More Unsettling Evidence on the Perfect Markets Hypothesis

David N. DeJong and Charles H. Whiteman

If you're so smart, why aren't you rich? While most people have no answer for this question, economists have a ready-made response—the perfect (or efficient, or rational) markets hypothesis. Loosely speaking, the hypothesis states that asset returns (on stocks, bonds, real estate, and the like) are not forecastable: The market aggregates information efficiently, or market participants use that information rationally, or the market is perfectly void of unexploited profit opportunities. If the hypothesis holds, trying to get rich in the stock market amounts to what Burton G. Malkiel, in the title of his famous book, called *A Random Walk Down Wall Street* (1973).

If the perfect markets hypothesis is not valid, the only defensible answer to the question posed above is "I'm not that smart." Ego preservation might be reason enough, but economists cling to the hypothesis for other reasons, primarily because the alternatives are unpalatable. For one thing, it is difficult to fathom the nature of the imagined economic profit opportunities. These would involve something other than the usual rewards to productive activity. Instead, they would be easily identifiable, costless, and riskless ways to get rich. Furthermore, even if these opportunities were to exist, how could they persist? In such a scenario there would, in effect, be dollar bills lying all over the main street and no one stooping to pick them up.

The possibility of persistent unexploited profit opportunities in any market poses a theoretical problem for all of economics. If economists are wrong about such a simple premise, why should more subtle predictions (for example, that a tax cut will spur economic activity) be trusted?

David N. DeJong is an associate professor of economics at the University of Pittsburgh. Charles H. Whiteman is a professor of economics at the University of Iowa and a visiting scholar in the Atlanta Fed's research department.

An Introduction to the Perfect Markets Controversy

The crack in the foundation of economics may already have appeared. The empirical validity of the perfect markets hypothesis, once thought to be settled in the 1960s and 1970s by regression tests of one of its implications, has recently been challenged in several studies. Robert J. Shiller (1981b) and Stephen F. LeRoy and Richard D. Porter (1981), for example, purport to show that aggregate annual stock prices are too volatile to reflect the underlying, fundamental values predicted by the hypothesis.

However, economists have been quick to point out that Shiller's evidence of "excess volatility" is not truly evidence or is not evidence of imperfection in the financial markets. Indeed, one counterattack, mounted by Marjorie A. Flavin (1983), LeRoy and C.J. LaCivita (1981), and Ronald W. Michener (1982), argues that the volatility tests of Shiller and others suffer from sampling and specification error and that the properties of the data highlighted in the studies are in fact consistent with the hypothesis. The work of Terry A. Marsh and Robert C. Merton (1986, 1987) and Allan W. Kleidon (1986a, 1986b) yielded a second counterattack, arguing that the perfect markets hypothesis can be resurrected by simply modifying an assumption Shiller made regarding the temporal stability of aggregate dividends.

The temporal stability issue arises in the perfect markets context because dividends grow over time. Chart 1 illustrates this fact for real (corrected for inflation) aggregate dividends associated with the Dow Jones Industrials, the New York Stock Exchange, and the Standard and Poor's 500 for the years from 1928 to 1978.[1] The question for the analyst concerns what manner of growth this is and how it influences share prices. Are dividends well described by fluctuations around a smooth trend? (Are they "trend-stationary"?) Or are dividends characterized by noisy "ratchet" growth—that is, does a random additional payout each year add to ("integrate with") what was paid out the previous year? (Are they "integrated"?)

The distinction between the two dividend assumptions involves how one revises one's forecasts of future dividends in light of surprisingly high or low current dividends. If dividends are trend-stationary, one would expect that today's dividend surprise will not persist but that, at some point in the not-too-distant future, dividends will revert to values predicted by the historical trend; a trend-stationary process never strays very far from its normal course of growth. On the other hand, an integrated process behaves like a drunken sailor on a random walk: regardless of the intended path, at any given point it is a coin toss as to whether the sailor will step right or left. If dividends are integrated, each new surprise is built into the level of dividends so that they may tend to grow, but they do not trend.

For interpreting the evidence on the perfect markets hypothesis the distinction matters because the two types of processes carry different implications for the variability of dividends. If dividends are integrated, actual dividend variability might be much less than what could have occurred but did not (to continue the metaphor—by sheer luck the drunken sailor might have stayed on the sidewalk), whereas under trend-stationarity, observed dividend variability is a good indicator of potential variability.

Shiller followed common practice in assuming trend-stationarity and found evidence against the perfect markets hypothesis. Marsh-Merton and Kleidon showed that under the integration assumption, investors might have been wary of potential dividend fluctuations that did not come to pass—for example, they might have held fears of a Soviet invasion of Poland in the 1980s, which of course were never realized. In this case, the interpretation of Shiller's evidence is exactly reversed, and the data thus appear to support the perfect markets hypothesis.

To this twofold counterattack has come a fourfold response. First, the sampling and specification considerations raised by Flavin, Michener, and LeRoy-LaCivita could account for some excess volatility but not for the amount Shiller had found. Second, N. Gregory Mankiw, David Romer, and Matthew D. Shapiro (1985) and Kenneth D. West (1988) developed excess volatility-type tests that are valid under the integration assumption about dividends; these studies found evidence against the perfect markets hypothesis, though it was much less dramatic than the results of Shiller and LeRoy-Porter. Third, John Y. Campbell and Shiller (1987) explicitly adopted the integration assumption about dividends and showed that restrictions implied by the perfect markets hypothesis were still generally rejected by the data. Finally, a new attack on the perfect markets hypothesis, one that does not hinge on the temporal stability issue, has been provided by the mean-reversion studies of Lawrence H. Summers (1986), James M. Poterba and Summers (1988), Eugene F. Fama and Kenneth R. French (1988), and Andrew W. Lo and A.C. MacKinlay (1988). These studies search for evidence that low returns are likely to be followed by higher ones—returns "revert" to their mean. If returns

Chart 1
Real Aggregate Dividends on Stocks Traded, 1928–78

Dow Jones Industrials

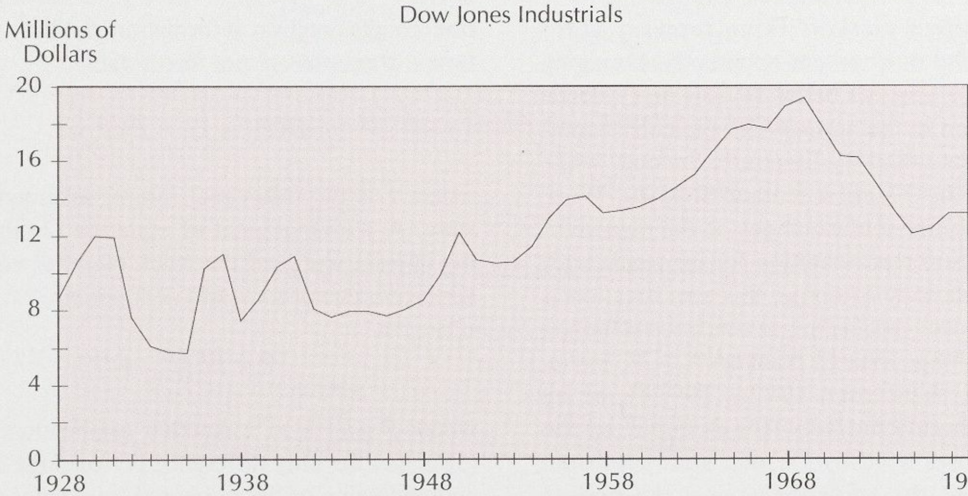

New York Stock Exchange

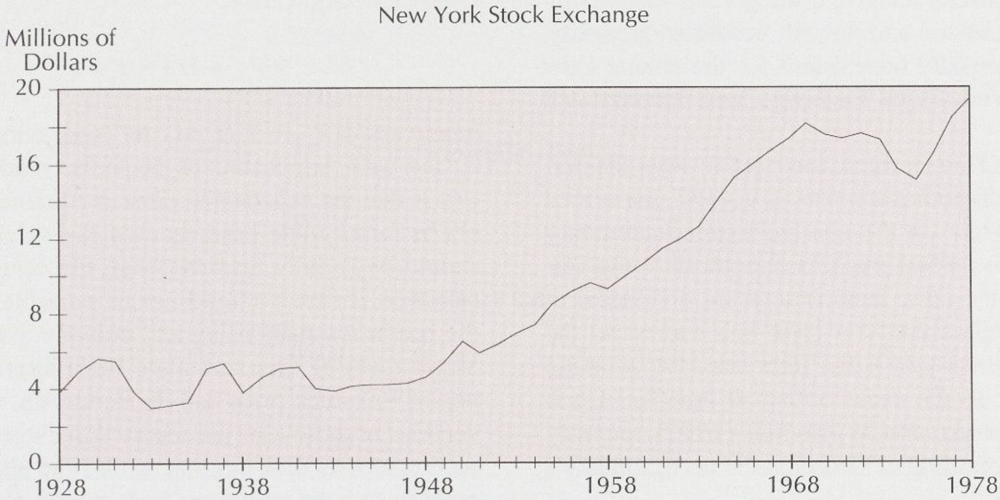

Standard and Poor's 500

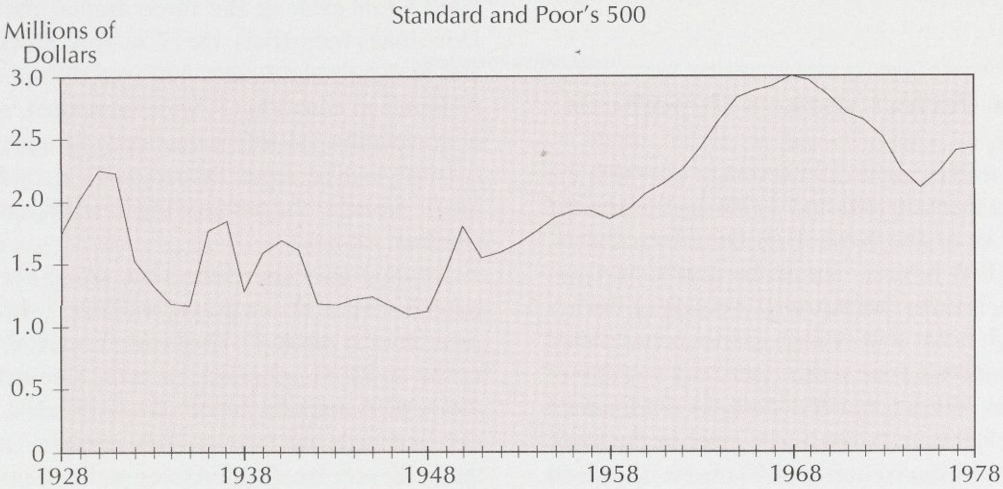

are mean-reverting, they are forecastable, and the perfect markets hypothesis is violated. Using aggregate monthly data, several of the studies did in fact find mean reversion and thus provided seemingly robust evidence against the perfect markets hypothesis.

Still, the perfect markets issue remains unresolved. Taking the most recent apparently damaging evidence first, the analysis below reveals no evidence of mean reversion in the annual data typically examined in the excess-volatility literature. Second, while the validity of the Mankiw-Romer-Shapiro, West, and Campbell-Shiller studies hinges on the plausibility of the integration representation for the stock price and dividend data they examine, this specification is at minimum suspect (according to evidence obtained by David N. DeJong, John C. Nankervis, N.E. Savin, and Charles H. Whiteman 1992a, 1992b), or by some measures highly unlikely (for example, by the Bayesian measures used by DeJong and Whiteman 1991). Moreover, the evidence against the perfect markets hypothesis under the integration assumption is not very dramatic and, as will be shown presently, is in fact practically nonexistent for the annual Dow Jones, New York Stock Exchange, and Standard and Poor's data.

As Herbert Hoover might have put it, what the perfect markets hypothesis needs is a good one-armed economist. While on the one hand (trend-stationary dividends) the perfect markets hypothesis looks implausible, on the other hand (integrated dividends), it is not very implausible. Yet there is a way to let the data decide which arm is the right one—by adopting an alternative to the usual statistical practice. Upon doing so, DeJong and Whiteman (1992) recently found more unsettling evidence on the perfect markets hypothesis.

Tests of the Perfect Markets Hypothesis

Numerous studies in the 1960s (many of which are summarized in Fama 1970) and 1970s had tested the "weak" version of the perfect markets hypothesis, which posits that returns cannot be predicted from their own past, and the "semistrong" version of the hypothesis, which holds that returns cannot be predicted using other economic time series. One way to think of these tests is as follows. Let P_t denote the ex-dividend price of a stock at time t (that is, the price of the stock just after the most recent dividend has been paid), and let the (end of) time t dividend be D_t; assume that P_t is known at time t, but that D_t is unknown. Then the expected one-period return to buying the stock (for price P_t), holding it long enough to receive the dividend payment D_t, and selling it (one period later for price P_{t+1}) is $E_t(P_{t+1} + D_t - P_t)/P_t$, where E_t denotes the expectation conditioned on information publicly available at time t. If returns are not forecastable,

$$E_t(P_{t+1} + D_t - P_t)/P_t = r, \tag{1}$$

where r is the mean real return, assumed to be constant. A rearrangement of equation (1) indicates that the current worth of the stock ought to equal the discounted value of its worth next period:

$$P_t = \beta E_t(P_{t+1} + D_t), \tag{2}$$

where $\beta = (1 + r)^{-1}$. Another implication of equation (1) is that the only difference between the realized discounted value of the stock and its current price results from expectation error:

$$\beta(P_{t+1} + D_t) - P_t = \epsilon_t, \tag{3}$$

where $\epsilon_t = \beta(P_{t+1} - E_t P_{t+1})$ is the expectation error.

The weak implication of the perfect markets hypothesis is that the expectation error in equation (3) should not be forecastable using its own past values (that is, it should be serially uncorrelated); the semistrong implication is that it should not be possible to forecast the return residuals using any publicly available data. Many tests of this sort have been applied to (and passed by) stock price and dividend data; the authors' versions of these tests are reported in Table 1. The first section reports statistics that result from testing the hypothesis that the return residuals are not serially correlated.[2] For each of the three annual data sets—the Dow Jones Industrials, the New York Stock Exchange, and Standard and Poor's 500 (see the appendix)—the residuals in equation (3) were computed, and the serial correlation statistic calculated. In no case is there reason to suspect that the residuals are serially correlated; weakly, the perfect markets hypothesis is not rejected.

Regression tests of the semistrong version of the hypothesis are also reported in Table 1. In each of the data sets the return residuals are forecastable by prices or the dividend-price ratio or both. Semistrongly, then, the perfect markets hypothesis is rejected, but the return residuals are not very forecastable—in each case the R^2 is very small. Moreover, as Shiller notes, "Such small correlations detected in long historical data are

Table 1
Tests of the Perfect Markets Hypothesis

	Dow Jones Industrials 1928–78	New York Stock Exchange 1926–81	Standard and Poor's 1871–1985
Are Return Residuals Serially Correlated? (Weak Test)			
Q-statistic	Q(21) = 11.10	Q(30) = 25.27	Q(21) = 12.29
Significance Level	0.96	0.71	0.93
Are Return Residuals Forecastable? (Semistrong Test)			
Regression on Lagged Price			
Coefficient	−0.241	−0.174	−0.141
t-statistic	−2.70	−2.36	−3.27
R Squared	0.13	0.10	0.09
Regression on Lagged Dividend-Price Ratio			
Coefficient	4252.568	0.148E+11	333.027
t-statistic	1.55	1.82	2.35
R Squared	0.05	0.06	0.05
Volatility Tests			
$\sigma(p)$	352.18	0.21E+08	52.37
$\sigma(p^*)$	116.81	0.55E+07	13.40
Variance Ratio Tests[a]			
$pVR(k)$	10.03	10.30	14.40
k	5	6	11
q	13	12	14
P-value	0.69	0.59	0.42

[a]Under the null hypothesis of no mean reversion, $pVR(k)$ is distributed as $\chi^2(q)$; k denotes the horizon over which returns are analyzed.

of questionable relevance to modern conditions and of minor interest given possible errors in data from remote times in history" (1981a, 294).

Weak though it may be, the forecastability evidence is evidence against the perfect markets hypothesis. Shiller (1981a, 1981b) argues that while this evidence might not convince anyone of anything, his volatility tests ought to. These tests, he suggests, have advantages over regression tests of "greater power in certain circumstances of robustness to data errors such as misalignment and of simplicity and understandability" (1981a, 292).

Shiller's tests are based on three assumptions:

- Stock prices reflect investor beliefs that are rational expectations of future dividends.
- The real expected rate of return on the stock market, r, is constant.
- Aggregate real dividends can be described by a trend-stationary process.

To derive the relationship on which the test is based, it is useful to begin by defining the ex post rational (or perfect markets) price per share, P. It is obtained by noting that equation (2) describes P_t in terms of D_t and P_{t+1}, which in turn must be describable by D_{t+1} and P_{t+2}, which must then be describable by D_{t+2} and P_{t+3}, and so on, and carrying out the "forward substitution," whereupon the current price can be expressed in terms of the entire future course of dividends:

$$P_t^* = \beta D_t + \beta^2 D_{t+1} + \beta^3 D_{t+2} + \ldots \equiv \Sigma_{k=0}^{\infty} \beta^{k+1} D_{t+k}. \quad (4)$$

In words, P is the present value of actual subsequent dividends. From the first assumption actual stock prices are ex ante rational—that is, they equal the current expectation of P,

$$P_t = E_t(P_t^*), \quad (5)$$

so that actual prices represent the expected present value of subsequent dividends. Equation (5) indicates that P_t is the optimal forecast of P_t^*; hence the forecast error u_t ($= P_t^* - P_t$) cannot be predicted using the actual price P_t. By implication, P^* must vary more than P:

$$\begin{aligned} variance(P^*) &= variance(P) \\ &+ variance(u) \geq variance(P). \end{aligned} \quad (6)$$

Chart 2 shows estimates of P and P^* for the three data sets, and sample estimates of *variance*(P^*) and *variance*(P) are reported in Table 1.[3] The calculations replicate Shiller's results for the Dow Jones Industrials and the Standard and Poor's and produce similar results in the New York Stock Exchange figures. Note that contrary to the prediction of the perfect markets hypothesis, the ex post rational stock price P^* is much less volatile than the price P itself: stock prices appear too variable to be accounted for by subsequent changes in dividends.

There are a number of reasons to view these calculations with some suspicion. First, the assumptions are questionable. On the constancy of the real rate, Michener (1982) and LeRoy and LaCivita (1981) have shown that in a theoretical model economy, risk-averse investors may behave in such a way that the real rate of return on the stock market fluctuates over time. This means that the r in equation (1) would not be the constant Shiller assumed. Further, if the return varies, the variance of P need not be less than the variance of P^*. But the amount by which the P^* bound could be exceeded was not as large as what Shiller had found. Indeed, Shiller (1981a) argued that for real rate variability to induce enough excess volatility to account for the P^* variance bound violations he found, the real rate would have had to vary from −8.16 to 17.27 percent, a range too large to be believable.

A second problem, discussed by Flavin (1983) and Kleidon (1986a), is that tests based on equation (6) are biased in favor of the inference of excess volatility. Their argument involves noting that the sample variances of P_t and P_t^* are biased estimates of the actual variances (because they are computed using deviations from estimated means rather than the true underlying means). Further, the greater the temporal dependence of a series, the greater the downward bias in its sample variance. Because P^* is a highly persistent series, its sample variance exhibits greater downward bias than that of P. However, Kleidon conceded that this bias does not seem to explain the dramatic violations Shiller reports.

Finally, a number of authors have attacked Shiller's third assumption. For example, Marsh and Merton argued that "his variance bound test results might be better interpreted as an impressive rejection of his model of the dividend process than as a rejection of stock market rationality" (1986, 485). To establish this point, Marsh and Merton maintained the first and second assumptions but used a different model to characterize how dividends, and hence rational prices, are determined. In so doing, they managed to reverse the inequality in equation (6) and establish a variance bound test that directly contradicts Shiller's.

Marsh and Merton began by assuming that managers who set dividend policies for business firms dislike frequent dividend changes, yet they try to achieve a constant dividend-price ratio.[4] The implication is that the managers would make the current dividend a weighted average of historical stock prices. But P^* is itself a weighted average of actual dividends, and these two facts together imply that P^* is a weighted average of actual prices. In fact, this simultaneity (dividends are set on the basis of past prices, and prices are the expected discounted value of future dividends) causes both prices and dividends to be integrated. As a result, Shiller's variance inequality should be reversed; intuitively, because an average is less volatile than the series from which it is computed, P^* should be "smoother" than P.[5]

Marsh and Merton (1987) provided empirical support for the dividend-smoothing idea by estimating what can be interpreted as the managerial equation for setting aggregate dividends as a function of stock prices. This equation, it turns out, explains much more of the variability of aggregate dividends than does the trend-stationary model attributed to Shiller.

Kleidon (1986b) has also argued that the apparent integration of stock prices may account for Shiller's findings. He failed to reject the null hypothesis of

Chart 2
Actual and Perfect Markets Prices for Portfolios of Stocks, 1928–78

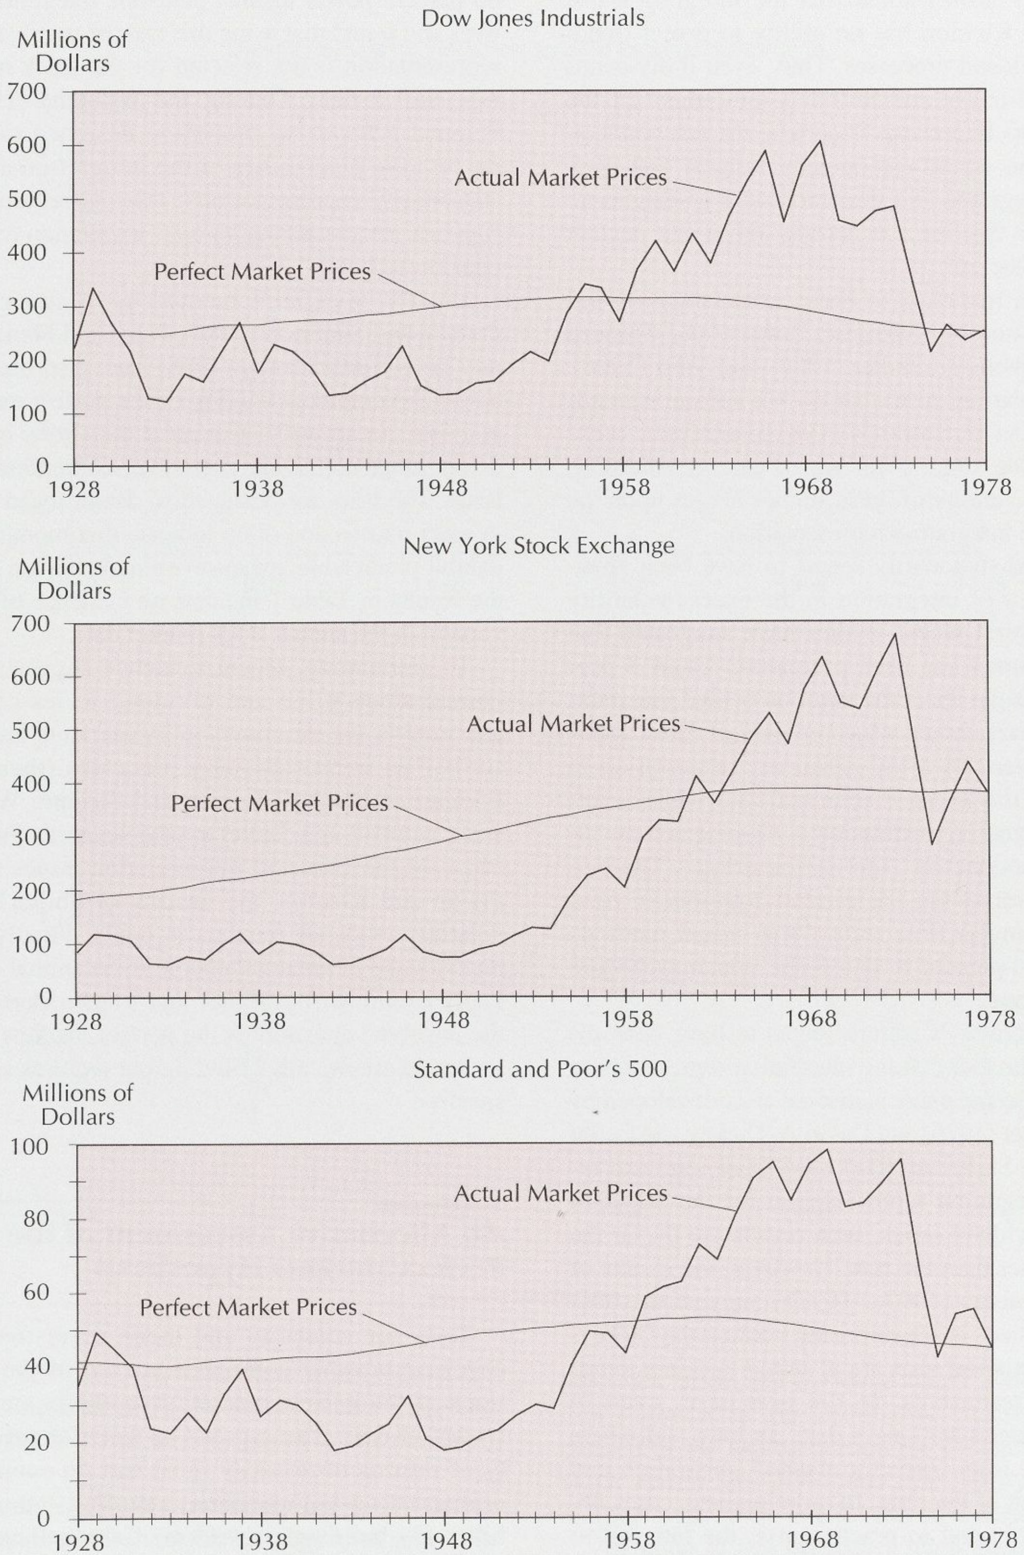

integration for the Standard and Poor's price and dividend series and used simulations to show that rationally determined prices can appear "excessively" volatile if dividends are integrated. However, West (1987) and Pierre Perron (1988) noted that the integration test employed by Kleidon has no ability to detect trend-stationary dividend processes. Thus, even if dividends are trend-stationary and Shiller's indictment of the perfect markets hypothesis is correct, the test conducted by Kleidon would erroneously suggest that dividends are integrated and thus mislead a researcher to conclude that Shiller was wrong about the perfect markets hypothesis.

Perhaps in response to these criticisms, several researchers, notably Mankiw, Romer, and Shapiro (1985), Campbell and Shiller (1987), and West (1988), developed volatility tests that do not subsume trend-stationarity. As in Shiller's (1981b) analysis, these studies provide some evidence of excess volatility. However, the validity of these studies hinges upon the validity of the integration representation.

While trend-stationarity seems to have been abandoned in favor of integration in the excess-volatility literature, several recent studies have suggested that this abandonment has been premature. Chart 3 provides some insight about the intuition behind this statement. The chart shows what remains of Dow Jones dividends when the trend-stationary assumption is adopted and the trend is removed and what is left when the integration assumption is adopted and the integration is undone by "first-differencing." The two methods for removing the growth component yield remarkably similar time series, and casual methods would be hard-pressed to determine which procedure is the correct one.

Trend-stationarity's stature seems to have been diminished by the fact that the integration representation is not often rejected using statistical tests developed by Wayne A. Fuller (1976) and David A. Dickey and Fuller (1981). While Shiller (1981a) used these tests and rejected the integration representation for the Standard and Poor's dividend series, such results are by far the exception rather than the rule. However, the results of DeJong and others (1992a, 1992b) suggest that this dearth of rejections is not surprising because the Dickey-Fuller type statistical tests are ill suited to detect trend-stationary alternatives. In the best case, even in experimental trend-stationary data, DeJong and others found that the tests correctly rejected integration less than 50 percent of the time. In more realistic cases designed to correspond to practical use, the rate fell to less than 20 percent. Hence, a failure to find convincing evidence against the integration hypothesis does not provide convincing evidence in its favor. Moreover, DeJong and others (1992a) developed a test of a trend-stationary representation that has approximately 60 percent power against plausible integrated alternatives and found that using this test the trend-stationarity representation is not rejected for the stock market series. Still further, DeJong and Whiteman (1991) developed a Bayesian procedure designed to evaluate the relative plausibility of the integration and trend-stationarity representations; this procedure strongly supports trend-stationarity over integration for the data considered here.

Finally, tests implemented by Summers (1986), Poterba and Summers (1988), Fama and French (1988), and Lo and MacKinlay (1988) have also been used to assess the empirical validity of the perfect markets hypothesis. An attractive feature of these tests is that they do not hinge on the integration versus trend-stationarity issue. The tests are designed to detect mean reversion in asset returns and often indicate that monthly returns exhibit predictable, mean-reverting behavior. However, the results in Table 1 indicate no evidence of mean reversion in the annual stock market data.

To summarize, recent evidence suggests that the annual stock price and dividend series commonly analyzed in the excess-volatility literature are well described as trend-stationary processes (therefore the findings of Mankiw, Romer, and Shapiro, West 1988, and Campbell and Shiller appear suspect). Further, because of the bias and interpretation issues raised by Flavin and Kleidon, the validity of Shiller's excess-volatility evidence remains in doubt. Finally, there is no evidence of mean reversion in the annual data originally examined by Shiller and Marsh-Merton. Thus, the unsolved question: is the perfect markets hypothesis sound empirically? Finding out requires a new perspective.

An Alternative Assessment of the Perfect Markets Hypothesis

Tests for volatility and mean reversion involve checking whether some, but not all, of the implications of the perfect markets hypothesis are violated by actual data. The regression tests for predictability of returns necessarily must rest on economically uninteresting assumptions regarding temporal stability. By building on work by Campbell and Shiller (1987) and DeJong and Whiteman (1991), DeJong and

Chart 3
Deviations of Real Dow Jones Dividends from Trend, 1928–78

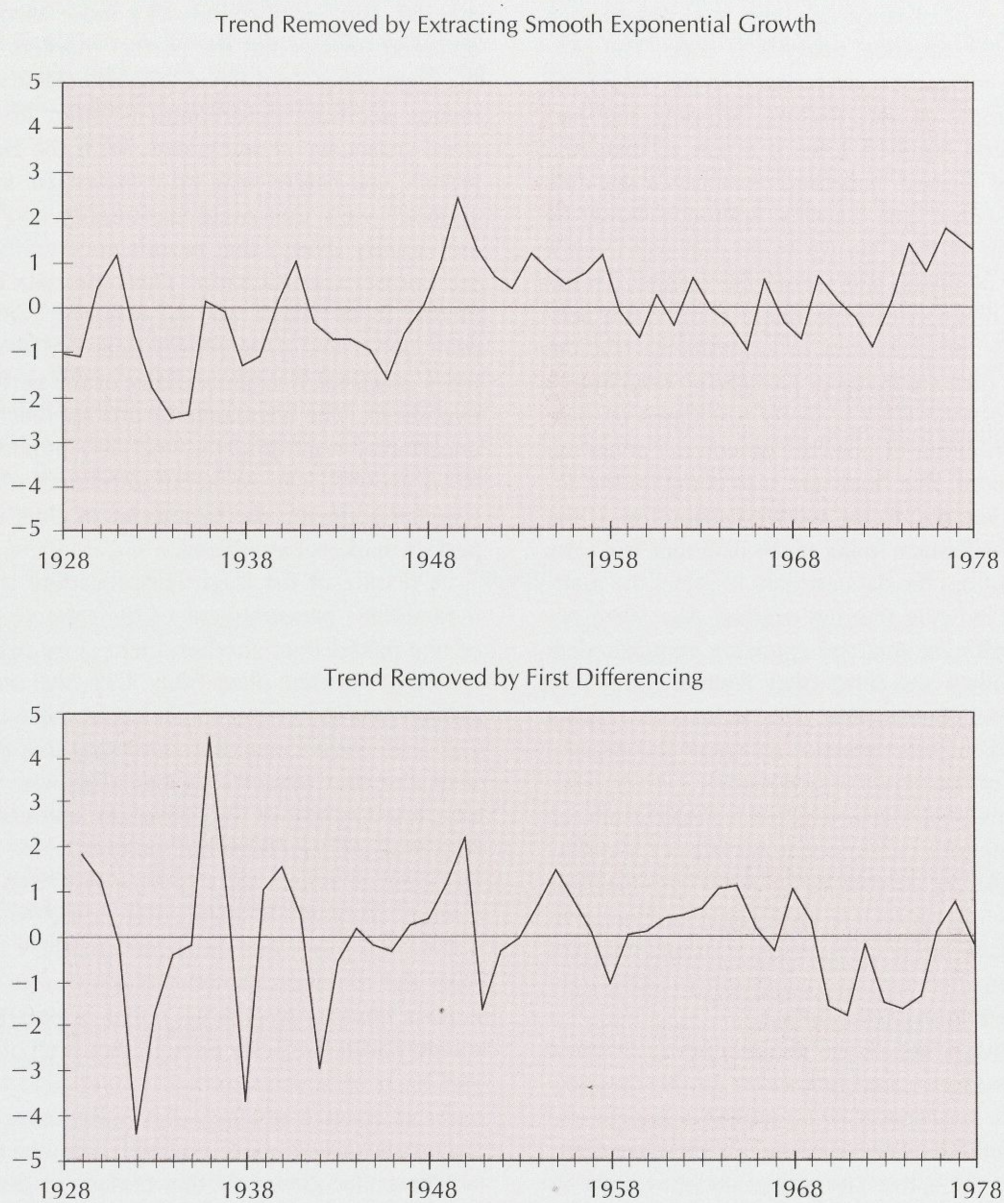

Whiteman (1992) have developed a procedure for testing all implications of the theory without needing to decide beforehand how to account for dividend growth.

The Campbell-Shiller procedure involves starting with an assumption about dividend growth and with a general, unrestricted statistical representation for the stock price and dividend data.[6] This representation summarizes the information in the data. In the context of the representation the connection between stock prices and dividends implied by the perfect markets hypothesis in equation (2) can be tested.

Adoption of this procedure highlights the sensitivity of inferences regarding the hypothesis to the dividend growth assumption. DeJong and Whiteman (1992, Table 1) report that the perfect markets hypothesis is resoundingly rejected under the trend-stationarity assumption and not rejected under the integration assumption.

That this sort of situation can arise is a consequence of how classical statistical analysis is done. The general approach is easiest to understand in the context of a simple example of dice-tossing. Suppose one had observed several hundred rolls of a pair of dice (the "sample") and wished to determine whether the dice were fair or loaded. The classical approach would be to locate a pair of dice known to be fair, roll them a large number of times, and record the results. (In actual practice, this sort of experiment is often not necessary because the properties of experiments like the rolling of fair dice are given in probability textbooks.) The results describe what is likely to happen on dice rolls (the distribution of possible outcomes) given the "fair-dice" model. The final step is to compare the actual data (the sample) to the model distribution. If the actual data distribution looks much different from the model distribution, the data are said to reject the model and would indicate that the original dice were not fair. It is significant that the approach requires data sufficient to sway a skeptic away from a belief in a particular model. Put another way, it involves asking whether the data look unusual or atypical from the point of view of the hypothesized model.

An alternative approach, which developed from the work of eighteenth-century statistician Sir Thomas Bayes, is to take the data rather than the hypothesis as given and ask what model is most likely to have generated those data. In the dice-tossing example, the data would be used to determine what sort of dice were most likely used in the actual tosses.

The contrast between the procedures is in what each views as known and unknown. In the classical procedure, one behaves as if the model is known and the data are unknown and random; in the alternative procedure, one takes the data as known (they have already been generated) and the model as unknown (because it is the object of study.)

Both approaches have strong foundations in probability theory, but they stem from different perspectives on the meaning of probability. The classical view associates probability with relative frequency: if, in a large number of identical experiments, a particular outcome (for example, the effectiveness of a drug) appears 75 percent of the time, the probability of that outcome is said to be 75 percent. The Bayesian view associates probability with the analyst's "degree of belief": the analyst would say that he or she was 75 percent certain that the drug works.

For many empirical questions, the two approaches yield identical answers, and no arcane issues involving the fundamental meaning of probability need be addressed. But for questions like those associated with dividend growth and the perfect markets hypothesis, the approaches do differ. Using the classical approach means deciding beforehand whether dividends are trend-stationary or integrated. With the Bayesian approach, one begins with prior views (for instance, the agnostic view that trend-stationarity and integration are equally likely) that permit uncertainty regarding the correct specification. Then the data are used to modify these beliefs, and the resulting combination of prior belief and data information can be used to address issues such as the validity of the perfect markets hypothesis. The advantage of this approach is that the uncertainty regarding the ancillary hypothesis (dividend growth) can be accommodated while the researcher addresses the hypothesis of chief interest (the perfect markets hypothesis).

A feature of the Bayesian procedure is that while it facilitates measurement of the relative plausibility of one model over another, there is no easily-defined notion of absolute plausibility. Classical and Bayesian statisticians argue constantly about these issues, with classical statisticians often claiming that they can assess absolute plausibility because their procedures take one hypothesis as given and ask whether actual data look unusual from that point of view. They claim that their tests can pit a given hypothesis against all comers. Bayesians argue that doing so always involves adopting questionable ancillary hypotheses (like the dividend growth assumption in testing the perfect markets hypothesis), that in practice the data record could be spotty enough that it would fail to reject each of several mutually exclusive hypotheses, and that science advances by finding new hypotheses that are relatively more plausible than the old ones.

The consequence of this feature of Bayesian reasoning is that, in testing a hypothesis like the perfect markets hypothesis, one must be willing to specify an explicit alternative—one must specify how stock prices would evolve if the perfect markets hypothesis were violated. There is freedom to try each of several alternatives, but each comparison must involve the perfect markets hypothesis and a specific alternative. For this reason, DeJong and Whiteman report tests of the perfect markets hypothesis against a variety of alternatives. Each of these tests accommodates uncertainty regarding dividend growth.

The DeJong-Whiteman results indicate that the perfect markets hypothesis appears relatively less plausible than a model designed specifically to exploit statistical regularities in forecasting prices and dividends. However, such a model has little theoretical foundation. DeJong and Whiteman show that researchers less certain of the nature of price-dividend interaction would find some, but not too much, reason to look beyond the perfect markets hypothesis.

Conclusion

Regression tests opened the controversy about the validity of the perfect markets hypothesis. Did the hypothesis pass? Were the tests powerful enough? When study after study either found no evidence against the perfect markets hypothesis or explained away any evidence found, the debate subsided. The controversy was reopened by applying volatility tests, which spurred an additional debate concerning the temporal stability of prices and dividends. While the stability issue seems to be settling in favor of trend-stationarity, the volatility tests suffer from bias, misspecification, and interpretational difficulties and have done little to settle the controversy over the perfect markets hypothesis.

Work that permits analysis of the perfect markets hypothesis without conditioning on the nature of dividend growth suggests that from at least one viewpoint something seems to be amiss. The problem with the perfect markets hypothesis may very well be the constancy of the real rate of return. Whatever the explanation, something other than the simplest version of the perfect markets hypothesis is needed to explain the aggregate dividend-price data. Indeed, economists are actively seeking better explanations (than the simple version of the perfect markets hypothesis) for why they are not rich.

Appendix

Data Set 1: Modified Dow Jones Industrial Average

Annual, 1928-78. Here, P_t and D_t refer to the real price and dividends of the portfolio of thirty stocks that made up the sample for the Dow Jones Industrial Average when it was created in 1928. However, because the stocks used to calculate this average are continually adjusted, Shiller has extensively modified it to control for these adjustments. These modifications are described in the appendix of Shiller (1981b).

Data Set 2: Value-Weighted New York Stock Exchange Index

Annual, 1926-81. P_t represents the January value-weighted New York Stock Exchange stock price divided by the January producer price index (PPI). D_t represents total dividends for the year accruing to the portfolio represented by the stocks in the index, divided by the annual average PPI. These data are contained in the Center for Research in Security Prices data set. In the Marsh and Merton (1987) analysis, P_t represents December rather than January prices, but January prices are considered here to remain consistent with the Shiller dating scheme.

Data Set 3: Modified Standard and Poor's 500 Series

Annual, 1871-1985. The price series is Standard and Poor's monthly composite stock price index for January, divided by the producer price index (January PPI starting in 1900, annual average PPI before 1900 scaled to 1.00 in the base year, 1979). The dividend series represents the total calendar year's dividends accruing to the portfolio, represented by the stocks in the index, divided by the average PPI for the year. These data differ slightly from Shiller's (1981b) numbers; some adjustments have been made to correct minor errors in the original data, and the series have been updated to 1985.

Notes

1. The Dow Jones Industrials and Standard and Poor's data were originally investigated by Shiller (1981a, 1981b), and the New York Stock Exchange data correspond to the series examined by Marsh and Merton (1987). The data are described in the appendix.
2. The statistic used was the Box-Pierce Q statistic. Further, the calculations in the table are carried out on data from which an exponential trend has been removed. Exponential detrending does not materially affect the implications of the perfect markets hypothesis: if P_t and D_t possess deterministic exponential trends, then equations like (1), (2), and (3) characterize the relationships between exponentially detrended prices and dividends but β is replaced by $\eta = (1 + g)/(1 + r)$, where g is the exponential growth rate.
3. The P series was calculated by backward iteration on equation (3) beginning with a terminal price $P = P_T$. Shiller used the average price over the sample for this terminal value; Flavin (1983), Mankiw, Romer, and Shapiro (1985), and Kleidon (1986a, 1986b) emphasize that this choice introduces biases into the calculations. Using the actual terminal price avoids (some of) this problem.
4. Lintner's (1956) interviews suggested that managers try to achieve a constant dividend-price ratio. Why they do it is another unsettled question.
5. Actually, the simultaneity complicates matters. It turns out that both P^* and P will have infinite variance under the dividend-smoothing assumption; sample estimates of variability will of necessity be badly biased downward, and P^* may appear to fluctuate less than P.
6. The representation is known in the econometrics literature as the vector autoregression.

References

Campbell, John Y., and Robert J. Shiller. "Cointegration and Tests of Present Value Models." *Journal of Political Economy* 95 (October 1987): 1062-88.

DeJong, David N., John C. Nankervis, N.E. Savin, and Charles H. Whiteman. "Integration versus Trend Stationarity in Time Series." *Econometrica* 60 (March 1992a): 422-33.

_____. "The Power Problems of Unit Root Tests in Time Series with Autoregressive Errors." *Journal of Econometrics* 53 (July-September 1992b).

DeJong, David N., and Charles H. Whiteman. "The Temporal Stability of Dividends and Stock Prices: Evidence from the Likelihood Function." *American Economic Review* 81 (June 1991): 600-617.

_____. "Modelling Stock Prices without Pretending to Know How to Induce Stationarity." University of Iowa manuscript, 1992.

Dickey, David A., and Wayne A. Fuller. "Likelihood Ratio Statistics for Autoregressive Time Series with a Unit Root." *Econometrica* 49 (July 1981): 1057-72.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25 (May 1970): 383-417.

_____, and Kenneth R. French. "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy* 96 (April 1988): 246-73.

Flavin, Marjorie A. "Excess Volatility in the Financial Markets." *Journal of Political Economy* 91 (December 1983): 929-56.

Fuller, Wayne A. *Introduction to Statistical Time Series*. New York: John Wiley and Sons, 1976.

Kleidon, Allan W. "Bias in Small Sample Tests of Stock Price Rationality." *Journal of Business* 59 (April 1986a): 237-61.

_____. "Variance Bounds Tests and Stock Price Valuation Methods." *Journal of Political Economy* 94 (October 1986b): 953-1001.

LeRoy, Stephen F., and C.J. LaCivita. "Risk Aversion and the Dispersion of Asset Prices." *Journal of Business* 54 (October 1981): 535-47.

LeRoy, Stephen F., and Richard D. Porter. "The Present Value Relation: Tests Based on Implied Variance Bounds." *Econometrica* 49 (May 1981): 555-74.

Lintner, John V., Jr. "Distribution of Incomes of Corporations among Dividends, Retained Earnings, and Taxes." *American Economic Review* 66 (May 1956): 97-113.

Lo, Andrew W., and A.C. MacKinlay. "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test." *Review of Financial Studies* 1 (1988): 41-66.

Malkiel, Burton G. *A Random Walk Down Wall Street*. New York: Norton, 1973.

Mankiw, N. Gregory, David Romer, and Matthew D. Shapiro. "An Unbiased Reexamination of Stock Market Volatility." *Journal of Finance* 40 (July 1985): 677-87.

Marsh, Terry A., and Robert C. Merton. "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices." *American Economic Review* 76 (June 1986): 483-98.

_____. "Dividend Behavior for the Aggregate Stock Market." *Journal of Business* 60 (January 1987): 1-40.

Michener, Ronald W. "Variance Bounds in a Simple Model of Asset Pricing." *Journal of Political Economy* 90 (February 1982): 166-75.

Perron, Pierre. "Trends and Random Walks in Macroeconomic Time Series: Further Evidence from a New Approach."

Journal of Economic Dynamics and Control 12 (June/September 1988): 297-332.

Poterba, James M., and Lawrence H. Summers. "Mean Reversion in Stock Returns: Evidence and Implications." *Journal of Financial Economics* 1 (October 1988): 27-60.

Shiller, Robert J. "The Use of Volatility Measures in Assessing Market Efficiency." *Journal of Finance* 36 (May 1981a): 291-304.

_____. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (June (1981b): 421-36.

Summers, Lawrence H. "Does the Stock Market Rationally Reflect Fundamental Values?" *Journal of Finance* 41 (July 1986): 591-601.

West, Kenneth D. "A Note on the Power of Least Squares Tests for a Unit Root." *Economics Letters* 24 (1987): 249-52.

_____. "Dividend Innovations and Stock Price Volatility." *Econometrica* 56 (January 1988): 37-61.