Stock Market Efficiency: An Autopsy?

This is the second in a series of three papers assessing the performance of the U.S. stock market. The first paper (Fortune 1989) dispelled the myth of increasing stock market volatility: it found that the monthly total rate of return on the Standard & Poor's 500 Composite Index has not been more volatile in the 1980s than in previous periods. Indeed, the peak of stock market volatility was in the 1930s. Others (for example, Schwert 1989) have reached the same conclusion using data going as far back as 1859. These observations suggest that investors do not face greater uncertainty about the returns on stocks than they have in the past, at least over periods of a month or longer. They also suggest that firms need not be concerned that the cost of equity capital has risen for risk-related reasons.

The present paper addresses the question of the efficiency of the stock market—do stock prices correctly reflect available information about future fundamentals, such as dividends and interest rates? Stated in another way, is the volatility of stock prices due to variation in fundamentals, or do other sources of volatility play a significant role?

The third paper will complete the trilogy by investigating the nature and consequences of very short-term (daily or intra-day) volatility. While the market's volatility over a period of a month or longer has not been increasing, rare—but prominent—daily spikes in stock price variation, which remain largely unexplained, have become the subject of public policy debate. Hence, the next paper will address episodes such as the October 19, 1987 crash and the break of October 13, 1989.

The present paper is structured as follows. Section I discusses the meaning of the Efficient Market Hypothesis (EMH) and draws out some of its implications for stock price behavior. The second section reviews the major stock market anomalies that cast clouds over the hypothesis of market efficiency, while the third section assesses "modern" evidence against the efficiency hypothesis. Section IV proposes an explanation for market inefficiency that is consistent with much of the evidence mar-

Peter Fortune

Professor of Economics at Tufts University and Visiting Scholar at the Federal Reserve Bank of Boston. *The author is grateful to Richard Kopcke and Joe Peek for constructive comments, and to David Zanger for his able research assistance.*

shalled in sections II and III. The paper concludes with a brief summary.

The purpose of this paper is not to draw out the policy implications of market inefficiency—that is the task of the next paper. However, the prominence of inefficiencies suggests a role for public policies that might be counterproductive in an efficient market. In short, this paper suggests that recent proposals for changes in margin requirements, introduction of trading halts, and other reforms might be productive.

I. The Efficient Market Hypothesis (EMH)

Practitioners are interested in the stock market because it is their bread and butter. Academic economists are interested for a very different reason: for them, the stock market provides an excellent laboratory for the evaluation of microeconomic theory. Common stocks are highly standardized products traded in an active auction market with very easy exit and entry of both producers (firms issuing equity) and consumers (investors purchasing shares); as a result, the prices of common stocks should conform to the implications of the theory of competitive markets.

The Efficient Market Hypothesis is the focus of the laboratory experiments, for it is the logical result of the application of microeconomic theory to the determination of stock prices. As Marsh and Merton point out (1986, p. 484):

> To reject the Efficient Market Hypothesis for the whole stock market . . . implies broadly that production decisions based on stock prices will lead to inefficient capital allocations. More generally, if the application of rational expectations theory to the virtually "ideal" conditions provided by the stock market fails, then what confidence can economists have in its application to other areas of economics where there is not a large central market with continuously quoted prices, where entry to its use is not free, and where short sales are not feasible transactions?

There were several reasons for the popularity enjoyed by the EMH in the 1960s and 1970s. First, it was rooted in a very strong theoretical foundation. This foundation began with Samuelson's work on the behavior of speculative prices, in which he showed that the prices of speculative assets should follow a random walk (1965). It was further buttressed by Harry Markowitz's theory of portfolio selection (1959) and by William Sharpe's construction of the Capital Asset Pricing Model (CAPM), which described the implications of optimal portfolio construction for asset prices in security market equilibrium (1964); both Markowitz and Sharpe won the 1990 Nobel Prize in Economics for their contributions. The final contribution was Robert Lucas's Rational Expectations Hypothesis (1978), which examined the implications of optimal forecasting for individual behavior and macroeconomic performance.

By the late 1970s, those who disputed the EMH found themselves facing an avalanche of sharply pointed and well-argued opposing positions. The theorists had apparently won, in spite of the paucity of supportive evidence, and the prevailing view was that no systematic ways exist to make unusual returns on one's portfolio. Practitioners were reduced to the position of ridiculing the EMH but could not make effective arguments against it.[1]

A second reason for the popularity of the EMH was the hubris of financial market practitioners during the 1960s. The "go-go" years had rested upon the notion that opportunities for unusual profits were abundant, and that it required only a reasonable person and a bit of care to sort out the wheat from the chaff in financial markets. The EMH provided an antidote to this hubris, for it argued that opportunities to make unusual profits were both rare and ephemeral: by their very nature, they were the result of temporary market disequilibria that are quickly eliminated by the actions of informed traders. Thus, the EMH counseled healthy skepticism in investment decisions. This skepticism about "beat the market" strategies has led to the popularity of index funds, which allow investors to hold "the market" without worrying about individual stocks.

"If the application of rational expectations theory to the virtually 'ideal' conditions provided by the stock market fails, then what confidence can economists have in its application to other areas of economics?"

The EMH, Definition 1: Prices Are Optimal Forecasts

The fundamental insight of the EMH is that asset prices reflect optimal use of all available information. A more formal statement is that the price of an actively traded asset is an optimal forecast of the asset's "fundamental value." To understand this notion, suppose that market agents think of each possible sequence of future events as a "state-of-the-world," and that there are N possible states of the world, to each of which a number s (s = 1,2,3, . . . , N) is assigned. For example, state-of-the-world 1 might be "dividends grow at 2 percent per year indefinitely, and a constant discount rate of 5 percent should be used," while s = 2 might be "dividends grow at 3 percent for two years, during which the discount rate is 7 percent, but thereafter dividends grow at 1 percent and the discount rate is 4 percent." Suppose also that the set of all the available information at time t is denoted by Ω_t and that $\pi(s|\Omega_t)$ is the probability that state s will occur, conditional on the information set available at time t.

Then for *each* state-of-the-world we can calculate a fundamental value of the asset, which we denote as $P^*(s)$. Hence, $P^*(4)$ is the fundamental value if state number 4 occurs. Note that there is no single fundamental value, rather there are N possible fundamental values, one for each state. If P_t is the market price and the *expected* fundamental value is $E(P_t^*|\Omega_t) = \Sigma_S P^*(s)\pi(s|\Omega_t)$, the EMH is embodied in the statement that the current price of the asset is equal to the expected fundamental value, or, more concisely, that the price of an asset is the best estimate of its fundamental value; that is,

$$(1) \qquad P_t = E(P_t^*|\Omega_t).$$

Consider the following simple example. There are three states-of-the-world: in state 1, the fundamental value is \$100, in state 2 it is \$75, and in state 3 the fundamental value is \$40. Investors do not know which state will materialize, but they have formed an assessment of the probability of occurrence of each state. Suppose that these probabilities are 0.25, 0.35 and 0.40, respectively. The market price under the EMH will be the expected fundamental value: P_t = 0.25(\$100) + 0.35(\$75) + 0.40(\$40) = \$67.25.

The EMH carries a number of strong implications about the behavior of asset prices. First, recall that Ω_t contains all the relevant information available at time t; this includes historical information (for example, past values of the asset price, the history of dividends, capital structure, operating costs, and the like), as well as current publicly available information on the firm's policies and prospects. Because P_t already incorporates all of the relevant information, the *unanticipated* component of the market price should be uncorrelated with *any* information available at the time the price is observed. A simple test of this proposition is to do a regression of P_t on a measure of the optimal forecast $E(P_t^*|\Omega_t)$ and upon any information that might be in Ω_t (say, the history of the stock price). This crude form of technical analysis[2] should result in a coefficient of 1.0 on $E(P_t^*|\Omega_t)$ and coefficients of zero on past stock prices, leading to the conclusion that all information in past stock prices has been embedded in the fundamental value.[3]

The fundamental insight of the EMH is that asset prices reflect optimal use of all available information.

A second implication bears on the sequence over time of prices under the EMH. Suppose we are at time t and we wish to forecast the price at time t + 1. If we knew the information that would be available at time t + 1, our forecast would be $E(P_{t+1}^*|\Omega_{t+1})$. But we do not know, at time t, the information available at time t + 1; we only know Ω_t. If r is the required rate of return on an asset with the risk level and other characteristics of the asset under consideration, the best forecast of P_{t+1} when we only know the elements in Ω_t is $E(P_{t+1}^*|\Omega_t)$.[4] Now, any new information arriving between time t and time t + 1 is, by definition, random so its effect on price creates a random deviation from today's best forecast.

From the optimal forecast definition of the EMH in (1), we see that the EMH implies the following sequence of prices:

$$(2) \quad P_{t+1} = (1 + r)P_t + \epsilon_{t+1}, \qquad \text{where } E(\epsilon_{t+1}) = 0$$

Equation (2) says that the sequence of prices will be a random walk with drift; price will vary randomly around a rising trend. Because new information is random, having no predictable components, ϵ_{t+1} has a zero mean and is without serial correlation.

This is the basis of the "random walk" tests of the EMH which estimate equations like (2) and search for serial correlations in the residuals. Appendix 1 uses time series analysis to determine whether daily changes in the closing value of the S&P500 index during the 1980s are consistent with a random walk. The answer appears to be "almost, but not quite." The results show a five-day trading cycle that can be used to predict stock price movements, but this cycle is a very small source of the total variation in the S&P500. Hence, it might not be strong enough to generate economic profits after transactions cost; this is, of course, consistent with the EMH.

The EMH, Definition 2: Risk-Adjusted Returns Are Equalized

The EMH can be restated in a different manner, which focuses on the rate of return on individual assets rather than their prices. The Capital Asset Pricing Model (CAPM), developed by Markowitz and Sharpe, states that in an efficient market the risk-adjusted *expected* returns on all securities are equal; any differences across assets in expected rates of return are due to "risk premia" arising from unavoidable (or "systematic") uncertainty.

The CAPM distinguishes between two types of risk: systematic risk, which affects all securities, each to a different degree, and unsystematic risk, which is unique to individual securities. Unsystematic risk can be avoided by an appropriate diversification of portfolios, based on the variances and covariances of security returns. Because unsystematic risk can be avoided without any sacrifice in the return expected from the investor's portfolio, it imposes no risk premium.

However, systematic risk, which affects all securities and is, therefore, unavoidable, will earn a risk premium. The CAPM defines a simple measure of the amount of systematic risk a security contains: its "beta coefficient." A security's beta coefficient measures the marginal contribution of that security to the market portfolio's risk: if $\beta = 0$, adding the security to the optimal portfolio does not affect portfolio risk; if $\beta < 0$, portfolio risk is reduced by adding the security to the portfolio: and if $\beta > 0$, the security adds to the portfolio's risk. The returns on a security with a beta of 1.0 move with the market—when the return on the market portfolio changes by 1 percent, the return on a $\beta = 1.0$ security moves in the same direction by 1 percent. Securities with betas greater than 1.0 have above-average risks, while securities with betas below 1.0 have below-average risks.

According to the CAPM, the realized return on a security is described by the following "characteristic line":

$$R_i = r_f + \beta_i(R_m - r_f) + \upsilon_i \qquad (3)$$

where R_i is the realized return on the specific security, r_f is the return on a risk-free asset (such as U.S. Treasury bills), R_m is the rate of return on the market portfolio (such as the S&P500), and υ_i is a zero-mean random variable whose variance measures the unsystematic risk of that security. The slope of this characteristic line, β_i, is the security's beta coefficient. The beta can be estimated by a bivariate regression of the excess return on an asset, $(R_i - r_f)$, on the excess return on the market, $(R_m - r_f)$.

Equation (3) describes the relationship between the *realized* return on an individual security and the *realized* return on the market. This provides the basis for answering the question, "What is the normal return on a security?" From the characteristic line we can see that because $E(\upsilon_i) = 0$, the *expected* return on an asset will be a linear function of the asset's risk level, as measured by its beta coefficient. For every security, the expected return will lie on the same straight line, called the Security Market Line (SML), which relates expected return to risk (beta). The SML is described by the following equation:

The Capital Asset Pricing Model states that in an efficient market the risk-adjusted expected returns on all securities are equal; any differences across assets in expected rates of return are due to "risk premia" arising from unavoidable uncertainty.

$$E(R_i) = r_f + (E_m - r_f)\beta_i \qquad (3')$$

where $E(R_i)$ is the *expected* return on the i*th* security, and E_m is the *expected* return on the market portfolio. This relationship represents the optimal forecast of a security's rate of return (rather than its price); under the EMH any deviations of the actual return, R_i, from this relationship must be random.

The term $(E_m - r_f)\beta_i$ is the "risk premium" for the *ith* security; it is the product of the market reward for a unit of risk, defined as the expected excess return on the market portfolio, and of the security's risk level, measured by its beta. The SML says that the optimal forecast of the return on a security is the risk-free interest rate plus a risk premium. Hence, on a risk-adjusted basis (when the risk premium is deducted from the expected return), all securities are expected to earn a return equal to the risk-free interest rate. Securities that have high betas, hence adding more to the portfolio risk, will carry higher expected returns because risk-averse investors will require higher returns on average to compensate them for the additional risks.

Each possible value of r_f and E_m will have a different SML describing security market equilibrium. Figure 1 shows the SML under the assumptions r_f = 8 percent and E_m = 12 percent. In this case the intercept will be 8 percent and the slope of the SML will be 4. The market portfolio would be at point "M," with a beta of 1.0 and an expected return of 12 percent; to verify this, simply substitute $\beta = 1.0$ into equation 3'. Discovery of a security whose *expected* return is above (or below) the SML is an indication of market inefficiency, for that security is expected to give a risk-adjusted return above (or below) the required level; that is, it is underpriced (or overpriced).

Thus, the SML can be used to describe the expected returns that are consistent with an efficient market. For example, point "HDL" on Figure 1 represents Handelman Corporation, a distributor of home entertainment media (records, video tapes, etc.). As of October 1990, Handelman's beta coefficient was 1.4, so the hypothetical SML predicts a return on HDL of 13.6 percent. Handelman's higher return is attributable solely to its above-average risk.

Figure 1

Efficient Markets:
The Security Market Line

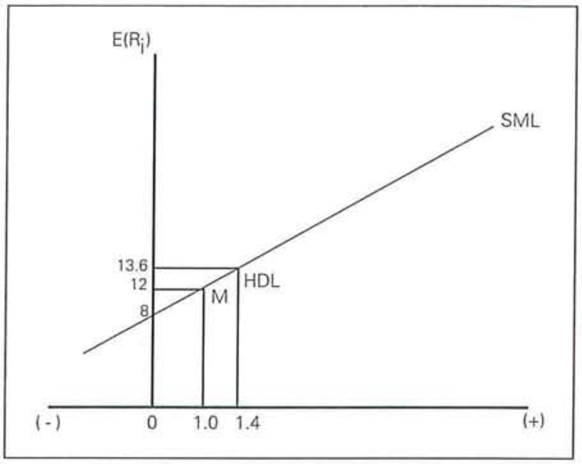

The equation for the SML is $E = r_f + (E_m - r_f)\ \beta_i$.

The SML drawn above assumes r_f = 8% and E_m = 12%.

Some Caveats

Both forms of the EMH rest on a strong assumption: the market equilibrium of asset prices is independent of the distribution across investors of the two basic raw materials of investment: information and wealth. In short, all those things that make different investors evaluate assets differently are treated as of negligible importance. Among these "irrelevant" factors are differences in probability assessments (more optimistic investors will invest a larger share in favored assets), differences in transactions costs (investors with low costs, such as financial institutions, will devote more resources to stocks than will high-cost investors, such as individuals), and differences in tax rates paid by investors.

If these factors can be ignored, the prices or returns on assets will be determined solely by fundamentals. But if they are important, prices can deviate, perhaps persistently, from fundamental values. Indeed most explanations of inefficiency in security markets rest on some form of heterogeneity among investors.

II. Stock Market Anomalies and the EMH

The theoretical victories of the EMH were not supported by empirical evidence. True, some studies did support the EMH; for example, numerous studies showed that stock prices were random walks in the sense that past stock prices provided no useful information in predicting future stock prices. But these studies do not represent the preponderance of the evidence for two reasons. First, gross inefficiencies can coexist with random walks in stock prices, as in the case of rational bubbles (which is discussed below). Second, and more important, by the 1980s a vast literature on stock market anomalies had developed. These anomalies, defined as departures from efficient markets that allow economic agents to enjoy unusually high (risk-adjusted) returns, appeared to lead to rejection of the EMH.

This section reviews some of the major anomalies in stock price determination that have been the traditional basis for rejecting the EMH. While a panel of coroners might not declare the EMH officially dead, by the late 1980s the burden of proof had shifted to the EMH adherents. The anomalies discussed in this section are among the list of causes that would appear on the death certificate.

Figure 2

Accumulated Value of $1 Invested in 1926

Source: Author's calculations, using Ibbotson (1990).

The Small-Firm Effect

Arguably the best-known anomaly in stock prices is the Small-Firm Effect: the common stocks of small-capitalization companies have, on average, exhibited unusually high rates of return throughout most of this century. This is shown clearly in Figure 2, which reports the accumulated values (assuming reinvestment of dividends) of an investment of one dollar in January of 1926 in two portfolios: the S&P500 and a portfolio of small-firm stocks. While small firms suffered more in the Great Depression, their growth since that episode has been far more dramatic than the growth in a portfolio represented by the S&P500.

According to the EMH, the small-firm effect should be due solely to higher beta coefficients for small stocks; in other words, the higher rate of return is solely due to higher risks. Any unwarranted growth should lead investors to restructure their portfolios to include more small-cap firms, thereby driving the price of small-cap stocks up relative to high-capitalization stocks and restoring a "normal" relationship in which all firms enjoy the same rate of return, after adjustment for risk. The evidence suggests, however, that the higher return on small-cap stocks is not attributable to higher risk. While in recent years the small-firm effect has disappeared, the puzzle is that it existed for so many years, in spite of general awareness that it was there.

The Closed-End Mutual Fund Puzzle

Another well-known anomaly involving a specific class of firms is the Closed-End Mutual Fund Puzzle, reflected in the discounts (and occasional premia) on closed-end mutual funds (Lee, Shleifer and Thaler 1990b; Malkiel 1977). Closed-end mutual funds differ from open-end mutual funds in that open-end funds keep the prices of their shares at the net asset value (NAV) by promising to buy or sell any amount of their shares at NAV. Closed-end funds, on the other hand, issue a fixed number of shares at inception, and any trading in those shares is between investors; this allows the closed-end fund share price to deviate from NAV, that is, closed-end funds can trade at either a discount or a premium. If the EMH is valid, then any sustained discount or premium on closed-end fund shares must be due to unique characteristics of the fund's assets or charter. In the absence of such distinguishing characteristics, any discounts or premiums would induce investors to engage in arbitrage that would eliminate the discount or premium. For example, an unwarranted discount would lead investors to buy the closed-end fund shares and sell short a portfolio of stocks identical to that held by the fund, thereby capturing a riskless increase in wealth equal to the discount. A premium, on the other hand, would induce investors to sell short the closed-end fund and buy an equivalent portfolio of stocks.

But closed-end fund shares typically sell at discounts, and the discounts are often substantial. Figure

3 shows the average year-end discount in the period 1970–89 for seven major diversified closed-end fund companies.[5] It is clear that the discounts move inversely to stock prices; periods of bull markets, such as 1968–70 and 1982–86, are associated with low discounts, while bear markets (the 1970s and 1987) are associated with high discounts. Thus, the price paid for a dollar of closed-end fund assets is procyclical.

Several reasons are offered for closed-end fund discounts. First, because of potential capital gains taxes on unrealized appreciation, a new buyer of closed-end fund shares faces a tax liability if the fund should sell appreciated securities; this potential tax liability justifies paying a lower price than the market value of the underlying securities.[6] Second, closed-end funds might have limited asset marketability if they buy letter stock or privately placed debt, which cannot be sold to the public without incurring the expense of obtaining Securities and Exchange Commission (SEC) approval or the restrictions on corporate policy often required by public market investors. Third, agency costs, in the form of high management fees or lower management performance, might explain the discounts.

Figure 3

Average Premium (+) or Discount (-) on Seven Closed-Ended Funds

Note: The seven companies are listed in text footnote 5.
Data from Barron's.

Malkiel (1977) found that the discounts were larger than could be accounted for by these factors, and other work has confirmed that this appears to be a true anomaly. To this should be added another puzzle: at inception, the initial public offering (IPO) of closed-end fund shares must incur underwriting costs and, as a result, the shares must be priced at a premium over NAV, after which the price of seasoned shares typically moves to a discount within six months. Why would informed investors buy the IPO, thereby paying the underwriting costs via capital losses as discounts emerge? Clearly, something irrational is going on!

Weekend and January Effects

Another class of anomalies focuses on specific time periods or seasonalities. Cross (1973) reported evidence of a Weekend Effect, according to which weekends tend to be bad for stocks; large market decreases tend to occur between the close on Friday and the close on Monday. Later work showed that the weekend effect really occurs between the Friday close and the Monday opening. In Appendix 1 a weekend effect is added to the time series model of stock prices for the 2,713 trading days in the 1980s. The result is resounding statistical support for a weekend effect. A plausible explanation of the weekend effect is that firms and governments release good news during market trading, when it is readily absorbed, and store up bad news for after the close on Friday, when investors cannot react until the Monday opening.

In recent years the January Effect has received considerable attention; the rate of return on common stocks appears to be unusually high during January. The primary explanation is the existence of tax-loss selling at year end: investors sell their losing stocks before year end in order to obtain the tax savings from deducting those losses from capital gains realized during the year. The selling pressure in late December is then followed by buying pressure in January as investors return to desired portfolio compositions. However, this explanation is not consistent with the EMH, according to which investors with no capital gains taxes, such as pension funds, should identify any tendency toward abnormally low prices in December and should become buyers of stocks oversold in late December. This means that tax-loss selling should affect the ownership of shares but not their price.

The January effect has been thoroughly investigated, and has been found to be more complicated than originally thought. Keim (1983) has shown that the January effect appears to be due largely to price behavior in the first five trading days of January; it is really an Early-January Effect. Also, Reinganum (1983) found that the January effect and the small-firm effect are commingled: the January effect appears to exist primarily for small firms and, in fact, much of the small-firm effect occurs in January.

In the time-series analysis reported in Appendix 1, a test was added for the January effect and for an early-January effect. The results do not support a January effect of either type in the 1980s, at least for the S&P500. The fact that it does not appear for large firms, which dominate the S&P500 and are the firms of primary interest to institutional investors, is consistent with the EMH. Arbitrage by well-informed institutional investors appears to prevent any late-December selling pressure from affecting the share prices of large-capitalization firms. This does not provide any conclusions about the January effect for small firms.

Figure 4

Annual Excess Return on Stocks, Classified by Value Line Rank, 1965 to 1989

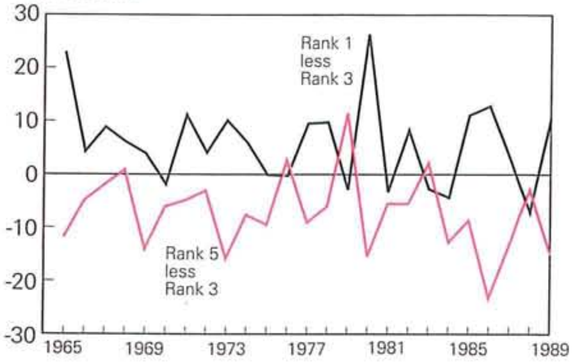

Source: Value Line Investment Survey.

The Value Line Enigma

Yet another well-known anomaly is the Value Line Enigma. The Value Line Investment Survey produces reports on 1700 publicly traded firms. As part of its service, Value Line ranks the common stocks of these firms in terms of their "timeliness," by which it means the desirability of purchasing the firm's shares. Value Line employs five timeliness ranks, from most timely (Rank 1) to least timely (Rank 5). Rank 3 is the designation for firms projected to increase in line with the market.

Figure 4 reports the annual average *excess* returns for Rank 1 and Rank 5 stocks. These returns are computed as the difference between the mean returns on stocks in the stated rank and the overall mean (Rank 3) returns. The computation assumes that the stocks are bought at the beginning of the year and sold at the end of the year, and transaction costs are not considered. It is clear that Rank 1 stocks generally perform better than average. In only five of the twenty-five years do Rank 1 stocks underperform the average; the probability of this happening by chance is only 0.00046.[7] Also, Rank 5 stocks tend to underperform. In only three of the twenty-five years do Rank 5 stocks perform better than average, and then the difference is small.

If the stock market is efficient, only one reason exists for higher rank stocks to generate higher returns: they have a higher level of market risk, that is, higher beta coefficients. Black (1971) found that the mean beta coefficients were roughly the same for stocks in each rank, concluding that the ranking system did have predictive value. However, Lee (1987) found that a stock's beta coefficient is inversely related to its Value Line rank: stocks for which purchase is timely tend to have higher betas. This suggests that the better performance of stocks ranked 1 and 2 is, at least in part, due to the higher average returns normally associated with higher risk.

Holloway (1981) examined the value of both active and passive trading strategies based on the Value Line Ranking System. An active trading policy was defined as purchasing Rank 1 stocks at the beginning of a year and holding them until the earliest of either the end of the year or a downgrade of the stock's Value Line rank, at which time the stock would be replaced by another Rank 1 stock to be held until year end. A passive, or buy-and-hold, strategy was defined as purchasing Rank 1 stocks at the outset of a year and selling them at year end. The active trading strategy generated higher returns than did the passive strategy when transaction costs were not considered, but was inferior to the buy-and-hold strategy when reasonable transaction costs were assessed. Hence, *active* trading using the Value Line ranking system is not a profitable strategy for investors.

However, Holloway found that even after adjustments for transactions costs and for risk, a passive strategy using Rank 1 stocks outperformed a passive strategy using Rank 3 stocks; the Value Line Ranking System did provide profitable information for those who are willing to buy and hold. It is noteworthy that this advantage existed even when adjustments were made for both transaction costs and risk (beta).

III. "Modern" Evidence of Inefficiency

The previous section reports the results of "traditional" approaches to assessing the EMH: examination of specific examples of departures from the EMH, called anomalies. During the 1980s several "modern" approaches were developed. These are the topic of this section.

Excess Volatility of Stock Prices

One of the more controversial "modern" tests of the EMH is based on the observed volatility of stock prices. Leroy and Porter (1981) and Shiller (1981) concluded that the observed amount of stock price volatility is too great to be consistent with the EMH. In order to understand this "excess volatility" argument, refer back to the first definition of an efficient market: a market is efficient if the price of the asset is an optimal forecast of the fundamental value, that is, if $P_t = E(P_t^* | \Omega_t)$.

The logic of the excess volatility argument is based upon a property of statistical theory: the optimal forecast of a random variable must, on average, vary by no more than the amount of variation in the random variable being forecasted. Thus, if the market price is an optimal forecast of the fundamental value—as the EMH implies—it should vary less than (and certainly no more than) the fundamental value.

A formal statement of the excess volatility argument is that the relationship between the fundamental price under the actual state (s) and the optimal forecast of the fundamental price is

$$P_s^* = E(P^* | \Omega_t) + \epsilon_s \tag{4}$$

where ϵ_s is a random variable that measures the deviation between the fundamental value for the state which actually occurs, P_s^*, and the optimal forecast of the fundamental value, $E(P^* | \Omega_t)$. If the forecast is optimal, these deviations must be random and uncorrelated with the forecast itself. Now, the EMH implies that $P = E(P_s^* | \Omega_t)$, which means

$$P_s^* = P + \epsilon_s. \tag{4'}$$

In other words, the correct price (conditional on knowing the true state) is equal to the market price plus a random term, denoted by ϵ_s, which measures the surprise resulting when the true state is known. This random term must be uncorrelated with P, because P is the optimal forecast and, therefore, already reflects any systematic information.

This provides the basis for the variance bounds tests of the EMH. Equation (4') shows that the variance of the fundamental price is equal to the variance of the market price *plus* the variance of the surprise. Turning this around produces the following relationship:

$$VAR(P_t) = VAR(P^*_t) - VAR(\epsilon_t). \tag{4''}$$

Because variances *must* be non-negative, if the EMH is valid the variance of the market price must be no greater than the variance of the fundamental value, or:

$$VAR(P_t) \le VAR(P_t^*). \tag{4'''}$$

Consider the following simple example, summarized in Table 1. Assume three states of the world, in each of which the dividend-price ratio is 10. In state 1 dividends paid at year end will be \$10 and the

Table 1
Example of Variance Bounds Tests
Three States—Three Years

State (s)	Fundamental Value (P_t^*)	Probability of State s in Year 1	2	3
1	\$100	.25	.50	.10
2	75	.50	.30	.80
3	40	.25	.20	.10
Market Price		\$72.50	\$ 80.50	\$74.00
Modal Fundamental Value*		\$75.00	\$100.00	\$75.00

*The modal fundamental value assumes that the most likely state occurs in each year.

fundamental value is $100, in state 2 dividends are $7.50 and the fundamental value is $75, and in state 3 dividends are $4 with a fundamental value of $40. These fundamental values are shown in column 2 of Table 1. At the beginning of each year the dividend to be paid at the end of the year is not known because the state that actually occurs is not known, but the *probability* of each state occurring is known. Therefore, at the beginning of each year investors know only the probability distribution of states and the dividend payment that each state entails. In this example, changes in the probability distribution across states correspond to the notion that new information is received by investors at the beginning of each year.

Under the null hypothesis of the EMH, the market price must vary by no more than the fundamental price. Any "excess" volatility is, therefore, a symptom of market inefficiency.

The "market's" problem is to determine a market price that best reflects that information. Table 1 assumes three years, with columns 3 to 5 showing the probability distribution of states in each year. The row marked "market price" shows the EMH market price, defined as the statistical expectation of fundamental values in each year. Thus, as time passes, the market price should increase from $72.50 to $80.50, then fall to $74.00; the sample standard deviation of the market price would be $4.25.

But the "correct" price, defined as the fundamental value associated with the realized state, would exhibit even larger movements. For example, if in each year the modal (most likely) state occurs, then the sequence of states is 2, 1, 2 and the fundamental values (row 5) would be $75 in year 1, $100 in year 2, and $75 in year 3. The sample standard deviation of these three "correct" prices would be $14.43, much greater than the sample standard deviation of the market price.[8] This result is consistent with the EMH.

Shiller's excess volatility tests were conducted as follows. He assumed a dividend valuation model in which the fundamental value is the present value of the perpetual stream of dividends resulting in each state-of-the-world. Using actual data on dividends paid over a very long period of time, and an assumption about the terminal price of shares, he calculated a time series for the fundamental value of the S&P500 index. He then compared the variance of that series with the variance of the observed values of the S&P500 and found that, if the discount rate was assumed to be constant, the variance of the market price was about six times the variance of the fundamental value—dramatic refutation of the EMH. However, if the discount rate was allowed to vary with interest rates (so that fundamental values exhibited greater variation), the market price had a variance about 1.5 times the variance of the fundamental price. In either case, the volatility of the stock market was greater than the upper bound implied by the EMH, leading Shiller to reject the EMH.

Under the null hypothesis of the EMH, the market price *must* vary by no more than the fundamental price. But Shiller's discovery implies either that the EMH is invalid or his test is invalid. This is a common problem of statistical tests: one must make assumptions about the world in order to construct any test, but one cannot know whether rejection of the null hypothesis is due to the invalidity of the hypothesis or to the invalidity of the assumptions.

The conclusion that excess volatility exists has been criticized for a number of reasons, each of which can be seen as a criticism of the test. Marsh and Merton (1986) disputed one of the assumptions underlying Shiller's test—that dividends are a stationary time series—and showed that if the process by which dividends are set is non-stationary, the EMH test is reversed: under the EMH, market prices *should* be more volatile than fundamental values. Kleidon (1986) has criticized the excess volatility test on statistical grounds, arguing that the Shiller test is an asymptotic test, assuming a very large sample of observations over time, and that the data available are necessarily finite, hence small-sample biases can weaken the test. In addition, the power of the test against reasonable alternative hypotheses is quite low, meaning that the test is not likely to reject the EMH when it should be rejected.

Whatever the validity of the excess volatility tests, they do provide an additional reason—other than observed anomalies—to doubt the validity of the EMH, and they have had a significant effect on the state of academic thinking about market efficiency.

Speculative Bubbles

It has long been a common practice to look back on dramatic collapses in asset prices and assign them to the bursting of a bubble. For example, following the October 1987 crash, many observers pointed out that stock prices had risen so rapidly in 1986 and 1987 that a bubble surely existed.

The notion of a "bubble" is a familiar one: a bubble reflects a difference between the fundamental value of an asset and its market price. Unfortunately, while the notion of a bubble has rhetorical force, it is a far more slippery concept than it appears. Clearly, a bubble is not merely a *random* deviation of price from value, for the law of large numbers suggests that purely random deviations will wash out over time without any necessity of collapse.

The bubble concept has been powerful because of the notion of self-fulfillment: bubbles are *self-fulfilling* departures of prices from fundamental values which continue until, for some reason, the conditions of self-fulfillment disappear. What do we mean by self-fulfilling bubbles? Recall that financial theory states that the market value of an asset (including dividends received) at the end of one period must be the market price at the end of the previous period, adjusted for growth at the required rate of return (r). That is, in equilibrium, where r is the required rate of return associated with the asset's risk level and E_t denotes an expectation conditional on information at time t:

$$E_t(P_{t+1} + D_{t+1}) = (1 + r)P_t. \quad (5)$$

This difference equation, when solved recursively, gives the following stock price model, which is the well-known present discounted value model:

$$P_t = \sum_{k=1}^{\infty} [(1 + r)^{-k} E_t D_{t+k}]. \quad (6)$$

One definition of a self-fulfilling speculative bubble is that its presence does not violate this description of asset prices. In that case, called a "rational bubble," market observers could not see the presence of a bubble and would not behave in ways that eliminate it. But how would a bubble remain invisible? To do this, it must be true that its existence does not violate the process shown in equation (5). If we define B_t as the size of the bubble, we can see that if the bubble is expected to grow at the required rate of return, that is, if $E_t B_{t+1} = (1 + r)B_t$, the bubble will be viable. In this case investors do not care if they are paying for a bubble because they expect to get the required return on that investment.

This definition of a rational bubble implies some very strong restrictions on bubbles. One is that bubbles cannot be negative: in order to be self-fulfilling, a negative bubble must become more negative at the geometric rate r, but the stock price will grow at a rate less than r because dividends are paid.[9] From (5) we can see that $E_t(P_t + 1) = (1 + r)P_t - E_t(D_t + 1)$. Hence, a negative rational bubble must ultimately end in a zero price, a result that, once acknowledged, must lead to the elimination of the negative bubble. Thus, while the market price might be below the fundamental value at a specific point in time, it cannot be the result of a rational bubble.

Because there is no upward limit on prices, a positive bubble can exist, although with some implausible consequences. First, as time passes a positive rational bubble must represent an increasing proportion of the asset's price. This is because the bubble must grow at the rate r, while the price grows at a rate less than r because of dividend payments.

Bubbles are self-fulfilling departures of prices from fundamental values, which continue until, for some reason, the conditions of self-fulfillment disappear.

But the idea that investors can project an indefinite increase in the relative size of the bubble undermines the existence of the bubble. Surely, if investors understand that a positive bubble means that the bubble must be an increasingly important component of price, they will imagine that at some time the bubble must burst. But as soon as they realize that it *must* burst, it *will* burst!

For example, suppose investors believe that a positive bubble exists but that it will not burst until the year 2091. They must, then, realize that in the year 2090 the market price must reflect only the fundamental value because the year 2090 investors will not pay for a bubble knowing that it will disappear. But if the year 2090 price is the fundamental value, no bubble can exist in 2089, and therefore the year 2089 price must be equal to the fundamental value. This chain of reasoning leads to the conclusion that a bubble cannot exist now. This will be true even if the collapse of a supposed bubble will not occur until after the passage of a very great (approaching infinite) time: as long as the resale price of the asset plays a negligibly small role in price determination, a bubble cannot exist![10]

If a rational bubble can never emerge, what is left of the notion of bubbles? Remember that a crucial assumption of the "bubbles cannot exist" paradigm is that investors behave as if they have an infinite time horizon. If investors have finite horizons, and plan to sell their shares before the present value of the sale becomes negligibly small, they will not project cash flows into the indefinite future, but will form judgments about the price at which the asset can be sold at the end of the horizon.[11] If, for example, the horizon is five years, the market price of the asset now will be described by the standard valuation equation $P_t = \Sigma_1^5[(1 + r)^{-k}E_tD_{t+k}] + (1+r)^{-5}E_tP_{t+5}$. If the expected resale price is simply the present value of expected dividends beyond that point, we are really back to the infinite-horizon model in which the ultimate resale price is irrelevant and bubbles cannot exist. For example, if $E_tP_{t+5} = \Sigma_6^\infty (1 + r)^{-k}E_tD_{t+k}$ we can see that the correct price will be described by equation (6).

Thus, the presence of a resale price whose expected value is not hinged to dividends beyond that point is necessary to the existence of rational bubbles. While it might be "rational" to use the infinite-horizon valuation model, it is not "realistic." Investors and traders do form judgments about the price at which they can sell assets, but they do not believe that the buyers are using an infinite-horizon model to decide the value of the asset.

Thus, rational bubbles are realistic descriptions of stock price performance; if the "market's" horizon is shorter than the time to the popping of a bubble, the bubble can continue. This is the essence of the "Greater Fool" explanation of speculative episodes: you will knowingly pay a price above fundamental value because you believe that someone later on will pay an even greater premium over fundamental value.

How should one go about testing the role of rational bubbles? This question is difficult to answer, for a rational bubble will not affect the sequence of prices until it breaks. The analysis of such low probability events is called the "peso problem": market prices will not reflect the effects of very low probability events even if they should have dramatic effects when they appear. Hence, it would be impossible to uncover a rational bubble as long as it exists. However, the disappearance of a bubble, such as a major decline in stock prices, can be examined to determine whether it was preceded by a speculative bubble in price.

Using the Ibbotson (1990) data for monthly returns on common stocks (S&P500) and one-month Treasury bills, a measure of stock price bubbles for the period February 1926 to December 1988 was constructed. This was done by computing, for each month, the difference between the total return on common stocks and the required return. The required return was computed as the one-month Treasury bill rate plus a risk premium. Denoting the actual return as R_t and the required return as $(r_t + \theta)$, where r_t is the one-month Treasury bill rate and θ is the risk premium, the bubble at time t is:

$$B_t = [1 + R_t - (r_t + \theta)]B_{t-1}. \qquad (7)$$

This approach assigns any difference between the observed return and the required return to

Figure 5

Estimated S&P 500 Stock Market Bubble

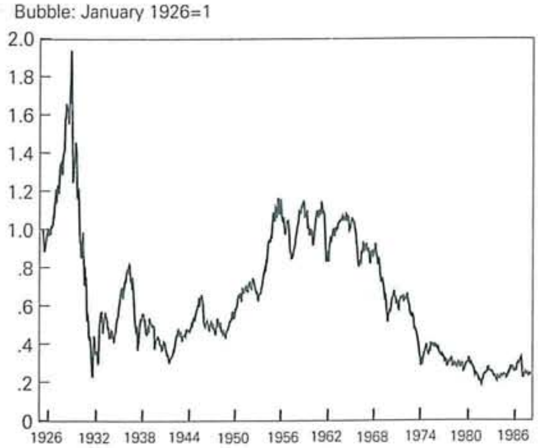

Source: Author's calculations, using Ibbotson (1990).

Table 2
Probability Model for Stock Market Crashes
Monthly Data 1926–88

	Size of Crash Over Next 12 Months					
	One Standard Deviation[a]		Two Standard Deviations		Three Standard Deviations	
CONSTANT	−1.5611 (−10.81)	−1.3795 (−6.34)	−2.4792 (−12.38)	−1.2698 (−4.50)	−3.7940 (−10.49)	−1.9569 (−3.92)
$\log BUBBLE_{t-1}$	+.9327 (+3.99)	+.8240 (+3.24)	+2.1689 (+4.52)	+1.3931 (+2.99)	+1.7837 (+2.23)	+.1897 (+.22)
TIME	n.a.	−.0007 (−1.12)	n.a.	−.0058 (−4.36)	n.a.	−.0139 (−2.46)
Number of Months	755	755	755	755	755	755
Proportion Predicted	.8808	.8808	.9603	.9603	.9881	.9881
Mean Probability	.7019	.7024	.8611	.8754	.9411	.9532
Number of Months Followed by Crashes	90	90	30	30	9	9

Note: Numbers in parentheses are t-statistics. The parameters are estimated using a logit model, according to which Prob(crash) = $1/1\{1 + \exp[-(a + bX)]\}$, where X is the list of explanatory variables (logBUBBLE and TIME).

n.a. = not applicable.

[a] In the 755 months in the period 1926:2 to 1988:12, the change in log bubble over the next 12 months had a mean of −0.0208 and a standard deviation of 0.2226. Hence, a "one-standard-deviation" crash is defined as a 12-month change in the logarithm of the bubble by an amount of −0.2434 or less. A two-standard-deviation crash is a change of −0.4660 or less.

changes in a bubble. Assigning any arbitrary positive value to the initial bubble makes it possible to trace out the path of the bubble using this difference equation. Note that the value assigned to the initial bubble is irrelevant since our interest is in movements in a bubble, not in its absolute size.

Figure 5 shows the path of our measure of the bubble over the period February 1926 to December 1988, using the assumptions that the initial bubble is 1.0 in January of 1926 and that θ is the average risk premium over the entire sample period.[12] The time series shown in Figure 5 indicates a very large speculative bubble in the late 1920s as the return on stocks sharply exceeded the required return. This was followed by the crashes of 1929–32. The next bubble emerged in the 1955–65 period, when the bubble appeared to remain high for a considerable period of time before a prolonged "crash" lasting through the 1970s. The bubble fell to a low point in late 1982 that matched the lows of the 1930s.

The Crash of 1987 has often been attributed to a speculative bubble emerging as prices rose dramatically in 1986 and the first nine months of 1987. However, our bubble measure does not support this interpretation. While stock returns were above the required return, creating an expanding bubble, the size of the bubble in September 1987 was so small that it could not be used to predict the crash.

Does our measure of a stock market bubble have any predictive value? Clearly, the concept of a bubble is intended to explain market crashes, not simply mild and temporary price declines. Therefore, we have employed a probabilistic model to determine whether the *probability* of a future crash is related to the size of the bubble. In an efficient market no relationship exists but, if speculative bubbles do exist, the probability of a crash should be a direct function of the size of the bubble.

Table 2 reports our results for several different definitions of a crash. A "one standard deviation crash" occurs if over the *next* twelve months the change in (the logarithm of) the bubble is more than one standard deviation *below* the mean change.[13] We also consider crashes of two and three standard deviations. The probability of a crash is assumed to be described by a logistic function, and the result is a logit[14] model in which the probability of a crash over the next twelve months is a function of the logarithm of the bubble at the end of the previous month. In order to correct for any trends in the relationship, we have also added a linear trend variable (TIME).

The results for 1926–88 suggest that the bubble does have predictive value: for both one- and two-standard-deviation definitions of a crash, the variable log $BUBBLE_{t-1}$ is statistically significant and has a positive sign: the bigger the bubble, the greater the probability of a crash in the next twelve months. The bubble size loses its predictive value if a crash is a three-standard deviation fall, but only 8 months fit this definition, and a model cannot predict events that occur so rarely.

Mean Reversion in Stock Returns

The phenomenon of mean reversion is the tendency for stocks that have enjoyed high (low) returns to exhibit lower (higher) returns in the future; that is, returns appear to regress toward the mean. A seminal test of mean reversion was undertaken by Fama and French (1988), who regressed the rate of return for a holding period of N months upon the rate of return during the previous N months. For example, they regressed the return over an 18-month period upon the return over the previous 18 months; a negative slope coefficient indicates mean reversion over an 18-month horizon. Fama and French found evidence of mean reversion for holding periods longer than 18 months.

While we have some skepticism about the Fama-French tests (the results appeared to be due primarily to the inclusion of the 1930s in the sample period), the phenomenon of mean reversion has been supported by other tests. Poterba and Summers (1988) found that the variances of holding-period returns do not increase in proportion to the length of the holding

Mean reversion can be thought of as another anomaly, but it is really much more: it is the common theme in most tests of stock market efficiency.

period, an indication of mean reversion.[15] Other variants of mean reversion tests also confirm the existence of mean reversion. For example, De Bondt and Thaler (1985) have found evidence of both a winner's curse and a loser's blessing in stock prices. Stocks that have experienced a recent reduction in their P/E ratios tend to have higher rates of return than equivalent stocks that have not been "losers," while stocks that have experienced increases in their P/Es tend to be losers in subsequent periods. The loser's blessing appears to be more dramatic than the winner's curse.[16]

Mean reversion can be thought of as another anomaly, but it is really much more: it is the common theme in most tests of stock market efficiency. For example, Shiller's excess volatility tests are an indirect test for mean reversion, and the analysis of bursting speculative bubbles is really an examination of sudden mean reversions.

IV. Why Is the Stock Market Inefficient?

Abundant evidence casts doubt on the Efficient Market Hypothesis. The natural next question is, "Why?" What aspects of investor behavior might account for these departures from the predictions of economic theory?

An important preliminary to answering this question is the observation that many of the anomalies shown in the previous section are really manifestations of one fundamental anomaly: the small-firm effect. For example, the January effect is primarily a characteristic of small firms (Keim 1983; Reinganum 1983), and the winner's curse and loser's blessing are also most prominent among small firms (De Bondt and Thaler 1985). Furthermore, the closed-end fund puzzle appears to be the result of the similarity of the markets for closed-end funds and for small-firm stock: Lee, Shleifer and Thaler (1990a) show that discounts on closed-end funds are highly correlated with performance of small-firm stocks, that institutions tend to shy away from both small-firm stocks and closed-end fund shares, and that the transaction size of both small-firm stocks and closed-end fund shares tends to be much lower than the transaction size for large-firm stocks.

This suggests a common denominator for many—but not all—of the departures from the EMH. They tend to be concentrated in stocks traded in relatively narrow markets where the "smart money" is not as likely to play. In short, inefficiencies might be associated with a form of market segmentation in which the EMH applies to stocks of large firms which are the province of financial institutions with access to research on fundamentals, while inefficiencies, and their associated profitable opportunities, appear to be concentrated among those who invest in the stocks of smaller firms, traded in less active markets with a lower quality of information.

Market Inefficiency and Market Segmentation

The claim that market participants can be segmented into highly informed and less informed investors, and that this fact is an important component of stock price determination, will create a sense of déjà vu: presenting it to market practitioners is a case of preaching to the choir. Indeed, the "smart money–dumb money" distinction has been around for as long as markets have existed, and it was enshrined in the work of early dissidents in the random walk debate. For example, Cootner (1964) argued that the profits of professional investors, who have low transactions costs, come from observing the random walk of stock prices produced by nonprofessionals and stepping in when prices wander sufficiently far from the efficient price.

Why has this view enjoyed a renaissance among academic economists? It is not merely because academics get to put notches on their guns when they disturb the conventional wisdom. Nor is it solely due to the more important reason that anomalies have become too numerous and well-documented to ignore. Each of these has played a role, but the fundamental reason is that only recently have economists provided a theoretical foundation for market segmentation.

EMH theorists rejected the market segmentation approach for several reasons. The first was that it clearly assumes irrational behavior, with the unsophisticated investors (henceforth called small investors) somehow driving prices of stocks (primarily of small firms) away from fundamental value. The second is that it assumes that the smart money allows this to happen and fails to step in when very small opportunities arise; in short, arbitrage is incomplete. The third problem is survival of the small investors; if small investors are buying high and selling low, as they must if they are giving large investors an opportunity for profits, then the population of small investors should diminish over time and the inefficiencies should disappear.

The possibility that some investors are "irrational" is not sufficient to induce inefficiency; nobody believes that all investors are rational, and so long as these investors are infra-marginal they are merely giving profits to large investors. True, theorists do not like irrationality, and might question why it should exist; but that is a question of psychology, not of economics. Furthermore, the proposition that irrationality is self-correcting because the irrational players incur losses and leave the game does not work well for two reasons. First, the mortality of investors, and the difficulty of transmitting wisdom and experience to the young, mean that a new crop of investors is always emerging which, if given sufficient endowments, can become players. The 1980s were an example of this, with young professionals having limited experience handling large amounts of money. Second, as we shall argue, it is possible that irrational investors do, in fact, get rewarded and are not eliminated.

It is possible that irrational investors do, in fact, get rewarded and are not eliminated.

The fundamental objection to market segmentation is the arbitrage objection. Large investors (it is argued) will ensure that prices must be nearly efficient. Because they care only about intrinsic value, if prices diverge from the efficient price, they will engage in arbitrage to restore the equality. The arbitrage objection can be dealt with by forgoing two implicit assumptions: that investors never plan to liquidate their stock positions, and that riskless arbitrage is possible. Each of these assumptions is addressed in turn.

As in the discussion of speculative bubbles, if investors have finite horizons, they will be concerned about the resale price of the security and form judgments about what that will be at the end of their horizon. One way of forming those judgments—the way proposed by traditional finance theory—is to estimate the future price as a present value of dividends received from that point on; that simply brings in the infinite life assumption through the back door. Another approach—which seems more plausible—is to recognize that investors are concerned with resale price and that their forecasts of resale price may well not reflect solely their judgments about future dividends; perhaps even more important will be their estimates of what other investors will be willing to pay. This was stated clearly by John Maynard Keynes (1964, pp. 155–56), who said of professional investors and speculators:

> They are concerned, not with what an investment is really worth to a man who buys it "for keeps", but with

> what the market will value it at, under the influence of mass psychology, three months or a year hence . . . professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole . . . We have reached the degree where we devote our intelligences to anticipating what average opinion expects average opinion to be.

But (an efficient market adherent might respond) even if many investors are forming judgments about resale prices that differ from fundamental value, a well-financed body of highly informed investors can prevent that from affecting market prices by engaging in riskless arbitrage. The response to this is that few opportunities for riskless arbitrage exist, and to the extent that arbitrage involves some risk, risk-averse investors will require a positive expected return (above opportunity costs), allowing inefficient pricing to continue.

An example of this is discounts on closed-end funds. Clearly, the market for closed-end fund shares must be dominated by investors who adopt a high probability that resale prices will deviate from fundamental values. If this were not true, investors would estimate that resale prices would be equal to future fundamental value, and riskless arbitrage would ensure that the current price reflects current fundamental value. For example, if current investors think that the future resale price will be above fundamental value, they will buy the closed-end fund shares and sell short a bundle of shares that replicate the closed-end fund portfolio. By doing this, they will enjoy profits, and their profit-seeking activities will ensure that current price is equal to current fundamental value.

Noise Trading

We have argued that prices *do* diverge systematically from fundamental values because prices *can* diverge systematically from fundamental values, because even well-informed investors are risk averse and will not engage in sufficient arbitrage activity to prevent this. This view has been formalized recently by a model of irrationality called "noise trading" by its proponents (Black 1986; Shleifer and Summers 1990). The noise trading model proposes that an important segment of the market consists of investors who bid prices away from fundamentals, thus introducing "noise" into stock prices. This noise, or "investor sentiment," is sufficiently broad in its impact, affecting many stocks, that investors cannot avoid it by diversification and must accept it as a source of systematic risk. Because it is systematic and undiversifiable, the noise affects the rate of return investors require on stocks and, therefore, market prices. Not all stocks are affected equally by noise risk—stocks of well-established firms, which are traded among informed investors, might not carry much noise risk, while stocks of small firms are more likely to bear this risk.

A simple model of noise trading is presented by DeLong, Shleifer, Summers and Waldmann (1990). This model, discussed in more detail in Appendix 2, assumes that young investors buy stocks and old investors sell stocks (to the young) to live on in their dotage. Sophisticated investors form optimal forecasts of the future price, but unsophisticated investors, called "noise traders," develop biased forecasts. Because sophisticated investors are risk averse and arbitrage is risky due to the possibility that the extent of price misperception by noise traders might change, sophisticated investors will not fully arbitrage away the influence of noise trading. Thus, noise traders can drive the market price away from the fundamental value.

The noise trading model proposes that an important segment of the market consists of investors who bid prices away from fundamentals, thus introducing "noise" into stock prices.

In the DeLong-Shleifer-Summers-Waldmann model, the degree of price misperception exhibited by noise traders—the difference between their forecasts and optimal forecasts—is assumed to be a random variable (denoted as ρ), which follows a normal probability distribution with mean ρ^* and variance σ^2. If $\rho^* = 0$, noise traders agree with sophisticated traders "on average," but noise trader forecasts will temporarily differ from sophisticated forecasts at any moment. In this case the equilibrium price of an asset will normally be below the fundamental value, the discount being necessary to compensate sophisticated traders for the risk that stock prices will deviate from fundamental value even more in the future. If noise traders are pessimistic ($\rho^* < 0$), the normal discount from fundamental value will be higher, while if noise traders are optimistic ($\rho^* > 0$), the normal discount will be lower or a premium might emerge. In addition to the normal discount arising from the average price misperception of noise traders, there is a temporary random discount due to temporary variations in optimism and pessimism ($\rho - \rho^*$).

The EMH adherent would ask why sophisticated investors do not dominate the market and, through arbitrage, force the market price to equal the fundamental value. If this did occur, the perceptions of noise traders would alter the ownership of stocks (sophisticated traders holding more when noise traders are pessimistic and less when noise traders are optimistic), but price misperceptions would not affect the market price of an asset.

The answer is, as noted above, that arbitrage is not riskless: no individual sophisticated trader can know that all sophisticated traders together will force equality of the market price with the fundamental value. There is always the possibility that noise traders will influence stock prices and that the sophisticated trader, when he arrives at the end of his horizon, will be forced to sell at a price even further below fundamental value than his cost.

This model is overly simple, designed for expository purposes and not as a strict representation of reality. But it does explain a number of important phenomena. For example, it explains the excess volatility of common stock prices found by Shiller: in the absence of noise trading, stock prices would always equal fundamental values, but with noise trading, stock prices will be more volatile than fundamental values because of the changing perceptions of noise traders. It also can explain the small-firm effect, and the related anomalies (for example, the January effect, the loser's blessing, the closed-end fund puzzle): the small-firm effect exists because the noise risk is higher among small firms, which are not as favored by sophisticated traders and in which noise traders play a larger role.

Furthermore, this simple model explains how the phenomenon of noise trading can persist. Friedman (1953) argued that, in the long run, prices must conform to fundamentals because speculators who paid incorrect prices would either go broke if they tended to buy high and sell low, or would force prices to equal fundamental value if they were sharp enough to buy low and sell high. In either case, what we now call noise trading would be a temporary phenomenon. In this model, however, noise traders create a more risky environment, but because their effects are pervasive and not idiosyncratic to individual stocks, the risk is not diversifiable and must earn a reward. Indeed, not only is noise trading consistent with an average return above the riskless rate, but it can be consistent with a higher average return for noise traders than for sophisticated investors if, as seems likely, noise traders tend to invest more heavily in noise-laden stocks, hence earning more of the risk premium associated with small stocks.

Noise Trading with Fads: A Simulation

The noise trading model can easily generate models of stock price movements that mimic the sharp breaks and apparent patterns visible in the real data. The crux of this is the possibility of "fads" that affect investor sentiment, measured by ρ. Small changes in these opinions can translate into very large changes in stock prices. For example, using equation (A2.1) in Appendix 2, we can calculate the effect of a change in investor sentiment of noise traders.

As noted above, noise trader price perceptions can be decomposed into two types: the normal perception is the average value of ρ, denoted by ρ^*, and is the degree of optimism that prevails over fairly lengthy periods; temporary price perceptions, denoted by $\rho - \rho^*$, prevail at any moment of time but do not affect the trend of stock prices.

Changes in the normal price perception, or ρ^*, can result in very large changes in stock prices. According to equation (A2.1), a change in ρ^* induces a change in stock price by the amount $(\mu/r)\ \rho^*$, where μ is the proportion of traders who are noise traders and r is the real interest rate. Assuming $\mu = 0.05$ and $r = 0.0042$ per month (5 percent per year), we calculate $(\mu/r) = 11.90$: a change in ρ^* by 0.01 (or 1 percent of fundamental value) will alter the stock price by 0.1190, or about 12 per cent of fundamental value.

Changes in stock prices due to variations in investor sentiment will be random so long as the normal price perception of noise traders is constant. However, if ρ^* is serially correlated—its current value depends on previous values—stock price changes can exhibit sharp breaks that do not conform to a model

Figure 6

Simulated 50-Year Stock Price History

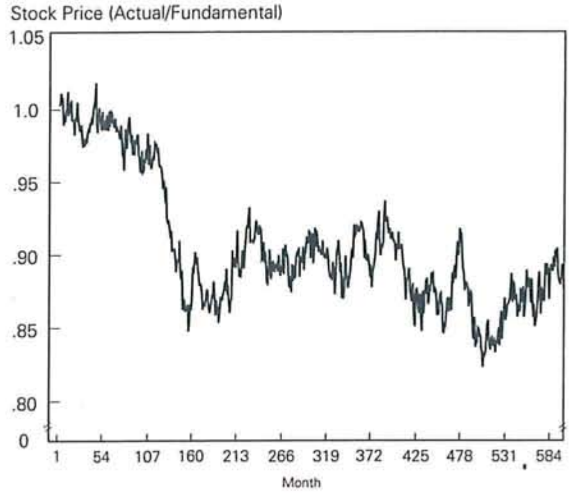

Source: See Appendix 2.

of simple random variation about an equilibrium level. This is likely to happen when there are fads in perceptions, as when optimism is reinforced by earlier phases of optimism.

To illustrate this, we have simulated monthly behavior of a hypothetical stock market using equation (A2.1). The settings for the crucial parameters are discussed in Appendix 2. In order to introduce the possibility of realistic results from a noise trading model we need to add the notion of "fads and fashions," in which investor interactions create waves of investor sentiment. Our way of introducing the possibility of contagious behavior is to assume that "opinion" follows a random walk shown by the autoregressive process:

$$\rho_t^* = \rho_{t-1}^* + \epsilon_t \quad (8)$$

where ϵ_t is white noise with mean $E(\epsilon_t) = 0$ and constant variance σ_ϵ^2. In this simple case, investors standing at a moment of time will forecast a constant value of ρ^* because zero is their optimal forecast of ϵ for every period. But the actual value of ρ^* will follow a path determined by equation (8). Because stock prices are very sensitive to ρ^*, this can result in realistic stock market cycles.

The results of one such experiment are shown in Figure 6. The results show what appear to be systematic patterns in the stock price imposed upon a bear market, in which the stock price falls from its fundamental value of 1.0 to about 0.88 at the end of the 50 years (600 months). Recall that the results are stock prices relative to fundamental value, so the figure does not mean that stock prices fall to 88 percent of the original value; an upward trend in the fundamental value could allow the stock price to rise even though, because of noise trading, it is not rising as fast as it should.

Repeated experiments will show essentially the same patterns of prolonged departures from fundamental value, punctuated by sharp breaks in price, though the sequence of bull and bear markets will be different in each simulation. It is clear that even with mild fads in noise trader misperceptions, there can be dramatic and apparently systematic cycles in stock prices. While this does not prove that noise trading is the source of the kind of stock market cycles we observe, it does show that noise trading, supplemented with contagious investor interactions or fads, is a plausible way of explaining observed stock price behavior.

V. Summary and Conclusions

This paper assesses the current state of the efficient market hypothesis, which was the conventional wisdom among academic economists in the 1970s and most of the 1980s. It reviews the empirical evidence and concludes that it provides an overwhelming case against the efficient market hypothesis. This evidence exists in the form of a number of well-established anomalies—the small firm effect, the closed-end fund puzzle, the Value Line enigma, the loser's blessing and winner's curse, and a variety of anomalies surrounding seasonality, such as the January effect and the weekend effect. Many of these anomalies are more pronounced among small-firm stocks, suggesting that the efficient market hypothesis might be more appropriate for stock of large firms, but analysis of the S&P500, which is dominated by large firms, also finds important anomalies such as a weekend effect, slow mean reversions in returns, and stock price volatility in excess of the amount predicted by fundamentals.

These anomalies can be explained by resorting to a model of "noise trading," in which markets are segmented with the "smart money" enforcing a high degree of efficiency in the pricing of stocks of large firms while less informed traders dominate the market for small firms. This model can explain many of the anomalies, and it can generate cycles in stock prices that are very similar to those observed in the real world.

Our fundamental conclusion is that the efficient markets hypothesis is having a near-death experience and is very likely to succumb unless new technology, as yet unknown, can revive it. This conclusion has a number of policy implications. The fundamental implication is that security market inefficiency provides

Our fundamental conclusion is that the efficient market hypothesis is having a near-death experience and is very likely to succumb unless new technology, as yet unknown, can revive it.

an economic foundation for public policy interventions in security markets. Clearly, if markets are efficient, hence conforming to the paradigm of pure competition, there is little reason for a security market policy: the market works to correct imbalances and to efficiently disseminate information.

However, if inefficiencies do abound, reflecting barriers to entry in transactions or inefficient collection, processing, and dissemination of information, there might be a role for public policy. For example, the existence of the Securities and Exchange Commission, whose primary function is to ensure equal access to relevant information, would be questionable in a world with an efficient market for information; in that case market prices would more accurately reflect all relevant information. As another example, in a world of efficient markets, sharp changes in prices, such as the October 1987 break, would reflect dramatic changes in fundamentals and should not elicit public policy responses. But in an inefficient market, in which investor sentiment clouds the influence of fundamentals, policies designed to mitigate price changes (daily price limits, market closings under certain conditions) might be appropriate.

The objective of this paper is not an examination of sharp and maintained price breaks such as October of 1987. But the noise trading model does suggest one reason why that break appeared to be lasting in its effect on stock prices. To the extent that the price break was not associated with changes in fundamentals (and it is widely agreed that it was not), it could have adversely affected investor sentiment, inducing prices to remain below fundamental values for prolonged periods. A useful analogy is the prevalence of discounts on closed-end mutual funds—these discounts appear to exist because investors recognize that they might get larger.

The purpose of the third paper in this trilogy will be to pursue the question of dramatic breaks in stock prices, and to investigate the wisdom and efficacy of policies designed to address potential problems of short-run stock price volatility that arise from stock market inefficiency.

Appendix 1: Time Series Analysis of Daily Stock Prices in the 1980s

In this appendix we report some results of tests for random walks and for specific anomalies using daily stock price data for the 1980s. Figure A-1 shows the data: the daily closing price of the S&P500 for each of the 2,713 trading days from January 2, 1980 to September 21, 1990. The chart reports the logarithm of the price index, rather than the index itself, for two reasons. First, analysis of stock prices is usually done on the logarithm of the price because desirable statistical properties are the result; in particular, the changes in log stock prices are close to normally distributed, a property that allows a broad range of statistical tools to be brought to bear. Second, a graph of the log price has the property that the slope of the line measures the percentage rate of change of the price; for example, the chart shows that the rate of increase in stock prices was particularly high from mid-1982 through mid-1983, then slowed somewhat until the October 19, 1987 break, and after that break, the rate of increase was lower than it had been in the earlier bull market periods.

Are Daily Stock Prices a Random Walk?

An attempt to fit this model to the daily S&P500 closing prices for 1980–90 failed to support the random walk hypothesis: while the intercept and slope coefficients were consistent with the EMH, the residuals were not white noise, but showed significant autocorrelation. Further experimentation using time series methods led us to conclude

Figure A-1

Logarithm of Daily S&P 500 Closing Price

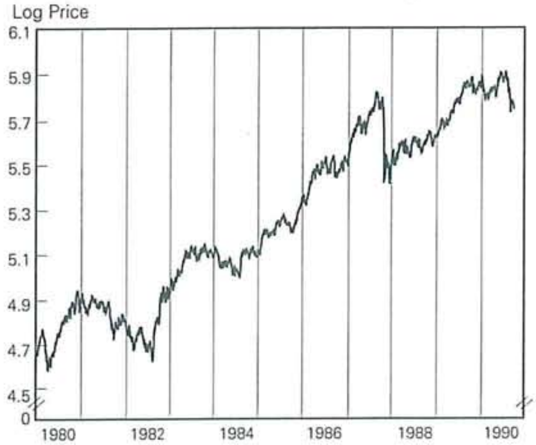

Source: Data Resources, Inc.

that the log of the daily closing value of the S&P500 corresponded to an Integrated Moving Average (IMA) model; we found that an IMA(1,5) model, using first differences and five daily moving average terms, was sufficient to eliminate autocorrelation. This equation, in which movements in the first difference of logP are described by a five-day moving average of white noise terms, is reported as equation 1 in Table A-1.

The moving average coefficients are statistically significant, and reveal the following pattern. If there is a downward shock ("crash") in the price change, the following adjustments will occur: on the following day the price change will be slightly more than the normal amount, after which it will increase at a slightly below-normal rate for three days, ending with a slightly above-normal increase on the fifth day. After five days, if no other shocks have occurred, the abnormal behavior is over. Given the short period of fluctuation, it is no surprise that longer-term data (such as monthly data) do not reveal departures from the EMH.

Thus, our data lead us to reject the random walk implications of the EMH for such short intervals of time as one day. However, the departure from a random walk is not of major economic significance; in short, it is not "bankable." First, the low coefficient of determination ($\overline{R}^2 = 0.01$) tells us that while the moving average terms are statistically significant, they explain only about 1 percent of the variation in the change in log price—there is a high probability that any potential profits from trading strategies based on the knowledge of the time series structure will be swamped by random variations. Second, even at its most profitable, the optimal trading strategy might not cover its costs: the optimal strategy would be to buy (sell) the day *after* a major fall (rise) in prices, then sell (buy back) the portfolio after it is held for four days. If we calculate the profits from doing this after a major price decline (defined as a decline greater than all but only 10 percent of price declines) the profits are only about 1.4 percent of the initial cost; this would not cover retail transactions costs, though it could cover institutional transaction costs.

Is There a Weekend Effect in the 1980s?

In order to examine the Weekend Effect during the 1980s we have re-estimated equation 1 of Table A-1 by adding two dummy variables: WKEND, which has a value of 1 if the trading day is a Monday, zero otherwise, and HOLIDAY, which has a value of 1 if the current trading day was preceded by a one-day holiday. The five-day moving average behavior reported in Table A-1, equation 1, is reproduced in equation 2. In addition, the WKEND dummy variable has a coefficient that is both negative and statistically significant. Thus, the daily data for the 1980s do contain a significant Weekend Effect. The HOLIDAY dummy is not statistically significant, indicating that one-day closings are not associated with systematic differences in price behavior.[17]

The January Effect in the 1980s

Equations 3 and 4 of Table A-1 incorporate dummy variables for the January Effect. Equation 3 includes JAN-

Table A-1
Tests of Random Walk Hypothesis
IMA(1,5) Model, Dependent Variable = Δ logP

Independent Variable	Equation 1	2	3	4
Constant	+.0004	+.0007	+.0007	+.0007
	(1.85)	(3.12)	(2.80)	(3.07)
MA(1)	+.0526	+.0546	+.0546	+.0546
	(2.74)	(2.84)	(2.84)	(2.84)
MA(2)	−.0370	−.0369	−.0369	−.0369
	(1.92)	(1.92)	(1.92)	(1.92)
MA(3)	−.0200	−.0193	−.0192	−.0193
	(1.04)	(1.00)	(.99)	(1.00)
MA(4)	−.0548	−.0544	−.0541	−.0544
	(2.85)	(2.83)	(2.81)	(2.83)
MA(5)	+.0575	+.0518	+.0518	+.0517
	(2.85)	(2.69)	(2.69)	(2.69)
WKEND	n.a.	−.0017	−.0017	−.0017
		(3.19)	(3.19)	(3.19)
HOLIDAY	n.a.	+.0007	+.0006	+.0006
		(.23)	(.19)	(.21)
JANUARY	n.a.	n.a.	+.0006	n.a.
			(.82)	
EARLYJAN	n.a.	n.a.	n.a.	+.0007
				(.19)
$\bar{R}^2$	.0089	.0120	.0119	.0117
SEE	.0109	.0109	.0108	.0108
Q(156)	169.7670	169.3580	169.1590	169.2850
p	.2179	.2248	.2281	.2260

Note: The sample period is January 2, 1980–September 21, 1990. Numbers in parentheses are absolute values of t-statistics. p is the probability level of the Q-statistic (156df).
n.a. = not applicable

UARY, defined as 1 for trading days in January and zero otherwise, and equation 4 includes EARLYJAN, a dummy variable defined as 1 in the first five trading days of January and zero otherwise. In neither case is the evidence consistent with a January Effect in the 1980s; neither coefficient is statistically significant. We cannot, however, conclude that the January Effect has disappeared; it is possible that it still exists for small firms, but that it does not exist for the large firms that are in the S&P500.

Our regressions do, however, allow us to conclude that the January Effect can no longer be used as a profitable strategy for a broad range of large firms listed on the major exchanges. The fact that it does not appear for large firms, which are the firms of primary interest to institutional investors, is consistent with the Efficient Markets Hypothesis: arbitrage by institutional investors prevents late-December selling pressure from affecting the share prices of large-capitalization firms, while small-cap stocks that have performed poorly do not have the attention of institutional investors and are oversold at year end.

Appendix 2: A Model of Noise Trading

The De Long-Shleifer-Summers-Waldmann model of noise trading results in a market price of common stock described by the following equation:

$$\text{(A2.1)} \quad P = 1 + \mu\rho^*/r + \mu(\rho - \rho^*)/(1 + r) - 2\gamma\mu^2\sigma^2/[r(1 + r)^2]$$

where μ is the proportion of investors who are noise traders, r is the real interest rate on riskless securities, and γ is the degree of absolute risk aversion, assumed to be the same for all investors. It is assumed that the degree of price misperception by noise traders, denoted by ρ, is normally distributed with mean ρ^* and variance σ^2. The price described by (A2.1) is the market price relative to the fundamental value, so P = 1.0 means that the price is equal to fundamental value.

The last three terms reflect the influence of noise trading; if $\mu = 0$, there are no noise traders and $P_t = 1$, that is, stock prices are determined by fundamentals alone. The second term reflects the influence of the mean amount of mispricing. If $\rho^* > 0$ ($\rho^* < 0$), noise traders are *normally* bullish (bearish), and to the extent that they are important in the market (measured by μ), this raises (reduces) stock prices; sophisticated traders will take opposing positions (such as reducing their stock holdings when prices are above fundamentals). The third term reflects the effect of *unusual* bullishness or bearishness of noise traders, measured by the random variable $(\rho - \rho^*)$; once again, the effect this has on stock prices depends on the relative numbers of noise traders.

Finally, the last term reflects the effect of uncertainty about the future degree of price misperception. The net effect of this is negative because the uncertainty imposed by noise traders will discourage investment by both sophisticated traders and noise traders, both of whom realize that price reversals can occur. This effect will be smaller, the less risk averse investors are (if $\gamma = 0$ nobody cares about risk so it will not affect prices) and the less important are noise traders (if $\mu = 0$ noise traders do not exist so they cannot affect prices). Given the values of γ and μ, the stock price will be negatively related to the size of the variance of noise trader misperceptions (σ^2).

If we measure the volatility of stock prices by the standard deviation conditional on information available in the previous period, for example, by $s = E_{t-1}(P_t - E_{t-1}P_t)^2$, we see that:

$$\text{(A2.2)} \quad s = \mu\sigma/(1 + r)$$

which says that volatility will be greater the larger the representation by noise traders, the larger the variability in their price misperceptions, and the lower the rate of interest.

If this were the end of the story, our noise trading model would predict that stock prices will deviate randomly around a normal value that is determined by the following equation:

(A2.3) $$P_t = 1 + \mu\rho^*/r - 2\gamma\mu^2\sigma^2/[r(1 + r)^2].$$

If ρ^* is zero, stocks will be chronically undervalued because all traders recognize that noise exists and that they might have to sell their assets at prices below their fundamental values. This "noise risk" cannot be eliminated by diversification, and results in a market price less than the fundamental value. Any deviations from this constant price level would be purely random, arising from temporary deviations of ρ from ρ^*. In short, the model would simulate volatility but not replicate the patterns of bull and bear markets we observe in the real world.

In order to introduce the possibility of realistic results from a noise trading model we need to add the notion of "fads and fashions," in which investor interactions create waves of investor sentiment. One way of introducing the possibility of contagious behavior is to assume that opinions follow a random walk shown by the autoregressive process:

(A2.4) $$\rho_t^* = \rho^*_{t-1} + \epsilon_t$$

where ϵ_t is white noise with mean $E(\epsilon_t) = 0$ and constant variance σ_ϵ^2. In this simple case, investors standing at a moment of time will forecast a constant value of ρ^* because zero is their optimal forecast of ϵ for every period. But the actual value of ρ^* will follow a path determined by equation (A2.4). Because stock prices are very sensitive to ρ^*, this can result in realistic stock market cycles.

In order to complete the simulation model, we assume $\mu = 0.05$, that is, noise traders represent 5 percent of the market. We also assume $\sigma_\epsilon = 0.005$, so that opinion (ρ^*) follows a simple random walk with the standard deviation of 0.005 in the steps.[18] Finally, we assume $\sigma = 0.01$, so that in 84 per cent of the trials (months) the random component in the degree of mispricing (for example, $\rho - \rho^*$) will be within ± 1 percent of the fundamental value of the stock.

Any movements in the simulated stock price must be due to either ϵ or $(\rho - \rho^*)$. We complete the simulation model by assuming that both ϵ and $(\rho - \rho^*)$ are normally distributed with zero means and the standard deviations assigned above. Using a random number generator to pick values of ϵ and $(\rho - \rho^*)$ for each of our "months," we can track the stock price. This experiment was done 600 times to simulate the stock price over a period of 50 years (600 months).

The results of one simulation are shown in Figure 6.

[1] One popular joke intended to discredit the EMH was about two economists walking along a street in Chicago (the bastion of the EMH): one observes a \$20 bill lying on the sidewalk and begins to bend down to get it, but the other tells him not to bother, for if the bill were really there it would already have been picked up!

[2] Technical analysis is the use of historical information on stock prices to forecast future stock prices. Perhaps the best example of technical analysis is the Dow Theory, which identifies specific patterns in stock prices and uses them to form forecasts.

[3] For example, a simple dividend valuation model with a constant growth rate for dividends (g) and a constant real discount rate (r) implies $E(P_t^*) = D/(r - g)$; knowing current dividends per share (D) and using estimates of r and g allows one to test the EMH by regressing P_t on a variable defined as $D/(r - g)$ and any other variable in the information set Ω. Adjustment of the regression method for measurement error will, of course, be necessary.

[4] For expositional convenience we are ignoring the payment of dividends. If dividends are included, the valuation equation must be modified to read

$$E(P_{t+1}^* + D_{t+1}) = (1 + r)E(P_t^* \mid \Omega_t).$$

Thus, dividends are assumed to be reinvested.

[5] The companies are Adams Express (NYSE), Baker Fentress (OTC), General American Investors (NYSE), Lehman Corporation (NYSE), Source Capital (NYSE), Tri-Continental (NYSE), and Niagara Shares (NYSE).

[6] Of course, an open-end fund with unrealized appreciation exposes the investor to the same tax liability. But the liability cannot affect the open-end fund's share price because it is always equal to the net asset value; the only effect is to induce investors with high tax rates to prefer closed-end funds over open-end funds.

[7] If the Value Line ranking system is of no use, the probability that Rank 1 stocks will outperform Rank 3 stocks is 0.50. Under this null hypothesis the probability of 5 or fewer "failures" in 25 years is only 0.00046. This provides a strong reason to believe that the ranking system does have merit.

[8] The keen-eyed reader will observe that one can construct an example from Table 1 that shows the fundamental value varying less than the market price. For example, if the same state occurred in each year, the fundamental price would not vary at all! This is an artifact of the example: with a large number of possible states and a large number of time periods, the optimal forecast *must* vary by no more than the fundamental price.

[9] From (5) we can see that $E_t(P_{t+1}) = (1 + r)P_t - E_t(D_{t+1})$. If no dividends are expected, the price will be expected to grow at the rate r. If dividends are paid, the rate of increase in price will be less than r, but the rational bubble must grow at the rate r.

[10] Thus, the imposition of the transversality condition $\lim_{t\to\infty} [(1 + r)^{-t}P_t] = 0$ is sufficient to crush a nascent bubble.

[11] Hence a necessary condition for existence of a speculative bubble is a finite time horizon. This is not, however, sufficient. Tirole (1982) has shown that even with a finite time horizon, a speculative bubble cannot exist if expectations are rational, that is, if investors' forecasts are optimal. Hence, bubbles require both finite horizons and non-optimal forecasting. Stated differently, rational bubbles require inefficient markets.

[12] Attempts to identify time-variation in the risk premium were not successful, so the value of θ was set at the sample average for $(R_t - r_t)$; this was 0.0070 per month or 0.0838 per year.

[13] For the 755 months in the period 1926:2 to 1988:12, the change in log bubble over the next 12 months had a mean of −0.0208 and a standard deviation of 0.2226. Hence, a "one-standard-deviation crash" is defined as a 12-month change in the logarithm of the bubble by an amount of −0.2434 or less. A two-standard-deviation crash is a change of −0.4660 or less.

[14] The logit model assumes that the probability of an event (π) is a logistic function of the following form:
$\pi_t = 1/\{1 + \exp[-(\alpha + \beta xt)]\}$.

[15] The logic of the Poterba-Summers test is simple. Suppose that the one-period rate of return on stocks is approximated by the change in the logarithm of the price. Suppose further that—as many financial studies assume—the change in the logarithm can be represented by a constant plus a random error term, so $\log P_{t+1} - \log P_t = \mu + \epsilon_{t+1}$. Then the average return over N periods is approximately $\log P_{t+N} - \log P_t = N\mu + (\epsilon_{t+1} + \epsilon_{t+2} \ldots + \epsilon_{t+N})$. If the EMH is correct, the ϵ's are identically and independently distributed. Denoting the variance of ϵ as σ^2, the variance of the N period return is $VAR(\log P_{t+N} - \log P_t) = N\sigma^2$: the variance of returns is proportional to the period over which the returns are experienced. If, as Summers and Poterba conclude, the variance increases less than in proportion to the period, the return on stocks is mean-reverting.

[16] The term "winner's curse" is used here in a different way than it is used in discussing the effects of mergers and acquisitions. In that context, the winner's curse is the tendency of those who outbid others to pay too high a price for the acquired firm.

[17] We also found that three-day weekends were no different from two-day weekends; the extra day makes no difference, just as a one-day holiday makes no difference.

[18] The other parameter in the model (γ) really plays no important role in the simulation. We set it at $\gamma = 0.10$.

References

Black, Fischer. 1971. "Yes, Virginia, There Is Hope: Tests of the Value Line Ranking System." Center for Research in Security Prices, Graduate School of Business, University of Chicago, May.

———. 1986. "Noise." *Journal of Finance*, July, pp. 529–543.

Cootner, Paul H. 1964. "Stock Prices: Random vs. Systematic Changes." In Paul H. Cootner, ed. *The Random Character of Stock Market Prices*. Cambridge, MA: The MIT Press.

Cross, Frank. 1973. "The Behavior of Stock Prices on Fridays and Mondays." *Financial Analysts Journal*, November/December, pp. 67–69.

De Bondt, Werner and Richard H. Thaler. 1985. "Does the Stock Market Overreact?" *Journal of Finance*, July, pp. 793–805.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy*, August, pp. 703–738.

Fama, Eugene F. and Kenneth R. French. 1988. "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy*, April, pp. 246–273.

Fortune, Peter. 1989. "An Assessment of Financial Market Volatility: Bills, Bonds and Stocks." *New England Economic Review*, November/December, pp. 13–28.

Friedman, Milton. 1953. "The Case for Flexible Exchange Rates." In Friedman, M., *Essays in Positive Economics*, pp. 157–302. Chicago, IL: University of Chicago Press.

Holloway, Clark. 1981. "A Note on Testing an Aggressive Investment Strategy Using Value Line Ranks." *Journal of Finance*, June, pp. 711–719.

Ibbotson Associates. 1990. *Stocks, Bonds, Bills and Inflation: 1990 Yearbook*. Chicago, IL: Ibbotson Associates.

Keim, Donald. 1983. "Size-Related Anomalies and Stock Return Seasonality." *Journal of Financial Economics*, June, pp. 13–22.

———. 1986. "The CAPM and Equity Return Irregularities." *Financial Analysts Journal*, May/June, pp. 19–34.

Keynes, John Maynard. 1964. *The General Theory of Employment, Interest and Money*. New York, NY: Harcourt, Brace and World.

Kleidon, Allan W. 1986. "Variance Bounds Tests and Stock Price Valuation Models." *Journal of Political Economy*, October, pp. 953–1001.

Lee, Cheng F. 1987. "Value Line Investment Survey Rank Changes and Beta Coefficients." *Financial Analysts Journal*, September/October, pp. 70–72.

Lee, Charles, Andrei Shleifer and Richard Thaler. 1990a. "Investor Sentiment and the Closed-End Fund Puzzle." Johnson Graduate School of Business, Cornell University, January.

———. 1990b. "Closed-End Mutual Funds." *Journal of Economic Perspectives*, Fall, pp. 153–164.

LeRoy, Stephen F. and Richard D. Porter. 1981. "The Present-Value Relation: Tests Based on Implied Variance Bounds." *Econometrica*, May, pp. 555–574.

Lucas, Robert E. 1978. "Asset Prices in an Exchange Economy." *Econometrica*, November, pp. 1429–1445.

Malkiel, Burton. 1977. "The Valuation of Closed-End Investment Company Shares." *Journal of Finance*, June, pp. 847–859.

Markowitz, Harry M. 1959. *Portfolio Selection: Efficient Diversification of Investments*. New York, NY: John Wiley and Sons.

Marsh, Terry A. and Robert C. Merton. 1986. "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Prices." *The American Economic Review*, June, pp. 483–498.

Poterba, James M. and Lawrence H. Summers. 1988. "Mean Reversion in Stock Prices: Evidence and Implications." *Journal of Financial Economics*, February, pp. 27–59.

Reinganum, Marc. 1983. "The Anomalous Stock Market Behavior of Small Firms in January: Empirical Tests for Tax-Selling Effects." *Journal of Financial Economics*, June, pp. 89–104.

Samuelson, Paul. 1965. "Proof That Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review*, Spring, pp. 41–49.

Schwert, G. William. 1989. "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance*, December, pp. 1115–1153.

Sharpe, William. 1964. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, September, pp. 425–442.

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

Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *The American Economic Review*, June, pp. 421–436.

Tirole, Jean. 1982. "On the Possibility of Speculation Under Rational Expectations." *Econometrica*, September, pp. 1163–1181.

Weisenbarger Investment Companies Service. 1988. *Investment Companies*, 48th Annual Edition. New York, NY: Weisenbarger Financial Services.