

NBER WORKING PAPER SERIES

MEAN REVERSION IN STOCK PRICES:
EVIDENCE AND IMPLICATIONS

James M. Poterba

Lawrence H. Summers

Working Paper No. 2343

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1987

We are grateful to Barry Perlstein, Changyong Rhee, Jeff Zweibel and especially David Cutler for excellent research assistance, to Ben Bernanke, John Campbell, Robert Engle, Eugene Fama, Pete Kyle, Greg Mankiw, Julio Rotemberg, William Schwert, Kenneth Singleton, and Mark Watson for helpful comments, and to James Darcel, Matthew Shapiro, and Ian Tonks for data assistance. This research was supported by the National Science Foundation and was conducted while the first author was a Batterymarch Fellow. The research reported here is part of the NBER's research program in Financial Markets and Monetary Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Mean Reversion in Stock Prices: Evidence and Implications

ABSTRACT

This paper analyzes the statistical evidence bearing on whether transitory components account for a large fraction of the variance in common stock returns. The first part treats methodological issues involved in testing for transitory return components. It demonstrates that variance ratios are among the most powerful tests for detecting mean reversion in stock prices, but that they have little power against the principal interesting alternatives to the random walk hypothesis. The second part applies variance ratio tests to market returns for the United States over the 1871-1986 period and for seventeen other countries over the 1957-1985 period, as well as to returns on individual firms over the 1926-1985 period. We find consistent evidence that stock returns are positively serially correlated over short horizons, and negatively autocorrelated over long horizons. The point estimates suggest that the transitory components in stock prices have a standard deviation of between 15 and 25 percent and account for more than half of the variance in monthly returns. The last part of the paper discusses two possible explanations for mean reversion: time varying required returns, and slowly-decaying "price fads" that cause stock prices to deviate from fundamental values for periods of several years. We conclude that explaining observed transitory components in stock prices on the basis of movements in required returns due to risk factors is likely to be difficult.

James M. Poterba
Department of Economics
Massachusetts Institute of Technology
Cambridge, MA 02139
(617) 253-6673

Lawrence H. Summers
Department of Economics
Harvard University
Cambridge, MA 02138
(617) 495-2447

This paper examines the evidence on the extent to which stock prices exhibit mean-reverting behavior. The question of whether stock prices contain transitory components is important for financial practice and theory. For example, consider the question of investment strategy. If stock price movements contain large transitory components then for long-horizon investors the stock market may be much less risky than it appears when the variance of single-period returns is extrapolated using the random walk model. Market folklore has long suggested that those who "take the long view" should invest more in equity than those with a short horizon. Although harshly rejected by most economists, this view is correct if prices exhibit mean-reverting behavior.¹ Furthermore, the presence of transitory price components suggests the desirability of investment strategies involving the purchase of securities that have recently declined in value.

Important transitory components in stock prices could also impart some logic to economic agents' reluctance to tie decisions to current market values. Corporate managers often assert that their common stock is misvalued and claim that it would be unwise to base investment decisions on its current market price. A common procedure among universities and other institutions that rely on endowment income is to spend on the basis of a weighted average of past endowment values. Harvard University spends out of endowment according to a preset trend line regardless of the market's value. Such rules are hard to understand if stock prices follow a random walk, but make sense if prices contain important transitory components.

As a matter of theory, evaluating the extent of mean-reversion in stock prices is crucial for assessing claims such as Keynes' (1936) assertion that

"all sorts of considerations enter into market valuation which are in no way relevant to the prospective yield (p.152)." If divergences between market and fundamental values exist, but beyond some limit are eliminated by speculative forces, then stock prices will exhibit mean reversion. Returns must be negatively serially correlated at some frequency if "erroneous" market moves are eventually corrected.² As Merton (1987) notes, reasoning of this type has been used to draw conclusions about market valuations from failures to reject the absence of negative serial correlation in returns. Conversely, the presence of negative autocorrelation may signal departures from fundamental values, though it could also arise from risk factors that vary through time.

The paper is organized as follows. Section 1 begins by evaluating alternative statistical procedures for testing for transitory components in stock prices. We find that variance ratio tests of the type used by Fama and French (1986a) and Lo and MacKinlay (1987) come close to being the most powerful tests of the null hypothesis of market efficiency cum constant required returns against plausible alternative hypotheses such as the "fads" model suggested by Shiller (1984) and Summers (1986). Nevertheless, these tests have little power, even with data spanning a sixty year period. They have less than a one in four chance of rejecting the random walk model in favor of alternative hypotheses that attribute most of the variance in stock returns to transitory factors. We conclude that a sensible balancing of Type I and Type II errors suggests use of critical values above the conventional .05 level.

Section 2 examines the evidence on the presence of mean reversion in stock prices. For the United States, we analyze monthly data on real and excess New York Stock Exchange returns since 1926, as well as annual returns data for the

1871-1986 period. We also analyze evidence from seventeen other equity markets around the world, and study the mean-reverting behavior of individual corporate securities in the United States. The results are fairly consistent in suggesting the presence of transitory components in stock prices, with returns exhibiting positive autocorrelation over short periods but negative autocorrelation over longer periods.

Section 3 uses our variance ratio estimates to gauge the substantive significance of transitory components in stock prices. For the United States we find that the standard deviation of the transitory price component varies between 15 and 25 percent of value, depending on what assumption we make about its persistence. The point estimates imply that transitory components account for more than half of the variance in monthly returns, a finding that is confirmed by the evidence from other countries.

Section 4 addresses the question of whether observed patterns of mean reversion and the associated movements in ex ante returns are better explained by fundamentals such as changes in interest rates or market volatility, or as byproducts of noise trading. We review several types of evidence indicating the difficulty of accounting for observed transitory components on the basis of changes in real interest rates or risk premia. Noise trading appears to be a plausible alternative explanation for transitory price components.

Section 5 concludes by discussing some implications of our results and directions for future research.

1. Methodological Issues Involved in Testing for Transitory Components

A vast literature dating at least back to Kendall (1933) has tested the efficient markets/constant required returns model by examining individual autocorrelations in security returns. This literature, surveyed in Fama (1970), generally found little evidence of patterns in security returns and is frequently adduced in support of the efficient markets hypothesis. Recent work by Shiller and Perron (1985) and Summers (1986) has shown that such tests have relatively little power against interesting alternatives to the null hypothesis of market efficiency with constant required returns.

Several recent studies using new tests for serial dependence, notably Fama and French (1986a), have nonetheless rejected the random walk model. This section begins by describing several possible tests for the presence of stationary stock price components, including those used in recent studies. We then present Monte Carlo evidence on the power of each test against plausible alternatives to the null hypothesis of serially independent returns. We find that even the best possible tests have little power against plausible alternatives to the random walk model when we specify the (conventional) size of .05. We conclude with a discussion of general issues involved in test design when the data can only weakly differentiate alternative hypotheses, addressing in particular the degree of presumption that should be accorded to our null hypothesis of serially independent returns.

1.1. Test Methods

Recent studies employ different but related tests for mean reversion. Fama and French (1986a) and Lo and MacKinlay (1987) compare the relative variability

of returns over different horizons using variance ratio tests. Fama and French (1987) use regression tests which also involve studying the serial correlation in many-period returns. Campbell and Mankiw (1987), studying transitory components in real output, use parametric ARMA models to gauge the importance of mean reversion. As we shall see, each of these approaches involves using a particular function of the sample autocorrelations to test the hypothesis that all autocorrelations equal zero.

The variance ratio test exploits the fact that if the logarithm of the stock price, including cumulated dividends, follows a random walk then the return variance should be proportional to the return horizon.³ We study the variability of returns at different horizons, relative to the variation over a one-year period. When we analyze monthly returns, the variance ratio statistic is therefore:

$$(1) \quad VR(k) = [\text{Var}(R_t^k)/k]/[\text{Var}(R_t^{12})/12]$$

where $R_t^k = \sum_{i=0}^{k-1} R_{t-i}$, R_t denoting the total return in month t .⁴ This statistic converges to unity if returns are uncorrelated through time. If some of the price variation is due to transitory factors, this will generate negative autocorrelations at some lags and yield a variance ratio below one.

The variance ratio is closely related to earlier tests based on estimated autocorrelations. Cochrane (1986) shows that the ratio of the k -month return variance to k times the one-month return variance is approximately equal to a linear combination of sample autocorrelations. Using his results, it is straightforward to show that (1) can be approximated by:

$$(2) \quad \begin{aligned} \text{VR}(k) &\cong 1 + 2 \sum_{j=1}^{k-1} \left(\frac{k-j}{k}\right) \hat{\rho}_j - 2 \sum_{j=1}^{11} \left(\frac{12-j}{12}\right) \hat{\rho}_j \\ &\cong 1 + 2 \sum_{j=1}^{11} j \left(\frac{k-12}{12k}\right) \hat{\rho}_j + 2 \sum_{j=12}^{k-1} \left(\frac{k-j}{k}\right) \hat{\rho}_j . \end{aligned}$$

The variance ratio statistic places increasing positive weight on autocorrelations up to and including lag 11, with declining positive weight thereafter.

The small sample distribution of the variance ratio can be inferred from its relationship to the sample autocorrelations. Kendall and Stuart (1976) show that under the null hypothesis of serial independence, the j th sample autocorrelation has (i) an expected value of $-1/(T-j)$, where T denotes sample size, (ii) an asymptotic variance of $1/T$, and (iii) zero covariance with estimated autocorrelations at other lags. The expected value of $\text{VR}(k)$ is therefore:

$$(3) \quad E(\text{VR}(k)) = \frac{12 + 5k}{6k} + \frac{2}{k} \sum_{j=1}^{k-1} \frac{T-k}{T-j} - \frac{1}{6} \sum_{j=1}^{11} \frac{T-12}{T-j} .$$

The variance ratio statistics reported below are bias-corrected by dividing the measured variance ratio by $E(\text{VR}(k))$.⁵

A second test for mean reversion, used by Fama and French (1987), involves regressing multi-period returns on lagged values of multiperiod returns. This test is also designed to exploit information on the high-order autocorrelations in returns. The test is based on whether

$$(4) \quad \hat{\beta}_k = \frac{\sum_{t=2k}^T (\tilde{R}_t^k \tilde{R}_{t-k}^k)}{\sum_{t=2k}^T (\tilde{R}_{t-k}^k)^2}$$

is significantly different from zero, where \tilde{R}_j^k denotes the de-measured k -period return. The probability limit of this statistic may be approximated as a linear combination of autocorrelations:

$$(5) \quad \text{plim } \hat{\beta}_k = \frac{\rho_1 + 2\rho_2 + \dots + k\rho_k + (k-1)\rho_{k+1} + \dots + 2\rho_{2k-2} + \rho_{2k-1}}{k + (2k-2)\rho_1 + (2k-4)\rho_2 + \dots + 2\rho_{k-1}}$$

$$\cong \frac{1}{k} (1+(3-2k)\rho_1 + \dots + (3(k-1)-2k)\rho_{k-1} + k\rho_k + (k-1)\rho_{k+1} + \dots + \rho_{2k-1}).$$

$\hat{\beta}_k$ applies negative weight to autocorrelations up to order $2k/3$, followed by increasing positive weight up to lag k , followed by decaying positive weights. The difficulties that affect the variance ratio also induce small sample bias in $\hat{\beta}_k$; Fama and French (1987) use Monte Carlo simulations to correct this problem.

A third method of detecting mean reversion involves computing a likelihood ratio test of the null hypothesis of serial independence against a particular alternative. A wide range of likelihood ratio tests could be developed for different alternative hypotheses. We present results for two such tests below.

1.2 Power Calculations

To analyze the power of alternative tests for mean reversion, we consider the class of alternative hypotheses to the random walk model that Summers (1986) suggests, where the logarithm of stock prices (p_t) embodies both a permanent (p_t^*) and a transitory (u_t) component. The transitory component might be due to variation in required returns, or to some type of pricing fads. We assume that

$$(6) \quad p_t = p_t^* + u_t.$$

If the stationary component is a first-order autoregression

$$(7) \quad u_t = \rho_1 u_{t-1} + v_t$$

then

$$(8) \quad \Delta p_t = \epsilon_t + (1-L)(1-\rho_1 L)^{-1} v_t$$

where ϵ_t denotes the innovation in the nonstationary component, $p_t^* - p_{t-1}^*$.

If ν_t and ϵ_t are independent, it is straightforward to show that Δp_t follows an ARMA(1,1) process since

$$(9) \quad (1-\rho_1 L)\Delta p_t = (1-\rho_1 L)\epsilon_t + (1-L)\nu_t.$$

This description of returns allows us to capture in a simple way the possibility that stock prices contain transitory, but persistent, components.⁶ The parameter ρ_1 determines the persistence of the transitory component, while its importance in return movements is determined by the relative magnitudes of σ_ϵ^2 and σ_ν^2 .

We perform Monte Carlo experiments by generating 25,000 sequences of 720 returns; the length of each series corresponds to the number of monthly observations in the Center for Research in Security Prices' data base. Each return sequence is generated by drawing 720 pairs of standard normal variates. We set $\sigma_\epsilon^2 = 1$, so that the variance of returns (Δp_t) equals $1 + 2\sigma_\nu^2/(1+\rho_1)$. The share of the return variance accounted for by the stationary component is:

$$(10) \quad \delta = \frac{2\sigma_\nu^2}{(1 + \rho_1) + 2\sigma_\nu^2}.$$

We parameterize the return generating process by choosing ρ_1 and δ ; these choices determine σ_ν^2 . We consider cases where δ equals .25 and .75. We set $\rho_1 = .98$ for both cases, implying that innovations in the transitory price component have a half-life of 2.9 years.

In evaluating power, we use the empirical distribution of the test statistic generated with $\delta = 0$ (no transitory component) to determine the critical region for a one-sided .05 test of the random walk null against the alternative

hypothesis of mean-reversion. The panels of Table 1 report the probability that each test rejects the null hypothesis when the data are generated by the process indicated at the column head. The mean value of the test statistic under the alternative hypothesis is also reported.

The first row in Table 1 analyzes a size .05 test based on the first-order autocorrelation coefficient. This test has minimal power against the alternative hypotheses we consider: .059 when one quarter of the variation in returns is from the stationary factor, and .076 when three quarters of return movements are due to transitory pricing factors. These results confirm the findings of Shiller and Perron (1985) and Summers (1986).

The next panel in Table 1 considers variance ratio tests with values of k ranging from 24 to 96 months. The results suggest that the variance ratio tests are much more powerful than tests based on the first-order autocorrelation coefficient, but still have relatively little power to detect mean reversion for models with fairly persistent transitory components. When one quarter of the variation in returns is due to transitory factors, the power of the variance ratio tests range between .06 and .075. Even when three quarters of the variance in returns is due to the stationary component, the power of the test never rises above .190. The variance ratio tests over long horizons have somewhat more power than the tests over short horizons. It appears however that the power gains from lengthening the horizon are largely exhausted after 48 months. It will be useful in considering the empirical results below to recall that in the price fads framework, even when the transitory component in prices has a half-life of less than three years and accounts for three-quarters of the variation in returns, the variance ratio at 96 months is .67.

Table 1: Power of Alternative Tests for Transitory Components

Test Statistic	Parameters of Return Generating Process			
	$\rho = .98$ $\delta = .25$		$\rho = .98$ $\delta = .75$	
	Power	Mean Value of Statistic	Power	Mean Value of Statistic
First-Order Autocorrelation	.059	-.002	.076	-.007
Variance Ratio				
24 months	.067	.973	.137	.927
36 months	.069	.952	.156	.867
48 months	.071	.935	.161	.815
60 months	.073	.920	.180	.771
72 months	.075	.906	.186	.733
84 months	.073	.894	.186	.700
96 months	.071	.884	.187	.670
Return Regression				
12 months	.067	-.044	.137	-.089
24 months	.071	-.080	.158	-.158
36 months	.071	-.112	.159	-.210
48 months	.066	-.141	.144	-.250
60 months	.066	-.167	.132	-.282
72 months	.059	-.194	.113	-.308
84 months	.059	-.221	.097	-.332
96 months	.057	-.250	.086	-.354
LR Test	.076	1.244	.240	4.497

Notes: Tabulations are based on 25000 Monte Carlo experiments using monthly returns generated by the indicated process, where δ indicates the share of return variation due to transitory components, while ρ describes the monthly serial correlation in the transitory component.

The second panel in Table 1 shows power calculations for the long-horizon regression tests. The results are similar to those for variance ratios. Regression tests appear to be less powerful than variance ratios against our alternative hypotheses, however, since the maximum power as the length of the regression horizon varies is below the maximum power for the variance ratio test. For example, the best variance ratio test against the $\delta = .25$ case has a power of .075, while the best regression test's power is .071. Similarly, in the $\delta = .75$ case, the best variance ratio test has a power of .187; the best regression test has a power of .159. It is interesting to note that the power of the regression tests is maximized with windows of about forty-eight months.

The final panel of Table 1 presents results on likelihood ratio tests.⁷ When the data are generated by an ARMA(1,1) model, the Neyman-Pearson lemma dictates that the likelihood ratio test is the most powerful test of the null of serial independence against this particular alternative. The power values associated with these tests are therefore upper bounds on the possible power that any other tests could achieve for a given size. In practice, even likelihood ratio tests will have lower power since we are unlikely to construct the test for the precise alternative hypothesis that generated the data.

Although the likelihood ratio tests have somewhat more power than the variance ratio tests, .078 in the $\delta = .25$ case and .240 in the $\delta = .75$ case, the absolute power levels are still low. This provides perhaps the most telling demonstration of the difficulty of distinguishing the random walk model of stock prices from alternatives that imply highly persistent, yet transitory, price components. Even the best possible tests have very low power.

1.3 Evaluating Statistical Significance

The preceding discussion highlights the low power of available tests for the presence of transitory components in stock prices. One dramatic way of making this point is to note that using the conventional 5% significance level in choosing between the random walk hypothesis and our two alternatives involves in the best case a 76% probability of Type II error. For most tests, the Type II error rate would be between .85 and .95. Leamer (1978) echoes a point made in most statistics courses, but rarely heeded in practice, when he writes that "the [popular] rule of thumb, setting the significance level arbitrarily at .05, is ... deficient in the sense that from every reasonable viewpoint the significance level should be a decreasing function of sample size (p.92)."

How should a significance level be set? This is obviously a matter of judgment. Figure 1 depicts the attainable tradeoff between Type I and Type II errors for the most powerful variance ratio and regression tests, as well as for the likelihood ratio test against the alternative hypothesis that the data are generated by an ARMA(1,1) process with three quarters of the monthly return variation due to transitory price components. As our previous discussion suggests, the power curve for the variance ratio test lies everywhere between the frontiers attainable using regression and likelihood ratio tests. For the variance ratio test, a .40 significance level is appropriate if the goal is to minimize the sum of Type I and Type II errors. In order to justify using the conventional .05 test, one would have to assign three times as great a cost to Type I as to Type II errors.

Unless one is strongly attached to the random walk hypothesis, significance levels in excess of .05 seem appropriate in evaluating the importance of tran-

Type II vs. Type I Error

Alternative Tests of Mean Reversion




Figure 1

sitory components in stock prices. We see little basis for strong attachment to the null hypothesis. Many plausible alternative models of asset pricing, involving rational and irrational behavior, suggest the presence of transitory components. Furthermore, since the same problems of statistical power which plague our search for transitory components also complicate the lives of speculators, it may be difficult for speculative behavior to eliminate these transitory components. The only real solution to the problem of "low power" is the collection of more data. In the next section, we try to bring to bear as much data as possible in evaluating the importance of transitory stock price components.

2. Statistical Evidence on Mean Reversion

This section uses variance ratio tests to analyze the importance of stationary components in stock prices. We focus primarily on excess and real returns rather than nominal returns. Fama and French (1986a) and Lo and MacKinlay (1987) work with nominal returns, so they are implicitly testing the hypothesis that nominal ex-ante returns are constant.⁸ It seems more natural to postulate that the required risk premium is constant, or that the required real return is constant.

We analyze four major data sets. The first consists of monthly returns on the New York Stock Exchange for the period since 1926. These data have been used in other studies of mean reversion and are presented in part to demonstrate our comparability with previous work. Our second data set includes annual returns on the Standard and Poor's stock price indices for the period since 1871. Although these data are less reliable than the monthly CRSP data, they are available for a much longer period. Third, we analyze post-war monthly stock returns for seventeen stock markets outside the United States. Finally, we consider data on individual firms in the United States for the post-1926 period to explore both mean reversion in individual share prices, and to study whether share prices tend to revert to a market average.

2.1 Monthly NYSE Returns, 1926-1985

We begin by analyzing monthly returns on both the value-weighted and equal-weighted NYSE indices from the Center for Research in Security Prices data base for the 1926-1985 period. We consider nominal returns on these indices, excess returns with the risk-free rate measured as the Treasury bill yield, as well as real returns measured using the CPI inflation rate.

The variance ratio statistics for these series are shown in Table 2. We confirm Fama and French's (1986a) finding that returns at long horizons exhibit negative serial correlation, as reflected in values of the variance ratio far below unity. The same findings obtain for both real and excess returns. Typically, the results indicate that the variance of eight year returns is about four rather than eight times the variance of one year returns. The point estimates thus suggest more mean reversion in stock prices than the examples of the previous section where transitory components accounted for three quarters of the variance in returns. Despite the low power of our tests, the null hypothesis of serial independence is rejected at the .08 level for value-weighted excess returns, and the .005 level for equal-weighted excess returns.⁹ Mean reversion is more pronounced for the equal-weighted than for the value-weighted index, but the variance ratios at long horizons are well below unity for both indices.

The variance ratios also suggest that at horizons shorter than one year, there is some positive autocorrelation in returns. The variance of the one month return on the equal weighted index is only .79 times as large as it would be predicted to be given the variability of twelve-month returns. A similar conclusion applies to the value-weighted index. This finding of first positive then negative serial correlation accounts for Lo and MacKinlay's (1987a) result that variance ratios exceed unity in their weekly data, while variance ratios fall below one in other studies concerned with longer horizons.¹⁰

An issue that arises in analyzing results for the CRSP sample is the sensitivity of the findings to inclusion or exclusion of the Depression years. A number of previous studies, such as Officer (1973), have documented the unusual behavior of stock prices during the early 1930s, and one could make a serious

Table 2: Variance Ratios for U.S. Monthly Data, 1926-1985

Data Series	Annual Return Standard Deviation	1 Month	24 Months	36 Months	48 Months	60 Months	72 Months	84 Months	96 Months
Value-Weighted Nominal Returns	20.1%	0.783 (0.150)	1.029 (0.108)	0.974 (0.177)	0.883 (0.232)	0.802 (0.278)	0.707 (0.320)	0.586 (0.358)	0.542 (0.394)
Value-Weighted Excess Returns	20.6%	0.776 (0.150)	1.030 (0.108)	0.971 (0.177)	0.874 (0.232)	0.786 (0.278)	0.686 (0.320)	0.559 (0.358)	0.509 (0.394)
Value-Weighted Real Returns	20.7%	0.764 (0.150)	1.036 (0.108)	0.989 (0.177)	0.917 (0.232)	0.855 (0.278)	0.781 (0.320)	0.689 (0.358)	0.677 (0.394)
Equal-Weighted Nominal Returns	29.0%	0.800 (0.150)	1.002 (0.108)	0.912 (0.177)	0.851 (0.232)	0.750 (0.278)	0.598 (0.320)	0.418 (0.358)	0.335 (0.394)
Equal-Weighted Excess Returns	29.6%	0.798 (0.150)	0.998 (0.108)	0.902 (0.177)	0.831 (0.232)	0.721 (0.278)	0.563 (0.320)	0.378 (0.358)	0.290 (0.394)
Equal-Weighted Real Returns	29.6%	0.785 (0.150)	1.010 (0.108)	0.925 (0.177)	0.878 (0.232)	0.786 (0.278)	0.649 (0.320)	0.487 (0.358)	0.425 (0.394)

Notes: Calculations are based on the monthly returns for the value-weighted and equal-weighted NYSE portfolios, as reported in the CRSP monthly returns file. Values in parentheses are Monte Carlo estimates of the standard deviation of the variance ratio, based on 25000 replications. Each variance ratio is corrected for small sample bias and has a mean of unity under the null hypothesis of no serial correlation.

argument for excluding these years from analyses designed to shed light on current conditions. The counter-argument, suggesting this period should be included, is that the 1930s by virtue of the large movements in prices contain a great deal of information about the persistence of price shocks. We explored the robustness of our findings by truncating the sample period at both the beginning and the end. Excluding the first ten years of the sample slightly weakens the evidence for mean-reversion at long horizons. The negative serial correlation in nominal returns is virtually unaffected by this sample change, and the results for both equal-weighted real and excess returns are also quite robust. The long-horizon variance ratios for real and excess returns on the value-weighted index rise substantially, however. The 96-month variance ratios are .94 and 1.07 for these two return series, compared with .71 and .54 for the real and excess returns on the equal-weighted index. Truncating the sample to exclude the last ten years of data has an opposite effect on the estimated variance ratios; the evidence for mean reversion is even more pronounced than for the full sample period. The postwar period, another subsample we analyzed, displays less mean reversion than the full sample or the post-1936 period.

2.2 Historical Data for the United States

The CRSP data are the best available for analyzing recent U.S. experience, but the low power of the available statistical tests suggests the value of examining other data as well. This also reduces the data-mining risks stressed by Merton (1987). We therefore consider returns based on the combined Standard and Poors'/Cowles Commission stock price indices that are available beginning in 1871. These data have recently been checked and corrected for errors by Jones

and Wilson (1986); we use the series they report for the pre-1926 period. We analyze annual return series for the period 1871-1925, as well as the longer 1871-1985 period. The S&P data have the advantage of being used relatively infrequently in studies of the serial correlation properties of stock returns.¹¹ We again consider nominal, excess, and real returns.

The results are presented in Table 3. For the pre-1925 period, the nominal and excess returns display pronounced negative serial correlation at long horizons. For the real returns, however, this pattern is much weaker. Although the explanation of this phenomenon is unclear, it appears to result from the jagged character of the Consumer Price Index series in the years before 1900. The ex post inflation rate may prove a particularly unreliable measure of expected inflation during this period. The three lower rows in Table 3 present results for the full 1871-1985 sample period. All three return series show negative serial correlation at long lags, but real and excess returns provide less evidence of mean-reversion than the monthly post-1925 CRSP data.

2.3 Equity Markets Outside the United States

Additional evidence on mean reversion can be obtained by analyzing the behavior of equity markets outside the United States. We analyze returns in Canada for the period since 1919, in Britain since 1939, and in fifteen other nations for shorter post-war periods. The Canadian data set consists of monthly capital gains on the Toronto Stock Exchange. The British data are monthly returns, inclusive of dividends, on the Financial Times-Actuaries Share Price Index.

Results for these two equity markets are shown in the first two rows of Table 4. Both markets display substantively important mean reversion at long

Table 3: Variance Ratios for Historical U.S. Data

<u>Data Series</u>	<u>Annual Return Standard Deviation</u>	<u>24 Months</u>	<u>36 Months</u>	<u>48 Months</u>	<u>60 Months</u>	<u>72 Months</u>	<u>84 Months</u>	<u>96 Months</u>
Nominal Returns 1871-1925	15.6%	0.913 (0.140)	0.607 (0.210)	0.582 (0.265)	0.571 (0.313)	0.449 (0.358)	0.400 (0.398)	0.441 (0.436)
Excess Returns 1871-1925	16.2%	0.915 (0.140)	0.612 (0.210)	0.591 (0.265)	0.582 (0.313)	0.464 (0.358)	0.410 (0.398)	0.441 (0.436)
Real Returns 1871-1925	17.2%	0.996 (0.140)	0.767 (0.210)	0.806 (0.265)	0.820 (0.313)	0.737 (0.358)	0.710 (0.398)	0.807 (0.436)
Nominal Returns 1871-1985	18.2%	1.035 (0.095)	0.895 (0.143)	0.885 (0.179)	0.826 (0.211)	0.756 (0.240)	0.688 (0.266)	0.689 (0.290)
Excess Returns 1871-1985	18.9%	1.047 (0.095)	0.922 (0.143)	0.929 (0.179)	0.900 (0.211)	0.856 (0.240)	0.807 (0.266)	0.833 (0.290)
Real Returns 1871-1985	19.0%	1.035 (0.095)	0.880 (0.143)	0.876 (0.179)	0.842 (0.211)	0.797 (0.240)	0.756 (0.266)	0.781 (0.290)

Notes: Each entry is a bias-adjusted variance ratio with a mean of unity under the null hypothesis. Values in parentheses are Monte Carlo standard deviations of the variance ratio, based on 25000 replications. The underlying data are annual returns on the Standard and Poor's Composite stock index, backdated to 1871 using the Cowles data as reported in Jones and Wilson (1987).

horizons. In the Canadian data, the 96 month variance ratio is .585, while for the London data it is .794. The pattern of variance ratios shows slow decay in the Canadian case, while for the British market the variance ratio falls to .74 at the four year horizon and does not rise significantly at longer horizons. The Canadian data show statistically significant positive serial correlation at lags of less than twelve months. The one month variance is only .718 times the value that would be predicted based on the twelve month variance. For the British data, the one month variance is .832 times the twelve-month variance, suggesting somewhat less positive serial correlation at short horizons.

Table 4 also reports results for fifteen other stock markets outside the United States. The variance ratios are calculated from monthly returns that are based on stock price indices in the International Monetary Fund's International Financial Statistics publication. Unfortunately, the IFS does not provide dividend yields, so the reported returns correspond to capital gains alone. To assess the importance of this omission, we re-estimated the variance ratios for dividend-exclusive CRSP monthly data, and dividend-exclusive British stock market data. The results of these variance ratio calculations, presented in Appendix Table A1, show only minor differences as a result of dividend omission. Yield-inclusive data would surely be superior to the monthly returns we use, but we suspect that our results would be affected in only minor ways.¹²

Although the variance ratios for individual countries, based typically on data starting in 1957, have larger standard errors than the results for the U.S., Britain, or Canada, they are remarkably similar. Most of the countries display negative serial correlation at long horizons. In Germany, for example, the 96-month variance ratio is .462; in France it is .438. Only three of the

Table 4: Variance Ratios for International Data on Real Monthly Returns

<u>Return Series</u> <u>Country/Sample</u>	<u>Annual Return</u> <u>Standard Deviation</u>	<u>1 Month</u>	<u>24 Months</u>	<u>36 Months</u>	<u>48 Months</u>	<u>60 Months</u>	<u>72 Months</u>	<u>84 Months</u>	<u>96 Months</u>
Canada/1919-1986 (Capital Gains Only)	20.1%	0.711 ¹ (0.141)	1.055 (0.102)	0.998 (0.166)	0.912 (0.218)	0.799 (0.261)	0.692 (0.301)	0.617 (0.336)	0.575 (0.370)
U. K./1939-1986	20.9%	0.832 (0.167)	0.987 (0.125)	0.868 (0.198)	0.740 (0.259)	0.752 (0.317)	0.807 (0.358)	0.806 (0.400)	0.794 (0.440)

Austria/1957-1986 (Capital Gains Only)	21.4%	0.663 (0.214)	1.205 (0.156)	1.206 (0.254)	1.132 (0.334)	0.864 (0.403)	0.666 (0.464)	0.582 (0.518)	0.502 (0.566)
Belgium/1957-1986 (Capital Gains Only)	17.0%	0.718 ¹ (0.214)	1.054 (0.156)	1.137 (0.254)	1.121 (0.334)	1.060 (0.403)	0.876 (0.464)	0.807 (0.518)	0.776 (0.566)
Colombia/1959-1983 (Capital Gains Only)	21.4%	1.223 (0.214)	0.822 (0.156)	0.743 (0.254)	0.724 (0.334)	0.583 (0.403)	0.477 (0.464)	0.386 (0.518)	0.180 (0.566)
Germany/1957-1986 (Capital Gains Only)	23.8%	0.610 (0.214)	1.309 (0.156)	1.251 (0.254)	0.987 (0.334)	0.747 (0.403)	0.687 (0.464)	0.581 (0.518)	0.462 (0.566)
Finland/1957-1986 (Capital Gains Only)	22.1%	0.504 (0.214)	1.141 (0.156)	1.262 (0.254)	1.396 (0.334)	1.463 (0.403)	1.381 (0.464)	1.215 (0.518)	1.014 (0.566)
France/1957-1986 (Capital Gains Only)	23.6%	0.874 ¹ (0.214)	0.961 (0.156)	0.914 (0.254)	0.865 (0.334)	0.607 (0.403)	0.460 (0.464)	0.433 (0.518)	0.438 (0.566)
India/1957-1986 (Capital Gains Only)	15.6%	0.752 (0.214)	0.985 (0.156)	0.974 (0.254)	0.942 (0.334)	0.823 (0.403)	0.767 (0.464)	0.619 (0.518)	0.596 (0.566)
Japan/1957-1986 (Capital Gains Only)	20.0%	0.870 (0.214)	1.135 (0.155)	1.015 (0.254)	0.927 (0.334)	0.803 (0.403)	0.691 (0.464)	0.595 (0.518)	0.538 (0.567)
Netherlands/1957-1986 (Capital Gains Only)	20.0%	0.710 (0.214)	1.238 (0.155)	1.263 (0.254)	1.217 (0.334)	1.083 (0.403)	1.047 (0.464)	0.894 (0.518)	0.741 (0.567)

Table 4: Variance Ratios for International Data on Real Monthly Returns (Continued)

Return Series Country/Sample	Annual Return Standard Deviation	1 Month	24 Months	36 Months	48 Months	60 Months	72 Months	84 Months	96 Months
Norway/1957-1986 (Capital Gains Only)	24.2%	0.601 (0.214)	1.033 (0.155)	0.961 (0.254)	0.926 (0.334)	0.844 (0.403)	0.825 (0.464)	0.840 (0.518)	0.784 (0.567)
Philippines/1957-1986 (Capital Gains Only)	29.7%	0.910 (0.214)	0.908 (0.155)	0.749 (0.254)	0.707 (0.334)	0.703 (0.403)	0.839 (0.464)	0.898 (0.518)	0.887 (0.567)
South Africa/1957-1986 (Capital Gains Only)	23.2%	0.767 (0.214)	1.151 (0.155)	1.063 (0.254)	0.963 (0.334)	0.980 (0.403)	1.090 (0.464)	1.131 (0.518)	1.151 (0.567)
Spain/1961-1986 (Capital Gains Only)	27.7%	0.603 (0.230)	1.289 (0.166)	1.584 (0.273)	1.831 (0.359)	2.008 (0.433)	2.246 (0.498)	2.347 (0.556)	2.373 (0.609)
Sweden/1957-1986 (Capital Gains Only)	21.1%	0.728 (0.214)	0.898 (0.155)	0.822 (0.254)	0.901 (0.334)	0.885 (0.403)	0.916 (0.464)	0.760 (0.518)	0.629 (0.567)
Switzerland/1957-1986 (Capital Gains Only)	21.5%	0.789 (0.214)	1.343 (0.155)	1.395 (0.254)	1.300 (0.334)	1.034 (0.403)	0.749 (0.464)	0.489 (0.518)	0.382 (0.567)
U.S./1957-1986 (Capital Gains Only)	16.6%	0.813 (0.214)	0.814 (0.155)	0.653 (0.254)	0.656 (0.334)	0.696 (0.403)	0.804 (0.464)	0.803 (0.518)	0.800 (0.567)
Average Value (Excluding U.S.)		0.757 (0.051)	1.089 (0.037)	1.071 (0.060)	1.035 (0.079)	0.943 (0.095)	0.895 (0.110)	0.824 (0.123)	0.754 (0.134)
Average Value (Excluding U.S., Spain)		0.766 (0.052)	1.077 (0.038)	1.039 (0.062)	0.985 (0.081)	0.877 (0.098)	0.811 (0.112)	0.728 (0.125)	0.653 (0.137)

Notes: Values in parentheses are Monte Carlo standard deviations of the variance ratio statistics. The data for all markets except Britain and Canada are drawn from the International Monetary Fund, International Financial Statistics. In all cases except those marked with a *, the data are monthly averages of daily or weekly values. The U.K. data are point-sampled but only at the end of each year. The variance ratios are corrected for the time aggregation bias induced by averaging closing values of the index within each month.

fifteen countries have variance ratios at 96 months that exceed unity, and many are substantially below. Evidence of positive serial correlation at short horizons is also pervasive. In only one country, Colombia, is the variance ratio at one month greater than unity. The short data samples make it extremely difficult to reject the null hypothesis of serial independence for any individual country. Nonetheless, the similarity of the results for the majority of nations supports our earlier conclusion of potentially important transitory price components.

The average variance ratios at each horizon are shown in the last two rows of the table. The mean 96-month variance ratio is .754 when all countries are aggregated, and .653 when we exclude Spain (which is clearly an outlier, probably because of the unusual pattern of hyperinflation followed by deflation that it experienced during our sample). By averaging across many countries, we also obtain a more precise estimate of the long-horizon variance ratios. The standard error of the Spain-exclusive average for the 96-month variance ratio is .142 assuming that the variance ratios for different countries are independent. If we assume that these statistics have a correlation of .25, however, the standard error rises to .326, again implying that the null hypothesis of serial independence would not be rejected at standard levels.¹³ The qualitative results on positive autocorrelation at short horizons and negative autocorrelation at long lags are, however, supportive of our qualitative findings using CRSP data.

2.4. Individual Firm Data

We also consider evidence on mean reversion for individual firms. It is much less plausible on a priori grounds to expect transitory components in the

relative prices of individual stocks than in the market as a whole. Arbitrageurs should find the task of trading in individual securities to correct mispricing far easier than taking positions in the entire market to offset persistent misvaluations. In spite of this, some previous work has suggested that individual stock returns may exhibit negative serial correlation. Miller and Scholes (1982), for example, show that regressing ex post returns on the reciprocal of the stock price yields a significant negative coefficient. Since the reciprocal price is close to the cumulative value of past returns, this indicates higher returns after periods of poor performance.

We examine the 82 firms in the CRSP monthly master file that have no missing return information between 1926 and 1985. There are a number of obvious biases in a sample of this type. It is weighted toward large firms that have been traded actively over the entire period. Firms that went bankrupt or began trading during the sample period are necessarily excluded. Since the value weighted NYSE index shows less mean reversion than the equal weighted index, our sample of 82 large firms might display less mean reversion than a sample of smaller stocks traded over shorter periods. For these 82 firms, however, we compute variance ratios using both nominal and real returns. Because the returns on different firms are not independent, we also examine the returns on portfolios formed by buying one dollar of each firm, and short-selling \$82 of the aggregate market. That is, we examine properties of the time series $R_{it} - R_{mt}$ where R_{mt} is the value-weighted NYSE return.¹⁴

The mean values of the individual firm variance ratios are shown in Table 5. They suggest some long-horizon mean reversion for individual stock prices relative to the overall market or relative to a risk-free return. Although the

Table 5: Average Variance Ratios for Individual Company Monthly Returns, 1926-1985

<u>Return Concept</u>	<u>1 Month</u>	<u>24 Months</u>	<u>36 Months</u>	<u>48 Months</u>	<u>60 Months</u>	<u>72 Months</u>	<u>84 Months</u>	<u>96 Months</u>
Nominal Returns	0.954 (0.017)	1.031 (0.012)	0.992 (0.020)	0.932 (0.026)	0.861 (0.031)	0.782 (0.035)	0.706 (0.040)	0.678 (0.044)
Excess Returns Relative to Risk-Free Rate	0.942 (0.017)	1.035 (0.012)	1.000 (0.020)	0.950 (0.026)	0.888 (0.031)	0.820 (0.035)	0.755 (0.040)	0.739 (0.044)
Excess Returns Relative to Value-Weighted NYSE	1.088 (0.017)	1.034 (0.012)	1.019 (0.020)	1.002 (0.026)	0.968 (0.031)	0.928 (0.035)	0.898 (0.040)	0.886 (0.044)

Notes: Each entry reports the average of the variance ratios calculated for the 82 firms on the monthly CRSP returns file with continuous data between 1926:1 and 1985:12. Values in parentheses are Monte Carlo standard deviations of the average variance ratio, computed under the assumption that the variance ratios for different firms are independent.

point estimates suggest that only twelve percent of the 8-year variance in firm excess returns is due to stationary factors, the increased precision gained by studying returns on many independent firms enables us to reject the null hypothesis that all of the return variation arises from non-stationary factors. However, there is also much less evidence than for the market aggregate of positive short-run serial correlation in excess returns, since the one-month variance ratios are close to unity.

2.5. Summary

The power calculations of the last section demonstrate the difficulty of detecting mean reversion in stock prices. Given the low power of available tests, our results are quite striking. The point estimates generally suggest that over long horizons return variance rises less than proportionally with time, and in many cases imply more mean reversion than our examples in the last section where transitory factors accounted for three-fourths of the variation in returns. Many of the results imply rejections of the null hypothesis of serial independence at the .15 level, a level that may not be inappropriate given our previous discussion of size vs. power tradeoffs. Furthermore, each of the different types of data we analyze provides evidence of some deviation from serial independence in stock returns. Taken together, the results are stronger than any individual finding.

It is interesting to note that there is a clear tendency for more mean reversion in less broad-based and sophisticated equity markets. The U.S. data before 1925 show greater evidence of mean-reversion than the post-1926 data, especially when we recognize that the appropriate comparison series for the

Standard and Poor's index is the value-weighted NYSE. The equal-weighted portfolio of NYSE stocks exhibits more mean reversion than the value-weighted portfolio. In recent years, mean reversion is more pronounced in foreign countries with less sophisticated equity markets than the United States.

3. The Substantive Importance of Transitory Components in Stock Prices

Our discussion so far has focused on the strength of the statistical evidence regarding transitory price components. This section uses our point estimates of the degree of mean reversion in stock prices to assess their substantive importance. One possible approach would involve calibrating models of the class considered in the first section. We do not follow this strategy because our finding of positive autocorrelation over short intervals implies that the AR(1) specification of the transitory component is inappropriate. Instead, we use an approach that does not require us to specify a process for the transitory component, but allows us to focus on its standard deviation and the fraction of the variance in one period returns that can be attributed to it.

We treat stock prices p_t as the sum of a permanent component and a transitory component. The permanent component evolves as a random walk and the transitory component follows a stationary process, $A(L)u_t = v_t$. This decomposition may be given two (not necessarily mutually exclusive) interpretations. First, u_t may reflect "fads", speculation-induced deviations of prices from fundamental values. Second, u_t may be a consequence of changes in required returns. In either case, describing the stochastic properties of u_t is a way of characterizing the part of stock price movements that cannot be explained on the basis of changing expectations about future cash flows.

Given our assumptions, the variance of T period returns is:

$$(11) \quad \sigma_T^2 = T\sigma_\epsilon^2 + 2(1-\rho_T)\sigma_u^2.$$

where σ_ϵ^2 is the variance of innovations to the permanent price component, σ_u^2 is the variance of the stationary component, and ρ_T is the T-period autocorrelation of the stationary component. Given data on the variance of returns over

two different horizons T and T' , and assumptions about ρ_T and $\rho_{T'}$, a pair of equations with the form (11) can be solved to yield estimates of σ_ϵ^2 and σ_u^2 . Using σ_R^2 to denote the variance of one period returns, and $VR(T)$ the T -period variance ratio as in (1), estimates of σ_ϵ^2 and σ_u^2 are given by:

$$(12a) \quad \sigma_\epsilon^2 = \frac{\sigma_R^2 [VR(T)(1-\rho_{T'})T - VR(T')(1-\rho_T)T']}{2(1-\rho_{T'})T - (1-\rho_T)T'}$$

and

$$(12b) \quad \sigma_u^2 = \frac{\sigma_R^2 T' [VR(T) - VR(T')] T}{2(1-\rho_{T'})T - (1-\rho_T)T'}$$

Many pairs of variance ratios and assumptions about the serial correlation properties of u_t could be analyzed using (12a-b). We begin by postulating that u_t is serially uncorrelated at the horizon of 96 months. For various degrees of serial correlation at other horizons, we can then estimate the variance of the transitory component, σ_u^2 and the share of the return variation due to transitory components, $1 - \sigma_\epsilon^2/\sigma_R^2$. We focus on different possible values of ρ_{12} , the twelve-month autocorrelation in u_t , and present estimates based on values of 0, .35, and .70. The findings were not especially sensitive to our choice of ρ_{96} ; our table reports values of 0, .15, and .30.

Table 6 presents estimates of the standard deviation of the transitory component in stock prices for the value weighted and equal weighted U.S. portfolios over the period 1926-1985 for various values of ρ_{12} . Particularly for the equal weighted portfolio, the transitory component in stock prices is of substantial importance. Depending on our assumption about its serial correlation properties, it accounts for between 43 and 99 percent of the variance in equal

Table 6: Permanent and Transitory Return Components, U.S. Monthly Data

CRSP Return Series	$\rho_{12} = 0.0$		$\rho_{12} = .35$		$\rho_{12} = .70$	
	σ_u	$1 - \frac{\sigma_u^2}{\sigma_R^2}$	σ_u	$1 - \frac{\sigma_u^2}{\sigma_R^2}$	σ_u	$1 - \frac{\sigma_u^2}{\sigma_R^2}$
Value-Weighted Excess Returns:						
$\rho_{96} = 0.00$	9.7%	0.369	12.5%	0.400	21.6%	0.554
$\rho_{96} = 0.15$	-----	-----	12.3%	0.386	20.5%	0.500
$\rho_{96} = 0.30$	-----	-----	12.1%	0.373	19.6%	0.456
Equal-Weighted Excess Returns:						
$\rho_{96} = 0.0$	16.8%	0.657	21.7%	0.712	37.7%	0.986
$\rho_{96} = .15$	-----	-----	21.4%	0.687	35.8%	0.890
$\rho_{96} = .30$	-----	-----	21.0%	0.664	34.2%	0.812

Note: Calculations assume that the transitory components exhibit zero auto-correlation at lags of 96 or 72 respectively. σ_u is the standard deviation of the transitory return component, while σ_u^2/σ_R^2 is the share of the variation in annual returns that is accounted for by permanent factors.

weighted monthly returns, and has a standard deviation of between 14 percent and 37 percent. Results for the value weighted portfolio similarly suggest that the transitory component accounts for a large, though smaller, portion of the variance in returns. Some estimates are as high as 56 percent.

As one would expect, Table 6 indicates that increasing the assumed persistence of the transitory component raises both its standard deviation and its contribution to the return variance. More persistent transitory components are less able to account for declining variance ratios at long horizons. Therefore, to rationalize any given long horizon variance ratio, increasing the persistence of the transitory component will require increasing the weight on the transitory component relative to the permanent component. Sufficiently persistent transitory components will be unable to account for low long horizon variance ratios, even if they account for all of the variation in returns. For example, a transitory component that is almost as persistent as a random walk will be unable to explain very much long-horizon mean reversion.

Which cases in Table 6 are most relevant? As an a priori matter, it is difficult to see a compelling argument for assuming that transitory components should die out extremely quickly. Previous suggestions that there are "fads" in stock prices have typically suggested half lives of several years, implying that the elements in the table corresponding to $\rho_{12} = .70$ are most relevant. If the decay is geometric, this suggests a half life of two years for the transitory component. Even greater values of ρ_{12} might be even more plausible. One other consideration supports this conclusion. For given values of σ_e^2 and σ_u^2 , equation (11) permits us to calculate ρ_T over any horizon. A reasonable restriction, that ρ_T not be very negative over periods of up to 96 months, is only satisfied

for cases where ρ_{12} is large. For example, with $\rho_{96} = 0$, when we impose $\rho_{12} = .35$ the value of the stationary component's implied autocorrelation is $-.744$ at 36 months, -1.27 at 60 months, and $-.274$ at 84 months. Even more negative values obtain assuming $\rho_{12} = 0$, and the results are also less satisfactory when ρ_{96} is positive. In contrast, when $\rho_{12} = .70$ and $\rho_{96} = 0$, the implied values of ρ_{36} and ρ_{60} are $.168$ and $-.173$, respectively. Similar results obtain for other large values of ρ_{12} . This is because variance ratios continue to decline substantially between long and longer horizons, and as equation (11) demonstrates, rationalizing this requires declining values of ρ_T . If ρ_T starts small, it therefore must become negative to account for the observed variance ratio pattern. Imposing larger autocorrelations at short horizons does not necessitate such autocorrelation patterns.

Since other countries and historical periods exhibit patterns of variance ratio decline that are similar to those in the American data, we do not present calculations similar to those in Table 6 for them. As one would expect, countries with 96-month variance ratios lower than those for the United States have larger transitory components than the U.S., and vice versa.

Insofar as the evidence in the first section and in Fama and French (1987) is persuasive in suggesting that transitory components in stock prices are present and statistically significant, this section's results confirm Shiller's (1981) conclusion that models assuming constant ex-ante returns cannot account for a large fraction of the variance in stock market returns. Stock market volatility is excessive relative to the predictions of these models.¹⁵ Since our analysis uses only returns data and does not exploit the present value relationship between stock prices and expected future dividends, it does not suffer from some of the problems that have been highlighted in the volatility test debate.

4. The Source of the Transitory Component in Stock Prices

If stock prices have a transitory component, ex-ante returns must vary.¹⁶ Any stochastic process for the transitory component can be mapped into a stochastic process for ex-ante returns, and any pattern for ex-ante returns can alternatively be represented by describing the associated transitory component of prices. The economically interesting issue is whether variations in ex-ante returns are better explained by "fundamentals" such as changes in interest rates or volatility¹⁷, or instead as byproducts of price deviations caused by noise traders.¹⁸ This section notes a number of considerations that incline us toward the latter view.

4.1 How Variable Must Risk Premia Be?

It is instructive to calibrate the amount of variation in expected returns that risk factors would have to generate in order for them to account for the observed transitory components in stock prices. To do this we assume for simplicity that the transitory component follows an AR(1) process as postulated in the "fads" example of Summers (1986). This has the virtue of tractability, although it is inconsistent with the observation that actual returns exhibit positive, then negative, serial correlation.

Changes over time in the required return on common stocks can generate mean-reverting stock price behavior. If required returns exhibit positive autocorrelation, then an innovation that raises required returns will reduce share prices. This will generate a holding period loss, followed by higher returns in subsequent periods. We show in the appendix that when required returns follow an AR(1) process¹⁹, then

$$(13) \quad R_t - \bar{R} \cong \frac{1+\bar{g}}{1+\bar{r}-\rho_1(1+\bar{g})} (r_t - \bar{r}) - \frac{(1+\bar{r})^{-1}(1+\bar{g})^2}{1+\bar{r}-\rho_1(1+\bar{g})} (r_{t+1} - \bar{r}) + \zeta_t$$

where ζ_t , a serially uncorrelated innovation that is orthogonal to innovations about the future path of required returns (ξ_t), reflects revisions in expected future dividends. The constants in this expression depend upon \bar{d} and \bar{g} , the average dividend yield and dividend growth rate respectively. In steady state $\bar{r} = \bar{d} + \bar{g}$.

If changes in required returns and profits are positively correlated, as is plausible given the importance of shocks to the perceived productivity of capital, then the assumption that ξ_t and ζ_t are orthogonal will understate the variance in ex-ante returns needed to rationalize mean reversion in stock prices. Although it is possible to construct theoretical examples where profits and interest rates are negatively related, as in Campbell (1986), the empirical finding that bond and stock returns are weakly correlated suggests positive correlation between shocks to cash flows and required returns.²⁰ Negative correlation between ζ_t and ξ_t would cause bond and stock returns to move together.

Our assumption that required returns are given by $(r_t - \bar{r}) = (1 - \rho_1 L)^{-1} \xi_t$ enables us to rewrite (13), defining $\tilde{\xi}_t \equiv -\xi_{t+1} (1+\bar{r})^{-1} (1+\bar{g})^2 / [1+\bar{r}-\rho_1(1+\bar{g})]$, as

$$(14) \quad (1-\rho_1 L)(R_t - \bar{R}) \cong \tilde{\xi}_t + \zeta_t - (1+\bar{d})\tilde{\xi}_{t-1} - \rho_1 \zeta_{t-1}.$$

The first order autocovariance of the expression on the right-hand side of (14) is nonzero, but all higher-order autocovariances equal zero. Ansley, Spivey, and Wroblewski (1977) show that this implies that the right hand side of (14) is an MA(1) process that can be represented as $(1+\theta L)w_t$. Provided $\sigma_\xi^2 > 0$, this implies that returns follow an ARMA(1,1) process; if $\sigma_\xi^2 = 0$, then returns are white noise.

The simple model of stationary and nonstationary price components summarized in equation (9) also yields an ARMA(1,1) representation for returns. This allows us to calculate the amount of variation in required returns that would be needed to generate the same time series process for observed returns as would be generated by "fads" of various sizes in equation (6). We measure the size of fads, or transitory factors more generally, by σ_u , the standard deviation of the transitory component. In the appendix we show that the required return variance corresponding to a given fad variance is:

$$(15) \quad \sigma_r^2 = \frac{[1+\bar{r}-\rho_1(1+\bar{g})]^2(1-\rho_1)^2(1+\bar{r})^2}{\{(1+\bar{d})(1+\rho_1^2)-\rho_1[1+(1+\bar{d})^2]\}(1+\bar{g})^2} \sigma_u^2.$$

This expression indicates the variation in required returns needed to generate transitory components of a given size.

Table 7 reports the standard deviation of required excess returns, measured on an annual basis, implied by a variety of different fad models. We calibrate the calculations using the average excess return (8.9% per year) on the NYSE equal-weighted share price index over the 1926-1985 period. The dividend yield on these shares averages 4.5%, implying an average dividend growth rate of 4.4%. We use our estimates of the variance ratio at 96 months (from Table 2) to calibrate the degree of mean reversion.

The findings suggest that a great deal of variability in required returns is needed to explain the degree of mean reversion in prices. For example, if we postulate that the standard deviation of the transitory price component is 20%, then even when required return shocks have a half life of 2.9 years, the standard deviation of ex ante returns (at an annual frequency) must be 5.8%. Even

Table 7: Time-Varying Return Models Needed to Account for Mean Reversion

	<u>Standard Deviation of Transitory Component</u>			
<u>Half Life</u>	<u>15.0%</u>	<u>20.0%</u>	<u>25.0%</u>	<u>30.0%</u>
1.4 Years	7.9%	10.6%	13.2%	15.8%
1.9 Years	6.1%	8.2%	10.2%	12.3%
2.9 Years	4.4%	5.8%	7.3%	8.7%

Notes: Each entry indicates the standard deviation of required returns, assuming that required returns follow an AR(1) process with the half life indicated in the left margin. The calculations are calibrated using data on excess returns for the equal-weighted NYSE index over the 1926-1985 period. The average excess return for this period is 8.9% per year, with a dividend yield of 4.5%.

larger amounts of required return variation are needed to explain the same size price fads when the persistence of required return shocks is lower. These estimates of the standard deviation of required returns are large relative to the mean of ex post excess returns. If ex ante returns are never negative, they imply that ex ante returns must exceed 20% fairly frequently.

It is difficult to think of risk factors that could account for such large variations in required returns. Campbell and Shiller's (1986) conclusion that stock price movements have no predictive power for changes in discount rates is especially relevant in this context. They reason that if stock price movements are caused by changes in future discount rates, then realized values of future discount rates should be Granger caused by stock prices. They find no evidence that this is the case using data on real interest rates and market volatilities. While they find evidence that stock prices Granger cause consumption, the sign is counter to the theory's prediction.

4.2 Negative Ex-Ante Returns

The principle restriction implied by homogeneous expectations models of financial markets is that ex-ante returns conditional on public information can never be negative. This is not a property of some models with noise traders. Sufficiently optimistic noise traders may drive prices high enough to make ex-ante returns negative, so risk averse speculators may not be willing or able to short the market to the point where ex-ante returns are driven to zero.

There is an obvious problem with evaluating whether or not ex-ante returns are ever negative. In estimating any model, there is a possibility of overfitting that may cause the spurious appearance of negative ex-ante returns.

This is especially likely when many parameters are present. Table 8 therefore presents the fraction of ex-ante returns that were negative when various autoregressive models were estimated using the CRSP returns data for 1926-1985 along with Monte-Carlo calculations of the share of negative ex-ante returns that resulted when similar models were estimated using serially independent returns.

The results indicate that negative ex-ante returns show up reasonably frequently, and to a greater extent than can be explained by statistical overfitting. For example, in a regression of monthly excess returns on 24 lags, using the equal weighted data, 31.2 percent of the fitted values are negative compared with 25 percent in the corresponding Monte-Carlo calculation. The corresponding values for the value-weighted index are 33.3% and 28.1%, respectively. Our Monte Carlo results show that the p-value associated with the outcome for the equal-weighted case is .172, and for the value-weighted case, .240. These results provide weak evidence against the risk factors hypothesis, since they suggest negative ex-ante returns in some periods.²¹ Since it is not possible to know in which periods ex ante returns are actually negative, they do not have strong implications for investment strategy.

4.3 The Difficulty of Accounting for the Observed Autocorrelogram

In section 2 we demonstrated that stock returns exhibited positive serial correlation over short periods and negative serial correlation over longer stretches. The AR(1) transitory components model treated in the previous subsection can rationalize the second but not the first of these observations. It is instructive to consider what type of behavior for expected returns is necessary to account for both observations.

Table 8: Incidence of Negative Predicted Values of Excess Returns

<u>Length of Autoregression</u>	<u>CRSP Value-Weighted</u>			<u>CRSP Equal-Weighted</u>		
	<u>Actual</u>	<u>Monte Carlo Mean</u>	<u>p-Value</u>	<u>Actual</u>	<u>Monte Carlo Mean</u>	<u>p-Value</u>
12 months	.282	.197	.202	.249	.168	.168
24 months	.333	.281	.240	.312	.250	.172
48 months	.353	.343	.412	.311	.322	.494

Notes: Each entry reports the fraction of predicted returns, computed as the fitted values from long autoregressive models for returns, that lie below zero. All calculations are based on 720 monthly observations from the CRSP monthly returns file spanning the 1926-1985 period.

Positive serial correlation in ex-post returns over short periods requires that increases in prospective required returns that reduce stock prices are followed by reduced returns. Shocks that have a large effect on the expected discounted present value of required returns must not have a large impact on required returns in the immediately succeeding periods. The impulse response function for required returns must cross the zero-axis; a positive current shock must lead to a reduction in required returns for a period in the immediate future, followed by higher required returns at a later date. The analysis in Poterba and Summers (1986) of the time series properties of volatility, as well as Litterman and Weiss' (1986) work on the time series properties of real interest rates, does not suggest that these series have the impulse response functions needed to account for the behavior of market returns.

5. Conclusions

The empirical results in this paper suggest that stock returns exhibit positive serial correlation over short periods and negative correlation over longer intervals. This conclusion emerges from the standard CRSP data on equal and value weighted returns over the 1926-1985 period. It is corroborated by data from the pre-1925 period, data for individual firms, and data on stock returns in seventeen foreign countries. While individual data sets do not consistently permit rejection of the random walk hypothesis at a high confidence level, the various data sets taken together establish a fairly strong case against its validity. Furthermore our point estimates generally suggest that the transitory component in stock prices is quantitatively important, accounting for the bulk of the variance in returns.

While the temptation to apply more sophisticated statistical techniques to stock return data in an effort to extract more information about the magnitude and structure of transitory components is ever present, we doubt that a great deal can be learned in this way. Even the relatively gross features of the data examined in this paper cannot be estimated precisely. As the debate over volatility tests has illustrated, sophisticated statistical results are often very sensitive to maintained assumptions that are difficult to evaluate. We have checked all the statistical procedures used in this paper by applying them to pseudo data conforming exactly to the random walk model. Our suspicion, supported by the dramatic illustrations in Kleidon (1986), is that a great deal of more elaborate work on stock price volatility could not meet this test.

We suggest in the paper's final section that noise trading provides a plausible account for the transitory components that are present in stock

prices. Pursuing this issue will involve constructing and testing theories of either noise trading or changing risk factors that can account for the characteristic stock return autocorrellogram documented here. Evaluating such theories is likely to require that information other than stock returns be studied. Such information might include information on fundamental values, proxies for noise trading such as odd-lot sales, and indicators of risk factors such as ex-ante volatilities implied from stock market options. Only by comparing the persuasiveness of models based on the presence of noise traders and models based on changing risk factors, can we come to a judgment about whether or not financial markets are efficient in the sense of rationally valuing assets, as well as in the sense of precluding the generation of easy excess profits.

Endnotes

1. This point is independent of Fischer's (1984) conclusion in analyzing the appropriate portfolio strategy for an investor concerned about terminal consumption but restricted in the frequency with which he can rebalance his portfolio. Our conclusion applies to investors who can rebalance continuously, as most investors probably can.

2. Stochastic speculative bubbles, considered by Blanchard and Watson (1982), could generate deviations between market prices and fundamental values without negative serial correlation in returns. In the presence of any limits on valuation errors set by speculators or real investment opportunities, however, such bubbles could not exist.

3. Testing the relationship between the variability of returns at different horizons has a long tradition: Osborne (1959) and Alexander (1961) apply tests similar to the variance ratio to much shorter data samples.

4. When in subsequent sections we examine annual returns data, the denominator of the variance ratio is simply $\text{Var}(R_t)$.

5. When the horizon of the variance ratio is large relative to the sample size, this bias can be substantial. For example, with $T=720$ and $k=60$, the bias is $-.069$. It rises to $-.160$ if $k=120$. Detailed Monte Carlo analysis of the variance ratio statistic may be found in Lo and MacKinlay (1987b).

6. The ARMA(1,1) parameters associated with this model, for the general ARMA $(1-\phi L)\Delta p_t = (1+\theta L)w_t$, are:

$$\phi = \rho_1$$

$$\theta = \{- (1+\rho_1^2) - 2\sigma_v^2 + (1-\rho_1)[4\sigma_v^2 + (1+\rho_1)^2]^{1/2}\} / (2\sigma_v^2 + 2\rho_1)$$

$$\sigma_w^2 = -(\rho_1 + \sigma_v^2) / \theta.$$

7. We compute the likelihood value under each hypothesis using the exact maximum likelihood method described in Harvey (1981). Because of the bias toward negative autocorrelations induced by estimating the mean return in each data sample, the mean likelihood ratios are actually above one for each of the hypotheses we consider.

8. Real returns are analyzed in Fama and French (1987), a revision of Fama and French (1986a).

9. These p-values are calculated from the empirical distribution of our test statistic, based on Monte Carlo results. They permit rejection at lower levels than would be possible using the normal approximation to the distribution of the variance ratio, along with the Monte Carlo estimates of the standard deviation of the variance ratio.

10. French and Roll (1986) apply variance ratio tests to daily returns for a sample of NYSE and AMEX stocks for the period 1963-1982. They find evidence of negative serial correlation especially among smaller securities. The divergence between their findings and those of Lo and MacKinlay (1987a) is presumably due to differences in the two data sets.

11. These data have been used in some studies of stock market volatility, such as Shiller (1981).

12. The monthly stock index data from the IFS also suffer from a second limitation. In many cases they are time averages of daily or weekly index values. Working (1960) showed that the first difference of a time-averaged random walk would exhibit positive serial correlation, with a first order autocorrelation coefficient of .25 as the number of observations in the average becomes large. This will bias our estimated variance ratios. For the countries with time aggregated data we therefore modify our small-sample bias correction. Instead of taking the expected value of the first-order autocorrelation to be $-1/(T-1)$ when evaluating $E(VR(k))$ we use $.25-1/(T-1)$. The reported variance ratios have been bias-adjusted by dividing by the resulting expected value.

13. The cross-country correlation of own-currency returns is typically below .25 in our data sample. Only five of the 171 pairwise correlations for the eighteen countries we consider (including the U.S.) exceed .50.

14. We also applied variance ratio tests to the residuals from the market model estimated for each firm, imposing a constant β for the entire 1926-1985 period. These residuals showed less correlation than the excess returns relative to the market, computed as $R_{it} - R_{mt}$.

15. Shiller's conclusion that market returns are too volatile to be reconciled with valuation models assuming constant required returns has been disputed by Kleidon (1986) and Marsh and Merton (1986). Mankiw, Romer, and Shapiro (1985) provide evidence in support of Shiller's conclusion.

16. Several recent studies have considered the extent to which equity returns can be predicted using various information sets. Keim and Stambaugh (1986) find that between eight and thirteen percent of the variation in returns for a portfolio of stocks in the bottom quintile of the NYSE can be predicted using lagged information. A much smaller share of the variation in returns to larger companies can be accounted for in this way. Campbell (1987) finds that approximately eleven percent of the variation in excess returns can be explained on the basis of lagged information derived from the term structure.

17. Market efficiency does not require constant ex ante returns, and the models of Lucas (1978) and Cox, Ingersoll, and Ross (1985) study the pricing of assets with time-varying required returns. Fama and French (1986b) show that the negative serial correlation in different stocks may be attributable to a common factor, and interpret this finding as supporting the view that time-varying returns account for mean-reversion in prices.

18. Several recent papers, including Black (1986), Campbell and Kyle (1986), DeLong et al. (1987), and Shiller (1984), have discussed the role of noise traders in security pricing.

19. The possibility of negative expected excess returns is an unattractive feature of the simple model we have analyzed. In principle the analysis could be repeated using Merton's (1980) model, which constrains the expected excess return to be positive. The exact parallel between the time-varying returns model and the fads model would not hold in this case, however.

20. Campbell (1987) estimates that the correlation between excess returns on long-term bonds and corporate equities was .22 for the 1959-1979 period, and .36 for the more recent 1979-83 period.

21. We perform Monte Carlo calculations holding the expected return on the market constant at its sample mean value. If required returns varied through time, but were always positive, it is possible that the fraction of negative predicted returns would be greater than our Monte Carlo results suggest.