



American Finance Association

An Unbiased Reexamination of Stock Market Volatility

Author(s): N. Gregory Mankiw, David Romer, Matthew D. Shapiro

Source: *The Journal of Finance*, Vol. 40, No. 3, Papers and Proceedings of the Forty-Third Annual Meeting American Finance Association, Dallas, Texas, December 28-30, 1984 (Jul., 1985), pp. 677-687

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2327790>

Accessed: 26/11/2009 08:41

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

An Unbiased Reexamination of Stock Market Volatility

N. GREGORY MANKIW, DAVID ROMER, and MATTHEW D. SHAPIRO*

ABSTRACT

Recent work demonstrates serious statistical problems with standard volatility tests. This paper proposes new tests that are unbiased in small samples and that do not require assumptions of stationarity. The new tests continue to find evidence against the model positing rational expectations and a constant required rate of return on equity.

IN HIS SEMINAL WORK on the volatility of the stock market, Robert Shiller [10] examines the view that changes in the level of the stock market reflect news about future dividends. In particular, he considers a model in which stock prices are the present discounted value, using a constant discount rate, of expected dividends.¹ He provides what appears to be overwhelming evidence against this model. Dividends appear a smooth time series, while stock prices are volatile. The standard deviation of stock prices is, he concludes, about five times greater than would be expected from the volatility of dividends.

Recent work by Flavin [2], Marsh and Merton [6], and others calls into question the validity of Shiller's test for a variety of reasons. Flavin examines the small sample properties of volatility tests and shows that they are extremely biased toward finding excessive volatility. Marsh and Merton argue that the excess volatility of stock prices relative to dividends is not surprising since firms, consciously smoothing dividends, make dividends a nonstationary series. They show that under a plausible dividend pay-out rule which makes dividends a random walk, the standard volatility test finds excess volatility in *every* sample, even though the efficient markets model is correct.

In this paper, we propose new volatility tests. Our goal is to examine the same question Shiller addresses: can stock prices, a very volatile series, be the expected present discounted value of dividends, a relatively smooth series? The new tests, however, do not suffer from the problems Flavin and Marsh and Merton discuss.

In Section I we briefly review Shiller's volatility test and then present our new ones. In Section II we consider the objections to Shiller's test and show that our tests are immune to these problems. We apply our tests to the data in Section III. We find that our new unbiased tests continue to provide evidence contra-

* Mankiw and Romer are from the Massachusetts Institute of Technology, and Shapiro is from Yale University. The authors are grateful to Fischer Black, Olivier Blanchard, Stanley Fischer, James Poterba, and Robert Shiller for helpful comments.

¹ While such a model is highly restrictive and need not arise in a general equilibrium setting in which the required rate of return may vary (Michener [7]), it is the model that pervades most casual discussions of the stock market.

dicting the model. In Section IV we discuss the implications of our results and outline some unsettled issues.

I. Old and New Volatility Tests

Consider the standard present value relation (1):

$$P_t = \sum_{k=0}^{\infty} \gamma^{k+1} E_t D_{t+k} \quad (1)$$

where

- P_t = the price of the stock at time t ;
- D_{t+k} = the dividend paid at time $t + k$;
- E_t = the expectation conditional on information available at time t ; and
- γ = the discount factor, or $1/(1 + r)$, where r is the required rate of return.

The present value model (1) has broad application. LeRoy and Porter [4] and Shiller [10] use it as a model of stock prices. For concreteness, the discussion here centers on this application. However, if D_t denotes the short-term interest rate and P_t denotes the long-term interest rate, then equation (1) expresses the expectations theory of the term structure.² The permanent income hypothesis can also be represented in the form of equation (1). There are thus numerous economic models to which volatility tests can be applied.

A. Shiller's Test

Define P_t^* as the "perfect foresight," or "*ex post* rational," stock price. That is,

$$P_t^* = \sum_{k=0}^{\infty} \gamma^{k+1} D_{t+k}. \quad (2)$$

P_t^* is the present value of actual, rather than expected, dividends. Since the expectations operator E_t is linear, we know

$$P_t = E_t(P_t^*). \quad (3)$$

The observed stock price P_t is, according to the model, the expectation of the perfect foresight price P_t^* conditional on information available at time t .

Define the error ν_t such that

$$P_t^* = P_t + \nu_t. \quad (4)$$

By (3), ν_t is the error in forecasting P_t^* . As a rational forecast error, it is uncorrelated with information available at time t . In particular, ν_t must be uncorrelated with P_t . Thus, we know

$$V(P_t^*) = V(P_t) + V(\nu_t) \quad (5)$$

where $V(x)$ is the variance of x . Of course, $V(\nu_t) \geq 0$, implying

$$V(P_t^*) \geq V(P_t). \quad (6)$$

² See Shiller [9] and Singleton [13].

The variance of P_t^* thus provides an upper bound to the variance of the observed stock price P_t .

B. New Tests

Let P_t^0 be some "naive forecast" stock price:

$$P_t^0 = \sum_{k=0}^{\infty} \gamma^{k+1} F_t D_{t+k} \quad (7)$$

where $F_t D_{t+k}$ denotes a naive forecast of D_{t+k} made at time t . This naive forecast need not be a rational one. It is important, however, that the rational agents at time t have access to this naive forecast.

We begin with the identity:

$$P_t^* - P_t^0 = (P_t^* - P_t) + (P_t - P_t^0). \quad (8)$$

Note that $(P_t^* - P_t)$ equals ν_t and thus is uncorrelated with information available at time t . In particular,

$$E_t[(P_t^* - P_t)(P_t - P_t^0)] = 0 \quad (9)$$

since P_t and P_t^0 are known at time t . Squaring both sides of (8) and taking the expectation thus implies

$$E_t(P_t^* - P_t^0)^2 = E_t(P_t^* - P_t)^2 + E_t(P_t - P_t^0)^2. \quad (10)$$

This equality in turn implies

$$E_t(P_t^* - P_t^0)^2 \geq E_t(P_t^* - P_t)^2 \quad (11)$$

and

$$E_t(P_t^* - P_t^0)^2 \geq E_t(P_t - P_t^0)^2. \quad (12)$$

Finally, the law of iterated projections allows us to replace expectations conditional on information available at time t with expectations conditional on information available prior to the beginning of the sample period. That is, letting E denote the expectation conditional on the initial conditions, we have

$$E(P_t^* - P_t^0)^2 = E(P_t^* - P_t)^2 + E(P_t - P_t^0)^2; \quad (10')$$

$$E(P_t^* - P_t^0)^2 \geq E(P_t^* - P_t)^2; \quad (11')$$

$$E(P_t^* - P_t^0)^2 \geq E(P_t - P_t^0)^2. \quad (12')$$

As long as the expectations are taken conditional on the information available a finite amount of time before date t , nonstationarity poses no difficulties for the existence of these conditional expectations.³ Expressions (11') and (12') are the volatility relations we examine.

Expression (11') states that the market price is a better forecast of the *ex post* rational price, in terms of mean squared error, than is the naive forecast stock

³ Likewise, nonstationarity poses no difficulties for the existence of the conditional expectations of the sample variances of P^* and P . Marsh and Merton's observation that nonstationarity causes problems for Shiller's test concerns not the existence of expectations of sample variances, but, as we discuss in Section II, the statistical properties of the sample variances.

price. If the naive forecast is better than the market forecast, then the inequality (11') is violated and the model (1) is rejected.

Expression (12') states that the *ex post* rational price is more volatile around P^0 than is the market price. The inequality (12') is thus analogous to Shiller's volatility test (6), where the variance is centered not around the mean but around the naive forecast price.⁴

II. Some Problems with Volatility Tests

Recent work suggests that the application of Shiller's test (6) is not as straightforward as it might first appear. In this section, we consider two related critiques of this volatility test and show that our tests are immune to them.

The relation in (6) concerns the distributions of the random variables P_t^* and P_t . But, of course, we observe a given P_t^* and P_t only once. A single observation does not provide information about the variance of a random variable. The usual procedure used to test the variance relation (6) is to assume that stock prices and dividends are stationary around a deterministic trend. Under this assumption, the sample variances of the detrended series for P^* and P will, as the number of observations becomes large, converge to the corresponding population variances. The sample variances are thus used to test the variance bounds.

Both Flavin and Marsh and Merton point out serious difficulties with this procedure. Flavin's criticism concerns the small sample properties of the test. Sample variances are downward-biased estimators of population variances because sample means are used instead of population means.⁵ The smoother (the more highly autocorrelated) a series is, the greater the downward bias. Suppose, to consider Flavin's example, that the fundamental series follows a first-order autoregressive process. Then the market price is proportional to dividends. Since P_t^* is a weighted sum of future dividends, it is a weighted sum of future prices. This effect tends to smooth P^* relative to P , and thus to make the bias in estimating the variance of P^* greater than the bias in estimating the variance of P . Flavin calculates the finite sample distributions of the sample variances for variance bounds tests of the expectations hypothesis of the term structure. She finds that for typical sample sizes the bias toward finding excess volatility is considerable.

Marsh and Merton demonstrate that the use of sample variances may be inappropriate even as the number of observations becomes arbitrarily large. They focus on the possibility that dividends do not follow a stationary process. Suppose,

⁴ West [15] uses a test analogous to (11') to reject the linear-quadratic model of firms' inventory policies. Specifically, he shows that under the assumptions of his model, firms' actual behavior is outperformed by a naive inventory policy. Flavin [2] and Frankel and Stock [3] discuss volatility tests involving comparisons of expressions of the form $\Sigma(P_t^* - Z_t)^2$ and $\Sigma(P_t - Z_t)^2$. However, Flavin assumes stationarity and discusses only cases in which Z_t is a constant, and Frankel and Stock use information from the entire sample (and, thus, information not available at time t) to compute their Z_t 's.

⁵ In an independent, identically distributed sample, the estimated variance must be corrected by a factor of $T/(T - 1)$. In a serially correlated sample, the correction for bias is larger. See T. W. Anderson [1, p. 448].

to take a simple example, that dividends follow a random walk. Then Marsh and Merton show that if the null hypothesis of market efficiency is correct, then for *any* sample the sample variance of P^* will be greater than or equal to the sample variance of P . Intuitively, here P^* is not simply a weighted sum of P 's, but a weighted sum with the sum of the weights equal to one. Marsh and Merton show further that for a broader class of nonstationary processes in which dividends depend on "permanent earnings," as the number of observations becomes large the sample variances violate the variance bounds almost surely under the null hypothesis.

Both difficulties with volatility tests can be illustrated by considering the case in which dividends follow a first-order autoregressive process,

$$D_t = \rho D_{t-1} + \varepsilon_t. \quad (13)$$

ε is white noise with variance σ_ε^2 . The expectation of the present discounted value of the flow of dividends exists as long as ρ is less than $1 + r$ in absolute value. If this condition holds, P_t is simply $\gamma\rho/(1 - \gamma\rho) D_{t-1}$, where $\gamma = 1/(1 + r)$.

Consider the test statistic

$$Q = V(P^*) - V(P), \quad (14)$$

where $V(x)$ is the sample variance of x ; i.e., it is the sum of squared deviations from the sample mean divided by T . Since we do not observe the infinite stream of dividends, we define P_t^* by

$$P_t^* = \sum_{k=0}^{(T-t)-1} \gamma^{k+1} D_{t+k} + \gamma^{T-t} P_T. \quad (15)$$

Both Shiller and Marsh and Merton use the sample mean of the detrended P 's in place of P_T in (15). Under this formulation, it is not true that rationality implies $P_t = E_t P_t^*$.⁶ Using our formulation, P_t^* corresponds to the "perfect foresight price" for the policy of holding the stock until time T and then selling it at the prevailing price. Thus, in this case $P_t = E_t P_t^*$. Note that because of definition of P_t^* includes not only future dividends but also the actual selling price of the stock at some future date, $P_t = E_t P_t^*$ even in the presence of speculative bubbles. This implies that the inequalities (11) and (12) hold even if there are bubbles.

Making the simplifying assumption that $D_{-1} = 0$, straightforward but tedious algebra allows us to compute the expectation of the test statistic Q as a function of ρ , γ , T , and σ_ε^2 under the null hypothesis. This allows us to learn about the biases, although not the distributional properties, of the volatility tests.

We can write Q as

$$\begin{aligned} Q &= V(P + \nu) - V(P) \\ &= V(\nu) + 2 \text{Cov}(P, \nu) \\ &= \left(\frac{1}{T} \sum_t \nu_t^2 - \frac{1}{T^2} [\sum_t \nu_t]^2 \right) + \left(\frac{2}{T} \sum_t P_t \nu_t - \frac{2}{T^2} (\sum_t P_t)(\sum_t \nu_t) \right). \end{aligned}$$

⁶ Flavin emphasizes that this method of obtaining the terminal perfect foresight price is an important cause of the bias that she finds in Shiller's test.




Figure 1. The expectation of the test statistic $Q = V(P^*) - V(P)$, normalized by $E[V(P^*)]$, as a function of ρ and T for the case in which dividends follow a first-order autoregressive process.

The expectations of the component pieces (ignoring the special cases of $\rho = 1$, $\rho = \gamma$, and $\gamma = 1$) are:⁷

$$E[\sum_t \nu_t^2] = \frac{\sigma_\varepsilon^2 B^2 \gamma^4}{1 - \gamma^2} \left[T - \gamma^2 \frac{1 - \gamma^{2T}}{1 - \gamma^2} \right];$$

$$E[(\sum_t \nu_t)^2] = \frac{\sigma_\varepsilon^2 \gamma^4 B^2}{(1 - \gamma)^2} \left[T - 2\gamma \frac{1 - \gamma^T}{1 - \gamma} + \gamma^2 \frac{1 - \gamma^{2T}}{1 - \gamma^2} \right];$$

$$E[\sum_t P_t \nu_t] = 0;$$

$$E[(\sum_t P_t)(\sum_t \nu_t)] = \frac{\sigma_\varepsilon^2 \gamma^3 B^2}{(1 - \gamma)(1 - \rho)} \cdot \left[(T - 1) - \gamma \frac{1 - \gamma^{T-1}}{1 - \gamma} - \rho \frac{1 - \rho^{T-1}}{1 - \rho} + \rho \gamma \frac{\gamma^{T-1} - \rho^{T-1}}{\gamma - \rho} \right];$$

where $B = \frac{1}{1 - \gamma\rho}$.

Figure 1 plots $E[Q]$, normalized by the expectation of the sample variance of P^* , as a function of T for various values of ρ . We set γ equal to 0.942, the value Shiller uses. A conventional volatility test involves checking whether the sample variance of P^* exceeds that sample variance of P —i.e., whether Q is positive. But

⁷ The expectations for the special cases can be found by applying l'Hôpital's rule to the formulas in the text.

Figure 1 shows that, under the null hypothesis, the expectation of Q is negative for small sample sizes. In addition, as ρ approaches one, the sample size required for the expectation of Q to be positive increases. With $\rho = 0.8$, roughly 20 observations are needed before $E[Q]$ is positive. With $\rho = 0.99$, several hundred observations are needed. And if ρ is greater than or equal to one, the “small sample bias” remains as the number of observations becomes arbitrarily large: for any sample size, the expectation of the test statistic is negative.

The bias of the test statistic Q arises from the inclusion of the sample means. Because the new tests do not use sample means, their sample counterparts are unbiased. Define

$$S_1 = \frac{1}{T} \sum_{t=1}^T (P_t^* - P_t^0)^2 - \frac{1}{T} \sum_{t=1}^T (P_t^* - P_t)^2; \quad (16)$$

$$S_2 = \frac{1}{T} \sum_{t=1}^T (P_t^* - P_t^0)^2 - \frac{1}{T} \sum_{t=1}^T (P_t - P_t^0)^2. \quad (17)$$

The expectations of these test statistics under the null hypothesis are:⁸

$$\begin{aligned} E[S_1] &= \frac{1}{T} \sum_{t=1}^T E[(P_t + \nu_t - P_t^0)^2] = \frac{1}{T} \sum_{t=1}^T E[(P_t + \nu_t - P_t)^2] \\ &= \frac{1}{T} \sum_{t=1}^T [E[\nu_t^2] + E[(P_t - P_t^0)^2]] - \frac{1}{T} \sum_{t=1}^T E[\nu_t^2] \\ &= \frac{1}{T} \sum_{t=1}^T E[(P_t - P_t^0)^2] \geq 0; \end{aligned} \quad (18)$$

$$\begin{aligned} E[S_2] &= \frac{1}{T} \sum_{t=1}^T [E[\nu_t^2] + E[(P_t - P_t^0)^2]] - \frac{1}{T} \sum_{t=1}^T E[(P_t - P_t^0)^2] \\ &= \frac{1}{T} \sum_{t=1}^T E[\nu_t^2] \geq 0. \end{aligned} \quad (19)$$

Both derivations utilize the fact that ν_t is uncorrelated with P_t and P_t^0 . Note that the expectations of the test statistics are positive even in the case of heteroskedasticity. That is, $E[(P_t - P_t^0)^2]$ and $E[\nu_t^2]$ need not be the same for all t .

Finally, our tests avoid two other difficulties that may arise in various forms of volatility tests. The first difficulty concerns detrending. The levels of stock prices and dividends are clearly nonstationary. Thus, to make a case for stationarity detrended data must be used. But a variety of evidence suggests that, in general, detrending may cause severe econometric difficulties (Nelson and Kang [8]). In the case of volatility tests in particular, Marsh and Merton [5] argue that Shiller’s procedures for detrending introduce important biases. But because the “naive forecast” P_t^0 can grow as dividends grow, our tests avoid the need for detrending.

The second difficulty arises in a potential remedy for small sample bias that follows the spirit of Flavin’s analysis. Flavin derives the small sample distribution

⁸ Recall that the operator E denotes the expectation conditional on the initial conditions.

of the test statistic $V(P^*) - V(P)$ under the assumption that fundamentals follow a particular process. At some cost in intuitive appeal—one is no longer simply checking which of $V(P^*)$ and $V(P)$ is larger—use of the small sample distribution appears to eliminate problems of bias and to allow precise statements of significance levels. In fact, however, the problem of bias remains. P_t is assumed to be the forecast of the present discounted value of future dividends given a particular information set. But agents' information sets may be larger than the particular information set considered. Because an optimal predictor based on more information has greater variance than one based on less information, P_t will be more variable than the econometrician expects. In the case of stationarity, for example, the sample variance of P will converge to the population variance of the estimator based on the larger information set; this full information population variance exceeds the population variance of the limited information estimator, to which the econometrician expects the sample variance of P to converge. The test is thus inconsistent. Intuitively, the source of this bias is that the investigator interprets movements in P_t caused by information outside the limited information set as "noise," or excess volatility. Again, the new tests that we consider are unbiased.⁹

III. Results

In this section we apply our proposed tests to the data. We use the same data Shiller [10] uses, which now extend from 1872 to 1983. The dividend and price series thus refer to a large portfolio of commonly traded stocks.

The first decision that needs to be made regards the choice of the naive forecast of dividends, $F_t D_{t+k}$. We let our naive forecast be completely myopic. That is, we set

$$F_t D_{t+k} = D_{t-1}. \quad (20)$$

Certainly such a myopic forecast is possible at any point in time: it does not require knowledge of parameters estimated over the entire sample. Recall that the naive forecast need not be efficient in any sense.

We can now compute P_t^* as in Equation (15) and P_t^0 by

$$P_t^0 = [\gamma/(1 - \gamma)]D_{t-1}. \quad (21)$$

We can then compute the sample mean squared errors in Equation (10'). Table I presents the results for various values of the required rate of return, r . Figure 2 displays the three time series for a required rate of return of six percent, which is approximately the average market return.

We find that our new inequalities are uniformly violated. Inequality (11') is tested by examining the first two columns. Contrary to the theory, the second column is uniformly greater than the first. That is, the naive forecast price P^0 is a better forecast of the perfect forecast price P^* than is the market price P .

⁹ Although Flavin does discuss the magnitude of the biases in previous tests of excess volatility, she does not in fact use the finite sample distributions to compute significance levels. Instead, a simple version of a noncentral test is utilized (see footnote 4).

Table I
Unbiased Volatility Tests (Not Weighted)

r (%)	$E(P^* - P^0)^2$	$=$	$E(P^* - P)^2$	$+$	$E(P - P^0)^2$
4	146		355		118
5	88		233		159
6	67		228		257
7	58		271		361
8	54		333		459
9	51		399		546
10	50		466		623




Figure 2. The Perfect Foresight Price (P^*), the Naive Forecast Price (P^0), and the Market Price (P) for the case $r = 6\%$.

Similarly, inequality (12') is violated for plausible values of r . Thus, the volatility of the market price around P^0 is greater than the volatility of the perfect foresight price around P^0 .

Examination of the series discussed above shows an upward trend. While this fact does not cast doubt on the above statistics, it does suggest that the relevant random variables might exhibit heteroskedasticity. The variance of the error, ν_t , may well grow through time. If so, undue weight may be accorded to recent experience. To correct for a possible inefficiency, we weight the errors by the inverse of the market price. The errors are thus similar to percent errors. It is straightforward to show that the proof of unbiasedness remains valid under this weighting. Given our particular sample, however, the weighting procedure makes the inequalities predicted by the theory especially likely to hold. The theory appears to perform worst during the 1960's and early 1970's, when the market

Table II
Unbiased Volatility Tests (Weighted)

r (%)	$E(P^* - P^0)^2$	=	$E(P^* - P)^2$	+	$E(P - P^0)^2$
4	0.387		0.840		0.176
5	0.221		0.336		0.081
6	0.139		0.156		0.090
7	0.095		0.109		0.132
8	0.070		0.116		0.182
9	0.055		0.146		0.233
10	0.046		0.186		0.281

price is dramatically higher than both P^* and P^0 ; weighting by the inverse of the market price assigns relatively low weight to these observations.

Table II presents the weighted mean squared errors. The results are somewhat less dramatic. In particular, the inequality (12') is not violated for low values of the required rate of return. Nonetheless, the weighted statistics continue to indicate that inequality (11') is violated. Our naive forecast, P^0 , is a better predictor of the perfect foresight price, P^* , than is the market price. Given that our naive forecast is *extremely* naive—dividends will never change—this result is surprising indeed.¹⁰

In summary, while our unbiased volatility tests do not find evidence as striking as that Shiller reports, we do find evidence contradicting the model. In particular, the naive prediction that dividends will never change outperforms the market as a forecast of the present value of *ex post* dividends.

IV. Conclusions

The original volatility tests of Shiller and LeRoy and Porter suggested that a considerable part of movements in stock prices cannot be attributed to information about future dividends. Recent work points out a variety of problems with these tests. In this paper, we present new tests that are unbiased in small samples, that do not require any assumptions of stationarity, and that avoid the need for detrending. These tests continue to find that the model positing no expected profit opportunities and a constant discount rate appears not to account for movements in the stock market. In particular, the naive prediction that dividends will never change seems to outperform the market price as a forecast of the *ex post* rational price. In addition, the theory performs only slightly better when confronted with the prediction that the market price should fluctuate less around the naive forecast price than does the perfect foresight price: only when we choose weights in a way that favors acceptance is the inequality not violated.

A variety of hypotheses could explain our results. The question of statistical significance remains open. One possibility is that the null hypothesis of market

¹⁰ Figure 2 shows that the theory performs particularly poorly in the 1960's. To see whether this fact is completely driving the results, we compute the test statistics ending in 1958, but still using January 1984 for the terminal price. This procedure also reduces the importance of the terminal price. The evidence we obtain, while somewhat less striking, again contradicts the model.

efficiency and a constant discount rate is correct, and that the violations of the inequalities predicted by the theory are not statistically significant.

It is also possible that observed stock market fluctuations are caused by changes in the discount rate. The discount rate could change either because of changes in the real interest rate or because of changes in the risk premium on equity. Shiller [11] reports mixed evidence on the question of whether changing discount rates could explain observed market fluctuations.

A final possibility is that the stock market simply does not accurately reflect underlying fundamentals. Shiller [12] argues that fads are an important cause of market movements. Summers [14] shows that the market could deviate substantially from fundamentals without the existence of discernible profit opportunities. If "animal spirits" are indeed an important force driving the market, then changes in the dividend-price ratio reflect neither news about future dividends nor changes in interest rates or in risk premiums.

REFERENCES

1. T. W. Anderson. *The Statistical Analysis of Time Series*. New York: Wiley, 1971.
2. Marjorie A. Flavin. "Excess Volatility in the Financial Markets: A Reassessment of the Empirical Evidence." *Journal of Political Economy* 91 (December 1983), 929-56.
3. Jeffrey A. Frankel and James H. Stock. "A Relationship between Regression and Volatility Tests of Market Efficiency." National Bureau of Economic Research Working Paper No. 1105, April 1983.
4. Stephen F. LeRoy and Richard D. Porter. "The Present-Value Relation: Tests Based on Implied Variance Bounds." *Econometrica* 49 (May 1981), 555-74.
5. Terry A. Marsh and Robert C. Merton. "Are Stock Prices Too Volatile?" Mimeo, Massachusetts Institute of Technology, May 1983.
6. ———. "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices." Sloan School Working Paper No. 1584-84, August 1984.
7. Ronald W. Michener. "Variance Bounds in a Simple Model of Asset Pricing." *Journal of Political Economy* 90 (February 1982), 166-75.
8. Charles Nelson and K. Kang. "Pitfalls in the Use of Time as an Explanatory Variable." NBER Technical Working Paper, 1983.
9. Robert J. Shiller. "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure." *Journal of Political Economy* 87 (December 1979), 1191-1219.
10. ———. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (June 1981), 421-36.
11. ———. "Consumption, Asset Markets, and Macroeconomic Fluctuations." *Carnegie-Rochester Conference on Public Policy* 17 (1982), 203-50.
12. ———. "Stock Prices and Social Dynamics." *Brookings Papers on Economic Activity* (No. 2, 1984), 457-98.
13. Kenneth J. Singleton. "Expectations Models of the Term Structure and Implied Variance Bounds." *Journal of Political Economy* 88 (December 1980), 1159-76.
14. Lawrence H. Summers. "Do We Really Know that Financial Markets Are Efficient?" National Bureau of Economic Research Working Paper No. 994, September 1982.
15. Kenneth David West. "A Variance Bounds Test of the Linear-Quadratic Inventory Model." In *Inventory Models and Backlog Costs: An Empirical Investigation*. Unpublished Ph.D. dissertation, Massachusetts Institute of Technology, May 1983.