

THE MYTH OF LONG-HORIZON PREDICTABILITY

Jacob Boudoukh^a, Matthew Richardson^b and Robert F. Whitelaw^{b*}

This Version: September 20, 2005

THE MYTH OF LONG-HORIZON PREDICTABILITY

Abstract

The prevailing view in finance is that the evidence for long-horizon stock return predictability is significantly stronger than that for short horizons. We show that for all practical purposes (i.e., with persistent regressors) the estimators are almost perfectly correlated across horizons under the null hypothesis of no predictability. For example, for the persistence levels of dividend yields, the analytical correlation is 99% between the 1- and 2-year horizon estimators and 94% between the 1- and 5-year horizons, due to the combined effects of overlapping returns and persistence of the predictive variable. Common sampling error across equations leads to OLS coefficient estimates and R^2 s that are roughly proportional to the horizon under the null of no predictability. This is the precise pattern found in the data. The asymptotic theory is corroborated and the analysis extended by extensive simulation evidence. We perform joint tests across horizons for a variety of explanatory variables, and there is little or no evidence of predictability in the data.

I. Introduction

Over the last two decades, the finance literature has produced growing evidence of stock return predictability though not without substantive debate. The strongest evidence cited so far is that of long-horizon stock returns regressed on variables such as dividend yields, term structure slopes, and credit spreads, among others. A typical view is expressed in the standard empirical textbook for financial economics: “The Econometrics of Financial Markets” by Campbell, Lo and MacKinlay (1997, p.268):

At a horizon of one month, the regression results are rather unimpressive: The R^2 statistics never exceed 2%, and the t-statistics exceed 2 only in the post-World War II subsample. The striking fact about the table is how much stronger the results become when one increases the horizon. At a two-year horizon the R^2 statistic is 14% for the full sample ... at a four-year horizon the R^2 statistic is 26% for the full sample.

However, there is an alternative interpretation of this evidence. Researchers should be equally impressed by the short- and long-horizon evidence for the simple reason that the regressions are almost perfectly correlated. In fact, for an autocorrelation of 0.953 for annual dividend yields, we show analytically that the 1-year and 2-year predictive estimators are 98.8% correlated under the null of no predictability. The correlations are even higher for longer horizons. For example, the correlation between the 4- and 5-year horizon estimators is 99.6%. This degree of correlation manifests itself in multiple-horizon regressions in a particularly unfortunate way. Since the sampling error that is

almost surely present in small samples shows up in every regression, both the estimators and R^2 s are proportional to the horizon.

The contribution of this paper is to provide analytical expressions for the correlation across multiple-horizon estimators and then show, through simulation evidence, that these expressions are relevant in small samples. The analytical expressions relate the correlation across these estimators to both the degree of overlap across the horizons and to the level of persistence of the predictive variable. Our findings are related to an earlier literature that looks at joint tests of the random walk hypothesis of stock prices using multiple-horizon variance ratios and autocorrelations, amongst other estimators (see, e.g., Richardson and Smith (1991) and Richardson (1993)). This earlier literature stressed the importance of taking into account the degree of overlap. The problem here is much more severe. In the previous univariate framework, the predictive variable, past stock returns, is approximately i.i.d. In this paper's framework, the predictive variable, such as dividend yields, is highly persistent.

We show through simulation that any sampling error present in the data under the null of no predictability will show up in the same manner in every multiple-horizon regression when the predictive variable is highly persistent. Using box plots and tables describing the relation across the multiple horizon estimates and R^2 s, we illustrate the exact pattern one should expect under the null hypothesis. In particular, the multiple-horizon estimates are monotonic in the horizon approximately two-thirds of the time and the mean ratios of the 2- to 5-year estimators to the 1-year estimator are 1.93, 2.78, 3.56 and 4.28, respectively. Consider, for example, the actual estimated coefficients for the regression of 1- to 5-year stock returns on dividend yields over the 1926-2004 sample period: 0.117,

0.236, 0.355, 0.424 and 0.469, which correspond to monotonically increasing estimates with corresponding ratios of 2.01, 3.03, 3.62 and 4.01. We show that these estimates lie in the middle of the distribution of possible outcomes under the null of no predictability.

The theoretical and simulation analyses stress the importance of interpreting the evidence jointly across horizons. We develop an analytical expression for a joint test based on the Wald statistic. While the level of high persistence implies that it can be dangerous to document and/or interpret regressions over multiple horizons, the joint tests show that this persistence may lead to powerful tests in an economy in which predictability exists. For example, the predictability may take a particular form in which the multiple-horizon coefficients are much less tied together than implied by the null of no predictability. In applying the joint tests to commonly used predictive variables, we point out various anomalies and show that, in contrast to the conclusions of the existing literature, the book-to-market ratio generates the strongest evidence of predictability. Interestingly, this variable is often cited as the variable with the greatest cross-sectional explanatory power. It is well known that factors in a multi-factor CAPM that capture changes in the investment opportunity set should both predict future stock returns and explain the cross-section of stock returns.

The paper proceeds as follows. In Section 2, we provide the expressions for analyzing multiple-horizon regressions and show that the basic findings carry through to small samples. The small sample results are especially alarming in the context of the existing literature. Section 3 applies the results to a number of data series and evaluates existing evidence. Of some interest, we provide joint tests of predictability. Section 4 concludes.

II. Multiple Horizon Regressions

A. The Existing Literature

The seminal paper in this area that first documents evidence of multivariate stock return predictability over multiple horizons is Fama and French (1988).¹ In brief, they regress overlapping stock returns of one month to four years on dividend yields, and report coefficients and R^2 s that increase somewhat proportionately with the horizon. The paper has over 250 citations to date, and it documents what has become one of the dominant stylized facts in empirical finance. As an illustration of the consensus view, consider John Cochrane's description of the three most important facts in finance in his survey "New Facts in Finance" (1999, p.37).

Now, we know that ...

[Fact] 2. Returns are predictable. In particular: Variables including the dividend/price (d/p) ratio and term premium can predict substantial amounts of stock return variation. This phenomenon occurs over business cycle and longer horizons. Daily, weekly, and monthly stock returns are still close to unpredictable...

This view is emphasized repeatedly in other surveys (see, for example, Fama (1988), Campbell (2000, 2003), and Barberis and Thaler (2003), among others). Furthermore, it is often an important fact to which to calibrate theoretical models (see, among others, Campbell and Cochrane (1999), Menzly, Santos and Veronesi (2004), and Lettau and Ludvigson (2005)). For example, Lettau and Ludvigson (2005, p.584) write:

¹ See also Campbell and Shiller (1988).

It is not necessary to delve far into recent surveys of the asset pricing literature to uncover a few empirical results that have come to represent stylized facts, part of the “standard view” of U.S. aggregate stock market behavior...

[Fact] 2. Returns on aggregate stock market indexes in excess of a short-term interest rate are highly forecastable over long horizons. The log-dividend price ratio is extremely persistent and forecasts excess returns over horizons of many years.

It is fairly well known since Fama and French (1988), and in particular from Campbell (2001), that the key determinants of long-horizon predictability are

- (i) the extent of predictability at short horizons, and
- (ii) the persistence of the regressor.

Thus, the R^2 at long horizons *relative* to the single period R^2 is a function of (ii). Everything else the same (in particular, the single period predictability) higher persistence results in a higher fraction of long-horizon returns that is explainable. As a function of the horizon, the R^2 first rises with the horizon but eventually decays. The effect diminishes eventually due to the exponential decline in the informativeness of the predictive variable. As we show below, (ii) also matters in the case of no predictability but in the presence of sampling error. Nevertheless, this important fact has not been used as the main line of attack against the evidence in support of multivariate predictability of stock returns.

There have been three principal alternative arguments put forward in the literature. The first is data snooping, which is perhaps best described by Foster, Smith, and Whaley

(1997). The basic idea is that the levels of predictability found at short horizons is not surprising given the number of variables that researchers can choose from ex post. Their findings are somewhat supported by a variety of papers, including Bossaerts and Hillion (1999), Cremers (2002), and Goyal and Welch (2003).

A second approach has been to look at the small sample biases of the estimators. In a seminal paper in the area, Stambaugh (1999) shows that the bias can be quite severe given the negative correlation between contemporaneous shocks to returns and the predictive variable, which usually involves some type of stock price deflator. His findings suggest much less predictability once the estimators are adjusted for the bias. However, Lewellen (2004) argues that the effect of the bias may be much smaller if one takes into account the persistence of the predictive variable. Lewellen's approach is similar to Stambaugh's (1999) Bayesian analysis of the predictability problem. While these papers certainly question the magnitude of the predictability, they do not address long-horizon predictability per se.

The third line of criticism of predictability, first explored by Richardson and Stock (1989) in a univariate setting, uses an alternative asymptotic theory in which the horizon increases with the sample size. In particular, Valkanov (2003) argues that long-horizon regressions have poor properties relative to the standard asymptotics. He shows that the estimators may no longer be consistent and have limiting distributions that are functionals of Brownian motions. In fact, the distributions are not normal and are not centered around the true coefficient. He shows that this alternative asymptotic theory works better in small samples. His results can be viewed as the theoretical foundations for earlier simulated

distributions by Kim and Nelson (1993) and Goetzmann and Jorion (1993), and for the intuition put forward by Kirby (1997), who used standard asymptotics.

These three criticisms of the predictability literature are important and have validity. Our paper focuses on a different aspect of predictability by examining the joint properties of the regression estimators across horizons. The conclusions here resemble closely the conclusions reached in Richardson and Smith (1991) and Richardson (1993) with respect to the long-horizon evidence against the random walk in Fama and French (1998) and Poterba and Summers (1988). In many ways, the arguments here are more damaging because we show that the degree of correlation across the multiple-horizon estimators is much higher than in the case of long horizon tests for the random walk. In fact, the null hypothesis of no predictability implies the exact pattern in coefficients and R^2 's found in papers presenting evidence in favor of predictability. We show these results in the next two subsections.

B. Multiple Horizon Regressions – Statistical Properties

We consider regression systems of the following type:

$$\begin{aligned}
 R_{t,t+1} &= \alpha_1 + \beta_1 X_t + \varepsilon_{t,t+1} \\
 &\vdots \\
 R_{t,t+J} &= \alpha_J + \beta_J X_t + \varepsilon_{t,t+J} \\
 &\vdots \\
 R_{t,t+K} &= \alpha_K + \beta_K X_t + \varepsilon_{t,t+K},
 \end{aligned} \tag{1}$$

where $R_{t,t+J}$ is the J -period stock return, X_t is the predictor, e.g., the dividend yield, and $\varepsilon_{t,t+J}$ is the error term over J periods. As is well known from Hansen and Hodrick (1980) and Hansen (1982), among others, the error terms are serially correlated due to

overlapping observations. Deriving through the standard GMM calculations under the null of no predictability and under conditional homoskedasticity (e.g., Richardson and Smith (1991)), one can derive the covariance matrix for any two horizons, J and J^* , for $\hat{\beta}_J$ and $\hat{\beta}_{J^*}$:

$$\text{var}(\hat{\beta}_J, \hat{\beta}_{J^*}) = \frac{\sigma_R^2}{\sigma_X^2} \begin{pmatrix} J+2 \sum_{l=1}^{J-1} (J-l) \rho_l & J + \left[\sum_{l=1}^{J-1} (J-l) (\rho_l + \rho_{l+(J^*-J)}) \right] + \sum_{l=1}^{J^*-J} J \rho_l \\ \dots & J^*+2 \sum_{l=1}^{J^*-1} (J^*-l) \rho_l \end{pmatrix}, \quad (2)$$

where $J^* > J$ and ρ_l is the l^{th} -order autocorrelation of X_t . The above expression for the covariance matrix of the estimators is not particularly intuitive though it is immediately apparent that, for J close to J^* , the estimators are almost perfectly correlated. Less obvious is the fact that, for $\text{cov}(X_t, X_{t-l}) \approx \sigma_X^2$, the estimators are also almost perfectly correlated irrespective of the horizon. Intuitively, the persistence of X_t acts in much the same way overlapping horizons do in terms of independent information across multiple horizons.

A popular simplification is to assume that X_t follows an AR(1) (see Campbell (2001), Boudoukh and Richardson (1994), Stambaugh (1993), and Cochrane (2001), among others). Under the AR(1) model, $\text{cov}(X_t, X_{t-l}) = \rho_l \sigma_X^2 \approx \rho^l \sigma_X^2$ where ρ is the autoregressive parameter for X_t . The covariance matrix in (2) reduces to a much simpler form:

$$\text{Var}(\hat{\beta}_J, \hat{\beta}_{J^*}) = \frac{\sigma_R^2}{\sigma_X^2} \begin{pmatrix} J + \frac{2\rho}{(1-\rho)^2} [(J-1) - \rho(J-\rho^{J-1})] & J + \frac{\rho}{(1-\rho)^2} \{ 2[(J-1) - \rho(J-\rho^{J-1})] + (1-\rho^J)(1-\rho^{J^*-J}) \} \\ \dots & J^* + \frac{2\rho}{(1-\rho)^2} [(J^*-1) - \rho(J^*-\rho^{J^*-1})] \end{pmatrix}. \quad (3)$$

For the special case $J=I$, the correlation between J and J^* is

$$\frac{(1-\rho)^2 + \rho(1-\rho)(1-\rho^{J^*-1})}{\sqrt{(1-\rho)^2} \sqrt{J^*(1-\rho)^2 + 2\rho[(J^*-1) - \rho(J^* - \rho^{J^*-1})]}}. \quad (4)$$

For example, for $J=1$ and $J^*=2$, we get $\sqrt{\frac{1+\rho}{2}}$. In our sample the autocorrelation of the dividend yield is 0.953, which yields a correlation of 0.988 between the 1- and 2-year estimators. It does not get much better as the horizon J^* increases to 3-, 4-, and 5-years, producing correlations of 0.974, 0.959, 0.945, respectively.² Even at a 10-year horizon, the correlation is over 87%. With the types of sample sizes faced by researchers in the field of empirical finance, these results suggest one has to be extremely cautious in interpreting the coefficients separately (as has been the case in the existing literature).

For the typical 1- through 5-year horizons looked at in the literature, we provide the analytical covariance matrix of the estimators under the null of no predictability and the dividend yield's ρ of 0.953:

$$\text{var}(\hat{\beta}_{1-5}) = \frac{\sigma_R^2}{\sigma_X^2} \begin{pmatrix} 1 & 0.988 & 0.974 & 0.959 & 0.945 \\ & 1 & 0.993 & 0.982 & 0.970 \\ & & 1 & 0.995 & 0.986 \\ & & & 1 & 0.996 \\ & & & & 1 \end{pmatrix}. \quad (5)$$

Several observations are in order. First, the high degree of correlation across the multiperiod estimators implies that, under the null hypothesis of no predictability, the regressions are essentially redundant. This is important because there is little doubt that the literature has not taken this viewpoint. Second, under the null, the estimators are asymptotically distributed as multivariate normal with mean zero. While this is clearly

² Of course, these correlations are lower bounds as the magnitude increases as J increases for a fixed J^* .

not true in small samples,³ consider using this distribution to understand the effect of sampling error across the equations. Specifically, conditional on $\hat{\beta}_1 = \bar{\beta}_1$, what do we expect $\hat{\beta}_{J^*}$ to equal? Using properties of a bivariate normal, we can write

$$E[\hat{\beta}_{J^*} | \hat{\beta}_1 = \bar{\beta}_1] = \left(1 + \frac{\rho(1 - \rho^{J^*-1})}{1 - \rho}\right) \bar{\beta}_1. \quad (6)$$

For ρ close to 1, the pattern in coefficients should basically be proportional to the horizon. As an example, for $\rho=0.953$, $\hat{\beta}_{J^*}$ as a multiple of $\hat{\beta}_1$ is $1.953 \hat{\beta}_1$, $2.861 \hat{\beta}_1$, $3.727 \hat{\beta}_1$ and $4.552 \hat{\beta}_1$ for the 2-, 3-, 4- and 5-year horizon regressions, respectively. Similarly, for the R^2 of the regression,

$$E[R_{J^*}^2 | R_1^2 = \bar{R}_1^2] = \frac{\left(1 + \frac{\rho(1 - \rho^{J^*-1})}{1 - \rho}\right)^2}{J^*} \bar{R}_1^2. \quad (7)$$

For ρ close to 1, the R^2 s also increases significantly with the horizon, i.e., the ratio of R^2 s is 1.907, 2.729, 3.472 and 4.143 for the 2-, 3-, 4- and 5-year horizon regressions, respectively.

The intuition is straightforward. Compare the regression of $R_{t,t+1}$ on X_t to that of $R_{t,t+K}$ on X_t . The former regression involves summing the cross product of the sequence of $R_{t,t+1}$ and X_t for all t observations. Note that for a persistent series X_t , there is very little information across the sequence of X_t s. Thus, when an unusual draw from $R_{t,t+1}$ occurs (denote it R_{t^*,t^*+1}) and this observations happens to coincide with the most recent

³ See Stambaugh (1999) for small sample bias and Valkanov (2003) for nonnormality of the distributions of the estimators.

value of the predictive variable, X_{t^*} , it will also coincide with all the surrounding X_t variables such as X_{t^*-1} , X_{t^*-2} , and X_{t^*-3} . Since R_{t^*,t^*+1} shows up K times in $R_{t,t+K}$ within the sample period (i.e., R_{t^*+1-K,t^*+1} , R_{t^*+2-K,t^*+2} , \dots , R_{t^*,t^*+K}), the impact of the unusual draw of R_{t^*,t^*+1} for $R_{t,t+K}$ will be roughly K times than that of $R_{t,t+1}$.

At first glance, the above results provide a fairly devastating verdict for the strategy of running multiple long-horizon regressions. However, this view is not necessarily an accurate one. Because the regressions are linked so closely under the null hypothesis of no predictability, joint tests may have considerable power under alternative models. What are these alternatives? The models must be such that the long horizons pick up information not contained at short horizons. Thus, the standard model in which short-horizon returns are linear in the current predictor and that predictor follows an ARMA process is clearly not a good candidate. One would be better to focus on estimating the short horizon in this case (see, for example, Campbell (1991), Hodrick (1992), and Boudoukh and Richardson (1994), among others). It should be noted though that the standard model is often chosen for reasons of parsimony rather than on an underlying theoretical basis. In the next section, we will look at the evidence associated with performing joint tests across horizons.

The above theoretical results are based on asymptotic properties of fixed-horizon estimators. There is a priori reason to be wary of these properties in small samples. In particular, there is considerable evidence of a small sample bias in the coefficient estimators and of non-normality in small samples. For example, as mentioned in Section II.A, Stambaugh (1999) derives a small sample bias of potentially large magnitude.

Therefore, it would be useful to evaluate some of the above conclusions in the context of this bias. In addition, Valkanov (2003) shows that, under alternative asymptotics in which the horizon length increases with the sample size, the estimators are no longer consistent and can have limiting distributions that deviate significantly from normality. This alternative asymptotic theory seems to behave better in small samples than the fixed-horizon one (at least for long-horizon estimators). As a result, the next section looks at small sample properties of the estimators and, in particular, the patterns in sampling error across equations. As a preview of the results to come, the basic tenet of equations (2) and (3), namely the dependence across equations, carries through to small samples.

C. Multiple Horizon Regressions – Simulation Evidence

We simulate the model in equation (1) under the assumption of no predictability, an AR(1) process on X_t , and 75 years of annual data. The analysis is performed over 1- to 5-year horizons with the AR parameter ρ , the standard deviation of X_t and $\varepsilon_{t,t+1}$, and the correlation between $\varepsilon_{t,t+1}$ and $u_{t,t+1} \equiv (X_{t+1} - \rho X_t)$ chosen to match the data.⁴ The simulations involve 100,000 replications each.

Table 1A reports the simulated 5x5 correlation matrix of the multiple horizon estimators. Consistent with the theoretical analytical calculations provided in Section II.B, the correlations tend to be high even for the most distant horizons. For example, the simulated correlations between the 1-year and 2- to 5-year horizon estimators are 0.966,

⁴ Specifically, for the regression of annual stock returns on the most commonly used predictive variable, namely dividend yields, we estimate $\rho = 0.953$, $\sigma_{\varepsilon} = 0.196$, $\sigma_u = 0.166$ and $\sigma_{\varepsilon u} = -0.683$. While the magnitudes of σ_u and σ_{ε} do not matter, this is not true for either the persistence variable ρ (Boudoukh

0.925, 0.883 and 0.841, respectively. These results show that the correlation calculations under the fixed-horizon asymptotics in II.B hold in small samples. This finding implies that the estimators' almost perfect cross-correlation leads to little independent information across equations. Thus, the sampling error that is surely present in small samples shows up in every equation in (1).

As shown in II.B, the persistence ρ is an important determinant of the magnitude of the correlation matrix of the multiple-horizon estimators. Figure 1A graphs the correlation between the 1-year and 2- to 5-year horizon estimators as a function of ρ , in particular, for values $\rho = 0.953, 0.750, 0.500, 0.250$ and 0.000 . Consistent with the theory in II.B, the correlations decrease as ρ falls. The drop-off can be quite large as the horizon increases. For example, as a function of the above ρ s, the 1- and 2-year estimators have a correlation of 0.966, 0.917, 0.849, 0.776, and 0.698, respectively, in contrast to the 1- and 5-year estimator's correlation of 0.841, 0.684, 0.545, 0.465, and 0.429. Even when the predictive variable has no persistence, the correlation can still be quite high due to the overlapping information across the multiple horizon returns.

However, the staggering result in Table 1A is that 65.2% of all the replications produce monotonic estimates as a function of the horizon. That is, almost two-thirds of the time the data produces coefficients increasing or decreasing with the horizon. This fact coincides with the predictions from the asymptotic theory given in II.B. To understand how unlikely monotonicity is, suppose that the five different multiple-horizon estimators were i.i.d. In this setting, the probability of a monotonic relation is 0.83%,

and Richardson (1994)) or the correlation $\sigma_{\epsilon u}$ (e.g., Stambaugh (1999)). Thus, we also investigate different values for these parameters.

approximately $1/78^{\text{th}}$ the true probability for the multiple-horizon estimators. Even with overlapping horizons, monotonicity drops sharply as ρ falls, i.e., from the 65.2% value to 36.7%, 19.7%, 10.8%, and 6.0% for $\rho = 0.953, 0.750, 0.500, 0.250,$ and $0.000,$ respectively. This result just further highlights the importance of persistence in the predictive variable for generating these patterns.

One possible explanation for this finding is that the small sample bias increases with horizon (e.g., Stambaugh (1999), Goetzmann and Jorion (1993) and Kim and Nelson (1993)). Table 1A confirms the increasing small sample bias with the means of the 1- to 5-year coefficients equal to 0.048, 0.092, 0.132, 0.169, and 0.203, respectively. To investigate whether the monotonicity is due to this bias, Table 1B duplicates Table 1A under the assumption that $\sigma_{\varepsilon u} = 0$. For this value, the small sample bias is theoretically zero, and the estimates are unbiased in our simulations as well. Interestingly, the monotonicity falls only slightly to 57.2%. Furthermore, Table 1B shows that the correlation matrix across the multiple horizon estimators is virtually identical to that of Table 1A. Thus, the monotonicity property is being driven by the almost perfect correlation across the estimators and the increasing horizon.

As described in Section II.A, much of the literature has argued for predictability by focusing on the increase in the coefficient estimates as a function of the horizon. Both theoretically and in simulation, we show that this is expected under the null of no predictability. An alternative measure of predictability also considered in the literature is the magnitude and pattern of R^2 s across horizons. While the R^2 is linked directly to the coefficient estimate, it is nonetheless a different statistic of the data. Table 2A reports the

simulated 5x5 correlation matrix of the multiple horizon R^2 s as well as their means, medians, standard deviations and monotonicity properties.

Similar to Table 1A, the R^2 s are all highly correlated across horizons. For example, the simulated correlation between the 1-year and 2- to 5-year horizon R^2 s are 0.948, 0.886, 0.823, and 0.762, respectively. This degree of correlation leads to R^2 s being monotonic in the horizon 51.3% of the time under the null of no predictability. This is the exact pattern documented in the literature. These results are not due to the Stambaugh (1999) small sample bias as both the degree of correlation and monotonicity also show up in simulated data without the bias (see Table 2B where the cross equation correlation is zero). Also, analogous to the evidence for the multiple-horizon coefficient estimators, the degree of cross-correlation and monotonicity depends crucially on the level of persistence ρ of the predictive variable.

Figures 1A and 1B provide the correlation coefficients between the one- and the J -period betas and R^2 s. The correlations are plotted for different persistence parameters and demonstrate the monotonic and nearly linear relation one would expect for both sets of correlations, and the dependence of this effect on the persistence parameter.

The theoretical calculations of Section II.B imply an even stronger condition than monotonicity. In particular, for ρ close to 1, the coefficients and R^2 s should be linearly increasing on a one-to-one basis with the horizon under the null hypothesis. Because this is the typical pattern found in U.S. data, it seems worthwhile investigating this implication via simulation under the null of no predictability. We compare the ratio of the 2- to 5-year coefficient and R^2 estimates to the 1-year estimates. Since there are numerical issues when using denominators close to zero, we condition on the 1-year

estimate having an absolute value greater than 0.01 or an R^2 greater than 0.5%. Approximately 87% and 61% of the simulations satisfy these criteria, respectively.

The results are provided in Table 3A. As predicted by the theory, the mean ratios of the estimates are 1.93, 2.78, 3.56, and 4.28 for the 2-, 3-, 4-, and 5-year horizons, respectively. The R^2 are equally dramatic with corresponding ratios of 1.95, 2.87, 3.75, and 4.59.⁵ Note that these simulations are performed under the null of no predictability – betas are zero but the other parameters are calibrated to match the joint distribution of stock returns and dividend yields in the data. How do these results compare with the estimated coefficients and R^2 s from the actual data? In the data the ratios for the 2-, 3-, 4-, and 5-year horizons are 2.01, 3.03, 3.62, and 4.00 for the beta estimates and 1.86, 3.01, 3.44, and 3.73 for the R^2 s. The similarities are startling.

Figures 2A and 2B plot the ratios as a function of the persistent parameter ρ . For large ρ , both the coefficient estimates and R^2 s increase linearly with the horizon with fairly steep slopes (albeit not one-for-one). As persistence drops off, the slope diminishes dramatically. For $\rho=0$, the ratio plot is actually flat. Nevertheless, given the high persistence of the predictive variables used in practice, the more relevant ratios would be those corresponding to steep slopes. These graphs refer to the mean of the ratio and do not address the distribution of these ratios. Understanding the distribution allows us to address whether the actual estimates fall within these empirical null distributions.

To better understand the statistical likelihood of the observed evidence in light of the distribution of the various relevant coefficients under the null, Figures 3A and 3B show

⁵ Similar to the earlier tables, Table 3b shows that these findings are not due to the Stambaugh bias and hold equally well for $\sigma_{\epsilon t} = 0$.

box plots of the distribution of the multiple horizon coefficient estimates and R^2 's conditional on the 1-year coefficient estimate and R^2 being close to the actual values (i.e., $\beta=0.117$ and $R^2=4.31\%$). The box plots shows the median, the 25th and 75th percentiles, and the more extreme 10th and 90th percentiles of the distribution. Several observations are in order. First, consistent with Figures 2A and 2B, the percentiles linearly increase at a fairly steep rate. Second, the actual values of the coefficients and R^2 's from the data (marked as diamonds in the graph) lie uniformly between the 25th and 75th percentiles. This explains why there is almost no hope of finding evidence against the null. That is, given some amount of sampling error, the null of no predictability produces precisely the pattern one would expect in the coefficients. Because the sample sizes are relatively small, the data almost guarantees sampling error is present. Third, the plots show what matters is the magnitude of the coefficient at short horizons. As evidenced by the voluminous literature on stock return predictability in finance, no researcher has ever pointed to the short horizon evidence as being in any way remarkable.

III. Evidence

The theory and corresponding simulation evidence provided in Section II suggests that it will be very difficult to distinguish between the null of no predictability and alternative models of time-varying expected returns that involve persistent autoregressive processes. The reason is that sampling error produces virtually identical patterns in both R^2 's and coefficients across horizons. However, this finding does not necessarily imply that joint tests will not distinguish the null from other alternatives. Recall that the null implies highly correlated regression coefficient estimators, which necessarily induce the

coefficient pattern. Even with unremarkable coefficient estimators, yet nonconforming coefficient patterns, one might find strong evidence against the null of no predictability.

In this section, we look at a number of commonly used variables to test for the predictability of stock returns. For stock returns, we use the value-weighted (VW) CRSP portfolio. For predictors, we use the log dividend yield on the S&P500, the log earnings yield on the S&P500, the default spread between Baa and Aaa yields, the term spread between long-term government bond yields and T-bill yields, and the log book-to-market ratio.⁶ In addition, we include the log dividend yield on the S&P500 and the CRSP value-weighted index calculated directly from the returns with and without dividends (available on CRSP).

The regression analysis corresponds to equation (1) and covers horizons of 1 to 5 years over the period 1926-2004. For each set of multiple-horizon regressions, we calculate the coefficient, its analytical standard error (using results in II.B), its asymptotic p-value and its simulated p-value (under an AR(1) with matching parameters). In addition, a joint Wald test is conducted across the equations using both asymptotic and simulated p-values. These results are reported in Table 4.

All the series show the much-cited pattern of increasing coefficient estimates and corresponding R^2 's. For the dividend yield on the S&P500, the dividend yield on the CRSP VW index, the dividend yield on the S&P500 (Goyal/Welch), the log earnings yield on the S&P500, the default spread, the term spread, and the log book-to-market ratio, the increases in R^2 's from the 1-year to the 5-year horizon are 4.31% to 16.11%, 4.98% to 17.24%, 1.80% to 16.05%, 3.31% to 10.65%, 0.03% to 3.31%, 3.30% to

14.43%, and 3.00% to 16.37%, respectively. However, the persistence levels of the associated variables are 95.3%, 93.3%, 89.5%, 76.2%, 79.2%, 61.5%, and 93.9%, respectively. Therefore, the fact that many of the series, especially the dividend yield, earnings yield, term spread and book-to-market, have significant coefficients using asymptotic p-values across most of the horizons should not come as a surprise. Under the null, the regressions at each horizon are virtually the same regression.

Table 5 is an alternative representation of the results in Table 4, i.e., the ratios of the coefficient estimates and R^2 s across horizons. With the exception of the default spread, the ratios for both quantities are similar to the simulated ratios under the null hypothesis of no predictability. In all cases, the ratios (and therefore the underlying coefficient estimates and R^2 s) increase with the horizon. Thus, the finding that some of the 1-year regressions are significant and the same variables produce virtually identical patterns at longer horizons is actually evidence that the annual regression results were due to sampling error.⁷ The joint tests confirm this phenomenon by producing much higher p-values, e.g., 0.14 for the dividend yield on the CRSP VW index, 0.09 for the dividend yield on the S&P500 (Goyal/Welch data), 0.39 for the earnings yield, 0.75 for the default spread, 0.31 for the term spread, and 0.03 for the book-to-market ratio, our only significant finding at the 5% level.

Several observations illustrate the nature of the joint tests. First, consider the regression results for the dividend yield on the CRSP VW index versus the S&P500 index (Goyal/Welch data). By almost any eyeball measure, the evidence appears stronger

⁶ See Goyal and Welch (2003). The data are available on Amit Goyal's website (<http://www.bus.emory.edu/AGoyal/>).

for the VW index. Four of the five horizons produce p-values less than 0.05 and the remaining one lies at 0.06 (using asymptotic p-values). In contrast, the S&P500 index produces only three coefficients with p-values less than 0.05 and also produces a p-value of 0.12 at the 1-year horizon. Nevertheless, the p-value of the joint test for the S&P500 dividend yield is lower than that for the CRSP VW index. Table 5 shows that the ratios across horizons for the S&P500 are less like those produced in simulations under the null, and thus they provide some evidence against the null.

Second, the one significant variable under the joint cross-horizon test, i.e., book-to-market, does not look any more impressive than the various dividend yield measures, earnings yield or term premium in terms of its individual coefficient estimates and corresponding R^2 s. Yet, the significance level of the joint test is much higher. Why? The pattern in the coefficients, while monotonic, is much less linear and one-to-one than implied by a series with a 94% persistence level. This example illustrates the power of the joint test to uncover seemingly innocuous differences across horizons.

Third, the simulated p-values in general show much less significance for both the individual and joint tests. This mirrors the small sample findings of Goetzmann and Jorion (1993), Kim and Nelson (1993) and Valkanov (2003). However, one interesting aspect of the simulated results is that, while book-to-market is no longer significant at the individual horizon level (with p-values increasing from the 5% level to the 20% level), the joint test remains significant at the 10% level. Again, this is because the pattern in the coefficient estimates is not consistent with that of the null with a persistent predictive

⁷ This has even greater support once the researcher takes into account the data snooping arguments of Foster, Smith and Whaley (1997).

variable. As shown in Tables 1A and 1B, the correlation pattern across multiple-horizon estimators is robust to small sample considerations.

IV. Conclusion

Long-horizon stock return predictability is considered one of the more important pieces of evidence in the empirical asset pricing literature over the last couple of decades (see, e.g., the textbooks of Campbell, Lo and MacKinlay (1997) and Cochrane (2001)). The evidence is set forth as a yardstick for theoretical asset pricing models and is slowly penetrating the practitioner community (see Brennan and Xia (2005) and Asness (2003) for two recent examples). In addition, long-horizon predictability has also been documented in other markets, which is perhaps not surprising given our analysis. For example, the highly cited work of Fama and Bliss (1987) and Mark (1995) document results similar in spirit to the ones discussed in this paper for bond returns and exchange rates, respectively. Both papers involve highly persistent regressors and document nearly linearly increasing betas and R^2 s.

In this paper we show that stronger long-horizon results, in the form of higher betas and increasing R^2 s, present little if any independent evidence over and above the short-horizon results. Under the null hypothesis of no predictability, sampling variation can generate small levels of predictability at short horizons. This result is well known. Our research shows that higher levels of predictability at longer horizons are to be expected as well.

References

- Asness, C., 2003, "Fight the Fed Model," *Journal of Portfolio Management*, 11-24.
- Barberis, N., and R. Thaler, 2003, "A Survey of Behavioral Finance," *The Handbook of the Economics of Finance*, ed. North Holland: Amsterdam.
- Bossaerts, P., and P. Hillion, 1999, "Implementing Statistical Criteria to Implement Return Forecasting Models: What Do We Learn?" *Review of Financial Studies*, 12 (2), 405-428.
- Boudoukh, J., and M. Richardson, 1993, "The Statistics of Long-Horizon Regressions," *Mathematical Finance*, 4 (2), 103-120.
- Brennan, M., and Y. Xia, 2005, "Persistence, Predictability, and Portfolio Planning," Wharton working paper.
- Campbell, J., 2000, "Asset Pricing at the Millennium," *Journal of Finance*, 55 (4), 1515-1567.
- Campbell, J., 2001, "Why Long Horizons? A Study of Power Against Persistent Alternatives," *Journal of Empirical Finance*, 8, 459-491.
- Campbell, J., 2003, "Consumption-Based Asset Pricing" in Chapter 13 of *The Handbook of the Economics of Finance*, ed. North Holland: Amsterdam, 803-881.
- Campbell J., and J. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107 (2), 205-251.
- Campbell, J., A. Lo, and C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.
- Campbell, J., and R. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1 (3), 195-228.
- Cochrane, J., 1999, "New Facts in Finance", *Economic Perspectives*, 23 (3), 36-58.
- Cremers, M., 2002, "Stock Return Predictability: A Bayesian Model Selection Perspective," *Review of Financial Studies*, 15 (4), 1223-1249.
- Fama, E., 1998, "Market Efficiency, Long-Term Returns, and Behavioral Finance," *Journal of Financial Economics*, 49, 283-306.

Fama, E., and R. Bliss, 1987, "The Information in Long Maturity Forward Rates," *American Economic Review*, 77, 680-92.

Fama E., and K. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3-25

Foster, D., T. Smith, and R. Whaley, 1997, "Assessing Goodness-of-Fit of Asset Pricing Models: The Distribution of the Maximal R-Squared," *Journal of Finance*, 52 (2), 591-607.

Goetzmann, W., and P. Jorion, 1993, "Testing The Predictive Power of Dividend Yields," *Journal of Finance*, 48, 663-679.

Goetzmann, W., and P. Jorion, 1995, "A Longer Look at Dividend Yields," *Journal of Business*, 68 (4), 483-508.

Goyal, A., and Welch I., 2003, "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 49 (5), 639-654.

Hodrick, R., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *Review of Financial Studies*, 5(3), 357-386.

Kim, M.J., and C.R. Nelson, 1993, "Predictable Returns – the Role of Small Sample Bias," *Journal of Finance*, 48 (2), 641-661.

Kirby, C., 1997, "Measuring the Predictable Variation in Stock and Bond Returns," *Review of Financial Studies*, 10 (3), 579-630

Lettau, M., and S. Ludvigson, 2005, "Expected Returns and Expected Dividend Growth," *Journal of Financial Economics*, 76 (3), 583-626.

Lewellen, J., 2004, "Predicting Returns with Financial Ratios," *Journal of Financial Economics*, 74 (2), 209-235.

Mark, N. C., 1995, "Exchange Rates and Fundamentals, evidence on Long Horizon Predictability," *American Economic Review*, 85 (1), 201-218.

Menzly, L., T. Santos and P. Veronesi, 2004, "Understanding Predictability," *Journal of Political Economy*, 112, 1-47.

Poterba, J., and L. Summers, 1988, "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics*, 22, 27-60.

Richardson, M., 1993, "Temporary Components of Stock Prices: A Skeptic's View," *Journal of Business and Economics Statistics*, 11 (2), 199-207.

Richardson, M., and T. Smith, 1991, "Tests of Financial Models in the Presence of Overlapping Observations," *Review of Financial Studies*, 4 (2), 227-254.

Richardson, M., and J. Stock, 1989, "Drawing Inferences from Statistics Based on Multi-Year Asset Returns," *Journal of Financial Economics*, 25, 323-348

Stambaugh, R.F., 1993, "Estimating Conditional Expectations When Volatility Fluctuates," NBER Technical Paper 140.

Stambaugh R.F., 1999, "Predictive Regressions," *Journal of Financial Economics*, 54 (3), 375-421.

Valkanov R., 2003, "Long-Horizon Regressions: Theoretical Results and Applications," *Journal of Financial Economics*, 68 (2), 201-232

Panel A: $\sigma_{eu} = -0.683$

Horizon	Coefficient Estimates			Correlations			
	Mean	SD	Median	Horizon			
				2	3	4	5
1	0.048	0.068	0.037	0.966	0.925	0.883	0.841
2	0.092	0.128	0.074		0.979	0.946	0.908
3	0.132	0.182	0.109			0.985	0.957
4	0.169	0.232	0.142				0.988
5	0.203	0.277	0.175				
% monotonic		65.24					

	Test Statistics			Size		
	Mean	SD	Median	Horizon		
				10%	5%	1%
Wald	7.733	119.281	5.536	20.75	13.15	5.42
P-value	0.396	0.294	0.354			

Panel B: $\sigma_{eu} = 0$

Horizon	Coefficient Estimates			Correlations			
	Mean	SD	Median	Horizon			
				2	3	4	5
1	0.000	0.063	0.000	0.960	0.913	0.867	0.823
2	0.000	0.121	0.001		0.977	0.940	0.900
3	0.000	0.175	0.001			0.984	0.954
4	0.000	0.227	0.002				0.987
5	0.000	0.276	0.001				
% monotonic		57.22					

	Test Statistics			Size		
	Mean	SD	Median	Horizon		
				10%	5%	1%
Wald	7.199	59.730	5.203	18.92	12.06	5.10
P-value	0.423	0.300	0.392			

Table 1: Distribution of Coefficient Estimates and Test Statistics

Panel A reports the mean, standard deviation, and median of the coefficient estimates from the predictive regression (equation (1)), and the correlations between these estimates, for horizons of 1 to 5 years across 100,000 simulations. “% monotonic” is the percentage of the simulations that produce coefficients that are monotonic in the horizon. Panel A also reports the mean, standard deviation, and median of the joint Wald test statistic (across horizons), the associated p-values, and the percentage of statistics that reject the null of no predictability at the 10%, 5%, and 1% levels. There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\rho = 0.953$, $\sigma_e = 0.196$, $\sigma_u = 0.166$, $\sigma_{eu} = -0.683$. Panel B reports the same statistics for $\sigma_{eu} = 0$ (all other simulation parameters are the same as in Panel A).

Panel A: $\sigma_{eu} = -0.683$

Horizon	Mean	SD	Median	Correlations			
				Horizon			
				2	3	4	5
1	1.791	2.343	0.885	0.948	0.886	0.823	0.762
2	3.386	4.280	1.746		0.968	0.916	0.858
3	4.845	5.948	2.564			0.976	0.932
4	6.182	7.412	3.326				0.981
5	7.415	8.708	4.089				
% monotonic		51.30					

Panel B: $\sigma_{eu} = 0$

Horizon	Mean	SD	Median	Correlations			
				Horizon			
				2	3	4	5
1	1.345	1.861	0.619	0.927	0.845	0.767	0.695
2	2.524	3.399	1.202		0.957	0.892	0.821
3	3.612	4.769	1.743			0.969	0.914
4	4.624	5.992	2.278				0.976
5	5.570	7.094	2.790				
% monotonic		42.50					

Table 2: Distribution of R^2 s

Panel A reports the mean, standard deviation, and median of the R^2 s from the predictive regression (equation (1)), and the correlations between these R^2 s, for horizons of 1 to 5 years across 100,000 simulations. “% monotonic” is the percentage of the simulations that produce R^2 s that are monotonic in the horizon. There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\rho = 0.953$, $\sigma_\varepsilon = 0.196$, $\sigma_u = 0.166$, $\sigma_{eu} = -0.683$. Panel B reports the same statistics for $\sigma_{eu} = 0$ (all other simulation parameters are the same as in Panel A).

Panel A: $\sigma_{\varepsilon u} = -0.683$

Horizon	Coefficient Estimate Ratios				R ² Ratios			
	Mean	SD	Median	# of Sim.	Mean	SD	Median	# of Sim.
2	1.930	0.854	1.916	87,166	1.953	0.826	1.868	61,467
3	2.784	1.842	2.731	87,166	2.872	1.752	2.597	61,467
4	3.564	2.963	3.461	87,166	3.753	2.792	3.209	61,467
5	4.276	4.183	4.119	87,166	4.595	3.891	3.739	61,467
% monotonic	69.85				% monotonic 59.40			

Panel B: $\sigma_{\varepsilon u} = 0$

Horizon	Coefficient Estimate Ratios				R ² Ratios			
	Mean	SD	Median	# of Sim.	Mean	SD	Median	# of Sim.
2	1.869	1.102	1.886	85,717	1.904	0.961	1.782	54,603
3	2.615	2.309	2.663	85,717	2.754	1.992	2.397	54,603
4	3.261	3.655	3.346	85,717	3.548	3.086	2.884	54,603
5	3.822	5.085	3.94	85,717	4.292	4.204	3.234	54,603
% monotonic	61.57				% monotonic 48.88			

Table 3: Distribution of Coefficient Estimate and R² Cross-Horizon Ratios

Panel A reports the mean, standard deviation, and median of the coefficient estimate and R² ratios (i.e., β_i/β_1 and R_i^2/R_1^2 , $i = 2, \dots, 5$) from the predictive regression (equation (1)) across the simulations out of the 100,000 for which $|\beta_1| > 0.01$ or $R_1^2 > 0.5\%$, respectively. “% monotonic” is the percentage of these simulations that produce coefficient estimates and R²s that are monotonic in the horizon. There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\rho = 0.953$, $\sigma_\varepsilon = 0.196$, $\sigma_u = 0.166$, $\sigma_{\varepsilon u} = -0.683$. Panel B reports the same statistics for $\sigma_{\varepsilon u} = 0$ (all other simulation parameters are the same as in Panel A).

	Horizon					Wald
	1	2	3	4	5	
log dividend yield, S&P500 ($\rho = 0.953$)						
β	0.117	0.236	0.355	0.424	0.469	9.559
Std. Err.	0.066	0.130	0.193	0.255	0.316	
Asym. P-value	0.037	0.034	0.033	0.048	0.069	0.089
Sim. P-value	0.144	0.129	0.116	0.137	0.166	0.190
R^2	4.315	8.019	13.009	14.856	16.107	
log dividend yield, CRSP VW ($\rho = 0.933$)						
β	0.136	0.273	0.404	0.474	0.526	8.405
Std. Err.	0.071	0.140	0.207	0.273	0.337	
Asym. P-value	0.027	0.025	0.025	0.041	0.059	0.135
Sim. P-value	0.137	0.121	0.113	0.138	0.164	0.226
R^2	4.977	9.181	14.357	15.776	17.244	
log dividend yield, S&P500 (Goyal/Welch) ($\rho = 0.895$)						
β	0.068	0.196	0.301	0.376	0.421	9.503
Std. Err.	0.058	0.114	0.167	0.218	0.268	
Asym. P-value	0.123	0.042	0.035	0.043	0.058	0.091
Sim. P-value	0.273	0.137	0.114	0.119	0.138	0.138
R^2	1.798	6.878	11.410	14.481	16.045	
log earnings yield, S&P500 ($\rho = 0.762$)						
β	0.098	0.217	0.315	0.370	0.366	5.241
Std. Err.	0.062	0.117	0.167	0.214	0.257	
Asym. P-value	0.058	0.032	0.030	0.042	0.077	0.387
Sim. P-value	0.127	0.078	0.068	0.081	0.128	0.417
R^2	3.308	7.396	10.907	12.315	10.647	
default yield spread ($\rho = 0.792$)						
β	0.429	3.299	4.967	7.209	9.113	2.673
Std. Err.	2.782	5.268	7.582	9.742	11.767	
Asym. P-value	0.439	0.266	0.256	0.230	0.219	0.750
Sim. P-value	0.561	0.392	0.383	0.353	0.340	0.778
R^2	0.032	0.855	1.364	2.347	3.313	
term yield spread ($\rho = 0.615$)						
β	2.694	4.046	6.233	9.681	11.707	5.966
Std. Err.	1.712	3.078	4.270	5.325	6.270	
Asym. P-value	0.058	0.094	0.072	0.035	0.031	0.310
Sim. P-value	0.076	0.111	0.085	0.041	0.035	0.329
R^2	3.300	3.394	5.669	11.175	14.432	

Table 4: Coefficient Estimates and R^2 s from Predictive Regressions

The table reports results from the regression of 1- to 5-year CRSP value-weighted returns on various lagged predictor variables (equation (1)) for the period 1926-2004 (75 observations). β is the estimated coefficient, with associated asymptotic standard error (equation (3)), p-value under the null hypothesis of no predictability, and the asymptotic Wald test and p-value for the joint hypothesis of no predictability across horizons. The table also reports simulated p-values (100,000 simulations) for both the individual coefficients and the Wald test. There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using parameter estimated from the data.

	Horizon					
	1	2	3	4	5	Wald
	log book-to-market ratio ($\rho = 0.939$)					
β	0.075	0.174	0.272	0.341	0.366	12.050
Std. Err.	0.050	0.099	0.147	0.193	0.239	
Asym. P-value	0.067	0.039	0.032	0.039	0.063	0.034
Sim. P-value	0.265	0.199	0.167	0.175	0.221	0.090
R^2	2.997	7.287	12.568	16.072	16.366	

Table 4 Cont'd

		Horizon			
		2	3	4	5
ln (D/P) (S&P500)	β_i/β_1	2.012	3.031	3.618	4.005
	R^2_i/R^2_1	1.858	3.015	3.443	3.733
ln (D/P) (CRSP VW)	β_i/β_1	2.005	2.965	3.471	3.858
	R^2_i/R^2_1	1.845	2.884	3.169	3.464
ln (D/P) (S&P500)	β_i/β_1	2.896	4.447	5.542	6.207
	R^2_i/R^2_1	3.826	6.346	8.054	8.924
ln (E/P) (S&P500)	β_i/β_1	2.214	3.205	3.768	3.728
	R^2_i/R^2_1	2.236	3.297	3.723	3.219
Default spread	β_i/β_1	7.686	11.572	16.796	21.232
	R^2_i/R^2_1	26.939	42.980	73.974	104.411
Term spread	β_i/β_1	1.502	2.314	3.594	4.345
	R^2_i/R^2_1	1.028	1.718	3.386	4.373
ln (B/M)	β_i/β_1	2.309	3.615	4.522	4.855
	R^2_i/R^2_1	2.431	4.193	5.362	5.460

Table 5: Coefficient Estimate and R²s Ratios from Predictive Regressions

The table reports the coefficient estimate ratios (β_i/β_1 , $i = 2, \dots, 5$) and the R² ratios (R^2_i/R^2_1 , $i = 2, \dots, 5$) from the regression of 1- to 5-year CRSP value-weighted returns on various lagged predictor variables (equation (1)) for the period 1926-2004 (75 observations).

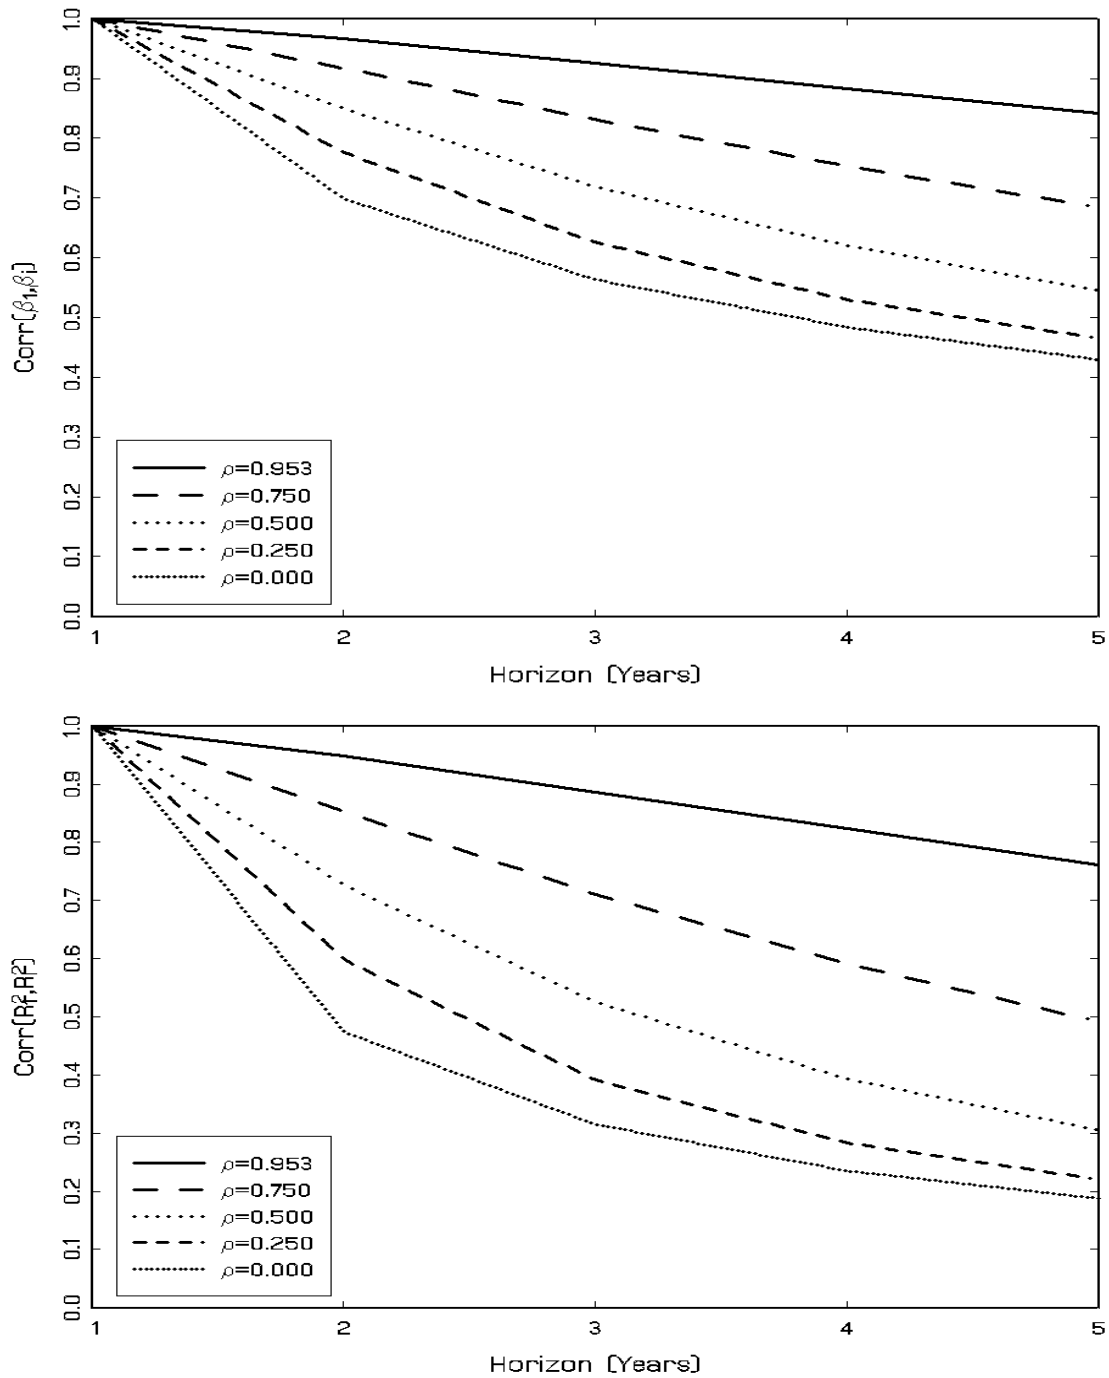


Figure 1: Cross-Horizon Correlations between Coefficient Estimates and R^2 s

The top panel plots the correlation between the coefficient estimate at the 1-year horizon and those at the 2- to 5-year horizons from the predictive regression (equation (1)) across 100,000 simulations for different values of ρ (the autocorrelation of the predictor variable). There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\sigma_\varepsilon = 0.196$, $\sigma_u = 0.166$, $\sigma_{\varepsilon u} = -0.683$. The bottom panel plots the analogous correlations for the predictive regression R^2 s.

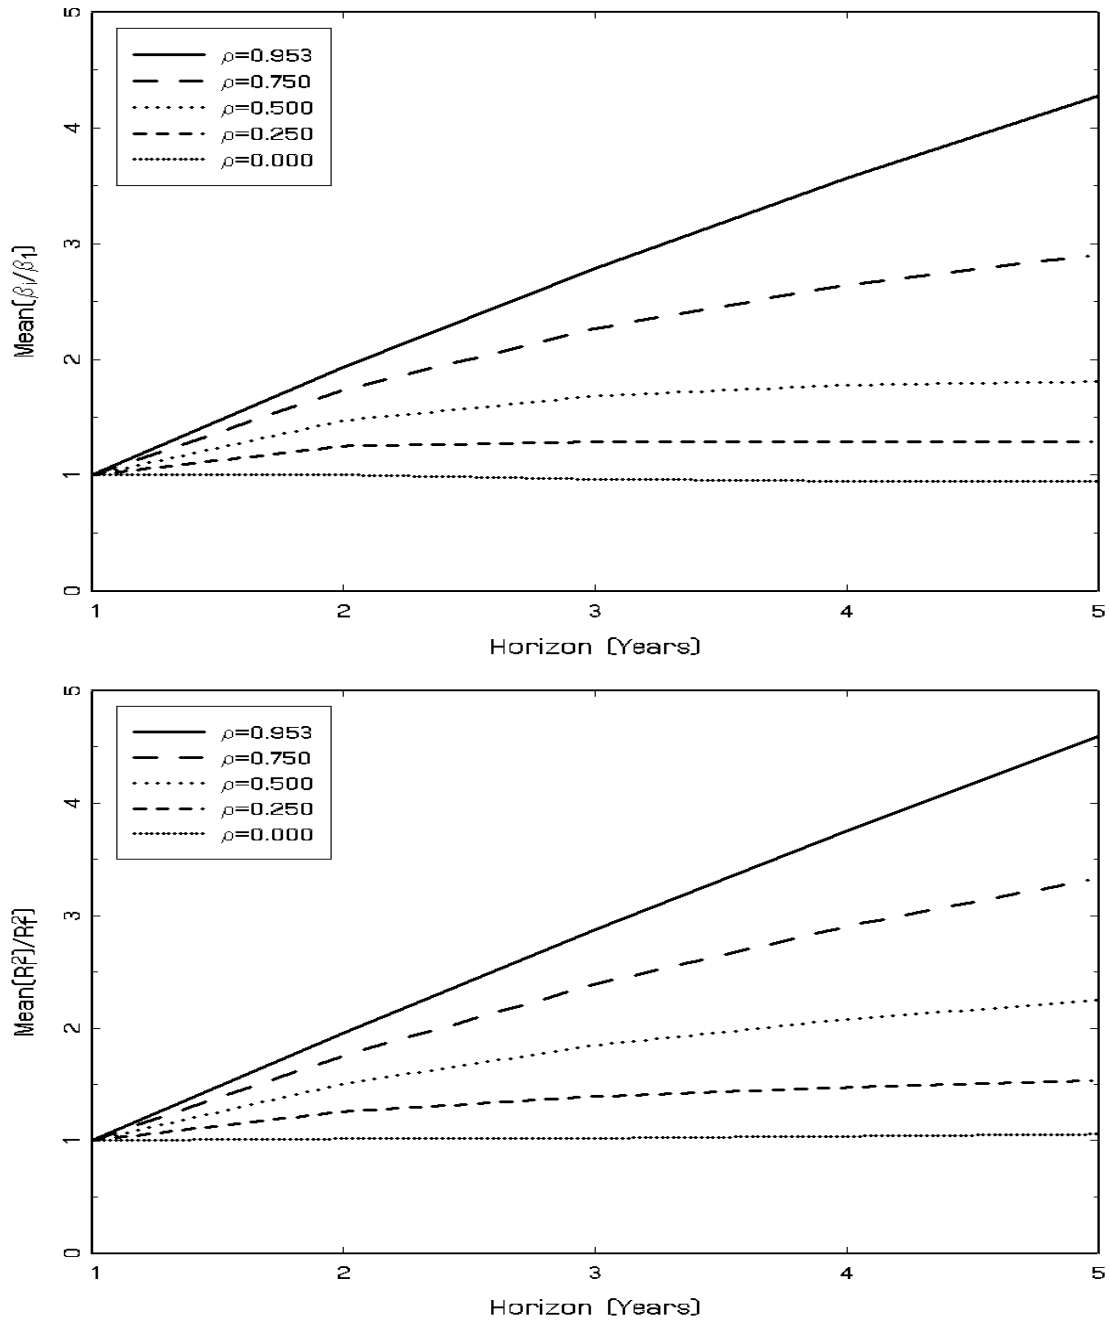


Figure 2: Mean Coefficient Estimate and R^2 Cross-Horizon Ratios

The top panel plots the mean coefficient estimate ratios (i.e., β_i/β_1 , $i = 2, \dots, 5$) from the predictive regression (equation (1)) across the simulations out of the 100,000 for which $|\beta_1| > 0.01$ for different values of ρ (the autocorrelation of the predictor variable). There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\sigma_\varepsilon = 0.196$, $\sigma_u = 0.166$, $\sigma_{\varepsilon u} = -0.683$. The bottom panel plots the means of the analogous R^2 ratios (i.e., R_i^2/R_1^2 , $i = 2, \dots, 5$) for simulations with or $R_1^2 > 0.5\%$.

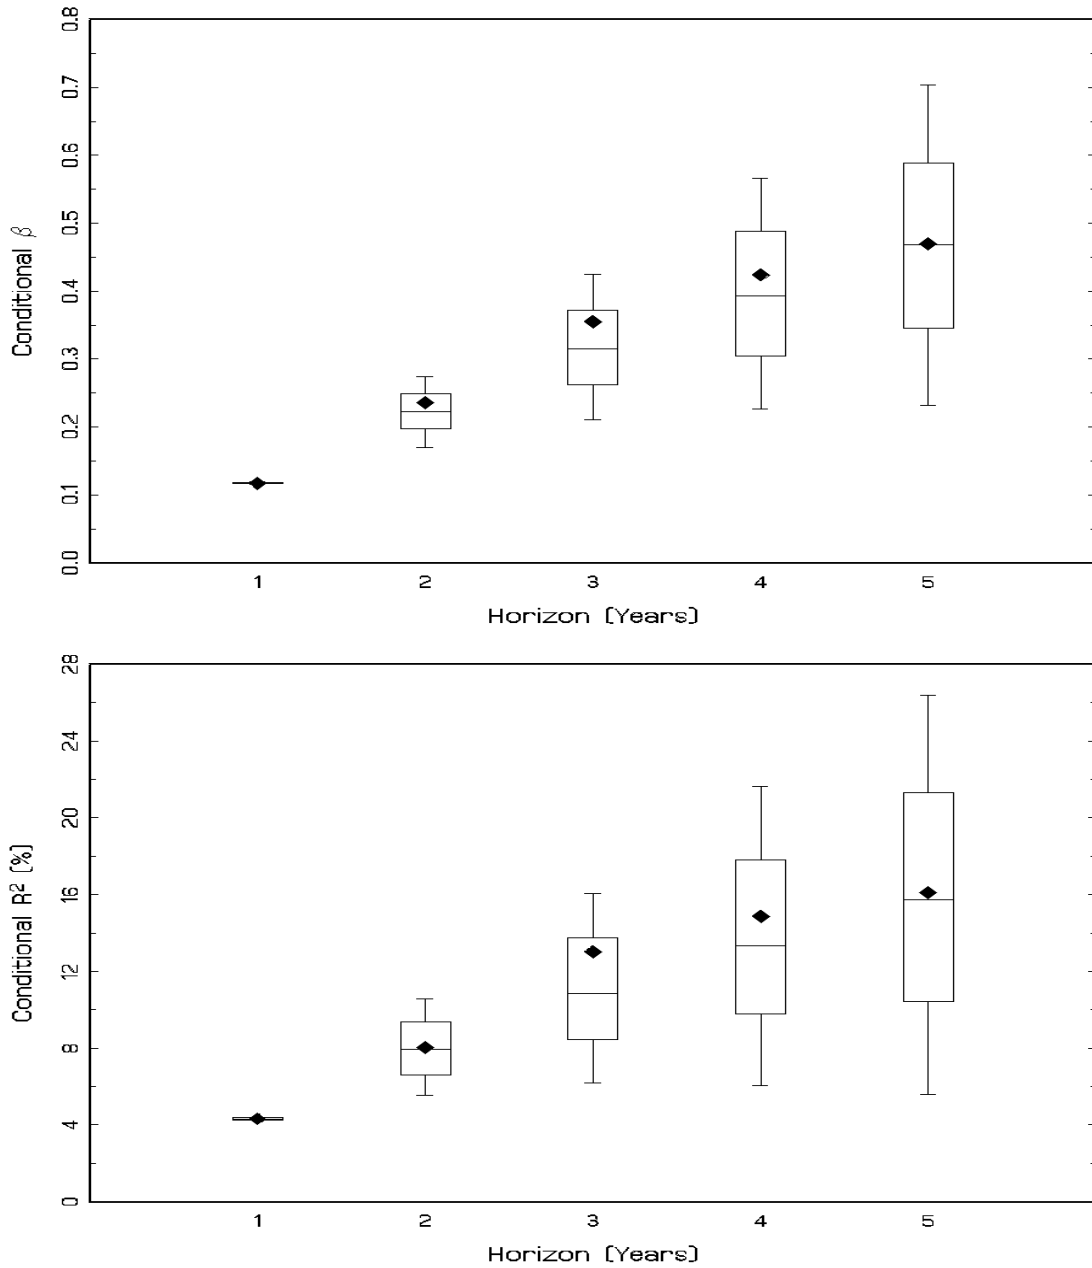


Figure 3: Conditional Distribution of Coefficient Estimates and R^2 s

The top panel provides a box plot of the simulated distributions of the coefficient estimates for horizons 2- to 5-years from the predictive regression (equation (1)) for the 971 out of 100,000 simulations for which $0.115 < \beta_1 < 0.119$. The boxes show the median, 25th/75th percentiles, and 10th/90th percentiles. The diamonds mark the actual coefficient estimates from the first regression in Table 4. There are 75 observations for each simulation, and simulations are performed under the null hypothesis of no predictability using the parameters $\rho = 0.953$, $\sigma_\varepsilon = 0.196$, $\sigma_u = 0.166$, $\sigma_{\varepsilon u} = -0.683$. The bottom panel plots the analogous simulated distributions for the predictive regression R^2 s for the 899 simulations for which $4.215\% < R^2_1 < 4.414\%$ and the corresponding actual R^2 s.