example, the variable-length moving-average strategy produced on average a daily return for buy periods of 0.042 percent, which is about 12 percent per year. In contrast, the corresponding daily return for the sell periods is -0.025 percent, or about -7 percent per year. This strategy results in a daily standard deviation of 0.89 percent for buy periods and a higher one, 1.34 percent, for sell periods.

The remainder of the paper is organized as follows: Section I describes the data and our technical trading rules; Section II presents the empirical results of the tests utilizing traditional techniques; Section III describes the bootstrap methodology, Section IV presents the empirical results from the bootstrap simulations, and Section V concludes and summarizes our results.

I. Data and Technical Trading Rules

A. Data

The data series used in this study is the Dow Jones Industrial Average (DJIA) from the first trading day in 1897 to the last trading day in 1986—a collection of 90 years of daily data. Stock price averages are available on a daily basis back to February 1985, but 1897 was the first full year for the industrial average. Before this date, Charles H. Dow, editor of the *Wall Street*

⁶ Our selection of trading strategies was clearly influenced by previous work in this area. However, all the strategies chosen have a very long history. Modern technical analysis probably originated in the work of Charles Dow near the turn of the century. Examples of the rules used in this paper can be found more than 60 years ago by influential market participants. For example, the ideas of trading ranges and resistance or support levels can be found in Wyckoff (1910). More “recent” references to these techniques can be found in Neill (1931) and Schabacker (1930). The use of moving averages was discussed by Gartley (1930). Further examples of the important early studies of these techniques have been carefully collected in Coslow (1966). The early use and popularity of these methods reduces the possibility that data-snooping biases are driving our results since we have at least 60 years of “fresh” data. Moreover, the early studies used just a few years in examining their trading strategies, thus contaminating little of the early data. A second issue here is the sensitivity of our results to the exact moving-average lengths used. Recent results in LeBaron (1990) for foreign exchange markets suggest that the results are not sensitive to the actual lengths of the rules used. We have replicated some of those results for the Dow index.

Journal, had occasionally published stock averages of various kinds, but not on a regular basis. No other index of U.S. stocks has been available for so long a period of time.⁷

The stocks included in the index have changed from time to time. Changes were more frequent in the earlier days. From the beginning, the list included large, well-known, and actively traded stocks. In recent years the 30 stocks in the index represent about 25 percent of the market value of all NYSE stocks. All the stocks are very actively traded and problems associated with nonsynchronous trading should be of little concern with the DJIA.

In addition to the full sample, results are presented for four subsamples: 1/1/97-7/30/14, 1/1/15-12/31/38, 1/1/39-6/30/62, and 7/1/62-12/31/86. These subsamples are chosen for several reasons. The first subsample ends with the closing of the stock exchange during World War I. The second subsample includes both the rise of the twenties and the turbulent times of the depression. The third subsample includes the period of World War II and ends in June 1962, the date at which the Center for Research in Securities Prices (CRSP) begins its daily price series. The last subsample covers the period that was extensively researched because of data availability.

B. Technical Trading Rules

Two of the simplest and most widely used technical rules are investigated: moving average-oscillator and trading range break-out (resistance and support levels). According to the moving average rule, buy and sell signals are generated by two moving averages of the level of the index—a long-period average and a short-period average. In its simplest form this strategy is expressed as buying (or selling) when the short-period moving average rises above (or falls below) the long-period moving average. The idea behind computing moving averages is to smooth out an otherwise volatile series. When the short-period moving average penetrates the long-period moving average, a trend is considered to be initiated. The most popular moving average rule is 1-200, where the short period is one day and the long period is 200 days. While numerous variations of this rule are used in practice, we attempted to select several of the most popular ones: 1-50, 1-150, 5-150, 1-200, and 2-200. The moving-average decision rule is often modified by introducing a band around the moving average. The introduction of a band reduces the number of buy (sell) signals by eliminating “whiplash” signals when the short and long period moving averages are close. We test the moving average rule both with and without a one percent band.

Our first rule, called the variable length moving average (VMA), initiates buy (sell) signals when the short moving average is above (below) the long

⁷ Pierce (1991) contains a brief history of the early Dow averages along with the series used in this study.

moving average by an amount larger than the band. If the short moving average is inside and band no signal is generated. This method attempts to simulate a strategy where traders go long as the short moving average moves above the long and short when it is below. With a band of zero this method classifies all days into either buys or sells. Other variations of this rule put emphasis on the crossing of the moving averages. They stress that returns should be different for a few days following a crossover. To capture this we test a strategy where a buy (sell) signal is generated when the short moving average cuts the long moving average from below (above). Returns during the next ten days are then recorded.⁸ Other signals occurring during this ten-day period are ignored. We call this rule a fixed-length moving average (FMA).

There are numerous variations of the moving average rule that we do not examine. We focus on the simplest and most popular versions. Other variants of the moving average rule also consider the slope of the long-period moving average in addition to whether the short-period moving average penetrated from above or below. In other versions changes in trading volume are examined before buy (sell) decisions are reached. Thus, numerous moving average rules can be designed, and some, without a doubt, will work. However, the dangers of data snooping are immense. We present results for all the rules examined and place emphasis on the robustness of the results over time.

Our final technical rule is trading range break-out (TRB). A buy signal is generated when the price penetrates the resistance level. The resistance level is defined as the local maximum. Technical analysts believe that many investors are willing to sell at the peak. This selling pressure will cause resistance to a price rise above the previous peak. However, if the price rises above the previous peak, it has broken through the resistance area. Such a breakout is considered to be a buy signal. Under this rule, a sell signal is generated when the price penetrates the support level which is the local minimum price. The underlying rationale is that the price has difficulties penetrating the support level because many investors are willing to buy at the minimum price. However, if the price goes below the support level, the price is expected to drift downward. In essence, technical analysts recommend buying when the price rises above its last peak and selling when the price sinks below its last trough.

To implement the trading range strategy, we defined rules in accordance with the moving average strategy. Maximum (or minimum) prices were determined based on the past 50, 150, and 200 days. In addition, the rule is implemented with and without a one percent band. As with the moving-average rule, numerous variations of the basic trading range strategy are being implemented in practice.

⁸ The selection of 10-day returns is arbitrary. For some rules we tried two-week returns and obtained essentially the same results.

II. Empirical Results: Traditional Tests

A. Sample Statistics

Table I contains summary statistics for the entire series and four subsamples for 1- and 10-day returns on the Dow Jones series. Returns are calculated as log differences of the Dow level. In Panel A the results for the daily returns are presented. These returns are strongly leptokurtic for the entire series and all the subsamples. All of the subperiods except one show some signs of skewness. Volatility is largest for the subperiod containing the Great Depression, and it appears to have declined in the most recent subperiods. Serial correlations are generally small with the exception of a few relatively

Table I
Summary Statistics for Daily and 10-Day Returns

Results are presented for the full sample and 4 nonoverlapping subperiods. Returns are measured as log differences of the level of the Dow index. 10-day returns are based on nonoverlapping 10-day periods. $\rho(i)$ is the estimated autocorrelation at lag i for each series. Numbers marked with * (**) are significant at the 5% (1%) levels for a two-tailed test. "Bartlett std." refers to the Bartlett standard error for the autocorrelation, $1/\sqrt{N}$.

Panel A: Daily Returns					
	Full Sample	97-14	15-38	39-62	62-86
<i>N</i>	25036	5255	7163	6442	6155
Mean	0.00017	0.00012	0.00014	0.00020	0.00020
Std.	0.0108	0.0099	0.0147	0.0075	0.0088
Skew	-0.1047**	-0.4804**	0.0193	-0.7614**	0.2707**
Kurtosis	16.00**	8.86**	12.75**	13.60**	11.57**
$\rho(1)$	0.033**	0.013	0.009	0.117**	0.079**
$\rho(2)$	-0.026**	-0.020	-0.029*	-0.068**	-0.001
$\rho(3)$	0.012*	0.041**	-0.006	0.036**	0.009
$\rho(4)$	0.046**	0.085**	0.055**	0.028*	-0.012
$\rho(5)$	0.022**	0.042**	0.027*	0.014	-0.011
Bartlett std.	0.006	0.014	0.012	0.012	0.013
Panel B: 10-Day Returns					
	Full Sample	97-14	15-38	39-62	62-86
Mean	0.0017	0.0012	0.0014	0.0019	0.0019
Std.	0.0351	0.0339	0.0486	0.0272	0.0296
Skew	-0.4583**	-0.1762	-0.9105**	-1.1551**	-0.0786
Kurtosis	7.91**	4.59**	8.51**	9.05**	3.91**
$\rho(1)$	0.037*	-0.004	0.065*	0.032	-0.002
$\rho(2)$	0.018	0.044	0.001	-0.090*	-0.041
$\rho(3)$	0.013	0.071	0.056	-0.037	0.007
$\rho(4)$	0.011	-0.125**	0.024	0.045	0.026
$\rho(5)$	0.032	0.094*	-0.022	0.018	-0.021
Bartlett std.	0.019	0.043	0.037	0.039	0.040

large values at the first lag in the two most recent subperiods. Panel B reports the values for 10-day nonoverlapping returns. A reduction in kurtosis for all subperiods is observed. Volatility shows a similar pattern to the daily returns with the largest value in the period containing the Great Depression, and the lowest values in the two most recent subperiods. For the 10-day returns the autocorrelations are generally quite small with no indication of an increase in the latter subperiods.

B. The Moving-Average Strategy

Results from trading strategies based on moving average rules for the full sample are presented in Panel A of Table II. The rules differ by the length of the short and long period and by the size of the band. For example (1, 200, 0) indicates that the short period is one day, the long period is 200 days, and the band is zero percent. We present results for the 10 rules that we examined. The moving-average rule is used to divide the entire sample into either buy or sell periods depending on the relative position of the moving averages. If the short moving average is above (below) the long, the day is classified as a buy (sell). This rule is designed to replicate returns from a trading rule where the trader buys when the short moving average penetrates the long from below and stays in the market until the short moving average penetrates the long moving average from above. After this signal the trader moves out of the market or sells short. We will refer to this rule as the "variable-length moving average" (VMA). In Table II we report daily returns during buy and sell periods and corresponding t -statistics.⁹

The results in Table II are striking. The last column lists the differences between the mean daily buy and sell returns. All the buy-sell differences are positive and the t -tests for these differences are highly significant, rejecting the null hypothesis of equality with zero. In every case the introduction of the one percent band increased the spread between the buy and sell returns. The first two columns in Panel A are the number of buy and sell signals generated. For each of the trading rules about 50 percent more buy signals are generated than sells, which is consistent with an upward-trending market.

⁹ The t -statistics for the buys (sells) are,

$$\frac{\mu_r - \mu}{(\sigma^2/N + \sigma^2/N_r)^{1/2}},$$

where μ_r and N_r are the mean return and number of signals for the buys and sells, and μ and N are the unconditional mean and number of observations. σ^2 is the estimated variance for the entire sample. For the buy-sell the t -statistic is,

$$\frac{\mu_b - \mu_s}{(\sigma^2/N_b + \sigma^2/N_s)^{1/2}},$$

where μ_b and N_b are the mean return and number of signals for the buys and μ_s and N_s are the mean return and number of signals for the sells.

Table II

Standard Test Results for the Variable-Length Moving (VMA) Rules

Results for daily data from 1897–1986. Rules are identified as (short, long, band) where short and long are the short and long moving averages respectively, and band is the percentage difference that is needed to generate a signal. “N(Buy)” and “N(Sell)” are the number of buy and sell signals reported during the sample. Numbers in parentheses are standard *t*-ratios testing the difference of the mean buy and mean sell from the unconditional 1-day mean, and buy-sell from zero. “Buy > 0” and “Sell > 0” are the fraction of buy and sell returns greater than zero. The last row reports averages across all 10 rules. Results for subperiods are given in Panel B.

Panel A: Full Sample								
Period	Test	N(Buy)	N(Sell)	Buy	Sell	Buy > 0	Sell > 0	Buy-Sell
1897–1986	(1, 50, 0)	14240	10531	0.00047 (2.68473)	-0.00027 (-3.54645)	0.5387	0.4972	0.00075 (5.39746)
	(1, 50, 0.01)	11671	8114	0.00062 (3.73161)	-0.00032 (-3.56230)	0.5428	0.4942	0.00094 (6.04189)
	(1, 150, 0)	14866	9806	0.00040 (2.04927)	-0.00022 (-3.01836)	0.5373	0.4962	0.00062 (4.39500)
	(1, 150, 0.01)	13556	8534	0.00042 (2.20929)	-0.00027 (-3.28154)	0.5402	0.4943	0.00070 (4.68162)
	(5, 150, 0)	14858	9814	0.00037 (1.74706)	-0.00017 (-2.61793)	0.5368	0.4970	0.00053 (3.78784)
	(5, 150, 0.01)	13491	8523	0.00040 (1.97876)	-0.00021 (-2.78835)	0.5382	0.4942	0.00061 (4.05457)
	(1, 200, 0)	15182	9440	0.00039 (1.93865)	-0.00024 (-3.12526)	0.5358	0.4962	0.00062 (4.40125)
	(1, 200, 0.01)	14105	8450	0.00040 (2.01907)	-0.00030 (-3.48278)	0.5384	0.4924	0.00070 (4.73045)
	(2, 200, 0)	15194	9428	0.00038 (1.87057)	-0.00023 (-3.03587)	0.5351	0.4971	0.00060 (4.26535)
	(2, 200, 0.01)	14090	8442	0.00038 (1.81771)	-0.00024 (-3.03843)	0.5368	0.4949	0.00062 (4.16935)
Average				0.00042	-0.00025			0.00067
Panel B: Subperiods								
1897–1914	(1, 150, 0)	2925	2170	0.00039 (1.19348)	-0.00025 (-1.48213)	0.5323	0.4959	0.00065 (2.30664)
1915–1938	(1, 150, 0)	4092	2884	0.00048 (1.16041)	-0.00045 (-1.82639)	0.5503	0.4941	0.00092 (2.59189)
1939–1962	(1, 150, 0)	4170	2122	0.00036 (1.06310)	-0.00004 (-1.26932)	0.5422	0.5151	0.00040 (1.98384)
1962–1986	(1, 150, 0)	3581	2424	0.00037 (0.94029)	-0.00012 (-1.49333)	0.5205	0.4777	0.00049 (2.11283)

The mean buy and sell returns are reported separately in columns 3 and 4. The buy returns are all positive with an average one-day return of 0.042 percent, which is about 12 percent at an annual rate. This compares with the unconditional one-day return of 0.017 percent from Table I. Six of the ten tests reject the null hypothesis that the returns equal the unconditional returns at the 5 percent significance level using a two-tailed test. The other

four tests are marginally significant. For the sells, the results are even stronger. All the sell returns are negative with an average daily return for the 10 tests of -0.025 percent which is about -7 percent at an annual rate. The t -tests for equality with the unconditional mean return are all highly significant, with all the t -values being less than -2.5 .

The fifth and sixth columns in Table II present the fraction of buys and sells greater than zero. For buys, this fraction ranges from 53 to 54 percent and, for sells, it is about 49 percent. Under the null hypothesis that technical rules do not produce useful signals the fraction of positive returns should be the same for both buys and sells. Performing a binomial test shows that all these differences are highly significant and the null hypothesis of equality can be rejected.

The negative returns in Table II for sell signals are especially noteworthy. These returns cannot be explained by various seasonalities since they are based on about 40 percent of all trading days.¹⁰ Many previous studies found as we did that returns are predictable. This predictability can reflect either: (1) changes in expected returns that result from an equilibrium model, or (2) market inefficiency. In general, it is difficult to distinguish between these two alternative explanations. Although rational changes in expected returns are possible it is hard to imagine an equilibrium model that predicts negative returns over such a large fraction of trading days.

In Panel B of Table II we repeat the results for several subperiods. To save space, results are presented for only the (1, 150, 0) rule. We found no evidence that the results are different across the subperiods.

The second moving average test, the fixed-length moving average (FMA) rule, examines fixed 10-day holding periods after a crossing of the two moving averages. Results are presented in Table III. For all the tests the buy-sell differences are positive. The average difference without a band is 0.77 percent while the average with a one percent band is 1.09 percent. These are quite substantial returns given that the unconditional 10-day return from Table I is only 0.17 percent. For 7 of these 10 tests the null hypothesis that the difference is equal to zero can be rejected at the 5 percent level. The remaining 3 tests are marginally significant. As before, in all cases the addition of a one percent band to the trading rule increases the buy-sell difference.¹¹ The table also reports the buy and sell returns separately. For all of the buys the returns are greater than the unconditional mean 10-day return with an average of 0.53 percent. The sells are all negative and fall below the unconditional mean 10-day return with an average of -0.40

¹⁰ The series used in this study does not contain dividends. Results in Lakonishok and Smidt (1988) suggest that our finding of negative returns during sell periods will not be altered with the inclusion of dividends. In addition, in the context of our trading rules it is very unlikely that the pattern of dividend payments would differ substantially between buy and sell periods.

¹¹ There is an unusual effect of the band on the number of buy and sell signals. For several of the rules an addition of the band actually increases the number of buy signals. In our testing method signals which occur within 10 days of a previous signal are blocked out, so the elimination of a few sell signals may actually increase the number of buys.

Table III

Standard Test Results for the Fixed-Length Moving (FMA) Rules

Results for daily data from 1897–1986. Cumulative returns are reported for fixed 10-day periods after signals. Rules are identified as (short, long, band) where short and long are the short and long moving averages respectively, and band is the percentage difference that is needed to generate a signal. “ $N(\text{Buy})$ ” and “ $N(\text{Sell})$ ” are the number of buy and sell signals reported during the sample. Numbers in parentheses are standard t -ratios testing the difference of the mean buy and mean sell from the unconditional 1-day mean, and buy-sell from zero. “Buy > 0” and “Sell > 0” are the fraction of buy and sell returns greater than zero. The last row reports averages across all 10 rules.

Test	$N(\text{Buy})$	$N(\text{Sell})$	Buy	Sell	Buy > 0	Sell > 0	Buy-Sell
(1, 50, 0)	340	344	0.0029 (0.5796)	-0.0044 (-3.0021)	0.5882	0.4622	0.0072 (2.6955)
(1, 50, 0.01)	313	316	0.0052 (1.6809)	-0.0046 (-3.0096)	0.6230	0.4589	0.0098 (3.5168)
(1, 150, 0)	157	188	0.0066 (1.7090)	-0.0013 (-1.1127)	0.5987	0.5691	0.0079 (2.0789)
(1, 150, 0.01)	170	161	0.0071 (1.9321)	-0.0039 (-1.9759)	0.6529	0.5528	0.0110 (2.8534)
(5, 150, 0)	133	140	0.0074 (1.8397)	-0.0006 (-0.7466)	0.6241	0.5786	0.0080 (1.8875)
(5, 150, 0.01)	127	125	0.0062 (1.4151)	-0.0033 (-1.5536)	0.6614	0.5520	0.0095 (2.1518)
(1, 200, 0)	114	156	0.0050 (0.9862)	-0.0019 (-1.2316)	0.6228	0.5513	0.0069 (1.5913)
(1, 200, 0.01)	130	127	0.0058 (1.2855)	-0.0077 (-2.9452)	0.6385	0.4724	0.0135 (3.0740)
(2, 200, 0)	109	140	0.0050 (0.9690)	-0.0035 (-1.7164)	0.6330	0.5500	0.0086 (1.9092)
(2, 200, 0.01)	117	116	0.0018 (0.0377)	-0.0088 (-3.1449)	0.5556	0.4397	0.0106 (2.3069)
Average			0.0053	-0.0040			0.0093

percent. For all the tests the fraction of buys greater than zero exceeds the fraction of sells greater than zero.

The profits that can be derived from these trading rules depend, among other things, on the number of signals generated. The lowest number of signals is for the (2, 200, 0.01) rule which generates an average of 2.8 signals per year over the 90 years of data. The largest number of signals is generated by the (1, 50, 0) rule with 7.6 signals per year. We explore the following strategy: upon a buy signal, we borrow and double the investment in the Dow Index; upon a sell signal, we sell shares and invest in a risk-free asset. Given that the number of buy and sell signals is similar we make the following assumptions: (1) the borrowing and lending rates are the same, and (2) the risk during buy periods is the same as the risk during sell periods. Under these assumptions such a strategy, ignoring transaction costs, should produce the same return as a buy and hold strategy. Using the (1, 50, 0.01) rule as an example, there are on average about 3.5 buy and sell signals per year. On the

buy side, because of leverage, we gain on average, 1.8 percent (3.5×0.0052). On the sell side, by not being in the market, we gain 1.6 percent (3.5×0.0046). This results in an extra return of 3.4 percent, before transactions costs, which is substantial when compared to the 5 percent annual return on the Dow Index (excluding dividends).

C. Trading Range Break

Results for the trading range break rule are presented in Table IV. With this rule buy and sell signals are generated when the price level moves above or below local maximums and minimums. Local maximums and minimums are computed over the preceding 50, 150, and 200 days. We also use a band technique where the price level must exceed the local maximum by one percent, or fall below the minimum by one percent. For the trading range break rule we compute 10-day holding period returns following buy and sell signals.

The results are presented in the same format as Table III. The average buy-sell return is 0.86 percent. Of the six tests, all reject the null hypothesis of the buy-sell difference being equal to zero. The buy return is positive across all the rules with an average of 0.55 percent. For 3 out of the 6 rules, the buy returns are significantly different from the unconditional 10-day return at the 5 percent level, and the remaining 3 rules are marginally significant. One

Table IV

Standard Test Results for the Trading Range Break (TRB) Rules

Results for daily data from 1897-1986. Cumulative returns are reported for fixed 10-day periods after signals. Rules are identified as (short, long, band) where short and long are the short and long moving averages respectively, and band is the percentage difference that is needed to generate a signal. "N(Buy)" and "N(Sell)" are the number of buy and sell signals reported during the sample. Numbers in parentheses are standard *t*-ratios testing the difference of the mean buy and mean sell from the unconditional 1-day mean, and buy-sell from zero. "Buy > 0" and "Sell > 0" are the fraction of buy and sell returns greater than zero. The last row reports averages across all 6 rules.

Test	N(Buy)	N(Sell)	Buy	Sell	Buy > 0	Sell > 0	Buy-Sell
(1, 50, 0)	722	415	0.0050 (2.1931)	0.0000 (-0.9020)	0.5803	0.5422	0.0049 (2.2801)
(1, 50, 0.01)	248	252	0.0082 (2.7853)	-0.0008 (-1.0937)	0.6290	0.5397	0.0090 (2.8812)
(1, 150, 0)	512	214	0.0046 (1.7221)	-0.0030 (-1.8814)	0.5762	0.4953	0.0076 (2.6723)
(1, 150, 0.01)	159	142	0.0086 (2.4023)	-0.0035 (-1.7015)	0.6478	0.4789	0.0120 (2.9728)
(1, 200, 0)	466	182	0.0043 (1.4959)	-0.0023 (-1.4912)	0.5794	0.5000	0.0067 (2.1732)
(1, 200, 0.01)	146	124	0.0072 (1.8551)	-0.0047 (-1.9795)	0.6164	0.4677	0.0119 (2.7846)
Average			0.0063	-0.0024			0.0087

possible reason for this relatively strong rejection when compared to the moving average rules is the fact that this rule generates more buy and sell signals. The sell returns are negative across all the rules with an average of -0.24 percent. For the individual rules, 1 out of the 6 are significantly different from the unconditional 10-day return. Results for the subperiods are similar and are not presented to save space.¹²

III. Bootstrap Methodology

The results of Section II are intriguing, but there are still some missing pieces that the bootstrap methodology can help to solve. First, we do not compute a comprehensive test across all rules. Such a test would have to take into account the dependencies between results for different rules. We now develop a joint test of significance for our set of trading rules. This is accomplished by utilizing bootstrap distributions for these tests. This is a major advantage of the bootstrap methodology, since constructing such tests using traditional statistical methods would require properly accounting for the complex dependencies across the different rules, which is an extremely difficult task.¹³

Second, the t -ratios reported earlier assume normal, stationary, and time-independent distributions. For stock returns there are several well-known deviations from this assumed distribution. Among these are: leptokurtosis, autocorrelation, conditional heteroskedasticity, and changing conditional means. These important aspects of the data will be addressed using distributions generated from simulated null models for stock prices.¹⁴ Using this strategy we can address all the issues brought up earlier. A third benefit of this methodology is that we can examine the standard deviations of returns during the buy and sell periods. This gives us an indication of the riskiness of the various strategies during buy and sell periods. Overall, these results can shed some light on whether predictability of stock returns is a result of market inefficiency or time varying equilibrium returns. Complete resolution of this debate is not possible, however.

¹² Our strong results in support of technical analysis differ from some of the earlier-mentioned studies. This might be because we utilized a much longer period and hence were in a position to obtain much stronger rejections of the null hypothesis. Moreover, we focused on the Dow Jones Index, whereas some of the previous studies looked at individual securities which could be another explanation for the observed differences.

¹³ While there are cases where it is possible to determine the distribution of statistics based on sums of correlated random variables, this in general is not a simple analytical problem. Moreover, in this case the distribution of these random variables is unknown.

¹⁴ While our hypothesis testing using parametric bootstraps is able to account for some forms of heteroskedasticity, it is limited to the specific functional forms used for the volatility process (i.e., the GARCH-M and EGARCH models). More general adjustments for heteroskedasticity would certainly be useful, but it is not clear how they could be implemented in the technical trading methodology of this paper. Finally, the parametric bootstrap gives us a specification test of these commonly used null models. More general nonparametric adjustments for heteroskedasticity may be useful, but are still at the experimental stage and beyond the scope of this paper.

The bootstrap methodology is described in detail in the Appendix. We provide an informal description here. The returns conditional on buy (sell) signals using the raw Dow Jones data are compared to conditional returns from simulated comparison series. The trading rules classify each day in our sample as either buy (b), sell (s), or neutral (n). Where the classification of day t is based on information up to and including day t . Define the h day return at t as

$$r_t^h = \log(p_{t+h}) - \log(p_t).$$

We will be interested in various conditional expectations based on the trading rule signals. For example, let

$$m_b = E(r_t^h | b_t)$$

be the expected h day return from t to $t + h$ conditional on a buy signal at time t , and

$$m_s = E(r_t^h | s_t)$$

be the expected h day return from t to $t + h$ conditional on a sell signal at time t . We also look at conditional standard deviations,

$$\left(E\left[(r_t^h - m_b)^2 | b_t \right] \right)^{1/2}, \left(E\left[(r_t^h - m_s)^2 | s_t \right] \right)^{1/2}$$

These conditional expectations will be estimated using the sample means. The values from the Dow series will then be compared with empirical distributions from the simulated null models for stock returns. Our methodology can be used not only to assess the profitability of various trading strategies, but also as a specification test for alternative models.

The distributions of the conditional moments under various null hypotheses for stock return movements will be estimated using the bootstrap methodology inspired by Efron (1982), Freedman (1984), Freedman and Peters (1984a, 1984b), and Efron and Tibshirani (1986). Computer simulations of time series designed to capture the properties of the various null models are performed using the estimation-based bootstrap of Freedman and Peters (1984a, 1984b). In this procedure each model is fit to the original series to obtain estimated parameters and residuals. We standardize the residuals using estimated standard deviations for the error process. The estimated residuals are then redrawn with replacement to form a scrambled residuals series which is then used with the estimated parameters to form a new representative series for the given null model.¹⁵ The standardized residuals

¹⁵ Freedman (1984) has adduced theoretical arguments as well as simulation evidence that the estimation-based bootstrap gives good estimates of standard errors for a class of linear models driven by error processes with unknown variance matrices that must be estimated from the data. Unfortunately, there is little published work on time series processes driven by heteroskedastic innovations beyond a rather brief treatment in Freedman (1984). Two examples dealing with heteroskedasticity in a cross-sectional context are Beran (1986) and Hardle (1990). By using the estimation-based bootstrap in the context of the GARCH-M and Exponential GARCH we are clearly taking the procedure beyond what has been proved in the bootstrap literature. In the absence of theoretical proof of convergence we present "computer proofs" of convergence in Figures 1 and 2 which will be explained in the next section.

are not restricted to a particular distribution, such as Gaussian, by this procedure.¹⁶

Each of the simulations is based on 500 replications of the null model.¹⁷ This should provide a good approximation of the return distribution under the null model. The null hypothesis is rejected at the α percent level if returns obtained from the actual Dow Jones data are greater than the α percent cutoff of the simulated returns under the null model. The methodology described in this section combines tests based on technical trading rules with bootstrap techniques for generating distributions of statistics under null models.

In this study representative price series are simulated from the following widely used processes for stock prices: a random walk with a drift, autoregressive process of order one (AR(1)), generalized autoregressive conditional heteroskedasticity in-mean model (GARCH-M), and Exponential GARCH (EGARCH). The random walk with drift series was simulated by taking the returns from the original Dow Jones series and “scrambling” them. The term “scrambling” refers to the formal bootstrap sampling process described in the Appendix. The “scrambling” procedure forms a new time series of returns by randomly drawing from the original series with replacement. This scrambled series will have the same drift in prices, the same volatility, and the same unconditional distribution. However, by construction the returns are independent, identically distributed.

The second model for the simulation is an AR(1),

$$r_t = b + \rho r_{t-1} + \varepsilon_t, \quad |\rho| < 1$$

where r_t is the return on day t and ε_t is independent, identically distributed. The parameters (b, ρ) and the residuals e_t are estimated from the Dow Jones series using OLS. The residuals are then resampled with replacement and AR(1)'s are then generated using the estimated parameters and scrambled residuals. To the extent that returns over short periods of time are positively correlated, the technical strategies might produce “abnormal” returns. Conrad and Kaul (1990) report first-order autocorrelation of 0.20 for a value-weighted portfolio of the largest companies during the period 1962–1985. They find that higher order autocorrelation, beyond a lag of one day, is essentially zero. Our data also reveal the presence of significant autocorrelation in some of the subsamples. Hence, in the presence of positive

¹⁶ While the bootstrap procedure allows deviations from the Gaussian distribution in the residuals series, estimation is performed for the nonlinear models using a Gaussian likelihood. For certain cases Bollerslev and Wooldridge (1990) have shown that estimators based on Gaussian likelihood will still be consistent under other error distributions.

¹⁷ The tests in this paper will present simulated p -values. The number of repetitions required for the estimation of p -values and confidence intervals using bootstrap techniques is quite large. Using 500 repetitions of estimates of objects such as $P(X > c)$ where X is a random variable and c is a constant will have a maximum standard error of $\sqrt{(0.5^2/500)} = 0.022$. While this number gives an upper bound on the precision of our estimates it probably greatly exceeds our actual standard errors. In Section IV experiments are performed to test the reliability of our estimates by extending the bootstrap replications to 2000.

autocorrelation, the “abnormal” returns from our trading strategy might be a result of an autoregressive process that generates stock returns.

The third model for simulation is a GARCH-M model:

$$\begin{aligned} r_t &= a + \gamma h_t + b\varepsilon_{t-1} + \varepsilon_t \\ h_t &= \alpha_0 + \alpha_1 e_{t-1}^2 + \beta h_{t-1} \\ \varepsilon_t &= h_t^{1/2} z_t \quad z_t \sim N(0, 1). \end{aligned}$$

In this model, the error, ε_t , is conditionally normally distributed and serially uncorrelated, and the conditional variance, h_t , is a linear function of the square of the last period's errors and of the last period's conditional variance. This specification of the conditional variance implies positive serial correlation in the conditional second moment of the return process—periods of high (or low) volatility are likely to be followed by periods of high (low) volatility. The conditional returns in this model are a linear function of the conditional variance (see Engle, Lilien, and Robins (1987)) and the past disturbance, e_{t-1} . Under this return-generating process, volatility can change over time and the expected returns are a function of volatility as well as past returns. This is a rich specification and is popular in financial economics literature.¹⁸ The parameters and standardized residuals are estimated from the Dow Jones series using maximum likelihood. Once again the residuals are resampled with replacement and used along with the estimated parameters to generate GARCH-M series. The GARCH-M might be consistent with the efficient market hypothesis; higher ex ante returns are expected when conditional volatility is high under the GARCH-M specification.

The fourth model that will be estimated and simulated is the Exponential GARCH (EGARCH) model developed by Nelson (1991):¹⁹

$$\begin{aligned} r_t &= a + \gamma e^{h_t} + b\varepsilon_{t-1} + \varepsilon_t \\ h_t &= \alpha_0 + g(z_{t-1}) + \beta h_{t-1} \\ g(z_t) &= \theta z_t + \omega (|z_t| - (2/\pi)^{1/2}) \\ \varepsilon_t &= e^{(1/2)h_t} z_t \quad z_t \sim N(0, 1) \end{aligned}$$

While this model also tries to capture persistent volatility as does the GARCH model, it differs in two important ways. First, the log of the conditional variance now follows an autoregressive process. Second, it allows previous returns to affect future volatility differently depending on their signs. This is clearly seen in the g function above. This is designed to capture a phenomenon in asset returns observed by Black (1976) where negative returns were generally followed by larger volatility than positive returns.

¹⁸ The specification used in this paper, the GARCH(1, 1)-M, was determined using the Schwarz (1978) model specification criterion. It is similar to the model specification used in French, Schwert, and Stambaugh (1987).

¹⁹ Nelson (1991) replaces the normal distribution used here with the generalized error distribution.

Once again the model is estimated and the standardized residuals and estimated parameters are used to generate simulated Exponential GARCH series.

Table V contains estimation results for the three models which will be used for comparison with the actual Dow series, the AR(1), the GARCH-M, and the EGARCH. Panel A presents the results from estimation of an AR(1) using OLS. The results reveal a significant first order autocorrelation for the Dow series. Panel B shows the results from estimation of a GARCH-M using maximum likelihood. The model estimated also contains an MA(1) to account for the short horizon autocorrelations. The results are consistent with previous studies, such as French, Schwert, and Stambaugh (1987) who utilized the Standard and Poor's index over the time period 1928–84. The estimates of

Table V
Parameter estimates for AR(1), GARCH-M, and Exponential GARCH models

Estimated on daily returns series 1897–1986. The AR(1) is estimated by OLS. The GARCH-M and Exponential GARCH are estimated using maximum likelihood. r_t is the continuously compounded return on day t and h_t is the conditional variance on day t . The numbers in parenthesis are t -ratios.

Panel A: AR(1) Parameter Estimates						
$r_t = a + br_{t-1} + \varepsilon_t$						
	a		b			
	0.00015		0.03330			
	(2.50)		(5.27)			
Panel B: GARCH-M Parameter Estimates						
$r_t = a + \gamma h_t + b\varepsilon_{t-1} + \varepsilon_t \quad \varepsilon_t = h_t^{1/2}z_t$						
$h_t = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta h_{t-1}$						
$z_t \sim N(0, 1)$						
	α_0	α_1	β	γ	b	a
	1.32e-6	0.094	0.896	2.09	0.09	2.24e-4
	(25.4)	(51.8)	(522)	(2.80)	(14.2)	(3.47)
Panel C: Exponential GARCH Parameter Estimates						
$r_t = a + \gamma e^{h_t} + b\varepsilon_{t-1} + \varepsilon_t \quad \varepsilon_t = e^{(1/2)h_t}z_t$						
$h_t = \alpha_0 + g(z_{t-1}) + \beta h_{t-1}$						
$g(z_t) = \theta z_t + \omega(z_t - (2/\pi)^{1/2})$						
$z_t \sim N(0, 1)$						
	α_0	β	γ	b	θ	ω
	-0.1867	0.9795	2.4457	0.1024	-0.0692	0.1758
	(-23.0)	(1224)	(3.11)	(16.2)	(-34.6)	(53.3)
						$a * 10^4$
						-0.0550
						(-0.080)

the GARCH-M model indicate that the conditional variance of stock returns is time varying and is autocorrelated. The conditional variance in $t - 1$, h_{t-1} , and the shock in $t - 1$, ε_{t-1}^2 , are highly significant in estimating the conditional volatility at time t , h_t . The estimated parameters also show a significantly positive relation between conditional variance and conditional mean, γ of 2.09. This is in close agreement with French, Schwert, and Stambaugh, who estimated a γ of 2.41 for their entire sample. The b parameter, capturing the first order autocorrelation in the series, is also significantly positive. Panel C presents the results for the exponential GARCH. They display a strong persistence in variance, a significant MA(1) term, and a positive relation between conditional variance and conditional mean as did the GARCH. The important difference from the exponential GARCH is the estimated parameter, θ , which is significantly negative. This indicates the inverse relation between returns today and future volatility mentioned previously.

IV. Empirical Results: Bootstrap Tests

A. Random Walk Process

In Table VI we display the results of random walk simulations. To save space we present only a subset of the rules used. In order to describe the format of Table VI we start with the first row of Panel A which presents the results for the (1, 50, 0) VMA rule. All the numbers presented in this row are the fractions of the simulated result which are larger than the results for the original Dow series. Results for returns are presented in the columns labeled, Buy, Sell, and Buy-Sell, while the results for standard deviations are presented in the columns labeled σ_b and σ_s . The number in the column labeled Buy, which is 0.00, shows that none of the simulated random walks generated a mean buy return as large as that from the original Dow series. This number can be thought of as a simulated “ p -value.” Turning to the Sell column the fraction is 1.00, showing that all of the simulated random walks generated mean sell returns larger than the mean sell return from the Dow series. The fraction in the Buy-Sell column, 0.00, reports that none of the simulated random walks generated mean buy-sell differences larger than the mean differences for the Dow series. In the column σ_b (or σ_s) the reported numbers are 1.00 (or 0.00) showing that all (or none) of the standard deviations for the simulated random walks were greater (or smaller) than those from the Dow series. The results for the returns are consistent with the traditional tests presented earlier. However, the results for the standard deviations are new. Not only do the buy signals select out periods with higher conditional means, they also pick periods with lower volatilities. This is in contrast to sell periods where the conditional return is lower, and the volatility is higher. An often-used explanation for predictability in returns is changing risk levels. Our results are not in accord with this explanation. Aside from the negative returns during sell periods being inconsistent with existing

equilibrium models, the higher returns for buys do not seem to arise during riskier periods. Similar results are obtained for the other VMA rules.

In Panel B of Table VI the results are summarized across all the rules. We construct a statistic to jointly test our set of trading rules. The bootstrap technique allows us to choose any function to aggregate results across the various rules. We decided to use a simple average over the ten rules. Thus, for each of our 500 simulations we compute an average over all the rules for both returns and standard deviations. The first row of Panel B, labeled Fraction > Dow, follows the same format as the results presented in Panel A. Not surprisingly, these results strongly agree with those for the individual rules. The second row, labeled Mean, presents the returns and standard deviations for the buys, sells, and buy-sells, averaged over the 500 simulated random walks. The returns and standard deviations for buys and sells are essentially the same and close to their unconditional values reported in Table I. The third row, labeled Dow, presents the same statistics for the original Dow series. For both returns and standard deviations large differences not seen in the bootstrap results are observed. The average difference between the buys and sells is 0.067 percent as presented previously in Table II. The standard deviations for buys and sells are 0.89 and 1.34 percent, respectively, indicating that the market is less volatile during buy periods relative to sell periods. The results for the FMA and TRB rules are also presented in Table VI. The results, although somewhat weaker, are similar to results from the VMA rule and therefore will not be discussed.

Most of the results in this table are in accord with the findings using traditional methods (Tables II–IV), although some differences can be observed. For example, focusing on sell returns for the TRB (1,200,0) rule, Table VI indicates that only 3.4 percent of the random walk simulations generated a value as low as that in the original series. The entry in Table IV, based on traditional tests, is a t -statistic of -1.49 . The probability for a standard normal being less than this value is about 7.7 percent. This suggests that the distributional assumptions of the standard tests may have an impact on statistical inferences.

The sensitivity of our inferences to the choice of 500 replications is examined in Figure 1. As an example we use the rule that generates the largest estimated p -value in Table VI. This is the (1,200,0) FMA rule. The estimated p -values for the buy, sell, and buy-sell returns using 500 replications are, respectively, 0.11, 0.91, 0.02. Figure 1 presents estimates of these values for replication sizes varying from 100 to 2000. To make the graph easier to read the sells are presented as $1 - (\text{Fraction} > \text{Dow})$. The figure shows that extending the replications beyond 500 adds little to the reliability of the estimated p -values.

B. AR(1) Process

Table VII repeats the previous results for a simulated AR(1) process by utilizing the estimated residuals from the original series. This experiment is

Table VI
Simulation Tests From Random Walk Bootstraps for
500 Replications

The log difference series is resampled with replacement and exponentiated back to a simulated price series. The rows marked "Fraction > Dow" refer to the fraction of simulations generating a mean or standard deviation larger than those from the actual Dow series. The trading rules are, variable-length moving average (VMA), fixed-length moving average (FMA), trading range break (TRB). Panel B presents results for the averages across all the rules for each reported statistic. "Dow" refers to the actual mean return or standard deviation from the Dow, and "Mean" refers to the mean from the simulated series.

Panel A: Individual Rules							
Rule		Result	Buy	σ_b	Sell	σ_s	Buy-Sell
(1, 50, 0)	VMA	Fraction > Dow	0.00000	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.25000	0.17200	1.00000	0.13000	0.00000
	TRB	Fraction > Dow	0.01000	1.00000	0.85800	0.00000	0.01000
(1, 50, 0.01)	VMA	Fraction > Dow	0.00000	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.02400	0.59800	1.00000	0.01000	0.00000
	TRB	Fraction > Dow	0.00000	0.09800	0.86800	0.00000	0.00000
(1, 150, 0)	VMA	Fraction > Dow	0.00400	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.02800	0.69600	0.88400	0.00000	0.00400
	TRB	Fraction > Dow	0.03000	1.00000	0.99000	0.00000	0.00000
(1, 150, 0.01)	VMA	Fraction > Dow	0.00200	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.00800	0.98200	0.99400	0.00000	0.00000
	TRB	Fraction > Dow	0.00000	0.94600	0.96600	0.00000	0.00000
(1, 200, 0)	VMA	Fraction > Dow	0.01000	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.11200	0.83800	0.91200	0.08400	0.03000
	TRB	Fraction > Dow	0.06000	1.00000	0.96600	0.00000	0.00600
(1, 200, 0.01)	VMA	Fraction > Dow	0.00600	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.04600	0.99800	1.00000	0.00000	0.00000
	TRB	Fraction > Dow	0.01200	0.99200	0.97800	0.00000	0.00200
Panel B: Rule Averages							
Rule average	VMA	Fraction > Dow	0.00200	1.00000	1.00000	0.00000	0.00000
		Mean	0.00016	0.01078	0.00017	0.01077	0.00000
		Dow	0.00042	0.00890	-0.00025	0.01342	0.00067
Rule average	FMA	Fraction > Dow	0.01000	0.99400	0.99800	0.00000	0.00000
		Mean	0.00167	0.03391	0.00150	0.03402	0.00016
		Dow	0.00531	0.03064	-0.00399	0.04156	0.00930
Rule average	TRB	Fraction > Dow	0.00400	0.99800	0.98800	0.00000	0.00000
		Mean	0.00169	0.03394	0.00160	0.03400	0.00009
		Dow	0.00633	0.03000	-0.00238	0.05398	0.00871




Figure 1.

designed to detect whether the results from the trading rules could be caused by daily serial correlations in the series. For all our trading rules the return on a day in which a buy (or sell) signal is received is expected to be large (or small). If the returns are positively autocorrelated we should also expect higher (or lower) returns on the following days. Indeed, the results reported in Table V document some degree of positive autocorrelation.

Table VII, Panel B, confirms that some differences between buys and sells do occur with an AR(1) process. For the VMA rules the average buy return from the simulated AR(1) is 0.019 percent, and the average sell return is 0.014 percent. This compares with an unconditional return of 0.017 percent for the entire sample. The AR(1) clearly creates a buy-sell spread as predicted, but the magnitude of this spread is not large when compared with the Dow series which produces a spread of 0.067 percent. The reported “*p*-value” of zero confirms that this difference cannot be explained by the AR(1) process. For the FMA and TRB rules the spreads produced by the AR(1) process are 0.10 and 0.12 percent, respectively. These spreads should be compared to the much larger spreads from the original Dow series of 0.93 and 0.87 percent.

C. GARCH-M Process

The next simulations use a GARCH-M process. In this model both conditional means and variances are allowed to change over time. A changing conditional mean can potentially explain some of the differences between buy and sell returns.

Checking the Buy-Sell column in Table VIII, Panel B, the VMA rule shows

Table VII
Simulation Tests From AR(1) Bootstraps for 500 Replications

The AR(1) residual series is resampled with replacement and simulated using the AR(1) estimated parameters. The rows marked "Fraction > Dow" refer to the fraction of simulations generating a mean or standard deviation larger than those from the actual Dow series. The trading rules are, variable-length moving average (VMA), fixed-length moving average (FMA), trading range break (TRB). Panel B presents results for the averages across all the rules for each reported statistic. "Dow" refers to the actual mean return or standard deviation from the Dow, and "Mean" refers to the mean from the simulated series.

Panel A: Individual Rules							
Rule		Result	Buy	σ_b	Sell	σ_s	Buy-Sell
(1, 50, 0)	VMA	Fraction > Dow	0.00400	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.35800	0.37000	1.00000	0.24600	0.00800
	TRB	Fraction > Dow	0.02400	1.00000	0.75200	0.00000	0.02600
(1, 50, 0.01)	VMA	Fraction > Dow	0.00000	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.04800	0.81400	0.99800	0.03600	0.00400
	TRB	Fraction > Dow	0.00000	0.27000	0.77000	0.00000	0.00000
(1, 150, 0)	VMA	Fraction > Dow	0.01400	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.05200	0.84000	0.81400	0.00000	0.03200
	TRB	Fraction > Dow	0.08200	1.00000	0.97800	0.00000	0.00600
(1, 150, 0.01)	VMA	Fraction > Dow	0.00800	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.03000	0.99400	0.97400	0.00000	0.00200
	TRB	Fraction > Dow	0.00400	0.99200	0.94200	0.00000	0.00400
(1, 200, 0)	VMA	Fraction > Dow	0.01400	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.17400	0.93800	0.85600	0.15200	0.05800
	TRB	Fraction > Dow	0.10200	1.00000	0.94400	0.00000	0.02200
(1, 200, 0.01)	VMA	Fraction > Dow	0.01200	1.00000	1.00000	0.00000	0.00000
	FMA	Fraction > Dow	0.11000	1.00000	0.99800	0.00000	0.00000
	TRB	Fraction > Dow	0.02200	0.99800	0.96600	0.00000	0.00600
Panel B: Rule Averages							
Rule average	VMA	Fraction > Dow	0.00600	1.00000	1.00000	0.00000	0.00000
		Mean	0.00019	0.01078	0.00014	0.01077	0.00006
		Dow	0.00042	0.00890	-0.00025	0.01342	0.00067
Rule average	FMA	Fraction > Dow	0.04000	0.99800	0.99800	0.00000	0.00000
		Mean	0.00216	0.03485	0.00113	0.03495	0.00104
		Dow	0.00531	0.03064	-0.00399	0.04156	0.00930
Rule average	TRB	Fraction > Dow	0.00600	1.00000	0.96000	0.00000	0.00200
		Mean	0.00223	0.03498	0.00106	0.03485	0.00118
		Dow	0.00633	0.03000	-0.00238	0.05398	0.00871

Table VIII

Simulation Tests From GARCH-M Bootstraps for 500 Replications

GARCH-M returns series are simulated using the estimated parameters and standardized residuals. The rows marked “Fraction > Dow” refer to the fraction of simulations generating a mean or standard deviation larger than those from the actual Dow series. The trading rules are, variable-length moving average (VMA), fixed-length moving average (FMA), trading range break (TRB). Panel B presents results for the averages across all the rules for each reported statistic. “Dow” refers to the actual mean return or standard deviation from the Dow, and “Mean” refers to the mean from the simulated series.

Panel A: Individual Rules							
Rule		Result	Buy	σ_b	Sell	σ_s	Buy-Sell
(1, 50, 0)	VMA	Fraction > Dow	0.11600	0.99400	0.99800	0.14200	0.01600
	FMA	Fraction > Dow	0.58000	0.76400	0.99800	0.67600	0.07800
	TRB	Fraction > Dow	0.12800	0.98600	0.73600	0.26400	0.14600
(1, 50, 0.01)	VMA	Fraction > Dow	0.05600	0.99400	0.99200	0.14400	0.01400
	FMA	Fraction > Dow	0.19600	0.88600	0.99800	0.65600	0.00400
	TRB	Fraction > Dow	0.14200	0.96800	0.72600	0.45400	0.14400
(1, 150, 0)	VMA	Fraction > Dow	0.10600	0.98200	0.99600	0.15600	0.01000
	FMA	Fraction > Dow	0.17000	0.95200	0.76000	0.26600	0.12600
	TRB	Fraction > Dow	0.28000	0.99600	0.92800	0.25800	0.09000
(1, 150, 0.01)	VMA	Fraction > Dow	0.11000	0.98200	0.99600	0.14600	0.01200
	FMA	Fraction > Dow	0.13000	0.98800	0.94800	0.28400	0.03600
	TRB	Fraction > Dow	0.19800	1.00000	0.83800	0.39800	0.12000
(1, 200, 0)	VMA	Fraction > Dow	0.11600	0.98600	0.99600	0.14800	0.01200
	FMA	Fraction > Dow	0.36600	0.96400	0.85200	0.78200	0.19000
	TRB	Fraction > Dow	0.32400	1.00000	0.85800	0.24800	0.14800
(1, 200, 0.01)	VMA	Fraction > Dow	0.12000	0.98000	0.99600	0.15200	0.01200
	FMA	Fraction > Dow	0.27400	1.00000	0.99600	0.39600	0.01400
	TRB	Fraction > Dow	0.29600	1.00000	0.85600	0.41000	0.14000
Panel B: Rule Averages							
Rule average	VMA	Fraction > Dow	0.08600	0.98600	0.99600	0.15400	0.01000
		Mean	0.00031	0.01136	0.00013	0.01210	0.00018
		Dow	0.00042	0.00890	-0.00025	0.01342	0.00067
Rule average	FMA	Fraction > Dow	0.23000	0.99400	0.99200	0.48600	0.01600
		Mean	0.00369	0.04255	0.00147	0.04342	0.00222
		Dow	0.00531	0.03064	-0.00399	0.04156	0.00930
Rule average	TRB	Fraction > Dow	0.19600	1.00000	0.86200	0.34600	0.10600
		Mean	0.00577	0.05312	0.00167	0.05242	0.00410
		Dow	0.00633	0.03000	-0.00238	0.05398	0.00871

that the GARCH-M generates an average spread of 0.018 percent, compared with 0.067 percent for the Dow series. Of the simulations only one percent generated buy-sell returns as large as those from the Dow series. The GARCH-M generates a positive buy-sell spread that is substantially larger than the spread under the AR(1), but this spread is still small when compared with that from the original Dow series.²⁰ These findings are repeated for the FMA and TRB rules, although the results are somewhat weaker.

The discrepancies for sell returns are in particular large. For example, for the VMA rule, the GARCH-M generates a sell average return of 0.013 percent which should be compared with an actual return of -0.025 percent for the Dow series. This difference is highly significant as indicated by the “ p -value” of 0.996. Results for the FMA and TRB rules strongly support the VMA results, and overall present strong evidence that the GARCH-M is incapable of generating returns consistent with the negative returns for the sell periods.

To test the reliability of our results for the 500 replications a test of convergence is performed similar to that done for the random walk. Results are again estimated for the (1, 200, 0) FMA trading rule and are presented in Figure 2. This figure shows the estimated p -values varying the number of replications from 100 to 2000. The results reveal that the estimated p -values are reliable after 500 replications. It should be noted that this test is crucial




Figure 2.

²⁰ The GARCH-M estimated here also contains an MA(1) component. Part of this returns spread could be coming from this element of the model.

since the asymptotic properties of the bootstrap for the GARCH-M are not known.

The GARCH-M model not only fails in replicating returns, but also is unable to match the results for volatility. For the Dow series standard deviations are lower for the buy periods than for the sell periods. Panel B of Table VIII for the VMA rule shows the GARCH-M average standard deviation for buys to be 1.14 percent, which should be compared with 0.89 percent for the Dow series. The “ p -value” of 0.986 supports the significance of this difference. Hence, the GARCH-M is substantially overestimating the volatility for buy periods. This provides a partial explanation for the high returns that the GARCH-M generates during buy periods. The GARCH-M does better at replicating volatility during sell periods. The model generates a standard deviation of 1.21 percent while the actual standard deviation is 1.34 percent, and the “ p -value” for this difference is 15.4 percent. The performance of the GARCH-M in predicting volatility during sell periods is even better for the FMA and TRB.

In summary, the GARCH-M model fails to replicate the conditional returns for the Dow Index. Moreover, the focal point of the GARCH models is to predict volatility where, as with returns, the GARCH-M model is also not able to match the results in the actual series.²¹

D. EGARCH

The next simulations examine the EGARCH model. These results are presented in Table IX. The important difference in the EGARCH model is that the conditional variance reacts differently to positive and negative shocks to the returns series. This may have the potential to explain some of our results for standard deviations from the simple GARCH-M model.

In Panel B of Table IX we see in the Buy, Sell, and Buy-Sell columns that the results are similar to those from the GARCH-M. For example, for the VMA rule the EGARCH generates a mean buy-sell difference of 0.002 percent and a p -value of 0.00. Checking the buy and sell returns separately, we see that for the buys the estimated p -value is 0.00 and for the sells it is 99.8 percent. The results for the other two tests are generally consistent with the VMA results.

Turning to volatility, we observe some changes in the results when compared to GARCH-M. While the volatility during sell periods is generally similar to that from the GARCH-M, volatility during buy periods is different. For all three rules, the conditional buy volatilities are lower for the EGARCH

²¹ Although it seems very unlikely, market microstructure issues could still explain the results that we see. To address this issue in the context of the AR(1) and GARCH-M models we experimented with skipping one day before using a specific trading rule signal. In other words, a buy signal on Wednesday would generate a purchase at Thursday's close rather than Wednesday as is used for all the tests in the paper. We performed this experiment for quite a few of the rules and obtained results that are similar to our previous findings without skipping a day.

Table IX
Simulation Tests From Exponential GARCH Bootstraps for
500 Replications

Exponential GARCH returns series are simulated using the estimated parameters and standardized residuals. The rows marked "Fraction > Dow" refer to the fraction of simulations generating a mean or standard deviation larger than those from the actual Dow series. The trading rules are, variable-length moving average (VMA), fixed-length moving average (FMA), trading range break (TRB). Panel B presents results for the averages across all the rules for each reported statistic. "Dow" refers to the actual mean return or standard deviation from the Dow, and "Mean" refers to the mean from the simulated series.

Panel A: Individual Rules							
Rule		Result	Buy	σ_b	Sell	σ_s	Buy-Sell
(1, 50, 0)	VMA	Fraction > Dow	0.00600	0.34000	0.99800	0.03400	0.00200
	FMA	Fraction > Dow	0.65000	0.40800	1.00000	0.35200	0.03600
	TRB	Fraction > Dow	0.00800	0.17600	0.87800	0.23600	0.02400
(1, 50, 0.01)	VMA	Fraction > Dow	0.00000	0.31200	0.99600	0.02600	0.00000
	FMA	Fraction > Dow	0.17200	0.55400	0.99800	0.38800	0.00200
	TRB	Fraction > Dow	0.01800	0.13000	0.86600	0.35600	0.01600
(1, 150, 0)	VMA	Fraction > Dow	0.00400	0.23000	0.99800	0.05600	0.00000
	FMA	Fraction > Dow	0.14600	0.78400	0.84000	0.07800	0.06800
	TRB	Fraction > Dow	0.04000	0.50400	0.96800	0.31000	0.01200
(1, 150, 0.01)	VMA	Fraction > Dow	0.00000	0.14800	0.99800	0.04600	0.00000
	FMA	Fraction > Dow	0.11000	0.96200	0.97000	0.11600	0.01600
	TRB	Fraction > Dow	0.02600	0.76200	0.92400	0.30800	0.01400
(1, 200, 0)	VMA	Fraction > Dow	0.00200	0.30400	0.99800	0.05000	0.00000
	FMA	Fraction > Dow	0.35000	0.88000	0.90400	0.77800	0.12800
	TRB	Fraction > Dow	0.07400	0.81400	0.92000	0.24600	0.03200
(1, 200, 0.01)	VMA	Fraction > Dow	0.00400	0.16600	1.00000	0.05000	0.00000
	FMA	Fraction > Dow	0.25000	0.99400	0.99400	0.28400	0.00400
	TRB	Fraction > Dow	0.08200	0.91000	0.93800	0.34600	0.03000
Panel B: Rule Averages							
Rule average	VMA	Fraction > Dow	0.00000	0.20600	0.99800	0.04800	0.00000
		Mean	0.00024	0.00867	0.00022	0.01218	0.00002
		Dow	0.00042	0.00890	-0.00025	0.01342	0.00067
Rule average	FMA	Fraction > Dow	0.15600	0.96400	0.99800	0.32600	0.00200
		Mean	0.00351	0.03470	0.00200	0.04031	0.00151
		Dow	0.00531	0.03064	-0.00399	0.04156	0.00930
Rule average	TRB	Fraction > Dow	0.01600	0.58800	0.94400	0.30000	0.00800
		Mean	0.00310	0.03070	0.00354	0.05150	-0.00043
		Dow	0.00633	0.03000	-0.00238	0.05398	0.00871

than for the GARCH-M. For example, for the VMA rule the volatilities for the EGARCH and GARCH-M are 0.87 percent and 1.14 percent, respectively.

To summarize, the EGARCH does not explain the differences between the buy and sell returns generated by the technical trading rules. The EGARCH, however, produces a large difference between the buy and sell volatility levels. This difference leads to closer agreement with the actual Dow series for the buys, but large volatility during sell periods is still not well matched by this model.

V. Conclusions

The recent studies on predictability of equity returns from past returns suggest that the conclusion reached by many earlier studies that found technical analysis to be useless might have been premature. In this paper we investigate two of the simplest and most popular trading rules—moving averages and trading-range breaks—by utilizing a very long data series, the Dow Jones Industrial Average index from 1897 to 1986. For statistical inferences, in addition to the standard tests, we apply the bootstrap methodology. The returns conditional on buy (or sell) signals from the actual Dow Jones data are compared to returns from simulated comparison series generated by a fitted model from the null hypothesis class being tested. The null models tested are: random walk with a drift, AR(1), GARCH-M, and EGARCH.

Overall our results provide strong support for the technical strategies that we explored. The returns obtained from buy (sell) signals are not likely to be generated by the four popular null models. Consistently, buy (sell) signals generate returns which are higher (or lower) than “normal” returns. A typical difference in returns over a 10-day period between a buy and a sell signal is about 0.8 percent, which is sizable when compared to a “normal” 10-day upward drift of about 0.17 percent. The small positive autocorrelation in returns cannot explain the observed patterns and accounts for less than 10 percent of the differences in returns between buys and sells. Furthermore, the difference in returns between buys and sells is not easily explained by risk. The results reveal that, following a buy signal, stock returns are substantially less volatile than following a sell signal. The most intriguing result is for the moving average rule with a variable-length holding period. Under this rule, an investor is continuously in the market. Following buy signals, the market increased at an annual rate of 12 percent. In contrast, following sells, a decrease of 7 percent is observed. This negative conditional return over a large fraction of trading days is an intriguing result because predictably negative returns are inconsistent with existing equilibrium models.

The popular GARCH-M model cannot explain the returns that our various rules generate. Moreover, our statistical design enables us to better understand the problems with the GARCH-M. For example, the GARCH-M predicts roughly the same volatility for buys and for sells. The actual results reveal a different picture. Following a buy signal the volatility is substan-

tially below the volatility following sell signals. Therefore, the GARCH-M model fails not only in predicting returns, but also in predicting volatility. The EGARCH model also displays similar problems in matching conditional means in the data. However, it performs better than the GARCH-M in predicting volatility, although it also fails in matching the volatility during sell periods.

Our results are consistent with technical rules having predictive power. However, transactions costs should be carefully considered before such strategies can be implemented. Of course, there are cases where the marginal transaction costs are zero, such as for pension funds that must reinvest dividends and funds contributed by sponsors. Opportunities also might exist in the futures markets where transactions costs are very small.

In sum, this paper shows that the returns-generating process of stocks is probably more complicated than suggested by the various studies using linear models. It is quite possible that technical rules pick up some of the hidden patterns. We would like to emphasize that our analysis focuses on the simplest trading rules. Other more elaborate rules may generate even larger differences between conditional returns. Why such rules might work is an intriguing issue left for further studies.

Appendix: Use of the Bootstrap to Estimate p -Values for Technical Trading Rules Under Null Models

We want to be able to estimate confidence intervals for trading profits under various null models whose parameters will be estimated on our data. To do this we briefly explain the bootstrap method we shall use. A more complete description can be had by writing the authors.

Following Efron (1982, chapters 5 and 10) and Singh (1981), let (E_1, \dots, E_n) be n IID draws, i.e., a random sample, from distribution F . Let $T(E_1, \dots, E_n; F)$ be a random variable of interest which may depend directly upon F . Let F_n be the empirical distribution function that puts mass $1/n$ on E_i , $i = 1, 2, \dots, n$. The bootstrap is a method that approximates the distribution of $T(E_1, \dots, E_n; F)$ under F by the distribution of $T(E_1, \dots, E_n; F_n)$ under F_n , where (E_1, \dots, E_n) denotes an IID sample from F_n . In other words the bootstrap estimates

$$\text{Prob}\{T(E_1, \dots, E_n; F) \in A\} \quad (\text{A1})$$

by

$$\text{Prob}\{T(E_1, \dots, E_n; F_n) \in A\} \quad (\text{A2})$$

Here, $\text{Prob}\{\cdot\}$ denotes the probability of event $\{\cdot\}$. Bootstrapping is done by taking B bootstrap samples, $Z_b = (E_{1,b}, \dots, E_{n,b})$, $b = 1, 2, \dots, B$ (each of which consists of n IID draws from F_n). Estimate (A2) by

$$\frac{1}{B} \sum_{b=1}^B I_A(T(Z_b; F_n)), \quad (\text{A3})$$

where I_A denotes the indicator function of event A which takes the value 1 when the event occurs and zero otherwise. By taking B to infinity with cheap computer time one can get as close to (A2) as one wishes. We shall show how to apply the bootstrap to estimate quantities based upon technical trading rules.

Case 1: $\{X_t\}$ is IID.

Assume $\{X_t\}$ is IID, drawn from F . Start with the random walk,

$$X_t = Y_t - Y_{t-1} = E_t, \quad E_t \sim F, \tag{A4}$$

where $Y = \log(P)$ where P is stock price, X is returns, and E_t is the innovation at date t .

We wish to test the adequacy of (A4) using a data set of $n = 1$ stock prices P_1, \dots, P_{n+1} . Build returns by setting

$$X_t = Y_t - Y_{t-1} = \log(P_t) - \log(P_{t-1}), \quad t = 2, 3, \dots, n + 1. \tag{A5}$$

Define date t to be a buy signal if

$$P_t > MA_{t,L}, \quad \text{and} \quad P_{t-1} < MA_{t-1,L}, \tag{A6}$$

where

$$MA_{t,L} = \frac{1}{L} \sum_{j=0}^{L-1} P_{t-j}. \tag{A7}$$

For each buy signal purchase one dollar's worth of stock, hold it for h periods, then sell it to obtain

$$R_t^h \equiv \frac{P_{t+h} - P_t}{P_t}. \tag{A8}$$

Let

$$R_n = \frac{1}{n} \sum_{t \in B} R_t^h, \tag{A9}$$

where B is the set of all buy signals.

The bootstrap simulations allow us to compare R_n with the empirical distribution for R_n from the simulated random walks, $R_{n;Z}$. We can then estimate objects such as $\text{Prob}(R_{n;Z} > x)$.

We will write $R_{n;Z}$ when we want to focus attention on R as a function of the random sample $Z = \{X_1, \dots, X_n\}$. Unfortunately, deriving an analytical expression for either a small sample or asymptotic distribution for $R_{n;Z}$ is very difficult. Fortunately, we can estimate it using the bootstrap. We will need to show a few properties for $R_{n;Z}$ and that it conforms to the standard bootstrap cases. First, under the null model

$$\lim_{n \rightarrow \infty} R_{n;Z} = E(R), \quad a.s.$$

Second, $E(R)$ can be written in the form $E(R) = t(F)$. Third, we will show $E(R_{n;Z}) = E(R)$ so that $E(R_{n;Z})$ is an unbiased estimator of $E(R)$. Fourth,

we will show how the bootstrap can be used to estimate any quantity of the form $\text{Prob}\{R_{n;Z} \in A\}$.

The first issue that must be dealt with is stationarity. The price process $\{P_t\}$ generated by (A4) is not a stationary process. This is because $P_t = e^{Y_t}$, and Y_t is stationary in first differences, not in levels. It turns out, however, that stationarity of $\{X_t\}$ allows us to write the probability of event (A6) and the quantities (A8) and (A9) as time stationary functions of a finite number of $\{X_t\}$. Stationarity and ergodicity of $\{X_t\}$ will allow us to prove that $R_{n;Z}$ converges to $E(R)$ almost surely as $n \rightarrow \infty$. It will also allow us to carry over central limit theorems and allow us to use the bootstrap to estimate quantities of the form $\text{Prob}\{R_{n;Z} \in A\}$. Central limit theory is needed by Bickel and Freedman (1981), Singh (1981), and Freedman (1984) to justify use of the bootstrap quantity (A2) and to approximate the quantity of interest (A1).

First, use the homogeneity of degree zero of (A6) in P_{t-L} to write it in the form

$$\exp(X_t + \dots + X_{t-L+2}) > \{\exp(X_t + \dots + X_{t-L+2}) + \dots + \exp(X_{t-L+2})\}/L,$$

and

$$\begin{aligned} &\exp(X_{t-1} + \dots + X_{t-L+1}) \\ &< \{\exp(X_{t-1} + \dots + X_{t-L+1}) + \dots + \exp(X_{t-L+1})\}/L. \end{aligned} \tag{A6'}$$

Note that the indicator function I_B , of event (A6'), which we shall call BUY, can be written as a time stationary function of (X_t, \dots, X_{t-L+1}) . Turn now to returns for a given t in B .

From (A8) returns can be written

$$\frac{P_{t+h} - P_t}{P_t} = \exp\left(\sum_{j=1}^h X_{t+j}\right) - 1, \tag{A10}$$

which is a time stationary function of a finite number of consecutive X 's. Hence R_n may be written as a time stationary function of a finite number of consecutive X 's. Hence,

$$R_{n;Z} = \frac{1}{n} \sum_{t \in B} \exp[X_{t+1} + \dots + X_{t+h}] - 1, \tag{A11}$$

where the sum is over all t that are indicated a buys. That is to say that (A11) can be written in the form, after defining the function $H(\cdot)$ and the vector $W = (X_{t-L+1}, \dots, X_{t+h})$ in the obvious way,

$$R_{n;Z} = \frac{1}{n} \sum H(W(t)). \tag{A12}$$

Equation (A12) is a time average of a time stationary function of a time stationary process. Therefore we would expect convergence almost surely of $R_{n;Z}$ to $E(R)$ and convergence in distribution of $n^{1/2}[R_{n;Z} - E(R)]$ to $N(0, V)$ where $N(0, V)$ denotes normal distribution with mean zero and variance V .

Routine application of the central limit theorems for weakly dependent processes in Gallant and White (1988) gives us central limit theory for (A12) for a much broader class of $\{X_t\}$ processes than the IID process (A4). This will be useful when we turn to the more general cases where the $\{X_t\}$ is a stable autoregressive process or a stable GARCH-M process. Turn now to the bootstrap.

By stationarity of the process $\{X_t\}$ we have

$$ER_{n;Z} = \frac{nE(H)}{n} = E(R) = E(H). \tag{A13}$$

In the case that $\{X_t\}$ is IID drawn from the distribution F the quantity EH may be written as a manifold integral against dF . Hence there is a mapping $J(F)$ from the space of distribution functions to the real line such that

$$EH = J(F), \tag{A14}$$

and, hence, we have the form appropriate for the bootstrap. In preparation for the bootstrap write (A12) thus

$$R_{n;Z} = \frac{1}{n} \sum H(W(t)) = T(Z; F), \tag{A15}$$

where Z denotes the entire sample (X_1, \dots, X_n) under the null hypothesis X_1, \dots, X_n is IID F . Recall that F_n is built from the data by placing mass $1/n$ on return $x_i, i = 2, 3, \dots, n + 1$ from (A5). We relabel $2, \dots, n + 1$ as $1, \dots, n$ in what follows. Note that F_n is the nonparametric maximum likelihood estimate of F and this fact coupled with (A14) plays a role in delivering the quality of bootstrap estimates (Efron (1982, chapter 5)). To estimate p -values we must estimate the cumulative distribution function,

$$\text{Prob}\{R_{n;Z} < x\} = \text{Prob}\{T(Z; F) < x\} = J(x; F). \tag{A16}$$

Do this by drawing B bootstrap samples $Z(b) = (X_{1,b}, X_{2,b}, \dots, X_{n,b})$ from $F_n, b = 1, 2, \dots, B$ and compute

$$\frac{1}{B} \sum I(T(Z_b; F_n) < x) = \hat{J}(x; F_n). \tag{A17}$$

The estimated p -value of the null model is now obtained from the data value r_n and

$$\hat{p} = 1 - \hat{J}(r_n; F_n). \tag{A18}$$

The estimated \hat{p} values are reported for two basic types of technical trading rule statistics for four classes of null models.

Case 2: $\{X_t\}$ is a stable autoregressive process.

Let the null hypothesis to be tested be the following. There are three parameters $a_0, a_1, |a_1| < 1$ such that,

$$X_t = a_0 + a_1 X_{t-1} + E_t, \quad X_0 = x_0, \quad \text{given } E_t \sim F \text{ IID}, \quad (\text{A19})$$

with zero mean and finite variance. Also, impose strong enough moment conditions on F so that the central limit theorem of White (1984, p. 119, see Exercise 5.17) can be applied to (A19) to show consistency and asymptotic normality of ordinary least squares estimates of a_0 and a_1 .

We follow Freedman (1984) to estimate the distribution of $R_{n;Z}$ as follows. First, estimate (A19) from the data by OLS and obtain the estimated residuals $\hat{e}_{t;n}$. Define F_n by placing mass $1/n$ on $\hat{e}_{t;n}$, $t = 1, \dots, n$. Second, obtain X_t^* by drawing E_t^* from F_n and generating using (A19) with the estimated parameters and $X_0^* = x_0$. Let Z^* represent the sequence of X_t^* . Now calculate $r_{n;Z^*}$ using this bootstrap sample and the trading rules. From these calculations the distribution for $R_{n;Z}$ can now be estimated.

Freedman (1984) extends the asymptotic theory of the bootstrap done by Bickel and Freedman (1981) to estimating the distribution function of parameter estimates for autoregressive models like (A19) under the stability condition $|a_1| < 1$ and some finiteness conditions on higher order moments of F . So we feel quite secure that the bootstrap should do a good job of estimating p -values for $R_{n;Z}$ under the null class (A19).

Case 3: GARCH-M and Exponential GARCH.

Both these models can be written in the form

$$X_t = G(I_t, A, E_t), \quad (\text{A20})$$

where A is a vector of parameters to be estimated, I_t is an information set which can include past X_{t-k} but no present or future X 's, and $\{E_t\}$ is an IID F process with mean zero and finite variance. In the section of our paper that uses the bootstrap to estimate confidence intervals for technical trading rules under these models we just used the bootstrap algorithm in the obvious way motivated by Case 2. That is to say, just estimate the parameter vector A and repeat the steps of the bootstrap algorithm replacing ordinary least squares estimation with the estimation method for A in (A20). We have not found any asymptotic justification for the bootstrap method like that of Freedman (1984) for general models like (A20). Therefore we follow Freedman and Peters (1984, p. 102) and do an evaluation experiment of the quality of our bootstrap estimates. These results are reported in Section IV of the text.

REFERENCES

- Alexander, S. S., 1961, Price movements in speculative markets: Trends or random walks, *Industrial Management Review* 2, 7-26.
- , 1964, Price movements in speculative markets: Trends or random walks, no. 2, in P. Cootner, ed.: *The Random Character of Stock Market Prices* (MIT Press, Cambridge, Mass.).
- Beran, R., 1986, Discussion of C. Wu's paper, "Jackknife, bootstrap, and other resampling methods in regression analysis," *Annals of Statistics* 14, 1261-1295.

- Bickel, P., and D. Freedman, 1981, Some asymptotic theory for the bootstrap, *Annals of Statistics* 9, 1196–1217.
- Bollerslev, T., 1987, A conditional heteroskedastic time series model for speculative prices and rates of return, *Review of Economics and Statistics* 69, 542–547.
- and Jeffrey M. Wooldridge, 1992, Quasi maximum likelihood estimation and inference in dynamic models with time varying covariances, *Econometric Reviews*, Forthcoming.
- Black, Fischer, 1967, Studies of stock price volatility changes, *Proceedings of the American Statistical Association, business and Economics Statistics Section* 177–181.
- Chopra, Navin, Josef Lakonishok, and Jay R. Ritter, 1992, performance measurement methodology and the question of whether stocks overreact, *Journal of Financial Economics* 31, 235–268.
- Conrad, Jennifer, and Gautam Kaul, 1989, Mean reversion in short-horizon expected returns, *The Review of Financial Studies* 2, 225–240.
- Coslow, Samson, and Harry D. Schultz, 1966, *A Treasury of Wall Street Wisdom* (Investors Press Inc., Pallisades Park, N.Y.).
- Cutler, D. M., J. M. Poterba, and L. H. Summers, 1991, Speculative dynamics, *Review of Economic Studies* 58, 529–546.
- De Bondt, W. F. M., and R. H. Thaler, 1985, Does the stock market overreact? *Journal of Finance* 40, 793–805.
- Efron, B., 1979, Bootstrap methods: Another look at the jackknife, *The Annals of Statistics* 7, 1–26.
- , 1982, *The Jackknife, The Bootstrap, and Other Resampling Plans* (Society for Industrial and Applied Mathematics, Philadelphia).
- and R. Tibshirani, 1986, Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy, *Statistical Science* 1, 54–77.
- Engle, Robert F., David Lilien, and Russel P. Robins, 1987, Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica* 55, 391–407.
- Fama, E. F., 1965, The behavior of stock market prices, *Journal of Business* 38, 34–105.
- and Marshall Blume, 1966, Filter rules and stock market trading profits, *Journal of Business* 39, 226–241.
- Fama, E. F., and K. R. French, 1986, Permanent and temporary components of stock prices, *Journal of Political Economy* 98, 246–274.
- Freedman, D., 1984, On bootstrapping two-stage least squares estimates in stationary linear models. *Annals of Statistics* 12, 827–842.
- and S. Peters, 1984a, Bootstrapping a regression equation: Some empirical results, *Journal of the American Statistical Society* 79, 97–106.
- , 1984b, Bootstrapping an econometric model: Some empirical results, *Journal of Business and Economic Statistics* 2, 150–158.
- French, Kenneth R., and Richard Roll, 1986, Stock return variances: The arrival of information and the reaction of traders, *Journal of Financial Economics* 17, 5–26.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Gallant, A. R., and H. White, 1988, *A Unified Theory of Estimation and Inference for Nonlinear Dynamic Models* (Basil Blackwell, New York, N.Y.).
- Gartley, H. M., 1930, *Profits in the Stock Market*, (H. M. Gartley Inc., New York, N.Y.).
- Goldberg, Michael D., and Stephan Schulmeister, 1988, Technical analysis and stock market efficiency, Working paper, New York University, C. V. Starr Center, New York.
- Hardle, W., 1990, *Nonparametric Regression* (Cambridge University Press, Cambridge).
- Jegadeesh, N., 1987, Predictable behavior of security returns and tests of asset pricing models, Ph.D. Dissertation, Columbia University.
- , 1990, Evidence of predictable behavior of securities returns, *Journal of Finance* 45, 881–898.
- Jensen, M. C., and G. Bennington, 1970, Random walks and technical theories: Some additional evidences, *Journal of Finance* 25, 469–482.
- , 1967, Random walks: Reality or myth—comment. *Financial Analysts Journal* 23, 77–85.

- Lakonishok, Josef, and Seymour Smidt, 1988, Are seasonal anomalies real? A ninety-year perspective *Review of Financial Studies* 1, 403-425.
- Leamer, E., 1978, *Specification Searches* (Wiley, New York).
- LeBaron, Blake, 1990, Technical trading rules and regime shifts in foreign exchange, Working paper, University of Wisconsin-Madison.
- Lehmann, Bruce N., 1990, Fads, martingales, and market efficiency, *Quarterly Journal of Economics* 105, 1-28.
- Levy, R. A., 1967a, Random walks: Reality or myth, *Financial Analysts Journal* 23, 69-77.
- , 1967b, Relative strength as a criterion for investment selection, *Journal of Finance* 22, 595-610.
- Lo, Andrew W., and Craig A. MacKinlay, 1990a, When are contrarian profits due to stock market overreaction? *Review of Financial Studies* 3, 175-206.
- , 1990b, Data-snooping biases in tests of financial asset pricing models, *Review of Financial Studies* 3, 431-468.
- Malkiel, B., 1981, *A Random Walk Down Wall Street*, 2 ed. (Norton, New York).
- Merton, R., 1987, On the state of the efficient market hypothesis in financial economics, in R. Dornbush, S. Fischer, and J. Bossons, eds.: *Macroeconomics and Finance: Essays in Honor of Franco Modigliani* (MIT Press, Cambridge, Mass.), pp. 93-124.
- Neill, Humphrey B., 1931, *Tape Reading and Market Tactics* (Forbes Publishing, New York, N.Y.).
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347-370.
- Pierce, Phyllis S., ed., 1991, *The Dow Jones Averages: 1885-1990* (Business One Irwin, Homewood, Ill.).
- Poterba, J. M., and L. H. Summers, 1988, Mean reversion in stock prices: Evidence and implications, *Journal of Financial Economics* 22, 27-59.
- Rosenberg, B., K. Reid, and R. Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management* 11, 9-16.
- Schabacker, Richard, 1930, *Stock Market Theory and Practice* (Forbes Publishing, New York, N.Y.).
- Schwarz, G., 1978, Estimating the dimension of a model, *Annals of Statistics* 6, 461-464.
- Sweeney, Richard, 1988, Some new filter rule tests: Methods and results, *Journal of Financial and Quantitative Analysis* 23, 285-300.
- Tomlinson, Lucile, 1953, *Practical Formulas for Successful Investing* (Wilfred Funk, New York, N.Y.).
- Van Horne, J. C., and G. G. C. Parker, 1967, The Random walk theory: An empirical test. *Financial Analysts Journal* 23, 87-92.
- White, H., 1984, *Asymptotic Theory for Econometricians* (Academic Press, New York, N.Y.).
- Wyckoff, Richard, 1910, *Studies in Tape Reading* (Fraser Publishing Company, Burlington, Vt.).