

Weak and Semi-strong Form Stock Return Predictability Revisited

WAYNE E. FERSON · ANDREA HEUSON · TIE SU

Boston College 140 Commonwealth Avenue, Chestnut Hill, MA. 02467
University of Miami 5250 University Drive, Coral Gables, FL. 33124
University of Miami 5250 University Drive, Coral Gables, FL. 33124

This revision: December 14, 2004

This paper makes indirect inference about the time-variation in expected stock returns by comparing unconditional sample variances to estimates of expected conditional variances. The evidence reveals more predictability as more information is used, and no evidence that predictability has diminished in recent years. Semi-strong form evidence suggests that time-variation in expected returns remains economically important.

© 1998-2004 by Wayne E. Ferson, Andrea Heuson and Tie Su. Ferson is a Research Associate of the National Bureau of Economic Research and the Collins Chair in Finance at Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA. 02467. ph (617) 552-6431, fax: 552-0431, email: wayne.ferson@bc.edu, <http://www2.bc.edu/~fersonwa>. Heuson and Su are Associate Professors of Finance at the University of Miami, 5250 University Drive, Coral Gables, FL. 33124. Heuson may be reached at (305) 284-1866, fax 284-4800, aheuson@miami.edu. Su may be reached at (305) 284-1885, fax 284-4800, tie@miami.edu.

1. Introduction

The empirical evidence for predictability in common stock returns remains ambiguous, even after many years of research. This paper makes indirect inference about the time-variation in expected stock returns by comparing unconditional sample variances of returns to estimates of expected conditional variances. The key to our approach is a sum-of-squares decomposition:

$$\text{Var}\{R\} = E\{\text{Var}(R|\Omega)\} + \text{Var}\{E(R|\Omega)\}, \quad (1)$$

where R is the rate of return of a stock and Ω is the public information set. $E(\cdot | \Omega)$ and $\text{Var}(\cdot | \Omega)$ are the conditional mean and variance and $\text{Var}\{\cdot\}$ and $E\{\cdot\}$, without the conditioning notation, are the unconditional moments. We are interested in the term $\text{Var}\{E(R | \Omega)\}$; that is, the amount of variation through time in conditionally expected stock returns. We infer this quantity by subtracting estimates of the expected conditional variance from estimates of the unconditional variance. We focus on the predictability in monthly stock returns. This is motivated by the empirical literature on asset pricing, which most commonly studies monthly returns.

We use two approaches to estimate the average conditional variances. These correspond to the classical description of increasing market information sets described by Fama (1970). Weak-form information considers only the information contained in past stock prices. This analysis, summarized in Table 1, builds on a comparison of daily and monthly sample variances, and is related to the variance ratios studied by Lo and MacKinlay (1988) and others. For this model we approximate the dynamics of expected returns with a step functions that changes once per month. Semi-strong form information relates to other lagged

variables that are clearly publicly available. Our analysis uses regressions for individual stock returns, on lagged firm-specific characteristics. Our approach is unique in that it relies on the covariation of the predictable components of individual stocks to draw inferences about index predictability. These results are reported in Table 2.

Studies of predictability in stock index returns typically report regressions with small R-squares, as the fraction of the variance in returns that can be predicted with lagged variables is small. The R-squares are larger for longer-horizon returns, because expected returns are considered to be more persistent than returns themselves.¹ However, because stock returns are very volatile, small R-squares can mask economically important variation in the expected returns. Stocks are long "duration" assets, so a small change in the expected return can lead to a large fluctuation in the asset value. To illustrate, consider the simple Gordon (1962) constant-growth model for a stock price: $P = kE/(r-g)$, where P is the stock price, E is the earnings per share, k is the dividend payout ratio, g is the future growth rate of earnings and r is the discount rate. The discount rate is the required or expected return of the stock. Consider an example where the price/earnings ratio, $P/E = 15$, the payout ratio, $k = 0.6$, and the expected growth rate, $g = 3\%$. The expected return is 7%. Suppose there is a shock to the expected return, ceteris paribus. In this example a change of one percent in r leads to approximately a 20% change in the asset value.

Of course, it is unrealistic to hold everything else fixed, but the example suggests that small changes in expected returns can produce large and economically significant changes in asset values. Consistent with this argument, studies such as Kandel and Stambaugh (1996), Campbell and Viceira (2001) and Fleming, Kirby and Ostdiek (2001) show that optimal

¹ Thus, the variance of the expected returns accumulates with longer horizons faster than the variance of returns, and the R-squares increase (see, e.g. Fama and French, 1988, 1989).

portfolio decisions can be affected to an economically significant degree by return predictability, even when the amount of predictability, as measured by R-squared, is small. Generalizing the Gordon model to allow for changes in growth rates, Campbell (1991) estimates that changes in expected returns through time may account for half of the variance of equity index values.

Our weak-form tests find no reliable evidence of predictability in modern data. Even so, a simulation study shows that the tests have the power to detect modest amounts of predictability. On the other hand, our semi-strong form tests find small but statistically and economically significant predictability. In contrast to recent studies that rely on aggregate predictor variables, we find no evidence that the predictability has diminished over time.

Section 2 discusses our approaches to measuring the variance of conditional expected stock returns. Section 3 presents the main results. Section 4 studies the power of our approach using simulations and the robustness of our empirical findings. Conclusions are offered in Section 5. Two Appendices discuss data, estimation issues and technical details.

2. Measuring Average Conditional Variances

2.1. Weak Form Information

Equation (1) says that we can estimate the variance of conditional expected stock returns by first estimating the unconditional variance and then subtracting an estimate of the average conditional variance. Thus, in order to use Equation (1) we need to estimate the average variance of the returns around the conditional mean, $E\{\text{Var}(R | \Omega)\} = E\{[R - E(R | \Omega)]^2\}$. The problem is that we don't know the conditional mean, $E(R | \Omega)$. Our approach in this section follows Merton (1980), who showed that while the mean of a stock return is hard to estimate, it is nearly irrelevant for estimating the conditional variance, when the time between

observations is short. We use high frequency returns to estimate the conditional variance, subtract its average from the monthly unconditional variance, and the difference -- according to Equation (1) -- is the variance of the monthly conditional mean.

Nelson (1990, 1992) develops Merton's idea. Ait-Sahalia and Mykland (2003) and Anderson, Bollerslev and Diebold (2004) discuss extensions. Suppose that the stock value can be approximated by a continuous process formed as a step function, with time intervals of length h between the steps. Take the interval $[T-h, T)$, chop it into D pieces, and consider the average of the D squared log price changes as an estimator for the conditional variance of the returns over the interval. Nelson proves the estimator is consistent, in the sense that it approaches the conditional variance in the "continuous record" limit, as h approaches zero and D becomes infinite. The intuition is that for small h , the conditional mean is effectively constant, so the sample variance approaches the conditional variance as D grows. By similar logic, Nelson (1992) shows that misspecification of the conditional mean function washes out as h gets small.

Evidence from Nelson (1991) supports the idea that for monthly stock returns, chopping the month into days should work well. He finds that daily returns measured with versus without dividends, or with versus without a simple adjustment for risk-free interest rates, produce virtually the same estimates of conditional variances. Similarly, Schwert (1990) finds that different dividend series have almost no effect on the measured daily variances of a long historical stock return series that we use in our analysis.

We estimate $E\{\text{Var}(R | \Omega)\}$ by the time series average of the daily return variances for each month. Using monthly returns data, we estimate the unconditional variance, $\text{Var}(R)$. Then, we infer the variance of the conditional expected returns by Equation (1). Let the return for month m be $R_m = \ln(V_m/V_{m-1}) = \sum_{j \in m} \Delta_j$, where V_m is the value of the stock at time

m and Δ_j is the daily log value change for day j . Assume that the conditional mean for month m is $\mu_m = E(\sum_j \Delta_j | j \in m)$, with $E(\Delta_j | j \in m) = \mu_m / D$, D being the number of days in the month. Thus, we approximate the dynamics of the conditional mean by a step function. The unconditional mean monthly return is $E(\mu_m) = \mu$, and we are interested in $\text{Var}(\mu_m)$, the variance of the monthly expected returns. Define the average daily variance, $\text{ADV} = E\{E[(\Delta_j - \mu_m / D)^2 | j \in m]\}$, and the unconditional monthly variance, $\text{MV} = E\{(R_m - \mu)^2\}$. Simple calculations show that $\text{Var}(\mu_m) = \text{MV} - D(\text{ADV})$.

2.2 Serial Dependence

The model uses the approximation that the means shift monthly, while daily returns fluctuate independently around the conditional means. However, there is weak serial dependence in daily stock returns. The question is whether or not to attribute this serial dependence to changes in the conditional expected return.

On the one hand, much of the literature on predictability allows that serial dependence may reflect changing conditional means. Fama and French (1988) use rate-of-return autoregressions to study predictability. Lo and MacKinlay (1988) and Conrad and Kaul (1988) model expected returns within the month as autoregressive processes. On the other hand, serial dependence in daily returns can arise from end-of-day price quotes that fluctuate between bid and ask (Roll, 1984) or from nonsynchronous trading of the stocks in an index. These effects should not be attributed to time-variation in the expected discount rate for stocks. We estimate $\text{Var}(\mu_m)$ with and without adjustments for serial dependence. To illustrate the adjustment, let $\theta = E\{E[(\Delta_j - \mu_m / D)(\Delta_{j'} - \mu_m / D) | j, j' \in m, |j - j'| = 1]\}$.

Assuming that the first order daily serial dependence reflects market microstructure effects unrelated to discount rates, we estimate $\text{Var}(\mu_m) = \text{MV} - D(\text{ADV} + 2\theta)$.

2.3 Finite Sample Biases

Biases in the finite sample variances and autocovariances arise due to estimation error in the sample means. To summarize, the objective is to estimate:

$$\text{Var}(\mu_m) = MV - D(ADV + 2\theta), \quad (2)$$

where

$$MV = E\{(R_m - \mu)^2\},$$

$$ADV = E\{E[(\Delta_j - \mu_m/D)^2 | j \in m]\}$$

and

$$\theta = E\{E[(\Delta_j - \mu_m/D)(\Delta_{j'} - \mu_m/D) | j, j' \in m, |j-j'| = 1]\}.$$

The finite sample variances, $\hat{M}V$ and $\hat{A}DV$, and the monthly average of the daily sample autocorrelations, have expectations that differ from the true values. Simple calculations show that

$$E(\hat{M}V) = MV - (D/M)(ADV + 2\theta), \quad (3)$$

$$E(\hat{A}DV) = [(D-1)/D] ADV - 2\theta/D,$$

$$E(\hat{\theta}) = \theta - (ADV + 2\theta)/D,$$

where M is the number of months in the sample. The system (3) provides three equations in three unknowns, and may be solved for unbiased estimators of MV , ADV and θ . These are given as:

$$\theta^* = \{D \hat{A}DV + D(D-1) \hat{\theta}\} / \{(D-1)(D-2) - 2\}, \quad (4)$$

$$ADV^* = \{2D \hat{\theta} + D(D-2) \hat{A}DV\} / \{(D-1)(D-2) - 2\},$$

$$MV^* = \hat{M}V + (D/M) ADV^* + (2D/M) \theta^*.$$

Using these estimators, we form the unbiased estimator of the variance of the monthly expected returns as:

$$\text{Var}(\mu_m)^* = MV^* - D(ADV^* + 2\theta^*). \quad (5)$$

In addition to the finite sample bias addressed above, there is a "finite record" bias, which arises because $h > 0$ and $D < \infty$. To address these biases we use Monte Carlo simulations. The calculations described above do not impose the requirement that variance estimates can't be negative. The Appendix B describes how this restriction is imposed. The restricted estimates involve numerical solution of a system of Kuhn-Tucker conditions. When a single parameter is involved, the solution sets the estimate to zero when the unrestricted variance estimate is negative. Section 4 further explores the power and robustness of the methods.

2.4. Semi-strong Form Information: Using Individual Stock Regressions

Much of the empirical literature on asset-return predictability uses regressions of stock-index returns on lagged, market-wide information variables. This approach raises two types of concerns. First, there are statistical problems associated with the regressions, especially when the data are heteroskedastic, the right-hand side variables are highly persistent or the left-hand side returns are overlapping in time.² The second issue is data mining. If the lagged variables arise from many researchers sifting through the same data sets, there is a risk of finding spurious predictability (Lo and MacKinlay 1990; Foster, Smith and Whaley 1997).

We use individual stocks to estimate the sum-of-squares decomposition in Equation

² Boudoukh and Richardson (1994) provide an overview of the statistical issues. Stambaugh (1999) and Ferson, Sarkissian and Simin (2003) provide more recent analyses and references.

(1), focusing on the aggregate predictability. Basic portfolio theory implies that individual-stock expected returns teach us about index predictability, only to the extent that they are correlated across the stocks. Consider the $N \times N$ covariance matrix of the conditional mean returns for N stocks, $\text{Cov}\{E(R | Z)\}$, where Z stands for the lagged, public information regressors. Letting $\underline{1}$ be an N -vector of ones, the variance of the conditional expected returns on an equally-weighted portfolio, R_p is:

$$\text{Var}\{E(R_p | Z)\} = (1/N^2) \underline{1}' \text{Cov}\{E(R | Z)\} \underline{1}. \quad (6)$$

Since there are $N(N-1)$ covariance terms, but only N variances in this expression, the expected return variance for the portfolio approaches the average of the firms' expected return covariances, while the individual stock predictability vanishes as N gets large.

To estimate the predictability of the index we model $\text{Cov}\{E(R | Z)\}$ from individual-stock regressions on lagged, firm-specific variables. We use Monte Carlo methods to handle the statistical issues, as described in Appendix B. There is some correlation between our firm-specific variables and the instruments selected in previous studies of aggregate predictability, so we are not completely immune to data mining bias. However, our measure does not rely on the direct index predictability that so many previous studies have explored, and the number of studies that examine individual-stock return predictability with time-series regressions is still relatively small. Using only firm-specific instruments we probably understate the correlations among the expected returns. If we use market-wide instruments for individual stocks, we are likely to overstate the correlations. Comparing the two cases we estimate a range of plausible values.

3. Empirical Results

3.1. Results using Weak-form Information

Table 1 presents estimates of predictability based on the comparison of monthly and daily return variances. Appendix A describes the data. Panel A presents results for the Standard and Poor's index over different subsamples and Panel B summarizes the results for the individual common stocks of twenty six large firms. Three estimators of predictability are shown. The estimator denoted as $\sigma(\mu_m)$ includes no adjustment for daily serial dependence, while the estimator $\sigma(\mu_m)^*$ adjusts for any first order autocorrelation, taking the view that daily serial correlation reflects microstructure issues unrelated to changes in discount rates. The estimator $\sigma(\mu_m)^\#$ takes the view that -- for individual stocks -- the main microstructure effect is the bid-ask bounce, which produces negative autocorrelation. Thus, negative autocorrelations are removed each month from the predictability calculation while no adjustment is made for positive ones. The average autocorrelation parameter, θ , is positive for the Standard and Poor's 500 index, and negative for 16 of the 26 firms. For comparison, the second column contains the unconditional standard deviations of the monthly returns, expressed as annual percentages.³

We first discuss the estimates in columns 3-5, labeled GMM estimates. These are the three estimates with analytical adjustments for finite sample biases. The estimates of predictability for the stock index range from 1.9% to 5.4% using Schwert's (1990) data for the 1885-1962 period. Over the 1962-2001 period where the CRSP daily data are available, the estimated volatility of the expected returns is 2%. After adjusting for positive serial dependence, which may arise from nonsynchronous trading of the stocks in the index, the

³ The monthly variance is multiplied by 12, then the square root of this result is multiplied by 100. All of the numbers in the tables are annualized this way.

estimator $\sigma(\mu_m)^*$ delivers a value of zero. Over the most recent 120 months of the sample all of the estimates of predictability are zero.

We estimate the finite sample biases in the predictability estimates using simulations. The simulations also provide empirical p-values for assessing the statistical significance of the results.⁴ The adjusted estimates are shown in columns 6-8 of Table 1. They tell a similar story. The estimates for the index range from zero to 5.1% for the 1885-1962 period, and the larger figures appear statistically significant. However, using the CRSP data for 1962-2001, the estimates are 1.44% or less, and none are statistically significant. Over the last 120 months the adjusted estimates are all equal to zero.

Panel B summarizes results for a sample of 26 large firms' individual common stocks. These cover the post-1962 period where the CRSP daily data are available. Averaged across the stocks, the predictability estimates range from zero to 2.2%, depending on the choice of estimator and sample period. The empirical p-values range from 0.23 to 0.35, thus providing no evidence of predictability. Even the extreme cases present no reliable evidence of predictability. Taking the stock with the maximum value of $\sigma(\mu_m)$, its empirical p-value is 0.07 or 0.08, depending on the subperiod. Accounting for the fact that this is the maximum of 26 cases, the results are insignificant.⁵

⁴ We resample from the actual data for a given stock or index, randomly with replacement. For each simulation trial we generate an artificial time series with the same number of daily observations as the original data series. The artificial data satisfy the null hypothesis that the expected return is constant. We compute the estimators on the artificial sample in exactly the same way as on the original samples. We repeat this for 1,000 trials. The average across the trials is the expected finite sample bias. We use the distribution of the simulated estimates to generate empirical p-values. These are the fraction of the simulations where the variance estimates are larger than the sample values. A small p-value means that the sample estimate is unlikely to occur if expected returns are constant.

⁵ Let u_i be the empirical p-value for the i -th experiment, $i=1,\dots,n$. Under the null hypothesis of no predictability the empirical p-value is uniformly distributed on $[0,1]$. Assuming independent experiments,

In summary, while the older historical data suggests economically significant predictability in the market index, there is little evidence of weak-form predictability in modern data. In particular, the evidence for the most recent ten years suggests that any weak-form predictability in the index has vanished.⁶

3.2. Semi-Strong Form Tests

Table 2 presents our estimates of predictability for an equally-weighted portfolio, based on the covariances of individual stock regressions on lagged variables. The regressions use monthly data from 1969 through 2001 and twenty large firms' common stocks. The data are described in Appendix A. The three rows for each sample period contain sample values of the implied predictability using either all the elements of the covariance matrix of the fitted expected returns, the off-diagonal elements only, or the diagonal elements only. The different cases are averaged according to the number of covariance elements involved. The "all elements" variances are the sum of the elements divided by N^2 . The "off-diagonal" variances are the sum of the off-diagonal terms divided by $N(N-1)$. The "Diagonal only" variances are the mean of the individual regression fitted variances. Thus, the all-elements

$$\begin{aligned} \text{Prob}\{\text{Min}_i u_i \leq q\} &= 1 - \text{Prob}\{u_i > q; i=1, \dots, n\} \\ &= 1 - \prod_i \text{Prob}\{u_i > q\} \\ &= 1 - (1-q)^n. \end{aligned}$$

For $q=0.07$ and $n=26$, the probability of finding the minimum p-value to be 0.07 or smaller in 26 trials is 85%. Of course, the trials are not independent, so the correct probability is somewhere between 7% and 85%.

⁶ Lo and MacKinlay (1999) present weak form tests with less evidence for predictability in more recent data, and suggest that the demise of such predictability may be related to "statistical arbitrage" trading by Wall Street firms. Nelson and Kim (1993) also find that weak-form evidence for stock index predictability is thin in post World War II data.

variance is a weighted average of the off-diagonal and diagonal-only variances, weighted according to the number of elements involved. The examples with diagonals only are unrealistic because they ignore the covariance across stocks, but they isolated the individual-stock predictability. That is, they show how the average stock would behave if it were treated like the index.

The estimates excluding the diagonals provide information on what would be expected to happen as the number of similar stocks in the portfolio becomes large. In this case using firm-specific regressors, the first column of figures shows that the implied predictability estimates fall in a narrow range, from 2.37% to 2.67%, depending on the subperiod.

The regressions behind Table 2 are subject to statistical biases, which we control via simulation as discussed in Appendix B. The bias-adjusted estimates are summarized in the second column of figures. Using only the off-diagonal terms, the average values are 1.63% to 1.84% annualized. Using the full covariance the estimates are 2.41% to 2.59%. These are statistically significant according to the empirical p-values.

The measures of predictability based on the diagonals only are summarized in the third row. The average diagonal elements are numerically larger than the average off-diagonal terms, but not significantly different from zero. Their right-tail p-values are larger than 20% in each experiment. This reflects the relatively large sampling variability of the regression estimates of expected returns, compared with their relatively small sampling covariability. This is one of the advantages of our approach, compared with previous studies that rely on direct regressions of stock indexes on lagged variables. Because the sampling covariability of the expected returns with firm-specific lagged variables is relatively small, we are able to estimate the predictability with relatively high precision. Thus, based on the

covariability, we can say that 1.8% to 2.6% is strongly statistically significant.

The estimates of predictability in Table 2 emphasizing the off diagonals, are similar whether we use the full sample or concentrate on subsamples of the last 120 months or the most recent 60 months.⁷ This is interesting in view of recent empirical studies that find index predictability, measured directly using lagged variables, has weakened in recent samples. It may be that the predictability was "real" when first publicized, but diminished as traders attempted to exploit it.⁸ Alternatively, the predictability may have been spurious in the first place, as a result of statistical biases and/or naive data mining. If the predictability is spurious we would expect lagged instruments to appear in the empirical literature, then fail to work with fresh data (e.g., Ferson, Sarkissian and Simin, 2003). But Table 2 presents no evidence that predictability is weaker in the recent subperiods.

As our firm-instrument-only measures do not rely on aggregate predictor variables, Table 2 provides new and interesting evidence that the underlying predictability has not diminished. This is consistent with the "efficient markets" view of predictability, as reflected in much of the conditional asset pricing literature (see reviews by Ferson (1985) and Cochrane, 2001). According to this view returns may be predictable if required expected returns vary over time in association with changing risk or risk aversion. If required expected returns vary over time there may be no abnormal trading profits and thus, no incentive to exploit the predictability. Predictability may therefore persist in an efficient market.

⁷ In additional experiments not reported in the table we include an additional lagged predictor, a measure of the firm's dividend yield. The results with this additional regressor (which is not available for all of the firms) are very similar to those in the table.

⁸ See Schwert (2003) for a review of this evidence and Goyal and Welch (2003) for a recent analysis.

The two right-hand columns of Table 2 show predictability estimates when each regression uses both firm-specific and market index ("Macro") characteristics as the lagged regressors, or alternatively, when only the Macro variables are used. These calculations show the effects of using the economy-wide lagged regressors. The estimated standard deviations of the monthly expected returns for the full sample period are between 4.2% and 8.6%, depending on whether or not we exclude the diagonals from the calculation. The economy-wide regressors increase the covariability of the expected return estimates. The "Macro only" examples in the far right column suggest that predictability is diminished somewhat in the more recent subperiods, which is consistent with the evidence cited earlier in studies that rely on economy-wide regressors.

Our semi-strong form estimates of predictability provide more reliable evidence of time-variation in monthly stock returns than our weak-form tests. This is expected if returns are more easily predicted using more information. But is 2-3% on an annual basis an economically significant effect? The simple Gordon model example from the introduction provides an illustration. Consider a month in which the required expected return jumps by roughly two standard deviations, say from 7% to 11%. Other things held fixed, the stock price would fall to half of its former value in response. Of course this overstates the effect, to the extent that a shock that changes the required return also changes the expected cash flows and future growth rates, but the example suggests the economic significance of predictability.

4. Power and Robustness

This section presents some results on the statistical power of our methods and on the robustness of our empirical findings. Concerns about power focus on the weak-form tests where we do not find significant predictability. Our approach contains several steps where

approximations are introduced or estimation is required. Each step is prone to some error, and the cumulative effect of the errors may result in low power. To evaluate the power we use simulations that include each of the steps. We also compare various estimators to isolate the impact of the different steps on power.

In the simulations we generate data following the model described in Section 2.1. The conditional means fluctuate each month as draws from a normal distribution, whose standard deviation controls the amount of variation in the expected returns. Daily returns fluctuate randomly around the monthly means, with variances chosen so that the first and second moments of the simulated returns match the sample returns. By setting the expected return variation to equal zero, we get critical values for a 5% test, defined such that 5% of the simulations produce statistics larger than the critical value. Setting the expected return variation to larger values, we trace out the power curves illustrated in Figures 1 and 2. The power of a particular test is the fraction of the simulation trials that produce a statistic larger than the critical value, given that an alternative hypothesis with time-variation in the expected returns generates the data. Figure 1 takes the variances of expected returns as the statistics. Figure 2 takes the empirical p-values of the expected return variances as the test statistics. The approach is the same, except that, because the empirical p-value is the result of a simulation, we now have to conduct simulations within the simulations.

Various estimators are displayed in figures 1 and 2. Auto is the sample autocorrelation statistic, and the others are the implied expected return estimators with various adjustments. We use the sample autocorrelation statistic as a "straw man" for comparison, because most of the statistics used in the literature on weak-form predictability are transformations of the sample autocorrelation (see, e.g. Cochrane, 1991). Of course, we expect the sample autocorrelation to perform poorly under our model of the return

dynamics, and the simulations bear this out. The power of the autocorrelation is nearly level at about 5%, independent of the amount of expected return variation. The estimator ABB refers to analytical and bootstrap bias adjustments, ABMB refers to analytical bias, microstructure bias and bootstrap adjustments. ABMNB refers to analytical bias, microstructure bias for negative autocorrelations only and bootstrap adjustments. Boot has only bootstrap bias adjustments and Unadj. refers to the estimators with no bias adjustments. The differences between the power of these estimators parse out the separate effects on power of the various steps in our estimation strategy. The figures show that the most complicated procedure, ABMB suffers lower power in some parts of the curve, but otherwise the various estimators give similar results.

Figures 1 and 2 suggest that our weak form tests have the power to detect modest amounts of predictability. For example, if the annual standard deviation of the expected returns is 3%, the power of the various statistics is 14% to 19%. Using 15% as the annual standard deviation of the monthly index returns, an expected return standard deviation of 3% means that a regression of the return on its expected return would produce an R-squared of only $(0.03/0.15)^2=4\%$. That is, the expected returns account for only 4% of the variance of the returns. With an annual standard deviation of 5%, the R-squared would be about 11%, which is similar to the values reported in some stock return regression studies. At this level the power of our tests is near 60%. At a standard deviation of 6% or more for the expected returns, the power of our tests is 80% or higher.

In tables 3 and 4 we explore the robustness of our empirical results. One of the questions to be further addressed relates to the treatment of the autocorrelations of daily returns. Panel A of Table 3 presents results for the Standard and Poor's index, replacing the stock index with index futures prices. Following Boudoukh et al. (1993) and Blume,

MacKinlay and Terker (1989), there should be no issues with microstructure-related autocorrelation in the index futures returns. The daily index futures are available from June of 1982 through September of 2002. The table shows that our main results are robust to the use of futures over this period, and over the more recent ten-year subsample. All of the point estimates of weak-form predictability are zero.

We use large firms' common stocks in the main experiments of tables 1 and 2, because such stocks should be representative of the market. However, small firms' stocks may have more serial correlation related to market microstructure effects. Small-firm stocks may also display more predictability related to market inefficiencies, as it is more costly to trade the stocks of smaller companies to exploit inefficiencies.

Panel B of Table 3 repeats the tests of Table 1, using 26 small-firm stocks in place of the large-firm stocks. The returns data are available from January of 1980 through September of 2002. The panel shows that the average weak-form predictability estimates and the point estimates for some of the small firms are somewhat higher than we found for the large firm stocks. The average bias-adjusted estimates are 0.25% to 1.35% over the full sample, and 1.25% to 1.32% for the most recent decade, depending on the estimator. The standard deviations of the monthly returns are also markedly higher for smaller firms. Thus, the average predictability estimate of 1.35% corresponds to only $(.0135 / .305)^2 = 0.2\%$ of the variance explained by the expected return. Given this low signal-to-noise ratio, the average right-tail p-values are all above 25%. The maximum values across the 26 firms seem more impressive; for example, bias-adjusted predictability values of 16-17% are found for one firm in the most recent decade. However, accounting for the multiple comparisons, the chances of finding numbers this large are as high as 7.5% to 18.8%.⁹

⁹ Following footnote 5, $1-(1-0.008)^{26}=18.8\%$ and $1-(1-0.003)^{26}=7.5\%$.

Table 4 presents the results of repeating the analysis of Table 2 using our sample of small firm stocks. For comparability with Table 2, the first twenty stocks are used. The available sample covers the 1981-2001 period. The results are very similar to those of Table 2. Using firm-specific lagged predictor variables the bias-adjusted estimates of semi-strong form predictability are between 1.78% and 3.62%, depending on how the diagonal elements of the expected return covariance matrix are treated. The diagonal elements on average, imply more predictability than the off-diagonal elements, as we saw before, but the covariability is somewhat lower among the small stocks. With the lower sampling covariability, the diagonal-only estimates of predictability become statistically significant. Again, there is no evidence that semi-strong form predictability has diminished in the most recent decade.

5. Conclusions

Small changes in expected returns can produce large and economically significant changes in asset values. This paper presents new estimates of time-variation in the expected returns of stocks, using indirect methods. Weak-form tests find no reliable evidence of predictability in modern data. Semi-strong form tests find small but economically significant predictability. In contrast to recent studies that rely on aggregate predictor variables, we find no evidence that the predictability has diminished in recent samples.

¹ The authors are grateful to Gurdip Bakshi, Hendrick Bessembinder, Charles Cao, John Cochrane, Pat Fishe, Bruce Grundy, Ravi Jagannathan, Herb Johnson, Avi Kamara, Terence Lim, Stewart Mayhew, Simon Pak, Mark Rubinstein, Robert Stambaugh, and William Ziemba for discussions and comments. The advice of the Editor, David Hsieh, and an anonymous

referee were especially helpful. Su acknowledges financial support from the Research Council at the University of Miami.

Appendix A: Semi-strong Variables

Our semi-strong form tests use data on lagged, firm-specific instruments and economy-wide, macro instruments for the index. The firm-specific instruments are obtained from CRSP and COMPUSTAT. They include for each stock month, (1) the average return over the previous twelve months; (2), the book-to-market ratio, defined as the most recently-available book value per share (quarterly Compustat data item #60 divided by item #61) divided by the one-month-lagged price per share; and (3) the earnings-to-price ratio, defined as the trailing four quarters' earnings (Compustat data item #11) divided by the one-month lagged stock price. The stock prices and returns are from CRSP.

The macro instruments include: (1) the lagged three-month Treasury Bill secondary market yield; (2) the lagged, one-month holding period return on a three-month Treasury Bill; (3) the spread between the Treasury Bill yield and the ten-year constant maturity Treasury Bond yield to maturity; and (4) the spread between Moody's Seasoned AAA and BAA corporate bond yields. All of the yield series are measured as the lagged monthly average of daily values, from the Federal Reserve. In addition, we use the one-month and twelve-month lagged holding period return on the Standard and Poor's Index, excluding dividends, from CRSP. We also use the dividend yield on the SPX index. The dividend yield is computed as the trailing 12-month dividends divided by the SPX index level. Monthly dividends are obtained from Bloomberg.

Appendix B: Estimation Issues

B.1 Imposing Positivity

Boudoukh, Richardson and Smith (1993) describe estimators for risk premiums, imposing the restriction that the expected risk premium is positive. Their tests involve the minimization of a quadratic form in the restricted and unrestricted estimates, which we adapt to the current setting as follows. Let $\hat{\phi}$ be an N-vector of unrestricted estimates of $\text{Var}(\mu_m)$ for N assets, whose sample values may be negative. Consider the estimator

$$\phi^* = \text{Arg Min}_{\phi} (\phi - \hat{\phi})' \text{Cov}(\hat{\phi})^{-1} (\phi - \hat{\phi}) \text{ subject to } \phi \geq 0, \quad (\text{B.1})$$

where the inequality in the constraint applies element-by-element. The Kuhn-Tucker conditions for this problem specify ϕ^* and an N-vector of multipliers $U = \{u_i\}_i$, which satisfy

$$\begin{aligned} (\phi^* - \hat{\phi})' \text{Cov}(\hat{\phi})^{-1} - U &= 0, \\ \phi_i^* u_i &= 0, i=1, \dots, N, \\ \phi_i^* &\geq 0, i=1, \dots, N, \text{ and} \\ u_i &\geq 0, i=1, \dots, N. \end{aligned} \quad (\text{B.2})$$

We solve the system (B.2) numerically to obtain the restricted estimators, ϕ^* . Note that when $N=1$ this amounts to setting $\phi_i^* = 0$ whenever the unrestricted value is negative.

B.2 Semi-strong Form Estimates

In our semi-strong form analysis, we use simulations based on a parametric form of the bootstrap to control statistical biases in the predictive regressions. The following regression

system is estimated for each asset:

$$\begin{aligned} R_{t+1} &= \mu + \epsilon_{t+1} \\ Z_{t+1} &= \mu_z + A (Z_t - \mu_z) + v_{t+1}, \end{aligned} \tag{B.3}$$

where A is an $L \times L$ matrix and there are L instruments in Z . The sample means estimate μ and μ_z , and the OLS coefficient estimates A . These estimates are taken as parameters of the simulation. Then, we resample from the vector of residuals $\{\epsilon_{t+1}, v_{t+1}\}$, randomly with replacement, and use these as the shocks in the simulation. We build up the time series of Z_t recursively in each simulated sample, along with the contemporaneous returns, which satisfy the null hypothesis of constant expected returns, μ . This approach preserves the first order autocorrelation of the instruments and accommodates the finite sample bias discussed by Stambaugh (1999), which arises when ϵ_{t+1} and v_{t+1} are correlated. We regress R_{t+1} on Z_t , using the simulated data samples, for each of the 1,000 simulation trials, and estimate the predictability exactly as in the original data. Since the covariance of the expected returns is zero, the true predictability is zero in the simulations. The average measured predictability, taken across the simulation trials, is our estimate of the bias, and the fraction of the simulations in which the actual data estimate exceeds the simulated value, is the empirical p-value.

References

Ait-Sahalia, Yacine, and Per A. Mykland, "The Effects of Random and Discrete sampling when estimating continuous-time diffusions," *Econometrica* 71 (2003) No. 2, 483-549.

Anderson, T.G., Tim Bollerslev and Francis X. Diebold, "Parametric and Nonparametric Volatility Measurement," forthcoming in L.P. Hansen and U. Ait-Sahalia (eds.) *Handbook of Financial Econometrics* (2004), Amsterdam, North Holland.

Boudoukh, Jacob, and Matthew Richardson, "The Statistics of Long-horizon Regressions," *Mathematical Finance* 4 (1994), 103-120.

Boudoukh, Jacob, Matthew Richardson and Tom Smith, "Is the Ex ante Risk Premium Always Positive? A New Approach to Testing Conditional Asset Pricing," *Journal of Financial Economics* 34 (1993), 387-408.

Blume, Marshall, A.C. MacKinlay and Bruce Terker, "Order Imbalances and Stock Price Movements on October 19 and 20, 1987," *Journal of Finance* 44 (1989), 827-848.

Campbell, John Y., "A Variance Decomposition for Stock Returns," *Economic Journal*, 101 (1991), 157-179.

Campbell, John Y. and Luis M. Viceira. "Who Should Buy Long-Term Bonds?," *American Economic Review*, 91 (2001), 99-127.

Cochrane, John H., "Asset Pricing," Princeton, N.J: Princeton University Press (2001). Pp. 530. ISBN 0-691-07498-4.

Cochrane, John H. "Volatility Tests And Efficient Markets: A Review Essay," *Journal of Monetary Economics*, 27 (1991), 463-486.

Conrad, Jennifer and Gautam Kaul. "Time-Variation In Expected Returns," *Journal of Business*, 61 (1988), 409-426.

Fama, Eugene F., "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25 (1970), 383-417.

Fama, E. and K. French, "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy* 96 (1988), 246-273.

Fama, Eugene F., and Kenneth R. French, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics* 25 (1989), 23-49.

Ferson, Wayne E., "Theory and Empirical Testing of Asset Pricing Models," Chapter 5 in *Finance, Handbooks in Operations Research and Management Science*, by Jarrow, Maksimovic and Ziemba (editors), Elsevier, (1995), 145-200.

FERSON, HEUSON AND SU

Weak and Semi-strong Form Stock Return Predictability Revisited

Ferson, Wayne E., Sergei Sarkissian and Timothy Simin, "Spurious regressions in Financial Economics?" *Journal of Finance* 58 (2003), 1393-1414.

Fleming, Jeff, Chris Kirby and Barbara Ost diek, "The Economic Value of Volatility Timing," *Journal of Finance* 61 (2001), 329-352.

Foster, F. Douglas, Tom Smith and Robert E. Whaley, "Assessing Goodness-of-fit of Asset Pricing Models: The Distribution of the Maximal R-squared," *Journal of Finance* 52 (1997), 591-607.

Gordon, Myron 1962, "The Investment, Financing and Valuation of the Firm," Irwin, Homewood, IL.

Goyal, Amit, and Ivo Welch, "Predicting the Equity Premium with Dividend Ratios," *Management Science* 49 (2003), No. 5 (March).

Hansen, Lars P., "Large Sample Properties of the Generalized Method of Moments Estimators," *Econometrica* 50 (1982), 1029-1054.

Kandel, Shmuel A., and Robert F. Stambaugh, " On the Predictability of Stock Returns: An Asset Allocation Perspective," *Journal of Finance* 51 (1996), 385-424.

Lo, Andrew and A. C. MacKinlay, "Stock Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies* 1 (1988).

Lo, Andrew and A. C. MacKinlay, "Data Snooping Biases in Tests of Financial Models," *Review of Financial Studies* 3 (1990), 431-468.

Lo, Andrew and A. C. MacKinlay, "A Nonrandom Walk Down Wall Street," Princeton University Press, Princeton, NJ (1999).

Lo, Andrew and A. C. MacKinlay, "The Econometrics of Financial Markets," Princeton University Press, Princeton, NJ (1997).

Merton, Robert, "On Estimating the Expected Return of the Market," *Journal of Financial Economics* 8 (1980), 323-361.

Nelson, Daniel B., "ARCH Models as Diffusion Approximations," *Journal of Econometrics* 45 (1990), 7-39.

FERSON, HEUSON AND SU

Weak and Semi-strong Form Stock Return Predictability Revisited

Nelson, Daniel B., "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica* 59 (1991), 347-370.

Nelson, Daniel B., "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," *Journal of Econometrics* 52 (1992), 61-90.

Nelson, C. and Myung J. Kim, "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance* 48 (1993), 641-661.

Roll, Richard., "A Simple Implicit Measure Of The Effective Bid-Ask Spread In An Efficient Market," *Journal of Finance*, 39 (1984), 1127-1139.

Schwert, G. W., "Indexes of US Common Stock Prices," *Journal of Business* 63 (1990), 399-426.

Schwert, G. W., 2003, in George M. Constantinides, Milton Harris and Rene M. Stulz, Editors: *Handbook of the Economics of Finance*, Elsevier Science Publishers, North Holland, pp. 743-800 ISBN: 0-444-5136-9.

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

Table 1
Return Predictability Based on Weak-Form Information

The predictability measures are the standard deviation of the expected returns, in percent per month, on an annualized basis.* The measures compare monthly unconditional return variances with average conditional variances, estimated from daily data. Panel A presents results for the Standard and Poor's stock index return over different subperiods. Panel B contains similar statistics for a sample of twenty six large firms' individual common stocks. For individual stocks the average statistics (Avg) are shown, along with the cases that produce the largest (Max) and smallest (Min) estimate of expected return variation. The columns present different estimators, and the monthly return standard deviations are included for comparison. The estimator $\sigma(\mu_m)$ is the annualized standard deviation of the monthly expected returns, in percent. The estimator $\sigma(\mu_m)^*$ is adjusted to remove the effects of all first order serial dependence in daily returns. The estimator $\sigma(\mu_m)^\#$ is adjusted to remove only the effects of negative first order serial dependence. The GMM estimates are found using the Generalized Method of Moments. The Finite Sample Adjusted estimates subtract a bootstrapped estimate of finite sample bias, under the null hypothesis of independent and identically distributed returns. Right-tail empirical p-values are on the second line of each case with bootstrapped finite sample adjustments.

Time Period	Monthly Return Standard Deviation	GMM Estimates:			Finite Sample Adjusted:			
		$\sigma(\mu_m)$	$\sigma(\mu_m)^*$	$\sigma(\mu_m)^\#$	$\sigma(\mu_m)$	$\sigma(\mu_m)^*$	$\sigma(\mu_m)^\#$	
Panel A: Standard and Poor's Stock Index								
1885-1962	17.8	5.40	1.86	5.40	5.09	1.44	0.00	
					0.00	0.11	1.00	
1962-2001	14.8	2.02	0.00	2.05	1.44	0.00	1.44	
					0.15	0.27	0.15	
1992-2001	13.5	0.00	0.00	0.00	0.00	0.00	0.00	
					0.35	0.37	0.37	
Panel B: Twenty-six Large-capitalization Common Stocks								
1962-2001	Avg	29.8	1.33	0.00	1.37	0.91	0.00	0.94
						0.23	0.30	0.24
	Max	37.9	7.83	0.00	7.87	6.24	0.00	6.23
						0.08	0.28	0.08
	Min	31.5	0.00	0.00	0.00	0.00	0.00	0.00
						0.29	0.35	0.30

table 1, continued...

Time Period	Monthly Return Standard Deviation	GMM Estimates: $\sigma(\mu_m) \sigma(\mu_m)^* \sigma(\mu_m)^\#$			Finite Sample Adjusted: $\sigma(\mu_m) \sigma(\mu_m)^* \sigma(\mu_m)^\#$			
Panel B: Twenty-six Large-capitalization Common Stocks, continued...								
1992-2001	Avg	32.3	0.89	2.20	1.15	0.66	0.97	0.68
						0.34	0.35	0.35
	Max	36.5	13.4	10.0	13.5	9.92	6.19	9.86
						0.07	0.18	0.07
	Min	29.6	0.00	0.00	0.00	0.00	0.00	0.00
						0.35	0.40	0.37

* Annualized figures are derived by multiplying the monthly decimal variance estimate of the expected returns by 12, taking the square root and multiplying the result by 100.

Table 2
Semi-strong Form Predictability

The implied predictability is the annualized standard deviation of the time-varying monthly expected returns, in percent, for an equally-weighted portfolio of 20 large common stocks. The figures are estimated from the covariance matrix of the individual stock regressions. The rows report alternative calculations where either all of the elements of the covariance matrix are used, or when the diagonal and off-diagonal elements are emphasized. The "All Elements" variance is a weighted average of the "off-diagonal" and the "Diagonal only" variance, weighted by the number of covariances involved. The columns report experiments where different predictor variables are in the regressions: Either firm-specific variables only, firm-specific and Macro variables, or Macro variables only. The Bias-adjusted estimates subtract a bootstrapped estimate of finite sample bias, under the null hypothesis of constant expected returns. Empirical right-tail p-values for this estimate are denoted as p-value.

Covariances Used	Firm Variables Only		Firm and Macro Variables	Macro Variables Only
	Implied Bias-Adjusted Predictability	P-value		

Panel A: October, 1969 - December 2001

All elements	3.18	2.41	0.000	6.67	5.16
Off-diagonal	2.37	1.82	0.000	5.93	4.88
Diagonal only	9.81	3.74	0.213	14.87	9.00

Panel B: January, 1992 - December, 2001

All elements	3.22	2.44	0.000	6.69	4.49
Off-diagonal	2.38	1.84	0.000	5.94	4.24
Diagonal only	10.02	3.78	0.222	15.01	7.89

Panel C: January, 1996 - December, 2001

All elements	3.85	2.59	0.001	8.60	4.76
Off-diagonal	2.67	1.63	0.050	7.64	4.44
Diagonal only	12.69	3.71	0.307	19.18	8.87

Table 3
Robustness of Weak-form Predictability Results

The predictability measures are the same as in Table 1. Panel A presents results for Standard and Poor's stock index futures returns over different subperiods. Panel B contains similar statistics for a sample of 26 small firms' individual common stocks.

Time Period	Monthly Return Standard Deviation	GMM Estimates: $\sigma(\mu_m) \sigma(\mu_m)^* \sigma(\mu_m)^\#$			Finite Sample Adjusted: $\sigma(\mu_m) \sigma(\mu_m)^* \sigma(\mu_m)^\#$			
Panel A: Stock Index Futures								
1986-2002	16.2	0.00	0.00	0.00	0.00	0.00	0.00	
					0.50	0.86	0.14	
1993-2002	16.2	0.00	0.00	0.00	0.00	0.00	0.00	
					0.45	0.77	0.20	
Panel B: Twenty-six Small-capitalization Common Stocks								
1980-2002	Avg	30.5	1.67	0.49	1.70	1.34	0.25	1.35
						0.25	0.31	0.25
	Max	31.8	9.58	0.00	9.82	7.77	0.00	7.94
						0.04	0.32	0.38
1993-2002	Avg	29.1	1.81	1.91	1.91	1.25	1.23	1.32
						0.31	0.33	0.32
	Max	40.2	20.5	20.0	20.8	16.9	16.1	17.0
						.004	.008	.003

Table 4
Semi-strong Form Predictability of Small-capitalization Stocks

The implied predictability is the annualized standard deviation of the time-varying monthly expected returns, in percent, for an equally-weighted portfolio of 20 small-firm common stocks. The methodology and symbology is otherwise the same as in Table 2.

Covariances Used	Firm Variables Only			Firm and Macro Variables	Macro Only
	Implied Predictability	Bias-Adjusted	P-value		
1981 - 2001					
All elements	2.25	2.04	0.000	4.82	4.16
Off-diagonal	1.90	1.78	0.000	4.55	4.07
Diagonal only	5.74	3.58	0.000	8.45	5.61
1992 - 2001					
All elements	2.39	2.04	0.000	4.75	4.16
Off-diagonal	1.81	1.65	0.000	4.25	3.98
Diagonal only	7.24	3.62	0.038	10.32	6.63

Fig. 1

**Power Curves:
Expected Return Variances**

Probability of Rejection, 5% Test




Fig. 2

**Power Curves
Empirical P-values**

Probability of Rejection, 5% Test

