



American Finance Association

Testing the Predictive Power of Dividend Yields

Author(s): William N. Goetzmann and Philippe Jorion

Source: *The Journal of Finance*, Vol. 48, No. 2 (Jun., 1993), pp. 663-679

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2328917>

Accessed: 25/11/2009 11:31

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

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

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

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



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

<http://www.jstor.org>

Testing the Predictive Power of Dividend Yields

WILLIAM N. GOETZMANN and PHILIPPE JORION*

ABSTRACT

This paper reexamines the ability of dividend yields to predict long-horizon stock returns. We use the bootstrap methodology, as well as simulations, to examine the distribution of test statistics under the null hypothesis of no forecasting ability. These experiments are constructed so as to maintain the dynamics of regressions with lagged dependent variables over long horizons. We find that the empirically observed statistics are well within the 95% bounds of their simulated distributions. Overall there is no strong statistical evidence indicating that dividend yields can be used to forecast stock returns.

A NUMBER OF RECENT studies appear to provide empirical support for the traditional use of the dividend-price ratio as a measure of expected stock returns. Rozeff (1984), for instance, finds that the ratio of the dividend yield to the short-term interest rate explains a significant fraction of movements in annual stock returns. Fama and French (1988) use a regression framework to show that the dividend yield predicts a significant proportion of multiple year returns to the NYSE index. They further observe that the explanatory power of the dividend yield increases in the time horizon of the returns; over four-year horizons, R^2 's range from a low of 19% to an astonishingly high value of 64%. Similar results are reported by Flood, Hodrick, and Kaplan (1987) and Campbell and Shiller (1988).

The apparent predictability of market returns from past values of dividend yields is regarded by Rozeff (1984) as support for the rejection of the random walk model of stock prices, and by Fama and French (1988) as support for the cyclical behavior of expected returns. Flood, Hodrick, and Kaplan (1987) interpret their results as support for time-varying expected returns to stocks. The direct, and somewhat disturbing, implication of most of these studies is that significant components of long-term stock returns may be predictable using combinations of past returns and macroeconomic variables.

There are a number of reasons, however, why these results should be regarded with caution. Given the persistent patterns of dividend payments, movements in dividend yields are essentially dominated by movements in prices. Therefore, the forecasting regressions suffer from biases due to the

* Columbia University and University of California at Irvine, respectively. We wish to thank Bob Hodrick and Stephen Ross for their advice. Useful comments were also received from Stephen Brown, John Campbell, John Hartigan, Jon Ingersoll, Ravi Jagannathan, Paul Kaplan, and Robert Shiller.

fact that the right-hand-side variables are correlated with lagged dependent variables, instead of being predetermined as assumed in standard statistical models. In addition, the usual GMM corrections to the standard errors are only valid asymptotically, and there is some question as to whether “asymptotic” should be measured in terms of years, decades or even centuries, especially for long-horizon forecasts.

Nelson and Kim (1992), in an independent study, analyze these biases in simulations of a VAR system under the null hypothesis of no predictability of returns. Using returns sampled annually, they report that the simulated distribution of t -statistics is displaced upward, but they still find some evidence of predictability at conventional significance levels. Hodrick (1992), making full use of the information available in monthly data, extensively simulates a VAR model applied to long-horizon returns, and also finds evidence of predictability in stock returns.

While the VAR model addresses the endogeneity of the predictor, it does not fully account for biases due to the fact that the regressor behaves like a lagged dependent variable. Our research illustrates a case where this problem still biases tests toward rejection of the null. We use the bootstrap methodology, as developed by Efron (1979), to model the distributions of regression statistics under the null hypothesis that stock returns are independently and identically distributed, and are not related to past dividends. Our approach differs from previous VAR simulations in that we explicitly incorporate the lagged price relation between returns and dividend yields. We find that this explicit specification of the null makes a substantial difference in hypothesis tests of the significance of the dividend yield regression.

In bootstrapped regressions of one- to four-year returns to the S & P stock return index on the preceding dividend yield, we fail to reject the null hypothesis that future returns are unrelated to past dividend yields at conventional significance levels. In addition, we find that the observed explanatory power of the model, as measured by the coefficient of determination, is only marginally significant when compared to bootstrapped distributions of R^2 's; OLS t -statistics over 18 and R^2 over 38% for all multiple year horizons are not unusual. Overall, when we explicitly model the null hypothesis as a random walk, we find that the observed regressions of returns on past dividend yields provide only marginal statistical evidence against the random walk model.

This paper is organized as follows. The basic dividend yield regressions are presented in Section I. Section II describes the modeling of the null hypothesis in a numerical analysis framework. The bootstrapping tests and results are presented in Section III. Finally, the last section contains some concluding observations.

I. Dividend Yields

Dividend yields have long been used to evaluate the expected return to investment in common stocks. If the stock price represents a claim to the

future stream of dividends, the price can be exactly determined assuming constantly growing dividends and a known discount rate. This is the model variously attributed to Williams (1938) or Gordon and Shapiro (1956). Campbell and Shiller (1988) more recently refer to this model as the “dividend-ratio” model in the absence of uncertainty:

$$P_t = \sum_{i=1}^{\infty} D_t(1+g)^i/(1+r)^i = \frac{D_{t+1}}{r-g}. \quad (1)$$

where P is the stock price, D is the dividend, r is the discount rate, and g is the constant growth rate of dividends. In the certainty model, the discount rate is the expected return on the stock. Although the model is not directly applicable to the case in which growth rates and discount rates vary through time, the model suggests that dividend yields should capture variations in expected stock returns.

If long-term market returns are predicted by the dividend yield, the following regression should produce a significant coefficient and a non-trivial R^2 :

$$R_{t,t+T} = \alpha_T + \beta_T Y_t + \epsilon_{t,t+T} \quad (2)$$

where $R_{t,t+T}$ is the compound total stock return from month t to month $t+T$, and Y_t is the ratio D_t/P_t , the annual dividend up to time t divided by the stock price as of time t .

The null hypothesis is that there is no relation between $R_{t,t+T}$ and Y_t , i.e. that $\beta_T = 0$. Following Fama and French (1988), we perform all tests using overlapping observations. As is well known, this procedure results in more powerful tests but induces a moving average process in the errors, which invalidates the usual OLS standard errors.¹ The standard error corrections are computed using the method proposed by Hansen and Hodrick (1980), in which the autocovariances are estimated from the data, with a modification due to White (1980) and Hansen (1982) that allows for conditional heteroskedasticity. These will be referred to as “Generalized Method of Moments” estimators. There are situations, however, where the variance-covariance matrix of the estimated coefficients is not positive definite. Therefore, standard errors are also reported using a correction due to Newey and West (1987) that ensures that the matrix is positive definite.²

To investigate the predictive power of dividend yields, we use data on the S & P 500 index over the period 1927 through 1990. Monthly total and income returns were used to construct a price series P , exclusive of dividends, from which monthly dividend payments were inferred. Because of seasonalities in monthly dividends, an annual dividend series D was computed by reinvesting the dividends at the monthly riskless rate. A price series P^R was also computed from the total returns R , assuming reinvestment of dividends; this series represents the value of a fund invested in the S & P 500 stocks with

¹ See for instance Richardson and Smith (1991).

² The covariances were weighted up to the number of overlaps.

Table I
Long-Horizon Dividend Yield Regressions

$$R_{t,t+T} = \alpha_T + \beta_T Y_t + \epsilon_{t,t+T}$$

where $R_{t,t+T}$ is defined as the total stock return from time t to time $t + T$, $Y_t = D_t/P_t$ is the annual dividend yield measured as of time t . The t -statistics presented are: $t(\text{OLS})$, from the classical OLS regression, $t(\text{GMM})$, which adjusts for heteroskedasticity and the moving average process induced by the overlapping observations (using the sample autocovariances), and $t(\text{NW})$, which in addition uses the Newey and West adjustment to ensure that the variance-covariance matrix of estimated coefficients is positive definite.

Horizon (Months)	Beta	<i>t</i> -Statistic			<i>R</i> ²
		OLS	GMM	NW	
Period: 1927–1990					
1	0.386	2.80	1.21	1.21	0.010
12	5.108	9.97	2.95	3.17	0.116
24	9.071	12.73	5.32	5.38	0.179
36	12.939	15.38	5.97	4.77	0.244
48	21.392	21.42	7.37	7.30	0.390
Period: 1927–1958					
1	0.415	1.89	0.85	0.85	0.009
12	5.846	6.97	2.53	2.61	0.116
24	10.778	8.97	5.70	4.64	0.183
36	16.167	11.30	5.16	3.74	0.269
48	28.987	18.06	11.20	7.05	0.493
Period: 1959–1990					
1	0.668	2.76	2.42	2.42	0.020
12	8.757	11.04	4.03	4.46	0.247
24	14.401	14.29	3.65	4.32	0.363
36	19.740	17.15	4.56	4.90	0.459
48	27.911	20.26	4.68	5.19	0.551

reinvested dividends. The appendix contains a complete description of the data.

Table I presents long-horizon forecasts for returns on the S & P index over the period 1927 to 1990, and two subperiods of equal length, 1927 to 1958 and 1959 to 1990. Over the total 64-year period, the results seem to suggest strong predictive ability for dividend yields. The slope coefficients increase from 0.39 to 21.39 when the horizon lengthens from 1 month to 4 years.³ The OLS t -statistic also increases uniformly with the horizon, up to 21.42; since the R^2 is a simple transformation of the OLS t -statistic, it increases to 39%, which appears to be substantial. These OLS t -statistics and R^2 's, however,

³ Part of this increase is explained by the fact that the variance of the dependent variable increases over longer horizons. In line with most previous research, we have chosen not to adjust returns for the length of the horizon.

are seriously biased upward because of overlapping observations. This is apparent from the much lower values of the GMM and Newey and West (NW) t -statistics, which are 7.37 and 7.30 at the 4-year horizon. These values, however, still indicate predictability.

The lower panels of the table reports estimates for the two subperiods. A similar pattern emerges: high values of β , increasing with maturity up to 29.0 and 27.9, and high R^2 's, going up to 49% and 55% in the two respective subperiods. The fact that the predictive power is consistently stronger over shorter sample periods to some extent suggests that the small-sample bias may be enhancing the apparent significance of the results.

II. Testing Predictive Power

In fact, dividend yield regressions are similar to regressions on a lagged dependent variable, which suffer from well-known biases. To illustrate this, assume as before that dividends grow at a constant rate g . For illustration purposes, it is slightly more convenient to use continuously compounded returns, and the series with reinvested dividends. Dividends then grow as $\ln D_t^R = \ln D_0^R + gt$, and regression (2) becomes:

$$\begin{aligned}\ln(P_{t+T}^T/P_t^R) &= \alpha + \beta \ln D_t^R - \beta \ln P_t^R + \epsilon_{t,t+T} \\ &= \alpha' + \beta(-\ln P_t^R) + \gamma t + \epsilon_{t,t+T}.\end{aligned}\quad (3)$$

In this setup, all of the economic time variation in the dividend yield derives from the time variation in the price level.

Equation (3) suffers from the classical bias due to the fact that a right-hand-side variable is a lagged dependent variable. When $T = 1$, Kendall (1973), for instance, shows that the OLS estimate, although consistent, is centered at values less than zero in finite samples, even when the slope is truly zero. Dickey and Fuller (1979) tabulate by simulations new values for the OLS t -statistic under the null. The downward bias is shown to be substantial in small samples, and is of the order of $(-4/n)$, where n is the sample size.

In overlapping regressions, however, it is harder to predict the extent of the bias for the slope coefficient. While the GMM corrections to the standard errors are valid asymptotically, even with 64 years of data, there are only 16 truly independent observations in the case of 4-year overlapping returns. Econometric theory offers little guidance as to whether 16 observations qualify for asymptotic status. For these reasons, it seems prudent to investigate the small-sample characteristics of these regressions by numerical analysis.

A. The Bootstrap

The bootstrap was proposed by Efron (1979) as a nonparametric randomization technique that draws from the observed distribution of the data to

model the distribution of a test statistic of interest. For example, suppose $X = (X_1 \dots X_n)$ are i.i.d. random variables drawn from the unknown distribution F . Define θ as some parameter of the population, and $\hat{\theta}(X)$ as an estimator of θ . Let $\hat{F}$ be the sample distribution, from the observed X_i , that assigns a mass $1/n$ to each X_i . The bootstrap estimates the distribution of $\hat{\theta}(F)$ by the sampling distribution $\hat{\theta}(\hat{F})$. This procedure is carried out by the following steps:

1. From the observed $(X_1 \dots X_n)$, compute the test statistic $\hat{\theta}(X)$.
2. Draw a bootstrap sample $(X_1^* \dots X_n^*)$, with replacement, from the empirically observed distribution $\hat{F}$.
3. Calculate $\hat{\theta}^* = \hat{\theta}(X^*)$ from the pseudodata.
4. Repeat steps 2 and 3 K times, obtaining $\hat{\theta}_k^*$, $k = 1, \dots, K$.

The empirical distribution of $(\hat{\theta}_k^* - \hat{\theta})$ can be used to approximate the distribution of $(\hat{\theta} - \theta)$. The asymptotic properties of the bootstrap for commonly used statistics such as the mean, median, variance, and distribution quantiles have been studied by Bickel and Freedman (1981) and Freedman (1981). Freedman (1981) demonstrates the validity of the bootstrap for the regression model, showing that, under the assumptions of predetermined variables and i.i.d. errors, the distribution of the coefficients converges to the normal. As we shall see, the bootstrapped distributions of coefficients may converge to a nonnormal distribution when these conditions are violated.

The bootstrap approach, it should be noted, has limitations. For small sample sizes, the bootstrapped distribution $\hat{\theta}(\hat{F})$ may be a poor approximation to $\hat{\theta}(F)$. Schenker (1985), for instance, shows a simple example where the bootstrapped distribution of standard errors would lead the researcher to underestimate the size of confidence intervals. For sufficiently large sample sizes, however, an important advantage of the bootstrap is that it allows the researcher to control for the presence of potentially biasing factors such as the use of overlapping return intervals, the lagged correlation between independent and dependent regression variables, and other idiosyncracies in the distribution of the returns or in the error structure.

B. Bootstrapping Dividend Yields

We adapt the bootstrapping methodology to the dividend yield problem by considering the following model. We want to specify a temporal relationship between returns, dividends, and prices that is analogous to their historical pattern. In particular, price levels—the denominator in the yield term—are dependent upon the capital appreciation return history to that point, and dividends are highly autocorrelated. To capture this relationship in the bootstrap model, we randomly sample total returns R^* from their true distribution, subtract off the contemporaneous income return R_I^* to create a pseudo-capital-appreciation-return series R_X^* . We compound these to calculate a pseudo-price-level series, P^* , which in turn is used to create a pseudo-dividend-yield, Y^* , where $Y_t^* = D_t/P_t^*$, in which D_t are the actual annual dividend flows, and P^* is the simulated price series.

Because total returns have been randomized, there is no relationship between returns and dividends. At the same time, the dividend series exhibits the high degree of persistence actually observed. Therefore, this setup introduces the possibility of autocorrelated errors in the regression, and potential bias in the estimated coefficients. This is a desired feature of the model, since it remains consistent with the null hypothesis.

Thus, our bootstrap procedure is the following:

1. Form the empirical distribution of monthly total stock returns and their associated income returns from the observed vectors of S & P returns.
2. Draw 718 of these observations ($R_1^* \dots R_{718}^*$), with replacement.
3. Compute the relevant statistics:
 - a. Form P^* from $R_X^* = R^* - R_I^*$ as described above. Use this to form $Y^* = D/P^*$.
 - b. Construct multiple horizon return vectors R_T^* for overlapping T-year returns.
 - c. Perform the regressions of future returns on dividend yields, for each horizon T, and save the resulting coefficients β_T^* , the R_T^{2*} and the t -statistics t_T^* .
4. Repeat steps 2 through 3 five thousand times.

III. Empirical Results

Table II reports the relevant quantiles, means, and standard deviations of the bootstrapped distributions of regression coefficients, t -statistics and R^2 's under the null hypothesis. The mean of the coefficient distribution increases uniformly in the return horizon, from 0.15 to 9.59, and the median of the coefficient differs increasingly from the mean, indicating right-skewed distributions. The extent of the skewness is apparent from Figures 1 and 2, which present histograms of the bootstrapped betas over 1-month and 4-year horizons.⁴ The critical value at the 5% level for the 4-year slope coefficient is 30.02, which is very high. Notice that the overlap is not the cause of the skewness; skewness occurs at the 1-month horizon as well as at the 4-year horizon.

The table also shows that all t -statistics increase in the time horizon. As other researchers have noted, the OLS t -statistics are grossly misleading since they fail to correct for the autocorrelated errors induced by the use of overlapping returns, as well as for the bias in the slope coefficient. At the 4-year horizon, OLS t -statistics above 21.16 are observed in 5% of all experiments. What is more worrisome, however, is the fact that the corrected GMM and NW t -statistics also appear to be seriously biased upward. At 4-year horizons, GMM t -statistics above 5.49 are observed in 5% of the

⁴ Similar results are found in empirical tests of the random walk assumption for real exchange rates, which involve regressions on lagged dependent variables. See for instance Abuaf and Jorion (1990).

Table II
Bootstrap on Total Returns: 1927 to 1990

The bootstrap uses 5000 replications, resampled from the actual distribution of total returns, under the null hypothesis of no linear relationship between returns and dividend yields. The table reports the distribution of β , t -statistics and R^2 from dividend yield regressions. The last two columns report the actual statistics observed over the 1927 to 1990 sample, as well as the empirical p -value, which is the proportion of times the observed statistic was exceeded under the null.

Statistic	Fractiles of Statistics							Mean	Std. Dev.	Observed Statistic	Empirical p-Value
	0.010	0.050	0.100	0.500	0.900	0.950	0.990				
1-month											
β	-0.071	-0.020	0.003	0.106	0.374	0.484	0.767	0.154	0.173	0.386	0.0920
$t(\text{OLS})$	-0.885	-0.298	0.073	1.284	2.436	2.753	3.453	1.270	0.928	2.795	0.0464
$t(\text{GMM})$	-0.936	-0.312	0.077	1.287	2.384	2.712	3.292	1.257	0.904	1.208	0.5358
$t(\text{NW})$	-0.936	-0.312	0.077	1.287	2.384	2.712	3.292	1.257	0.904	1.208	0.5358
R^2	0.00000	0.00003	0.00012	0.00217	0.00771	0.00981	0.01533	0.003	0.003	0.010	0.0466
12-month											
β	-0.972	-0.280	0.022	1.329	4.847	6.391	9.968	1.972	2.200	5.108	0.0878
$t(\text{OLS})$	-2.926	-0.995	0.117	4.298	8.613	9.816	12.220	4.373	3.311	9.969	0.0462
$t(\text{GMM})$	-1.031	-0.344	0.040	1.369	2.749	3.188	4.182	1.398	1.092	2.948	0.0740
$t(\text{NW})$	-1.207	-0.411	0.049	1.628	3.271	3.707	4.752	1.653	1.269	3.167	0.1146
R^2	0.00001	0.00029	0.00120	0.02431	0.08948	0.11317	0.16513	0.037	0.038	0.116	0.0462
24-month											
β	-2.190	-0.0642	0.044	2.936	10.520	13.468	20.894	4.222	4.661	9.071	0.1348
$t(\text{OLS})$	-3.768	-1.341	0.171	5.904	12.497	14.348	18.006	6.167	4.813	12.730	0.0932
$t(\text{GMM})$	-1.074	-0.342	0.046	1.466	3.157	3.861	5.715	1.587	1.425	5.319	0.0140
$t(\text{NW})$	-1.233	-0.414	0.053	1.710	3.539	4.226	5.746	1.785	1.435	5.379	0.0138
R^2	0.00062	0.00054	0.00210	0.04531	0.17368	0.21696	0.30380	0.069	0.073	0.179	0.0932
36-month											
β	-3.793	-1.102	0.026	4.776	16.482	21.567	32.303	6.745	7.417	12.939	0.1620
$t(\text{OLS})$	-4.542	-1.585	0.138	7.069	15.489	18.061	22.920	7.508	6.050	15.375	0.1034
$t(\text{GMM})$	-1.165	-0.388	0.022	1.555	3.644	4.659	7.178	1.679	3.643	5.972	0.0214
$t(\text{NW})$	-1.292	-0.435	0.038	1.779	3.890	4.794	6.824	1.917	1.634	4.773	0.0514
R^2	0.00003	0.00077	0.00281	0.06459	0.24710	0.30854	0.41814	0.098	0.102	0.244	0.1034
48-month											
β	-6.426	-1.564	0.050	6.872	23.174	30.018	46.231	9.589	10.546	21.391	0.1166
$t(\text{OLS})$	-5.349	-1.734	0.134	7.945	18.068	21.157	27.167	8.636	7.123	21.419	0.0460
$t(\text{GMM})$	-1.246	-0.406	0.030	1.632	4.149	5.492	9.333	1.972	2.579	7.367	0.0222
$t(\text{NW})$	-1.349	-0.456	0.034	1.822	4.235	5.283	8.107	2.045	1.849	7.300	0.0158
R^2	0.00005	0.00093	0.00354	0.08115	0.31226	0.38369	0.50653	0.124	0.127	0.390	0.0460




Figure 1. Distribution of beta: 1-month horizon. Distribution of slope coefficient from 1-month horizon dividend yield regression under the null. The histogram describes the bootstrap distribution, and the solid line describes the VAR distribution. The actual beta measured over 1927 to 1990 is also reported.

sample.⁵ The mean of the R^2 's also increases uniformly in the return horizon, from near zero to a mean value of 12% for 4-year forecasts. The 5% critical level for the 4-year R^2 is above 38%.

The right-hand-side columns in Table II report the observed statistics from the actual regressions of returns on lagged dividend yield. Note that, while the slope coefficients all exceed the median values for their bootstrapped distributions, none exceeds the relevant 95% fractile. Thus, it is difficult to conclude that the regression statistics differ significantly from those obtained under the null hypothesis as specified, using conventional significance levels. While some of the t -statistics and the R^2 barely exceed the 95% fractiles, these statistics may not be directly applicable in view of the asymmetry in the distribution of slope coefficients.

⁵ There were a few instances where the GMM covariance matrix was not invertible (5 cases at 3-year horizons, 1 at 4-year); the frequencies reported are out of the total number of cases for which GMM t -statistics could be calculated.




Figure 2. Distribution of beta: 4-year horizon. Distribution of slope coefficient from 4-year horizon dividend yield regression under the null. The histogram describes the bootstrap distribution, and the solid line describes the VAR distribution. The actual beta measured over 1927 to 1986 is also reported.

As noted in Richardson and Smith (1991) and Goetzmann (1990), it is insufficient to examine the coefficients separately. Following Hodrick (1992), we assess joint significance of all five slope coefficients by formulating a statistic similar to the Hotelling T^2 : $\hat{\beta} V(\hat{\beta}^*)^{-1} \hat{\beta}$, where $V(\hat{\beta}^*)$ is the covariance matrix of all 5,000 jointly bootstrapped betas. The observed statistic of 11.40, compared to its bootstrapped distribution, is exceeded in 9.5% of the experiments, which again indicates that the slope coefficients are only marginally significant.

Table III presents the results of the bootstrap applied to the two subperiods, 1927 to 1958 and 1959 to 1990. The empirical p -values of the actual 4-year slope coefficients are 0.130 and 0.317 in the two respective samples. Therefore, there is little statistical evidence that the slope coefficients are different from what could be expected under the null of no forecasting ability for dividend yields.

While the previous analysis is distribution free, it may be of some interest to investigate the sensitivity of the results to the distributional assumptions. Therefore, the numerical analysis was also performed using total returns R^*

Table III
Bootstrap on Total Returns: Actual Distribution
(Monthly Data: Subperiods)

The bootstrap uses 1000 replications. The "Observed Statistic" column reports the actual statistic over the subsample. The empirical p -value is the proportion of times the observed statistic was exceeded under the null.

Statistic	Period: 1927–1958				Period: 1959–1990			
	Bootstrap Betas		Observed Statistic	Empirical p -Value	Bootstrap Betas		Observed Statistic	Empirical p -Value
	Mean	Std. Dev.			Mean	Std. Dev.		
1-month								
β	0.247	0.271	0.415	0.195	0.380	0.352	0.668	0.160
$t(\text{OLS})$	1.219	0.933	1.888	0.239	1.411	0.843	2.761	0.044
$t(\text{GMM})$	1.195	0.899	0.855	0.658	1.425	0.849	2.424	0.122
$t(\text{NW})$	1.195	0.899	0.855	0.658	1.425	0.849	2.424	0.122
R^2	0.00606	0.00680	0.00924	0.239	0.00698	0.00636	0.01956	0.044
12-month								
β	3.030	3.436	5.846	0.158	4.768	4.148	8.757	0.137
$t(\text{OLS})$	4.152	3.542	6.972	0.201	4.982	2.972	11.041	0.022
$t(\text{GMM})$	1.424	1.244	2.530	0.170	1.686	1.079	4.028	0.021
$t(\text{NW})$	1.644	1.405	2.614	0.224	1.976	1.236	4.456	0.024
R^2	0.06750	0.07369	0.11586	0.201	0.07783	0.06669	0.24733	0.022
24-month								
β	6.278	7.450	10.778	0.202	10.063	8.381	14.401	0.243
$t(\text{OLS})$	5.838	5.305	8.970	0.265	7.170	4.479	14.291	0.057
$t(\text{GMM})$	1.806	1.991	5.700	0.025	2.095	1.572	3.651	0.126
$t(\text{NW})$	1.883	1.776	4.642	0.058	2.287	1.581	4.322	0.101
R^2	0.12186	0.13017	0.18311	0.265	0.14600	0.11933	0.36259	0.057
36-month								
β	9.757	12.150	16.167	0.205	15.657	12.758	19.740	0.309
$t(\text{OLS})$	7.089	6.762	11.302	0.248	8.845	5.799	17.154	0.081
$t(\text{GMM})$	1.772	7.805	5.163	0.111	2.337	6.092	4.562	0.145
$t(\text{NW})$	2.142	2.151	3.740	0.185	2.609	1.952	4.904	0.105
R^2	0.16414	0.17266	0.26908	0.249	0.20485	0.15875	0.45889	0.081
48-month								
β	13.462	17.845	28.987	0.130	21.882	17.763	27.911	0.317
$t(\text{OLS})$	8.153	8.138	18.061	0.126	10.353	7.051	20.258	0.082
$t(\text{GMM})$	2.605	7.862	11.204	0.040	2.459	10.371	4.677	0.218
$t(\text{NW})$	2.445	2.695	7.049	0.055	2.976	2.388	5.188	0.149
R^2	0.19912	0.20547	0.49333	0.126	0.25728	0.19112	0.55056	0.082

drawn from a normal distribution with mean and variance equal to those of the original sample.⁶ Table IV reports summary statistics for the simulation under normally distributed returns.

The table is substantially in agreement with the results in Table II; none of the slope coefficients is significant at the 5% level. The empirical p -values, however, are systematically lower than those reported from the bootstrap.

⁶ Because a pair of total and income returns is not available in this setup, the dividend yield D^R is constructed from the total price series P^R , as explained in the Appendix. The simulation first generates total returns R^* , from which the pseudoseries P^{R*} is obtained and combined with the actual annual dividend series D^R , to form a pseudo-dividend-yield $Y^* = D^R/P^{R*}$.

Table IV
Alternative Models
(Monthly Data: 1927 to 1990)

Five-thousand replications. Alternative models are: (1) a simulation assuming normally distributed returns, modelling the dependence between returns, prices and yields, (2) a bootstrap experiment assuming predetermined yields, (3) a bootstrap experiment based on a VAR model that allows for endogenous yields. For each model, the mean and standard deviation of the statistic are reported, as well as the empirical p -value, which is the proportion of times the observed statistic, reported in the second column, was exceeded under the null.

Statistic	Observed Statistic	Simulations: Normal Distribution			Bootstrap: Fixed Yield			Bootstrap: VAR		
		Bootstr. Betas			Bootstr. Betas			Bootstr. Betas		
		Mean	Std. Dev.	Empirical <i>p</i> -Value	Mean	Std. Dev.	Empirical <i>p</i> -Value	Mean	Std. Dev.	Empirical <i>p</i> -Value
1-month										
β	0.386	0.135	0.141	0.0628	0.003	0.145	0.0062	0.064	0.149	0.0322
$t(\text{OLS})$	2.796	1.307	0.856	0.0356	0.018	0.992	0.0034	0.338	0.986	0.0068
$t(\text{GMM})$	1.208	1.321	0.867	0.5628	0.017	1.003	0.1164	0.340	0.992	0.1908
$t(\text{NW})$	1.208	1.321	0.867	0.5628	0.017	1.003	0.1164	0.340	0.992	0.1908
R^2	0.0101	0.003	0.003	0.0356	0.0013	0.0019	0.0056	0.0014	0.0020	0.0070
12-month										
β	5.108	1.738	1.856	0.0588	0.025	1.821	0.0052	0.779	1.792	0.0166
$t(\text{OLS})$	9.969	4.500	3.103	0.0386	0.052	3.259	0.0026	1.202	3.191	0.0036
$t(\text{GMM})$	2.948	1.447	1.039	0.0686	0.001	1.159	0.0088	0.434	1.117	0.0192
$t(\text{NW})$	3.167	1.710	1.208	0.1064	-0.002	1.346	0.0116	0.496	1.291	0.0250
R^2	0.1163	0.037	0.036	0.0386	0.0135	0.0189	0.0030	0.0147	0.0204	0.0038
24-month										
β	9.071	3.722	4.006	0.0940	0.022	3.834	0.0114	1.610	3.647	0.0264
$t(\text{OLS})$	12.730	6.320	4.553	0.0822	0.038	4.324	0.0024	1.729	4.228	0.0066
$t(\text{GMM})$	5.319	1.618	1.286	0.0104	-0.010	1.245	0.0010	0.559	1.287	0.0022
$t(\text{NW})$	5.379	1.836	1.383	0.0134	-0.021	1.391	0.0012	0.583	1.375	0.0018
R^2	0.1791	0.070	0.068	0.0822	0.0235	0.0312	0.0028	0.0259	0.0349	0.0072
36-month										
β	12.939	5.959	6.477	0.1230	0.023	6.074	0.0200	2.539	5.584	0.0338
$t(\text{OLS})$	15.375	7.691	5.746	0.0908	0.027	4.969	0.0020	2.156	4.821	0.0046
$t(\text{GMM})$	5.972	1.754	2.586	0.0172	-0.022	1.296	0.0004	0.616	3.505	0.0052
$t(\text{NW})$	4.774	1.955	1.564	0.0442	-0.030	1.405	0.0016	0.669	1.438	0.0060
R^2	0.2444	0.099	0.097	0.0908	0.0309	0.0396	0.0020	0.0344	0.0451	0.0050
48-month										
β	21.392	8.534	9.363	0.0878	0.002	8.668	0.0092	3.606	7.621	0.0068
$t(\text{OLS})$	21.419	8.842	6.787	0.0402	0.006	5.422	0.0000	2.526	5.152	0.0002
$t(\text{GMM})$	7.367	1.958	2.483	0.0142	-0.030	1.332	0.0000	0.382	7.412	0.0084
$t(\text{NW})$	7.300	2.067	1.764	0.0110	-0.033	1.404	0.0000	0.755	1.482	0.0006
R^2	0.3895	0.126	0.122	0.0402	0.0368	0.0459	0.0000	0.0405	0.0522	0.0002

For instance, the 4-year p -value is 0.0878, instead of 0.1166 previously; the 1-month p -value is 0.0628, instead of 0.0920 previously. Differences appear to be due to the fact that there are fewer extreme observations in the tails of the normal distribution than actually observed over the sample; the bootstrap more accurately reflects the actual distribution of stock returns. In spite of these slight differences, the general conclusions are not affected by the distributional assumptions behind the returns simulations.

To understand why these results are so different from traditional regressions, consider another bootstrap experiment where the right-hand-side vari-

ables are exogenous. In this setup, the independent variables are taken to be the actual dividend yield $Y = D/P$, and monthly returns are randomized to generate the dependent variables and multiple horizon returns. The middle panel in Table IV reports information on the bootstrapped statistics in this traditional regression framework. The table shows that the bootstrapped GMM t -statistics correspond much more closely to their expected distribution than in Table II. For instance, the empirical one-tailed 5% level for the GMM t -statistic at the 4-year horizon is 2.13, which is close to what was expected from traditional regression theory. Using the empirical distributions in Table IV would lead to the misleading conclusion that all of the multiyear statistics are strongly significant. This further indicates the need to specify a framework allowing the right-hand-side variables to be endogenous.

As Nelson and Kim (1993) and Hodrick (1992) point out, the issue of endogeneity could be analyzed by modelling a first-order VAR process,

$$Z_{t+1} = AZ_t + u_{t+1} \quad (4)$$

where the columns of Z represent monthly stock returns and dividend yields. In order to simulate Z_{t+1} under the null, we set the slope coefficients in the return equation equal to zero. Then we bootstrap the sample distribution of errors.⁷

This procedure has two desirable features: it models the dividend yield as a highly autocorrelated series, with a first-order autocorrelation of 0.96, and also as an endogenous variable, with a contemporaneous correlation with returns of -0.90 . The bootstrapped distributions are summarized in the right-hand-side panel of Table IV. The VAR approach corrects for the small-sample bias due to the use of an endogenous variable, as well as for biases in the t -statistics, but only indirectly models the serial dependence resulting from the lagged price effect in dividend yield regressions.

As a result, the distributions of slope coefficients appear much closer to normal than under the bootstrap. The curves labelled "VAR" in Figures 1 and 2 clearly show that the VAR model only partially capture the skewness in slope coefficients. In terms of inference, Table IV might indicate predictability in dividend yields: for all horizons, the empirical p -value for the observed coefficient is below 5%. Our results suggest that these rejections may be misleading, because they do not explicitly incorporate the dynamics of regressions with lagged dependent variables.

The results of the dividend yield regressions, for which the price process is endogenous, bear a close resemblance to the well-known simulations performed by Granger and Newbold (1974) and further analyzed by Phillips (1986). Granger and Newbold regressed two independent random walks, and found rejection of the null the rule, rather than the exception. Indeed, their

⁷ Nelson and Kim (1993) randomize on the observed errors, while Hodrick (1992) generates the errors from a multivariate distribution following a GARCH process.

paper has frequently been cited as justification for the need to use differenced price series in econometric studies (see Plosser and Schwert (1978), for example). These results may help us understand the spuriously high R^2 's obtained in the preceding tests. The greater the overlap in the return series, the more closely the return series resembles a price level series, rather than a return series. Although each successive return may be independent (in fact, the bootstrapped returns are independent by construction), the series comprised of a rolling sum of returns is not. Likewise, dividends also resemble random walks. It would thus not be surprising to find that the combination of these two series in a regression could result in spurious conclusions regarding both significance and explanatory power.

To present another perspective on the results, Table V reports estimates of equation (3), which explains returns by the logarithm of the lagged price plus a time trend. As shown before, this specification assumes that dividends grow at a constant rate. Table V shows that the R^2 's obtained in this specification

Table V
Long Horizon Regressions on Lagged Prices and a Time Trend

$$\ln(P_{t,t+T}/P_t) = \alpha_T + \beta_T(-\ln P_t) + \gamma T^t + \epsilon_{t,t+T}$$

where $\ln(P_{t,t+T}/P_t)$ is the continuously compounded total stock return from time t to time $t + T$, P_t is the price as of time t , t is a time trend. The t -statistics presented are: $t(\text{OLS})$, from the usual OLS regression, $t(\text{GMM})$, which adjusts for heteroskedasticity and the moving average process induced by the overlapping observations.

Horizon (Month)	Beta	<i>t</i> -Statistic		Gamma	<i>t</i> -Statistic		R^2
		OLS	GMM		OLS	GMM	
Period: 1927–1990							
1	0.0164	2.83	2.22	0.0001	2.90	2.34	0.011
12	0.2217	11.04	2.23	0.0011	11.37	1.96	0.147
24	0.4461	17.60	2.77	0.0022	18.44	2.53	0.315
36	0.6418	25.69	6.28	0.0033	27.38	7.17	0.507
48	0.7973	33.25	7.04	0.0041	35.94	9.08	0.643
Period: 1927–1958							
1	0.0145	1.63	1.30	0.0001	1.97	1.96	0.011
12	0.2289	7.14	1.90	0.0011	7.41	1.41	0.155
24	0.4966	11.71	2.56	0.0023	12.46	1.96	0.351
36	0.7186	16.84	4.69	0.0036	20.21	6.93	0.584
48	0.8731	21.82	4.52	0.0046	29.04	9.35	0.749
Period: 1959–1990							
1	0.0210	2.01	1.55	0.0001	2.19	1.93	0.012
12	0.2383	6.77	2.88	0.0012	7.72	3.27	0.012
24	0.3336	7.36	3.03	0.0017	9.41	3.55	0.207
36	0.3612	7.24	3.18	0.0020	10.67	3.47	0.280
48	0.3915	6.27	3.28	0.0024	11.02	3.69	0.342




Figure 3. Comparison of 4-year returns and fitted values using yields and using lagged prices. Total 4-year returns are compared with the fitted values from regressions of returns on yields and on lagged prices plus a time trend.

range from 1.1% for a 1-month horizon to 64% for a 4-year horizon. These values, as well as *t*-statistics, are typically higher than those found for the equivalent dividend yield regressions. Figure 3 compares the 4-year returns with fitted values of dividend yield and price regressions. Apparently, the regression on lagged prices fits the returns even better than the regression on dividend yield. These results imply that the economic time variation in the ratio of dividends to past prices may be due primarily to the time variation in the price series, and that long-term forecasting ability has little to do with variation in dividends.

IV. Conclusions

Regression tests of long-horizon returns on dividend yields have previously been interpreted as providing strong evidence of predictability in stock returns. These studies, however, have failed to recognize the serious biases arising from regressions on lagged dependent variables. To illustrate how inference may be affected, we have used bootstrapping techniques, and have modelled the null hypothesis that returns conform to a random walk while at the same time preserving the actual patterns of dividends.

The results of our tests show that, in a simple setting with no linear relationship between future returns and the dividend-price ratio, the OLS procedure, even with standard errors corrected for overlapping data, will often yield results that suggest otherwise. Indeed the coefficients, t -statistics and R^2 's from such regressions are shown to be misleading in the sense that, even when generated by data conforming to the null hypothesis, they yield what might normally be interpreted as strongly significant results. The biases are much stronger than previously thought. As an example, consider the GMM t -statistics for 4-year returns measured over 64 years of monthly data. In a conventional setting, the 5% upper tail critical value is 2.1; using a VAR approach, this value increases to 3.9; our approach yields a value of 5.5.

Our findings argue for a different formulation of such tests, and caution against drawing inferences from usual regression statistics without a thorough understanding of their underlying distributions. The implications of these results extend far beyond tests of the predictive power of dividend yields. Time series studies of returns conditioned upon any ratio involving price levels are also susceptible to the biases reported here. While GMM adjustments and VAR simulations clearly help to adjust for overlaps and small sample biases, this study shows that in some situations they may not be adequate for the purposes of hypothesis testing.

Appendix: Data Construction

The data series were constructed as follows: monthly total, capital, and income returns on the S & P 500 index were obtained from Ibbotson Associates. These are defined, respectively, as

$$R_{t,t+1} = R_{t,t+1}^X + R_{t,t+1}^I = (P_{t+1} - P_t)/P_t + d_{t+1}/P_t,$$

where P_t is a price series that excludes the reinvestment of dividends. Setting P_0 at 100, we recursively compute P_t , as well as the monthly dividend d_t from these series.

A monthly annualized dividend series was computed from compounding twelve monthly dividends at the 1-month Treasury bill rate r_t :

$$D_t = d_t + (1 + r_t)d_{t-1} + (1 + r_t)(1 + r_{t-1})d_{t-2} + \dots$$

Then annual dividend yield is then defined as $Y_t = D_t/P_t$.

For the simulation using total returns, we computed a "total" price that represents the value of a fund mimicking the index with monthly reinvestment of dividends, as $R_{t,t+1} = (P_{t+1}^R - P_t^R)/P_t^R$, with $P_0^R = 100$. the actual income on such a fund is computed recursively from $R_{t,t+1}^I = d_{t+1}^R/P_t^R$, and D^R is computed from d^R as explained above.

REFERENCES

- Abuaf, N., and P. Jorion, 1990, Purchasing power parity in the long run, *Journal of Finance* 45, 157-174.

- Bickel, P., and D. Freedman, 1981, Some asymptotic theory for the bootstrap, *Annals of Statistics* 9, 1196–1271.
- Campbell, J., and R. Shiller, 1988, The dividend price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Dickey, D., and W. Fuller, 1979, Distribution of estimators for autoregressive time series with a unit root, *Journal of the American Statistical Association* 74, 427–431.
- Efron, B., 1979, Bootstrap methods: Another look at the jackknife, *Annals of Statistics* 7, 1–26.
- Fama, E., and K. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- Flood, K., J. Hodrick, and P. Kaplan, 1987, An evaluation of recent evidence on stock market bubbles, Working Paper 1971, National Bureau of Economic Research, Cambridge, Mass.
- Freedman, D., 1981, Bootstrapping regression models, *Annals of Statistics* 9, 1218–1228.
- Goetzmann, W., 1990, Bootstrapping and simulation tests of long-term patterns in stock market behavior, Ph.D. thesis, Yale University.
- Gordon, M., and E. Shapiro, 1956, Capital Equilibrium analysis: The required rate of profit, *Management Science* 3, 102–110.
- Granger, C. J. W., and P. Newbold, 1974, Spurious regressions in econometrics, *Journal of Econometrics* 2, 111–120.
- Hansen, L., 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.
- and R. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates, *Journal of Political Economy* 88, 829–853.
- Hodrick, R., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–386.
- Kendall, M. G., 1973, *Time-Series* (Griffin, London).
- Nelson, C., and M. Kim, 1993, Predictable stock returns: The role of small sample bias, *Journal of Finance* 48, 641–661.
- Newey, W., and K. West, 1987, A simple, positive definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Phillips, P., 1986, Understanding spurious regressions in econometrics, *Journal of Econometrics* 33, 311–340.
- Plosser, C. I., and W. Schwert, 1978, Money income and sunspots: Measuring economic relationships and the effects of differencing, *Journal of Monetary Economics* 4, 637–660.
- Richardson, M., and T. Smith, 1991, Tests of financial models in the presence of overlapping observations, *Review of Financial Studies* 4, 227–254.
- Rozeff, M., 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11, 68–75.
- Schenker, N., 1985, Qualms about bootstrap confidence intervals, *Journal of the American Statistical Association* 80, 360–361.
- White, H., 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817–838.
- Williams, J. B., 1938, *The Theory of Investment Value* (Harvard University Press, Cambridge, Mass.).