The Time-Variance Relationship of Security Returns: Implications for the Return-Generating Stochastic Process
Author(s): Philip R. Perry
Source: *The Journal of Finance*, Vol. 37, No. 3 (Jun., 1982), pp. 857-870
Published by: Blackwell Publishing for the American Finance Association
Stable URL: http://www.jstor.org/stable/2327713
Accessed: 27/11/2009 08:19

THE JOURNAL OF FINANCE • VOL. XXXVII, NO. 6 • JUNE 1982

The Time-Variance Relationship of Security Returns: Implications for the Return-Generating Stochastic Process

PHILIP R. PERRY*

ABSTRACT

Using a test statistic which specifically allows for parameter shifts over time, we investigate the time-variance relationship of security returns. The null hypothesis of stationary and independent increments is rejected, and the existence of a complex short-term reversal phenomenon is reported.

THIS PAPER DEALS WITH the compounding of security returns over time: every returns-generating process that has been proposed implies some specific relationship among returns as the time horizon increases, and these implications have not been adequately investigated. We focus specifically on the relationship between the variance of return and the time period over which that return is measured, and conduct a test of the common assumption that security returns are generated by a stochastic process which has stationary and independent increments. The test statistic which is used has been designed to allow for parameter shifts, and hence the results are quite general. The primary conclusions are that (1) this stochastic process does not have stationary and independent increments; (2) the observed discrepancies cannot be explained by any simple mechanism such as the presence of first-order serial correlation; and (3) there exists a complex short-term reversal phenomenon (which we quantify) which is unexplained by all of the existing models for this process.

The organization of the rest of the paper is as follows: Section I contains a brief review of previous research; Section II describes the hypothesis and test statistic; Section III discusses the data which were used; Section IV reviews the empirical results; and Section V contains a brief summary.

I. Previous Research

The topics studied in this paper are not new ones, although previous work has not been very extensive. This section briefly reviews the literature in this field, with emphasis on the most recent work.

* Assistant Professor of Finance, State University of New York at Buffalo.

I would like to thank Professors Robert Oliver, Mark Rubinstein, and especially Barr Rosenberg, for their advice and encouragement. I am also indebted to Jim Boness, Seymour Smidt, and Mark Weinstein for helpful comments and suggestions. Particular thanks are due to Michael Brennan and an unknown referee of this Journal. Responsibility for any remaining errors is, of course, my own. Financial support from the National Science Foundation through NSF Grant No. SOC 77-18087 is gratefully acknowledged.

Fama [8] was one of the earliest researchers to report the presence of serial correlation, although the statistical significance was considered marginal. Alexander [1] and Steiger [25] were concerned with a related phenomenon, the presence of trends, while Larson [12] found some evidence that returns are positively correlated over short (roughly, six days) periods of time, and negatively correlated over longer ones. Their results, however, were based on relatively small samples and not all used common stock returns. More recently, Young [27] used 35 years of monthly closing prices for 233 NYSE common stocks and found some evidence of a weak negative dependence in the return series which became significant over a long period (6–12 months). He interpreted this as evidence of a weak negative serial correlation that became increasingly important as the time interval increased in length. Greene and Fielitz [10] also report the existence of a long-term dependency in security returns.

The most recent and most extensive study of the time-variance relationship, that of Schwartz and Whitcomb [23; hereafter SW], was published in 1977. They used weekly data for 20 securities (randomly) drawn from the S & P 500; the time period covered was 1962–68. To analyze the time-variance relationship, SW postulated a simple autoregressive process for security returns, and showed that the autocorrelation coefficient is an important parameter in the relationship between the variance of return over a short period of time and the variance over a longer period. They also showed that the sign of the autocorrelation coefficient is crucial: if Var R_T is the variance of return over time period T, ρ the autocorrelation coefficient, and we define

$$\gamma = \frac{T}{\text{Var}\, R_T} \frac{d\, \text{Var}\, R_T}{dT} = \frac{d\, ln\, \text{Var}\, R_T}{d\, ln\, T}$$

then

$$\gamma > -1 \quad \text{if} \quad \rho > 0 \quad \text{and} \quad \gamma < -1 \quad \text{if} \quad \rho < 0$$

Having derived this relationship, SW estimated the parameter γ. They did this by breaking their data into intervals of different lengths and calculating the variance of return for each of these different time intervals; SW thus assumed that the variance of the underlying stochastic process is constant through time. Regression analysis was then used to estimate the relationship between variance and time. This was done for both total and residual returns; SW (quite properly) emphasized the results for residual returns. They found a γ which was significantly less than -1, and inferred from it a weekly autocorrelation coefficient of $-.09$, and a daily autocorrelation of $-.45$ (pp. 49, 50). Their "principal empirical finding is that, for the overwhelming majority of stocks in our sample, market model residuals are negatively autocorrelated" (p. 50).

The primary difficulty with the SW paper is its assumption that all observed deviations from an idealized random walk (i.e., all deviations from the ideal of independent returns and constant variance of return) may be attributed to the existence of a simple autoregressive process which generates the return sequence. That such an assumption is unwarranted is clearly shown by the results presented below, which indicate that the deviations from the expected (idealized) variance pattern cannot be explained by a simple serial correlation scheme. While the SW methodology is useful for demonstrating that security returns do not follow an idealized random walk, it is not useful for making inferences concerning the observed deviations. One example of this inference problem is their deduction that daily residual returns have a (first order) serial correlation coefficient of roughly −.45. Not only is this implausibly large (in absolute value), it also differs in sign from what is generally observed.

II. The Hypothesis and Test Statistics

Despite the large number of models that have been proposed for the stochastic process which generates security returns, a limited amount of general empirical testing has been done. This literature is well-known, and is not reviewed here.[1] It is important to note, however, that virtually all of the models that have been proposed for the security returns generating process assume that these returns are independent and identically distributed over time. The primary exceptions are the model proposed by Oldfield et al. [18], which posits a serial dependence in the transactions return sequence, and the Cox and Ross [6] models, which assume that the distribution is a stationary function of price; since prices change relatively slowly, these models effectively imply that returns are independently and identically distributed over short time periods.

This common assumption of independence and stationarity is a strong one, and has an implication with respect to the joint distribution of security returns over time that may be easily tested. Carrying out such an empirical test will enable us to establish some general restrictions on the nature of the returns-generating process in addition to evaluating the proposed models.

One of the difficulties in analyzing the stochastic process generating security returns is that the underlying process may be obscured by the effects of parameter shifts. The potential presence of such shifts requires that the properties of the statistical test used be carefully considered. In particular, the use of spectral analysis, which assumes a constant generating process through time, could easily yield misleading results. The test statistic which we use is developed below. This statistic is based on the behavior of the variance over time, explicitly allows for parameter shifts, and provides a measure of the magnitude of the variance compounding discrepancy as well as tests for its presence.

While we report empirical results pertaining to different compounding intervals, for ease of exposition the discussion here is in terms of a monthly interval. Define, for a given security,

r'_{tm}, the arithmetic return for day t in month m;
r_{tm}, the logarithmic return for day t in month m;
r'_m, the arithmetic reurn for month m;
r_m, the logarithmic return for month m; and
N_m, the number of (trading) days in month m.

[1] See, for example, Bachelier [2], Barnea and Downes [3], Blattberg and Gonedes [4], Clark [5], Fama [7, 8], Hsu, Miller, and Wichern [11], Mandelbrot [13, 14], Merton [15], Moore [16], Officer [17], Osborne [19], Praetz [21], Press [22], and Teichmoeller [26].

By definition,

$$\ln(1 + r'_m) = \sum_{t=1}^{N_m} \ln(1 + r'_{tm})$$

or

$$r_m = \sum_{t=1}^{N_m} r_{tm}$$

r_{tm} may be rewritten as $v_{tm}u_{tm}$, where u_{tm} is the standardized daily return, with mean zero and variance one, and v_{tm}^2 is the variance of r_{tm}. Then,

$$r_m = \sum_{t=1}^{N_m} v_{tm}u_{tm} = v_{\cdot m} \sum_{t=1}^{N_m} u_{tm} = v_{\cdot m}u_m$$

where $v_{\cdot m}$, the standard deviation for month m, is assumed constant across the compounding interval.

Two quantities are of interest.[2] The first is an estimate of the monthly variance for month m obtained from daily data, assuming that the daily returns are generated by a process which has stationary and independent increments:

$$V_{md} = \sum_{t=1}^{N_m} r_{tm}^2 = \sum_{t=1}^{N_m} v_{tm}^2 u_{tm}^2 = v_{\cdot m}^2 \sum_{t=1}^{N_m} u_{tm}^2$$

The second quantity is an estimate of this same variance obtained by squaring the monthly return:

$$V_{mn} = [\sum_{t=1}^{N_m} r_{tm}]^2 = v_{\cdot m}^2[\sum_{t=1}^{N_m} u_{tm}]^2$$

V_{mn} is, of course, an estimate of the monthly variance obtained without the assumption of stationary and independent increments. Under the assumptions of independence and stationarity,

$$E(V_{md}) = E(V_{mn})$$

and hence a joint test of these assumptions may be conducted by estimating the empirical relationship between V_{md} and V_{mn}.

The test statistic we use is the average of the ratios of V_{mn} to V_{md}, that is,

$$R1 = \frac{1}{M}\sum_{m=1}^{M}\left(\frac{V_{mn}}{V_{md}}\right) = \frac{1}{M}\sum_{m=1}^{M}\frac{(\sum_{t=1}^{N_m} r_{tm})^2}{\sum_{t=1}^{N_m} r_{tm}^2} = \left(\frac{1}{M}\right)\sum_{m=1}^{M}\left[\frac{(\sum_{t=1}^{N_m} u_{tm})^2}{\sum_{t=1}^{N_m} u_{tm}^2}\right]$$

where M is the total number of months for which data are available.

Under the null hypothesis of stationarity and independence, the probability limit (plim) of $R1$ as M increases is one. Hence, the null hypothesis of independent and stationary daily return distributions may be rejected if we can reject the hypothesis that the test statistic is equal to one. $R1$ has the advantage that it allows us to investigate the underlying stochastic process, assuming that the variance is constant only across the compounding interval used.

Because this test statistic involves a ratio of dependent variables, its sampling properties are unknown. In an attempt to clarify these properties, a Monte Carlo

[2] For each of these quantities, we make use of the empirical fact that for security returns, while the true mean is undoubtedly nonzero, the squared mean return is negligibly small compared to the mean squared return, and hence the former may be taken to be zero for our purposes.

study was performed, the results of which are reported in the Appendix and will be referred to below.

III. The Data

The basic data are daily returns obtained from the Center for Research in Security Prices (CRSP) at the University of Chicago; these returns are adjusted for stock splits and include dividends. The data were from 2 July 1962 to 30 December 1977, a total of 3,881 daily returns for each security. Since the companies were selected based on their 1976 characteristics, and the data extended back to 1962, there was a danger of introducing ex post selection bias. If this were a study of average returns, this bias would likely be strong enough to invalidate any conclusions we might have reached. We are, however, studying variances, not average returns, and for such a study this sample bias is not a serious problem.

It should be noted that it is standard CRSP practice, when dealing with a time period in which a particular security did not trade, to make the first return reported after such a period a multiday return covering the entire nontraded time period. Trading is likely to be suspended during a period of intense investor interest, and consequently multiday returns reported by CRSP are likely to be large (in absolute value) and hence not representative of daily returns. These multiperiod returns were few, and were eliminated from the data.

There are, of course, several well-known difficulties with using daily data:

(1) Bid-ask spreads. At any point in time, there are two different stock prices applicable—one to buy (bid) and one to sell (ask). Since only the price of the last transaction is known, and not whether it is the bid or ask price, some noise is introduced into the data. This is potentially serious since the typical bid-ask spread for the New York Stock Exchange is at least one-eighth, and this is sometimes comparable to the price change over a single day. This problem is analyzed in Perry [20], where it is shown that the consequences of the bid-ask spread are not serious for our hypothesis tests and data base.[3]

(2) Weekends. The problem here is that the return from Friday to Monday is taken to be a daily return when in fact it is a multiday return. One possibility is to consider this a three-day return since it encompasses three calendar days; another possibility is to consider this a one-day return since it encompasses one trading day. The published empirical evidence is quite scanty [8, 9] and, for our purposes, difficult to interpret, but seems to imply that a weekend is "worth" roughly 1.2–1.5 days. This problem is analyzed in Perry [20] where empirical evidence is presented that, at least for our data base and hypothesis tests, weekend returns may be treated as single day returns.[4]

[3] Utilizing some fairly mild assumptions, it can be shown that the presence of the bid-ask spread implies that our variance estimates are biased. The extent of this bias, however, is small: less than four percent of our estimates require an adjustment of five percent or more, and less than one-half of one percent require an adjustment of ten percent.

[4] The ratio of weekend variance to daily variance ranges from .89 to 1.32; the average is 1.03, as is the median. While the mean is marginally significantly different from one, with a t-statistic of 2.30, a "typical" value of 1.03 implies a correction of less than three-quarters of one percent to our variance estimates.

Table I

Sample Companies

1.	American Telephone & Telegraph
2.	Dow Chemical
3.	Eastman Kodak
4.	Exxon
5.	General Electric
6.	General Motors
7.	IBM
8.	Minnesota Mining & Manufacturing
9.	Proctor & Gamble
10.	Sears, Roebuck
11.	American Home Products
12.	Atlantic Richfield
13.	Burroughs
14.	Caterpillar Tractor
15.	Coca Cola
16.	Continental Oil
17.	Du Pont
18.	Ford Motor
19.	General Telephone & Electronics
20.	Getty Oil
21.	Gulf Oil
22.	Halliburton
23.	Johnson & Johnson
24.	Merck
25.	Mobil
26.	Phillips Petroleum
27.	Royal Dutch Petroleum
28.	Schlumberger
29.	Shell Oil
30.	Standard Oil (California)
31.	Standard Oil (Indiana)
32.	Texaco
33.	Union Carbide
34.	U.S. Steel
35.	Weyerhaeuser
36.	Xerox
37.	K Mart

(3) Timing. Trading is not continuous, and hence the closing prices reported for a group of securities are not all applicable to the same point in time. This is true of all price data but is particularly serious for daily data. To minimize this problem, we chose to work with the common stock returns of those forty companies in the S & P 500 which had the largest (equity) capitalization.[5] For companies as large as these, trading is virtually continuous and hence timing is not a problem for our data base. We do, however, make use of a market index. While there is no good way around the problem here, we have attempted to minimize its effects by using a value-weighted index.

Of the forty companies chosen, three were eliminated because of an insufficient

[5] As of 8 July 1976. These data were chosen for the sake of convenience: a list of the fifty largest, as of that date, is readily available [24].

number of data points (less than 3,000). The companies whose common stock returns were used in this study are listed in Table I.

Descriptive statistics for each logarithmic return series are given in Table II. As expected, the skewness is generally positive but small, and the kurtosis is greater than three. It is particularly interesting that all estimated serial correlation coefficients are positive. These results are consistent with those of Fama [8], who found positive daily serial correlation for 22 of the 30 Dow Jones Industrials over

Table II

Descriptive Statistics for Daily Logarithmic Returns.[a]

Sec	Mean	Std Dev	Min	Max	Skew	Kurt	Ser Cor
1	.000248	.00927	−.0680	.0716	.41	7.75	.071
2	.000447	.0137	−.0671	.0782	.19	5.70	.12
3	.000295	.0146	−.129	.0939	.051	6.66	.027
4	.000371	.0110	−.0468	.0630	.073	4.85	.17
5	.000234	.0141	−.0882	.0749	.049	5.68	.086
6	.000281	.0126	−.0744	.0864	.15	5.68	.078
7	.000373	.0131	−.0913	.0786	.12	5.35	.051
8	.000253	.0141	−.0958	.0872	.19	6.00	.10
9	.000347	.0118	−.0678	.0947	.27	7.04	.094
10	.000228	.0128	−.0830	.0815	.19	6.64	.13
11	.000411	.0138	−.0832	.0993	.18	6.56	.16
12	.000452	.0157	−.123	.090	.11	6.81	.11
13	.000582	.0202	−.148	.182	.36	7.66	.052
14	.000531	.0158	−.0721	.0785	.022	5.02	.097
15	.000434	.0140	−.118	.0872	−.17	9.05	.089
16	.000396	.0160	−.139	.0870	−.057	7.00	.12
17	.000120	.0123	−.0677	.0654	.23	5.37	.097
18	.000269	.0135	−.0840	.107	.25	5.97	.074
19	.000300	.0133	−.0703	.0927	.15	6.11	.078
20	.000698	.0186	−.115	.159	.32	6.55	.088
21	.000299	.0131	−.106	.0851	.19	6.78	.076
22	.000813	.0163	−.0844	.100	.13	4.73	.15
23	.000596	.0150	−.0756	.148	.40	8.23	.12
24	.000494	.0142	−.106	.0768	.091	5.70	.13
25	.000441	.0145	−.0796	.0950	.14	5.96	.11
26	.000411	.0165	−.146	.105	.0096	8.28	.14
27	.000465	.0128	−.0840	.0789	.099	6.76	.099
28	.000792	.0159	−.0935	.0850	.074	4.71	.17
29	.000377	.0137	−.0681	.0758	.18	5.10	.18
30	.000328	.0134	−.0677	.0693	.084	5.49	.13
31	.000544	.0127	−.106	.0765	.11	6.01	.16
32	.000234	.0133	−.0774	.0673	.12	5.54	.057
33	.000141	.0139	−.0718	.0930	.23	6.18	.12
34	.000217	.0151	−.0923	.156	.67	8.30	.076
35	.000428	.0163	−.0892	.108	.27	6.15	.15
36	.000503	.0195	−.101	.0833	−.12	4.76	.093
37	.000803	.0179	−.221	.0937	−.85	13.7	.14

[a] The number of data points used was at least 3877 (out of a total possible 3881), except for security 35, which had 3521 data points. Skewness is here defined as the ratio of the third moment to the square root of the cubed second moment; kurtosis is defined as the ratio of the fourth moment to the square of the second moment (for the normal distribution, this is 3).

the period September 1957 through September 1962. Since the presence of measurement error, in particular, the bid-ask spread, tends to bias the estimated correlation coefficients downwards, the true values are probably somewhat greater than indicated.

IV. Empirical Results

In this study, we are interested in the relationship between V_{md} and V_{mn} both for individual securities and in general. In order to investigate the question of compounding for the data base as a whole, the security returns would have to be pooled, that is, the test statistics would have to be calculated using all the data simultaneously. Such pooling of the total returns would, however, give misleading results since the observations are not independent: at any given point in time, the returns of all common stocks will be strongly influenced by a common factor, the market. Once the effect of the market is removed, there will be a time series of residual returns for each security. While these residuals are not perfectly independent across securities (there are, for example, industry effects), the primary source of cross-sectional dependence has been removed, and hence these data may be pooled with some confidence. Section A below discusses the estimation of the residual returns, and Section B reports the empirical results for both total and residual returns.

A. Estimation of residual returns

To remove the common effect of the market, we made use of the market model,

$$r_{it} = \alpha_i + \beta_i R_{Mt} + \epsilon_{it}$$

where:

r_{it} is the (logarithmic) return for security i in time period t;
R_{MT} is the (logarithmic) return for the market in time period t (we have used the CRSP RETV index, which is value-weighted, includes dividends, and encompasses all NYSE and AMEX securities);
α_i is a constant;
β_i is defined to be the covariance of r_i and R_M divided by the variance of R_M; and
ϵ_{it} is the random component of unsystematic return for security i in time period t.

For each security, the parameter β was estimated by regression analysis. Since our goal is to remove the common effect of the market, we have defined the residual return to be

$$\text{res}_{it} = r_{it} - (\hat{\beta}_i R_{Mt})$$

These residual returns have much smaller means and only slightly smaller variances than do the total returns, and hence the assumption that the mean return is zero is still valid.

To allow for possible nonstationarity in beta, these coefficients were estimated in three different ways: a single "overall" beta was calculated, using all available data; a sequence of monthly betas was calculated, using the returns in each month to estimate that month's beta; a sequence of quarterly betas was calculated, using the returns in each quarter to estimate that quarter's beta. A residual return series was calculated for each beta estimation method; however, the choice of residual series did not affect the test statistic significantly, and hence only the results from the residual returns derived from quarterly betas are reported below.

B. The test statistic

Values of $R1$ were calculated for each security's total return and residual return sequence for weekly, monthly, and quarterly compounding intervals; these results are shown in Table III.

For total returns, virtually all of the securities have compounding coefficients greater than one for the weekly compounding interval. A majority of coefficients are also greater than one for monthly and quarterly compounding although, based on the 95% confidence intervals established in the Appendix, only a few are significant. In no case is a coefficient significantly less than one. For the sample as a whole, the mean coefficients are significantly greater than one, although these figures are difficult to interpret because of a lack of independence across the different securities.

For residual returns, the mean weekly coefficient is significantly greater than one, the mean monthly coefficient is significantly less than one, and the mean quarterly coefficient is marginally significantly less than one. Not surprisingly, the pooled sample results for residual returns are very similar: for the weekly compounding interval, $R1 = 1.035$; for monthly compounding, $R1 = .950$; and for quarterly compounding, $R1 = .934$. Since it can be shown that the effect of assuming that the mean return is zero is to bias $R1$ upwards, these values of $R1$ are, if anything, too high. (Note that in Section II we have argued that this bias is in fact negligible.)

The effect of first-order serial correlation is analyzed in the Appendix, where it is shown that, assuming a simple first-order autoregressive process and positive daily serial correlation, the weekly coefficient should be approximately 1.19 and the compounding coefficient should increase as the compounding interval increases. The results given in Table III are even more striking, then, since they indicate that the compounding coefficient tends to decrease as the compounding interval lengthens from a week to a month.

The finding that the weekly compounding coefficient is less than expected indicates the presence of some type of reversal phenomenon that operates on a time scale of a week or less. The fact that these coefficients decrease as the compounding interval increases indicates that its effect becomes more pronounced over time. The presence of such a phenomenon, whatever its origin, is evidence against the assumption of stationary and independent increments in the return-generating process, and is also evidence that a simple first-order autoregressive process is inadequate to describe the stochastic process which generates security returns.

Table III

Values of R1

Security	Total Returns Compounding Interval			Residual Returns Compounding Interval		
	Weekly	Monthly	Quarterly	Weekly	Monthly	Quarterly
1	1.045	.973	.861	1.045	.938	.777
2	1.159	1.231	1.098	1.012	1.009	.873
3	1.039	.982	1.143	.958	.982	1.166
4	1.088	.920	1.019	1.101	1.021	1.057
5	1.023	1.062	1.074	.920	.855	.696
6	1.097	.997	1.144	1.015	.826	.836
7	1.016	.945	1.040	.930	.952	1.127
8	1.134	1.047	1.013	.987	.929	.864
9	1.011	.932	.862	.989	.877	.835
10	1.103	1.018	1.161	1.055	1.032	.921
11	1.140	1.069	1.148	1.092	.975	1.011
12	1.099	1.046	.918	1.078	.968	.952
13	1.066	.861	1.009	1.063	.998	1.213
14	1.103	1.135	1.014	.991	.904	.638
15	1.055	.848	.983	.941	.687	.749
16	1.091	1.009	1.011	1.015	.926	.847
17	1.087	1.024	1.243	1.169	1.195	1.371
18	1.049	1.143	1.093	1.038	1.010	.866
19	1.073	1.105	1.075	.933	.744	.690
20	1.134	1.067	1.202	1.034	.989	1.025
21	1.118	1.042	.912	1.053	.829	.691
22	1.108	1.014	1.198	1.008	.931	1.004
23	1.084	.998	1.382	1.076	.987	1.167
24	1.172	1.033	.878	1.092	.940	.730
25	1.123	.995	.928	1.067	.776	.741
26	1.089	1.027	.925	1.031	.941	.969
27	1.098	1.114	.935	1.103	.942	.661
28	1.190	1.112	1.210	1.079	1.017	1.123
29	1.207	1.123	1.104	1.138	.969	.900
30	1.148	1.011	.967	1.010	.979	.846
31	1.127	1.123	1.114	1.087	1.116	1.146
32	.958	.963	.932	.891	.853	.720
33	1.070	1.100	.877	.996	.938	.911
34	1.054	.982	.934	1.093	1.037	1.092
35	1.167	1.133	1.294	1.109	.952	.970
36	1.054	.951	1.241	1.047	1.180	1.226
37	1.137	1.133	1.525	1.052	1.041	1.170
Mean	1.095	1.034	1.067	1.035	.953	.935
Std dev	.053	.082	.153	.063	.104	.188
SE of mean	.0087	.013	.025	.010	.017	.031

V. Conclusion

The primary conclusion is that the null hypothesis of stationary and independent increments may be rejected, and that this rejection cannot be explained by any simple mechanism. Further, our results indicate the presence of some complex reversal mechanism: estimates of the weekly variance, using daily data, typically understate the true weekly variance, while similar estimates of the monthly and quarterly variance overstate the true variance. That this type of dependence exists over short time horizons has not been previously demonstrated. In addition, we have provided a measure of the size of the resulting variance-compounding discrepancy. What we have not provided is an economic rationale for this reversal phenomenon; this we leave to future research.

APPENDIX

A. The Effect of First-Order Serial Correlation

Assume that the (logarithmic) return series follows a first-order autoregressive process. Define:

r_t, the (logarithmic) return for period t;
ρ, the autocorrelation coefficient of the process; and
e_t, a random variable, uncorrelated with e_s for $s \neq t$.

Then, the model for r_t is

$$r_t = \rho r_{t-1} + e_t \tag{A.1}$$

It follows that the expectation of V_{mn} is

$$E[V_{mn}] = E[\textstyle\sum_{s=1}^{T} \sum_{t=1}^{T} r_s r_t]$$

Further,

$$E[V_{mn}] = \mathrm{Var}(r) \textstyle\sum_{s=1}^{T} \sum_{t=1}^{T} \rho_{s,t}$$

assuming stationarity and a mean of zero, where $\rho_{s,t}$ is the correlation coefficient between r_s and r_t. Thus,

$$E[V_{mn}] = \mathrm{Var}(r)[T + 2 \textstyle\sum_{i=1}^{T-1} (T-i)\rho_{1,1+i}]$$

Since the model given in Equation (A.1) implies that

$$\rho_{1,1+i} = \rho^i$$

we have that

$$E[V_{mn}] = \mathrm{Var}(r)[T + 2 \textstyle\sum_{i=1}^{T-1} (T-i)\rho^i]$$

$$E[V_{mn}] = E[V_{md}]\left[1 + \frac{2}{T} \textstyle\sum_{i=1}^{T-1} (T-i)\rho^i\right]$$

$$\frac{E[V_{mn}]}{E[V_{md}]} = 1 + \frac{2}{T} \textstyle\sum_{i=1}^{T-1} (T-1)\rho^i \tag{A.2}$$

From Equation (A.2), it is clear that the presence of positive serial correlation in the daily returns sequence implies that the test statistic $R1$ should be greater than one. Estimated values of this coefficient for quarterly, monthly, and weekly compounding intervals are given in Table A.I. The two values of ρ used in this table are the average values for daily total returns ($\rho = .107$) and residual returns

(ρ = .051). [As shown in Table II, all first-order serial correlation coefficients for daily total returns are positive. For the residual returns, all but five are positive, and those that are negative are quite small (the largest of these, in absolute value, is .019).]

Table A.I

Coefficients Implied by First-Order Serial Correlation

T	$\rho = .107$	$\rho = .051$
5 (weekly)	1.186	1.085
20 (monthly)	1.226	1.102
60 (Quarterly)	1.239	1.106

Note that if the simple autoregressive model is correct, the presence of positive serial correlation implies that the compounding coefficients should increase as the time horizon lengthens.

B. A Monte Carlo Study of the Properties of the Test Statistic R1

The Monte Carlo study was conducted assuming the null hypothesis was correct, that is, the test statistic $R1$ was calculated using sequences of independently generated and identically distributed random numbers; these numbers were each drawn from an $N(0, 1)$ distribution. One thousand simulations were performed with 3,840 data points each. Compounding intervals of weeks, months, and quarters were simulated using 5, 20, and 60 data points respectively. The estimated mean and standard deviation of the test statistic for each compounding interval are given in Table B.I.

Table B.I

Monte Carlo Results for the Test Statistic $R1$

	Compounding Interval		
	Weekly	Monthly	Quarterly
Mean	1.0007	1.0004	.9954
Std Dev	.0397	.0992	.1781

The distribution of the test statistic was approximately symmetric, with the two standard deviation range encompassing between 94.9 and 95.9% of the data. The two standard deviation range was thus used to establish approximate 95% confidence intervals; these intervals are given in Table B.II. (Note that these ranges are derived from the Monte Carlo results and are not based on an assumption of normality.) These intervals are indeed approximate, and are probably too small since the Monte Carlo study was done with $N(0, 1)$ variables whereas security returns characteristically exhibit "long tails" relative to the normal distribution.

Table B.II

95% Confidence Intervals for $R1$

Compounding Interval		
Weekly	Monthly	Quarterly
.9214–1.0801	.8019–1.1989	.6392–1.3516

REFERENCES

1. Sidney S. Alexander. "Price Movements in Speculative Markets: Trends or Random Walks" and "Price Movements in Speculative Markets: Trends or Random Walks, No. 2." In Paul H. Cootner, ed., *The Random Character of Stock Market Prices.* Cambridge, Mass.: M.I.T. Press, 1964, pp. 199–218 and 338–372.
2. Louis Bachelier. "Theory of Speculation." Ph.D. Thesis, Academy of Paris (1900). Translated and reprinted in Paul H. Cootner, ed., *The Random Character of Stock Market Prices.* Cambridge, Mass.: M.I.T. Press, 1964, pp. 17–78.
3. A. Barnea and D. Downes. "A Re-examination of the Empirical Deviations of Stock Price Changes." *Journal of the American Statistical Association* 68 (1973), 348–50.
4. Robert C. Blattberg and Nicholas J. Gonedes. "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices." *Journal of Business* 47 (1974), 244–80.
5. Peter Clark, "A Subordinated Stochastic Process Model With Finite Variance for Speculative Prices." *Econometrica* 41 (1973), 135–55.
6. John C. Cox and Stephen A. Ross. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics* 3 (1976), 145–66.
7. Eugene F. Fama. "Mandelbrot and the Stable Paretian Hypothesis." *Journal of Business* 36 (1963) 420–29.
8. Eugene F. Fama. "The Behavior of Stock-Market Prices." *Journal of Business* 38 (1965), 34–105.
9. Clive W. J. Granger and Oskar Morganstern. *Predictability of Stock Market Prices.* Lexington, Mass.: Heath Lexington Books, 1970.
10. Myron T. Greene and Bruce D. Fielitz. "Long-term Dependence in Common Stock Returns." *Journal of Financial Economics* 4 (1977), 339–49.
11. D. Hsu, R. Miller, and D. Wichern. "On the Stable Paretian Behavior of Stock Market Prices." *Journal of the American Statistical Association* 69 (1974), 108–13.
12. Arnold B. Larson. "Measurement of a Random Process in Futures Prices." In Paul H. Cootner, ed., *The Random Character of Stock Market Prices.* Cambridge, Mass.: M.I.T. Press, 1964, pp. 219–30.
13. Benoit Mandelbrot. "The Variation of Certain Speculative Prices." *Journal of Business* 36 (1963), 394–419.
14. Benoit Mandelbrot. "Comments on 'A Subordinated Stochastic Process Model With Finite Variance For Speculative Prices' by Peter K. Clark." *Econometrica* 41 (1973), 157–59.
15. Robert C. Merton. "Option Pricing When Underlying Stock Returns are Discontinuous." *Journal of Financial Economics* 3 (1976), 125–44.
16. Arnold B. Moore. "Some Characteristics of Changes in Common Stock Prices." In Paul H. Cootner, ed., *The Random Character of Stock Market Prices.* Cambridge, Mass.: M.I.T. Press, 1964, pp. 139–61.
17. R. R. Officer. "The Distribution of Stock Returns." *Journal of the American Statistical Association* 67 (1972), 807–12.
18. George S. Oldfield, Richard J. Rogalski, and Robert A. Jarrow. "An Autoregressive Jump Process for Common Stock Returns." *Journal of Financial Economics* 5 (1977), 389–418.
19. M. F. M. Osborn. "Brownian Motion in the Stock Market." *Operations Research* 7 (1959), 145–73.
20. Philip R. Perry. "Properties of Extreme Security Returns." Unpublished Ph.D. thesis, Graduate School of Business Administration, University of California at Berkeley (December 1979).

21. Peter D. Praetz. "The Distribution of Share Price Changes." *Journal of Business* 45 (1972), 49–55.
22. S. James Press. "A Compound Events Model of Security Prices." *Journal of Business* 40 (1967), 317–35.
23. Robert A. Schwartz and David K. Whitcomb. "The Time Variance Relationship: Evidence on Autocorrelation in Common Stock Returns." *Journal of Finance* 32 (1977), 41–55.
24. William F. Sharpe. *Investments.* Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1978.
25. William Steiger. "A Test of Nonrandomness in Stock Price Changes." In Paul H. Cootner, ed., *The Random Character of Stock Market Prices.* Cambridge, Mass.: M.I.T. Press, 1964, pp. 253–61.
26. John Teichmoeller. "A Note on the Distribution of Stock Price Changes." *Journal of the American Statistical Association* 66 (1971), 282–84.
27. W. E. Young. "Random Walk of Stock Prices; A Test of the Time Variance Function." *Econometrica* 39 (1971), 797–812.