Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?

Citation

Campbell, John Y. and Samuel B. Thompson. 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies* 21(4): 1509-1531.

Published Version

<http://dx.doi.org/10.1093/rfs/hhm055>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:2622619>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?

John Y. Campbell and Samuel B. Thompson*

* Campbell: Corresponding author. Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, USA, and NBER. Tel 617-496-6448. Email john_campbell@harvard.edu. Thompson: Arrowstreet Capital, LP, 44 Brattle Street, 5th floor, Cambridge MA 02138, USA. Email sthompson@arrowstreetcapital.com.

Abstract

Goyal and Welch (2006) argue that the historical average excess stock return forecasts future excess stock returns better than regressions of excess returns on predictor variables. In this paper, we show that many predictive regressions beat the historical average return, once weak restrictions are imposed on the signs of coefficients and return forecasts. The out-of-sample explanatory power is small, but nonetheless is economically meaningful for mean-variance investors. Even better results can be obtained by imposing the restrictions of steady-state valuation models, thereby removing the need to estimate the average from a short sample of volatile stock returns.

Towards the end of the last century, academic finance economists came to take seriously the view that aggregate stock returns are predictable. During the 1980s, a number of papers studied valuation ratios, such as the dividend-price ratio, earnings-price ratio, or smoothed earnings-price ratio. Value-oriented investors in the tradition of Graham and Dodd (1934) had always asserted that high valuation ratios are an indication of an undervalued stock market and should predict high subsequent returns, but these ideas did not carry much weight in the academic literature until authors such as Rozeff (1984), Fama and French (1988), and Campbell and Shiller (1988a,b) found that valuation ratios are positively correlated with subsequent returns and that the implied predictability of returns is substantial at longer horizons. Around the same time, several papers pointed out that yields on short- and long-term Treasury and corporate bonds are correlated with subsequent stock returns [Fama and Schwert (1977), Keim and Stambaugh (1986), Campbell (1987), and Fama and French (1989)].

During the 1990s and early 2000s, research continued on the prediction of stock returns from valuation ratios [Kothari and Shanken (1997) and Pontiff and Schall (1998)] and interest rates [Hodrick (1992)]. Several papers suggested new predictor variables exploiting information in corporate payout and financing activity [Lamont (1998) and Baker and Wurgler (2000)], the level of consumption in relation to wealth [Lettau and Ludvigson (2001)], and the relative valuations of high- and low-beta stocks [Polk, Thompson, and Vuolteenaho (2006)]. At the same time, several authors expressed concern that the apparent predictability of stock returns might be spurious. Many of the predictor variables in the literature are highly persistent: Nelson and Kim (1993) and Stambaugh (1999) pointed out that persistence leads to biased coefficients in predictive regressions if innovations in the predictor variable are correlated with returns (as is strongly the case for valuation ratios, although not for interest rates). Under the same conditions, the standard t -test for predictability has incorrect size [Cavanagh, Elliott, and Stock (1995)]. These problems are exacerbated if researchers are data mining, considering large numbers of variables, and reporting only those results that are apparently statistically significant [Foster, Smith, and Whaley (1997) and Ferson, Sarkissian, and Simin (2003)]. An active recent literature discusses alternative econometric methods for correcting the bias and conducting valid inference [Cavanagh, Elliott, and Stock (1995), Mark (1995), Kilian (1999), Lewellen (2004), Torous, Valkanov, and Yan (2004), Campbell and Yogo (2006), Jansson and Moreira (2006), Polk, Thompson, and Vuolteenaho (2006), and Ang and Bekaert (2007)].

A somewhat different critique emphasizes that predictive regressions have often performed poorly out-of-sample [Goyal and Welch (2003, 2006) and Butler, Grullon,

and Weston (2005)]. This critique had particular force during the bull market of the late 1990s, when low valuation ratios predicted extraordinarily low stock returns that did not materialize until the early 2000s [Campbell and Shiller (1998)]. Goyal and Welch (2006) argue that the poor out-of-sample performance of predictive regressions is a systemic problem, not confined to any one decade. They compare predictive regressions with historical average returns and find that historical average returns almost always generate superior return forecasts. They conclude that “the profession has yet to find some variable that has meaningful and robust empirical equity premium forecasting power.”

While it is not clear how much weight should be placed on out-of-sample statistics in judging the predictability of stock returns [Inoue and Kilian (2004)], in this paper we take up Goyal and Welch’s (2006) challenge and ask whether standard variables could have been used in real time to forecast twentieth and early twenty-first century stock returns. We show that simple restrictions on predictive regressions, suggested by investment theory, improve the out-of-sample performance of key forecasting variables and imply that investors could have profited by using market timing strategies.

We begin in Section 2 by comparing the in-sample and out-of-sample forecasting power of standard predictor variables through the end of 2005. We use at least 20 years of data to obtain initial coefficient estimates and restrict the forecast evaluation period to the period since 1927 when high-quality total returns data are available from CRSP. We calculate an out-of-sample R^2 statistic that can be compared with the usual in-sample R^2 statistic. Like Goyal and Welch (2006), we find unimpressive in-sample results and poor out-of-sample performance for many of the usual linear regressions.

Our contribution is to show that restricted regressions perform considerably better than these unrestricted regressions. We begin by considering two alternative restrictions that can be imposed on any theoretically motivated forecasting regression: first, that the regression coefficient has the theoretically expected sign; and second, that the fitted value of the equity premium is positive. We impose these restrictions sequentially and then together, and find that they substantially improve the out-of-sample performance of predictive regressions.

Next we look more closely at valuation ratios, for which theory restricts not only the sign but also the magnitude of the regression coefficients. We find that out-of-sample predictability can be improved even more by restricting the magnitude of the coefficients to the values implied by a simple steady-state model, such as the Gordon

(1962) growth model in the case of the dividend-price ratio.

Section 1 shows that several commonly used forecasting variables do have some ability to predict stock returns out-of-sample. The out-of-sample R^2 statistics are positive, but small. This raises the question of whether the predictive power is economically meaningful. In Section 2, we show that even very small R^2 statistics are relevant for investors because they can generate large improvements in portfolio performance. In a related exercise, we calculate the fees that investors would be willing to pay to exploit the information in each of our forecasting variables. Section 3 briefly concludes.

1 Theoretical Restrictions on Predictive Regressions

In this section, we conduct an out-of-sample forecasting exercise inspired by Goyal and Welch (2006), with modifications that reveal the effectiveness of theoretically motivated restrictions. We use monthly data and predict simple monthly or annual stock returns on the S&P 500 Index. This immediately creates a tradeoff between the length of the data sample and the quality of the available data. High-quality total return data are available monthly from CRSP since 1927, while total monthly returns before that time are constructed by interpolation of lower-frequency dividend payments and therefore may be suspect. Accordingly, we use the CRSP data period as our out-of-sample forecast evaluation period, but use earlier data to estimate an initial regression.

Table 1, whose format is based on the tables in Goyal and Welch (2006), reports the results. We begin by discussing monthly return forecasts reported in Panel A, and then discuss annual forecasts reported in Panel B. Each row of the table considers a different forecasting variable. The first four rows consider valuation ratios: the dividend price ratio, earnings price ratio, smoothed earnings price ratio, and book-to-market ratio. Each of these ratios is measured in levels, rather than logs, with some accounting measure of corporate value in the numerator, and market value in the denominator. The smoothed earnings price ratio, proposed by Campbell and Shiller (1988b, 1998) is the ratio of a 10-year moving average of real earnings to current real prices. Campbell and Shiller argue that this ratio should have better forecasting

power than the current earnings price ratio because aggregate corporate earnings display short-run cyclical noise; in particular, earnings drop close to zero in recession years, such as 1934 and 1992, creating spikes in the current earnings price ratio that have nothing to do with stock market valuation levels.¹

The next row of Table 1 uses a smoothed measure of accounting real return on equity (ROE) as the predictor variable. Real ROE is measured as the level of real earnings divided by lagged real book equity and the gross one-year inflation rate. This gives a measure of the total resources available to be divided between real payouts and real growth in book equity.² We smooth real ROE over ten years to remove short-run cyclical noise in corporate profitability. We include this variable in Table 1 not because it has been claimed to predict aggregate stock returns by itself, but because we will use it in combination with valuation ratios in Table 2.

The next five rows of Table 1 consider nominal interest rates and inflation: the short-term interest rate, the long-term bond yield, the term spread between long- and short-term Treasury yields, the default spread between corporate and Treasury bond yields, and the lagged rate of inflation.

The last two rows of Table 1 evaluate forecasting variables that have been proposed more recently: the equity share of new issues proposed by Baker and Wurgler (2000) and the consumption-wealth ratio of Lettau and Ludvigson (2001). Lettau and Ludvigson's variable is based on a cointegrating relationship between consumption, aggregate labor income, and aggregate financial wealth. Rather than estimate a separate cointegrating regression, we simply include the three variables directly in the forecasting equation for stock returns.³

The first column of Table 1 reports the first date for which we have data on each forecasting variable. For dividends, earnings, stock returns, long-term bond yields, and inflation we have data, originally assembled by Robert Shiller, back to the early 1870s. Other data series typically begin shortly after the end of World War I, although book equity does not become available until 1926, which means that the book-to-market ratio begins then and the ten-year smoothed ROE begins in 1936. All data series continue to the end of 2005. The second column reports the date at which we begin the out-of-sample forecast evaluation. This is the beginning of 1927, when accurate data on total monthly stock returns become available from CRSP, or 20 years after the date in column 1, whichever comes later.

The third and fourth columns of Table 1 report the full-sample t -statistic for

the significance of each variable in forecasting stock returns, and the adjusted R^2 statistic of the full-sample regression.⁴ It is immediately obvious from the column of t -statistics that many of the valuation ratios and interest-rate variables are statistically insignificant in predicting stock returns over the long sample periods considered here. The most successful variables are the earnings yield, the Treasury bill rate, and the term spread. Of the two recently proposed variables, net equity issuance is modestly successful and the consumption-wealth ratio is strikingly successful in-sample.

The remaining columns of Table 1 evaluate the out-of-sample performance of these forecasts, using an out-of-sample R^2 statistic that can be compared with the in-sample R^2 statistic. This is computed as:

$$R_{OS}^2 = 1 - \frac{\sum_{t=1}^T (r_t - \hat{r}_t)^2}{\sum_{t=1}^T (r_t - \bar{r}_t)^2}, \quad (1)$$

where $\hat{r}_t$ is the fitted value from a predictive regression estimated through period $t-1$, and $\bar{r}_t$ is the historical average return estimated through period $t-1$. If the out-of-sample R^2 is positive, then the predictive regression has lower average mean squared prediction error than the historical average return.⁵ We use the entire available history of stock returns, back to 1871, to estimate the historical average return. This gives the historical mean an advantage over predictive regressions with variables that have become available more recently, because more data are available to estimate the historical mean than to estimate such predictive regressions. However, this is a real-world advantage of the historical mean that should be taken into account in our tests.

The out-of-sample performance of the predictor variables is mixed. The fifth column of Table 1 shows that only two out of four valuation ratios (the earnings yield and smoothed earnings yield) and two out of five interest-rate variables (the Treasury bill rate and term spread) deliver positive out-of-sample R^2 statistics. The interest rate results are consistent with the conclusion of Ang and Bekaert (2007) that the Treasury bill rate and term spread are robust return predictors. The performance of these variables would be stronger if we started the sample period later, because the interest rate process changed dramatically at the time of the Federal Reserve-Treasury Accord in 1951. Of the two recently proposed variables, net issuance performs reasonably well but the consumption-wealth ratio does not. The difficulty of estimating coefficients in a short sample is particularly severe for this series because it includes three separately estimated components.⁶

All the regressions we have reported predict simple stock returns rather than log stock returns. The use of simple returns makes little difference to the comparison of predictive regressions with historical mean forecasts, but all forecasts tend to underpredict returns when log returns are used. The reason is that high stock market volatility in the 1920s and 1930s depressed log returns relative to simple returns in this period. Thus the gap between average stock returns in the late twentieth century and the early twentieth century is greater in logs than in levels.

1.1 Sign restrictions on slope coefficients and forecasts

Despite the mixed performance of unrestricted predictive regressions out of sample, it is premature to conclude with Goyal and Welch (2006) that predictive regressions cannot profitably be used by investors in real time. A regression estimated over a short sample period can easily generate perverse results, such as a negative coefficient when theory suggests that the coefficient should be positive. Since out-of-sample forecast evaluation begins as little as 20 years after a predictor variable first becomes available, while the historical mean return is estimated from the beginning of our data set, this can be an important problem in practice. For example, in the early 1930s the earnings-price ratio was very high, but the coefficient on the predictor was estimated to be negative. This led to a negative forecast of the equity premium in the early 1930s and subsequent poor forecast performance. In practice, an investor would not use a perverse coefficient but would likely conclude that the coefficient is zero, in effect imposing prior knowledge on the output of the regression.

In the remaining columns of Table 1, we explore the impact of imposing sensible restrictions on the out-of-sample forecasting exercise. In the sixth column, we set the regression coefficient to zero whenever it has the “wrong” sign (different from the theoretically expected sign estimated over the full sample). In the seventh column, we assume that investors rule out a negative equity premium, and set the forecast to zero whenever it is negative. In the eighth and final column, we impose first the sign restriction on the coefficient, and then the sign restriction on the forecast.

These restrictions never worsen and almost always improve the out-of-sample performance of our predictive regressions. With no restrictions, as we noted above, only earnings-based valuation ratios have a positive out-of-sample R^2 , but the slope restriction delivers a positive out-of-sample R^2 for the dividend yield and the sign restriction brings the out-of-sample R^2 close to zero for the book-to-market ratio.

The sign restriction also delivers a positive out-of-sample R^2 for the long-term bond yield and the consumption-wealth ratio.

Figure 1 illustrates the effect of the restrictions for the dividend-price ratio. The top panel shows annualized excess return forecasts based on the full-sample OLS regression coefficient, the rolling (“out-of-sample”) OLS regression coefficient without restrictions, and the out-of-sample coefficient with both coefficient and forecast sign restrictions. The bottom panel shows the cumulative out-of-sample R^2 for these three forecasts. The coefficient sign restriction significantly improves the forecasts in the 1930s, when the coefficient was estimated to be negative. The forecast restrictions bind periodically during the 1960s and 1990s. Valuation ratios were unusually low during these periods, leading to unprecedentedly low forecasts. Campbell and Shiller (2001) also noted the unusually low valuation ratios of the 1990s, and wrote “We do not find this extreme forecast credible; when the independent variable has moved so far from the historically observed range, we cannot trust a linear regression line.” Our forecast restrictions are a simple way to avoid such incredible forecasts.

Panel B of Table 1 reports comparable results for annual regressions, estimated using overlapping monthly data. In-sample t -statistics are corrected for serial correlation and are even lower than those reported in Panel A for monthly data. Despite this weak in-sample predictive power, these regressions perform quite well out-of-sample. When both our theoretical restrictions are imposed, all four valuation ratios and the three variables based on the Treasury yield curve have out-of-sample R^2 statistics of at least 2%. Net equity issuance and the consumption-wealth ratio also benefit from our restrictions but do not beat the historical mean return at the annual frequency.

1.2 Quantitative coefficient restrictions for valuation ratios

So far we have used theory only to impose weak sign restrictions on coefficients and excess return forecasts. We now go further and impose the restrictions of a simple steady-state model on the coefficients for valuation ratios. This approach can be interpreted as a conditional version of the procedure advocated by Fama and French (2002) to estimate the unconditional average mean return. It can also be interpreted as a way to combine the information from multiple data sources, for example, on market valuations, corporate profitability, and payout policy, without increasing the number of free coefficients that must be estimated.

We start from the Gordon (1962) growth formula,

$$\frac{D}{P} = R - G, \quad (2)$$

which describes the dividend-price ratio in a steady state with a constant discount rate and dividend growth rate. We combine this formula with the steady-state relation between growth and accounting return on equity,

$$G = \left(1 - \frac{D}{E}\right)ROE, \quad (3)$$

where D/E is the payout ratio, to obtain a growth-adjusted return forecast:

$$\hat{R}_{DP} = \frac{D}{P} + \left(1 - \frac{D}{E}\right)ROE. \quad (4)$$

The same approach can be used to obtain growth forecasts from the earnings yield. Using the fact that $D/P = (D/E)(E/P)$ we have:

$$\hat{R}_{EP} = \left(\frac{D}{E}\right)\frac{E}{P} + \left(1 - \frac{D}{E}\right)ROE, \quad (5)$$

a payout-ratio-weighted average of the earnings yield and the accounting return on equity. When return on equity equals the expected return, as might be the case in long-run equilibrium, then this implies that $\hat{R}_{EP} = E/P$.

Finally, since $E/P = (B/M)ROE$, we have:

$$\hat{R}_{BM} = ROE \left[1 + \frac{D}{E} \left(\frac{B}{M} - 1 \right) \right]. \quad (6)$$

To use these formulas in practice, one must decide how to combine historical and contemporaneous data on the right-hand side variables. We follow Fama and French (2002) by using historical average data on payouts and profitability, but differ from them by using current rather than historical average data on valuation ratios to obtain a return forecast conditional on the market's current valuation level. This procedure assumes that movements in valuation ratios, relative to historical cash flows, are explained by permanent changes in expected returns. It is a compromise between the view that valuation ratios are driven by changing forecasts of profitability, in which

case the implied movements in returns would be smaller, and the view that valuation ratios are driven by *temporary* changes in discount rates, in which case the implied return movements would be larger [Campbell and Shiller (1988a)].⁷

Table 2 reports regression forecasts using both the unadjusted valuation ratios from Table 1, and two variants of growth-adjusted ratios. Our first set of growth-adjusted ratios uses Equations (4), (5), and (6) with current data on the valuation ratios and historical data on the payout ratio (an average from the beginning of the sample) and the return on equity (a ten-year moving average as in Table 1).⁸ Our second set of growth-adjusted ratios in addition subtracts the historical average real interest rate from the beginning of the sample period in order to convert a theoretical real return forecast into a theoretical excess return forecast. The historical average real interest rate is extremely stable, so this final step is close to an intercept adjustment.⁹

We use these ratios in four different ways. First, we report unrestricted regressions of returns on the ratios. Second, we report regressions restricted as in Table 1 to have positive slope coefficients and positive return forecasts. Third, given that valuation ratios are positive, we can obtain reasonable return forecasts merely by bounding the intercept above zero and the slope coefficient to lie between zero and one. This is similar to the second approach but can be implemented by coefficient restrictions rather than by restricting forecasts directly. Finally, we impose the restrictions of the steady-state theory by restricting the intercept of the regression to be zero and the slope coefficient to be one. That is, we estimate the return forecast directly from the data without using any historical information on the covariance between returns and valuation ratios.¹⁰

Table 2 shows that the last and most restrictive approach delivers the best out-of-sample performance in monthly data. The out-of-sample R^2 statistics range from -0.66% to 0.32% when no restrictions are imposed, from -0.45 to 0.43% when the restrictions of Table 1 are imposed, and from 0.24% to 0.97% when the zero-intercept and unit-slope restrictions are imposed. These restrictions improve the out-of-sample monthly forecasting power of every predictive regression we consider. The out-of-sample R^2 statistics are also reliably positive in annual regressions that impose zero-intercept and unit-slope restrictions, ranging from 1.85% to 7.99%, but here the theoretical restrictions worsen out-of-sample predictive power in a few cases. The most striking example is the dividend-price ratio with no growth adjustment, where the zero-intercept restriction is the least theoretically appealing as it effectively assumes

zero real growth in dividends.

Goyal and Welch (2006) emphasize that out-of-sample performance has been particularly poor for predictive regressions in the last few decades. To investigate this issue, Table 3 repeats Table 2 for three subsamples: 1927-1956, 1956-1980, and 1980-2005. We choose the first subsample to end twenty years after we are first able to calculate ten-year smoothed ROE, so that growth adjustments in the second and third subsamples are entirely based on ROE and not on historical earnings growth. We omit the book-to-market ratio in the first subsample, as the series does not become available until 1926 and thus we cannot estimate coefficients before the first subsample begins.

As Goyal and Welch (2006) have noted, the ability of valuation ratios to forecast stock returns is strongest in the first subsample, which includes the Great Depression and World War II, slightly weaker in the second subsample, which includes the post-war boom and the oil shocks of the 1970s, and weakest in the third subsample, which includes the great equity bull market at the end of the twentieth century. Even in the third subsample, however, earnings-based valuation ratios outperform the historical mean return in forecasting future monthly stock returns. The outperformance is particularly convincing when the ratios are adjusted for growth, in which case it can be achieved either by imposing a zero intercept and a unit slope or by merely imposing a positive intercept and a slope between zero and one. Results are more mixed in annual data during the third subsample, but here too earnings-based valuation ratios perform reasonably well. The relatively weak performance of the dividend-price ratio in this period is consistent with the observation of Boudoukh, Michaely, Richardson, and Roberts (2007), who find that in recent years share repurchases have played a more important role in total payouts to shareholders.

Figure 2 compares the historical mean stock return with four forecasts based on the smoothed earnings yield: the unrestricted rolling regression forecast, the zero-intercept unit-slope forecast based on the smoothed earnings yield with no adjustments, and the zero-intercept unit-slope forecasts based on the two growth-adjusted versions of the smoothed earnings yield. The theoretically restricted forecasts are more stable over time and have superior out-of-sample R^2 statistics as illustrated in the bottom panel of the figure. Although there is a very gradual decline over time in the out-of-sample R^2 statistics of these forecasts, this does not mean that the forecasts are not working in the later years of the sample; it simply means that they are not beating the historical mean forecast as decisively as they did in earlier years.

2 How Large an R^2 Should We Expect?

In the previous section, we showed that many of the forecasting variables that have been discussed in the literature do have positive out-of-sample predictive power for aggregate U.S. stock returns, when reasonable restrictions are imposed on the predictive regression. However, the R^2 statistics are small in magnitude. This raises the important question of whether they are economically meaningful.

To explore this issue, consider the following example:

$$r_{t+1} = \mu + x_t + \varepsilon_{t+1}, \quad (7)$$

where r_{t+1} is the excess simple return on a risky asset over the riskless interest rate, μ is the unconditional average excess return, x_t is a predictor variable with mean zero and constant variance σ_x^2 , and ε_{t+1} is a random shock with mean zero and constant variance σ_ε^2 . For tractability, consider an investor with a single-period horizon and mean-variance preferences. The investor's objective function is expected portfolio return less $(\gamma/2)$ times portfolio variance, where γ can be interpreted as the coefficient of relative risk aversion.¹¹ If the investor does not observe x_t , the investor chooses a portfolio weight in the risky asset:

$$\alpha_t = \alpha = \left(\frac{1}{\gamma}\right) \left(\frac{\mu}{\sigma_x^2 + \sigma_\varepsilon^2}\right) \quad (8)$$

and earns an average excess return of:

$$\left(\frac{1}{\gamma}\right) \left(\frac{\mu^2}{\sigma_x^2 + \sigma_\varepsilon^2}\right) = \frac{S^2}{\gamma}, \quad (9)$$

where S is the unconditional Sharpe ratio of the risky asset.

If the investor observes x_t , the investor sets:

$$\alpha_t = \left(\frac{1}{\gamma}\right) \left(\frac{\mu + x_t}{\sigma_\varepsilon^2}\right), \quad (10)$$

where the denominator is now σ_ε^2 rather than $\sigma_x^2 + \sigma_\varepsilon^2$ because the variation in the predictor variable x_t is now expected and does not contribute to risk. The investor earns an average excess return of:

$$\left(\frac{1}{\gamma}\right) \left(\frac{\mu^2 + \sigma_x^2}{\sigma_\varepsilon^2}\right) = \left(\frac{1}{\gamma}\right) \left(\frac{S^2 + R^2}{1 - R^2}\right), \quad (11)$$

where

$$R^2 = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_\varepsilon^2} \quad (12)$$

is the R^2 statistic for the regression of excess return on the predictor variable x_t .

The difference between the two expected returns is:

$$\left(\frac{1}{\gamma}\right) \left(\frac{R^2}{1 - R^2}\right) (1 + S^2), \quad (13)$$

which is always larger than R^2/γ , and is close to R^2/γ when the time interval is short and R^2 and S^2 are both small. The *proportional* increase in the expected return from observing x_t is:

$$\left(\frac{R^2}{1 - R^2}\right) \left(\frac{1 + S^2}{S^2}\right), \quad (14)$$

which is always larger than R^2/S^2 and is close to R^2/S^2 when the time interval is short and R^2 and S^2 are both small.

This analysis shows that the right way to judge the magnitude of R^2 is to compare it with the squared Sharpe ratio S^2 . If R^2 is large relative to S^2 , then an investor can use the information in the predictive regression to obtain a large proportional increase in portfolio return. In our monthly data since 1871, the monthly Sharpe ratio for stocks is 0.108, corresponding to an annual Sharpe ratio of 0.374. The squared monthly Sharpe ratio is $S^2 = 0.012 = 1.2\%$. This can be compared with the monthly out-of-sample R^2 statistic for, say, the smoothed earnings-price ratio of 0.43% in the last column of Panel A of Table 1. A mean-variance investor can use the smoothed earnings-price ratio to increase average monthly portfolio return by a proportional factor of $0.43/1.2 = 36\%$. The absolute increase in portfolio return depends on risk aversion, but is about 43 basis points per month or 5.2% per year for an investor with unit risk aversion, and about 1.7% per year for an investor with a risk aversion coefficient of three. The calculation can also be done in the annual data shown in Panel B of Table 1, comparing the squared annual Sharpe ratio of 11.8% to the annual out-of-sample R^2 statistic for, say, the smoothed earnings-price ratio of 7.9%. A mean-variance investor can use the smoothed earnings-price ratio in annual data to increase average annual portfolio return by a proportional factor of $7.9/11.8 = 67\%$. The absolute increase in portfolio return is 9.6% per year for an investor with unit risk aversion, and 3.2% per year for an investor with a risk aversion coefficient of three.

The investor who observes x_t gets a higher portfolio return in part by taking on greater risk. Thus the increase in the average return is not pure welfare gain for a risk-averse investor. To take account of this, in Table 4 we calculate the welfare benefits generated by optimally trading on each predictor variable for an investor with a relative risk aversion coefficient of three. We impose realistic portfolio constraints, preventing the investor from shorting stocks or taking more than 50% leverage: that is, confining the portfolio weight on stocks to lie between 0% and 150%. The investor's optimal portfolio depends on an estimate of stock return variance at each point in time, and we assume that the investor estimates variance using a rolling five-year window of monthly data.¹² We report the change in utility, relative to investing with the historical mean forecast of the equity premium, for each of the models considered in Table 3. These utility differences have the units of expected annualized return, so they can be interpreted as the transactions costs or portfolio management fees that investors would be prepared to pay each year to exploit the information in the predictor variable.

The utility gains in Table 4 are broadly consistent with the out-of-sample R^2 statistics reported in Table 3. Utility gains are reliably positive using annual returns and coefficients fixed at theoretical levels: The only utility loss occurs for the unadjusted dividend-price ratio in the 1980-2005 sample, when share repurchases become important. The results are somewhat more mixed for monthly returns, but here too are generally positive. The utility gains tend to be greatest in the 1927-1956 subsample, but by contrast with Table 3, are not always smaller in 1980-2005 than in 1956-1980.

The utility gains reported in Table 4 are limited by the leverage constraint, together with the high average equity premium. Predictable variations in stock returns do not generate portfolio gains when there is a binding upper limit on equity investment. Utility gains would be larger if we relaxed the portfolio constraint or included additional assets, with higher average returns than Treasury bills, in the portfolio choice problem. On the other hand, Table 4 does not take any account of transactions costs. Modest gains from market timing strategies could be offset by the additional costs implied by those strategies. Optimal trading strategies in the presence of transactions costs are complex, and so we do not explore this issue further here. We note that even the baseline strategy based on a historical-mean forecast incurs rebalancing costs and that utility gains of 50 basis points per year, which are commonly achieved in Table 4, are sufficient to cover substantial additional costs.

Since small R^2 statistics can generate large benefits for investors, we should expect

predictive regressions to have only modest explanatory power. Regressions with large R^2 statistics would be too profitable to believe. The saying “If you’re so smart, why aren’t you rich?” applies with great force here, and should lead investors to suspect that highly successful predictive regressions are spurious. Since the squared Sharpe ratio and average real interest rate increase in proportion with the investment horizon, however, larger R^2 statistics are believable at longer horizons. Authors such as Fama and French (1988) have found that R^2 statistics increase strongly with the horizon when the predictor variable is persistent, a finding that is analyzed in Campbell, Lo, and MacKinlay (1997, Chapter 7) and Campbell (2001). This behavior is completely consistent with our analysis and empirical results.

3 Conclusion

A number of variables are correlated with subsequent returns on the aggregate U.S. stock market in the twentieth century. Some of these variables are stock market valuation ratios, others reflect the levels of short- and long-term interest rates, patterns in corporate finance or the cross-sectional pricing of individual stocks, or the level of consumption in relation to wealth. Goyal and Welch (2006) argue that in-sample correlations conceal a systematic failure of these variables out of sample: None have been able to beat a simple forecast based on the historical average stock return.

In this paper, we have shown that most of these predictor variables perform better out-of-sample than the historical average return forecast, once weak restrictions are imposed on the signs of coefficients and return forecasts. The out-of-sample explanatory power is small, but nonetheless is economically meaningful for investors. We have also explored theoretical restrictions on the coefficients relating valuation ratios to future returns. Although these restrictions are based on steady-state models, which do not fully describe asset prices in an economy with changing growth rates and discount rates, they improve the out-of-sample performance of valuation forecasts by removing the need to estimate coefficients from short samples of volatile stock returns. We find that even in the period 1980-2005, where return prediction is most challenging, theoretically restricted valuation models often outperform return forecasts based on the long-run historical mean of stock returns. These models also generate meaningful utility gains for mean-variance investors. Importantly, the performance of the models does not depend sensitively on the particular valuation ratio that is used or the manner in which it is adjusted for long-run growth. We inter-

pret this result as illustrating the well-known econometric principle that even false theoretical restrictions can be helpful in forecasting if they reduce the variance of a predictor more than they increase its bias.

We have deliberately kept our methods simple to avoid the charge that we have tested a wide variety of methods and have only reported those that succeed. The results suggest that further improvements in performance can be achieved by exploring alternative methods to combine theory with historical data. One obvious possibility is a Bayesian approach with priors centered either on a model with a constant expected stock return, as in Wachter and Warusawitharana (2006), or on a valuation model with constant or slowly changing profitability and payout ratios of the sort we have explored in this paper. Of course, given the quantity of historical data that are now available, optimal forecasts of stock returns going forward may place greater weight on the data, and less weight on theoretical restrictions, than those methods that most successfully predicted stock returns during the twentieth century.

Endnotes

Acknowledgements: We are grateful to Jan Szilagyi for able research assistance, to Amit Goyal and Ivo Welch for sharing their data, and to Malcolm Baker, Lutz Kilian, Martin Lettau, Sydney Ludvigson, Rossen Valkanov, the editor, and two anonymous referees for helpful comments on an earlier draft entitled “Predicting the Equity Premium Out of Sample: Can Anything Beat the Historical Average?” This material is based upon work supported by the National Science Foundation under Grant No. 0214061 to Campbell.

1. Goyal and Welch (2006) consider these variables, along with the ratio of lagged dividends to lagged prices (the “dividend yield” in Goyal and Welch’s terminology). We drop this variable as there is no reason to believe that it should be a better predictor than the ratio of lagged dividends to current prices (the “dividend price ratio”).

2. If earnings obey the clean-surplus relation, the growth rate of real book equity equals real ROE minus the payout ratio. ROE can alternatively be measured from the growth rate of book equity rather than from reported accounting earnings, as in Cohen, Polk, and Vuolteenaho (2006).

3. In 2003, the BEA revised the definitions of several variables used by Lettau and Ludvigson (2001) to construct their series. We use the updated data available on Martin Lettau’s website, not the original series used in Lettau and Ludvigson (2001). Data revisions raise the deeper problem that the series may not have been available to investors in real time, but we do not try to deal with this issue here.

4. The adjustment of the R^2 statistic for degrees of freedom makes only a very small difference in samples of the size used here. The adjustment is about 5 basis points for a regression starting in 1871, and about 10 basis points for a regression starting in 1927.

5. Clark and West (2005) point out that if the return series is truly unpredictable, then in a finite sample the predictive regression will on average have a higher mean squared prediction error because it must estimate an additional coefficient. Thus the expected out-of-sample R^2 under the null of unpredictability is negative, and a zero out-of-sample R^2 can be interpreted as weak evidence for predictability. We do not pursue this point here because, like Goyal and Welch (2006), we ask whether predictive regressions or historical average return forecasts have delivered better out-of-sample

forecasts, not whether stock returns are truly predictable.

6. Earlier drafts of this paper reported better out-of-sample performance for this series. BEA data revisions in 2003 and the addition of recent years to the forecast evaluation period are responsible for the change in results.

7. Cochrane (2006) also emphasizes the importance of the fact that valuation ratios do not forecast growth rates of cash flows. The ability of the theoretical models in this section to predict stock returns is consistent with Cochrane's results.

8. Before 1926, we do not have data on ROE and thus we cannot calculate ten-year smoothed ROE until 1936. Before that date, we use real earnings growth to estimate G .

9. One could consider forecasting the short-term real interest rate using recent data on inflation and nominal interest rates. There are two reasons not to pursue this approach. First, the volatility of inflation makes it difficult in practice. Second, the steady-state growth model delivers a long-run forecast of stock returns that should be matched to a long-run forecast of real bond returns, such as the TIPS yield or (since this is not available before 1997) the long-run historical average real interest rate. If the short-term real interest rate exceeds such a long-run real rate forecast, then it is quite possible that the short-term real stock return shifts upwards in parallel, leaving the excess stock return forecast unchanged.

10. A subtle but important issue is that we use the steady-state model to forecast the arithmetic average stock return, and take arithmetic averages of ROE and other historical data. Some authors, such as Siegel (1994), take geometric averages of historical data and forecast the geometric average stock return. These two approaches are equivalent if the volatility of dividend growth and stock returns are the same, as implied by the steady-state model, but are different in the data because stock returns are much more volatile than dividend growth and ROE. A full exploration of the two approaches is beyond the scope of this paper, but we note that the Siegel approach would generate higher forecasts of arithmetic average stock returns, and thus would perform even better in the late twentieth century than the approach we use here.

11. Merton (1969) presents the analogous portfolio solution for the case where the investor has power utility with relative risk aversion γ , asset returns are lognormally distributed, and the portfolio can be continuously rebalanced. Campbell and Viceria (2002, Chapter 2) use a discrete-time approximate version of Merton's solution.

Sentana (2005) also explores the relation between regression forecasts and optimal portfolio construction.

12. In an earlier draft we reported similar results for the case where the investor estimates variance using all available historical data. If these variance estimates are incorrect, for example because the predictor variable forecasts variance as well as return, this will reduce the utility generated by trading on the predictive variable.

References

- Amihud, Y., and C. Hurvich, 2004, "Predictive Regressions: A Reduced-Bias Estimation Method," *Journal of Financial and Quantitative Analysis*, 39, 813–841.
- Ang, A., and G. Bekaert, 2007, "Stock Return Predictability: Is It There?," *The Review of Financial Studies* 20, 651–707.
- Baker, M., and J. Wurgler, 2000, "The Equity Share in New Issues and Aggregate Stock Returns," *Journal of Finance*, 55, 2219–2257.
- Bollerslev, T., 1990, "Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model," *Review of Economics and Statistics*, 72, 498–505.
- Bollerslev, T., and J. Wooldridge, 1992, "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews*, 11, 143–172.
- Boudoukh, J., R. Michaely, M. Richardson, and M. Roberts, 2007, "On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing," *Journal of Finance*, 62, 877–916.
- Butler, A. W., G. Grullon, and J. P. Weston, 2005, "Can Managers Forecast Aggregate Market Returns?," *Journal of Finance*, 60, 963–986.
- Campbell, J. Y., 1987, "Stock Returns and the Term Structure," *Journal of Financial Economics*, 18, 373–399.
- Campbell, J. Y., 2001, "Why Long Horizons? A Study of Power Against Persistent Alternatives," *Journal of Empirical Finance*, 8, 459–491.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton.
- Campbell, J. Y., and R. J. Shiller, 1988a, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *The Review of Financial Studies*, 1, 195–228.
- Campbell, J. Y., and R. J. Shiller, 1988b, "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 43, 661–676.

- Campbell, J. Y., and R. J. Shiller, 1998, "Valuation Ratios and the Long-Run Stock Market Outlook," *Journal of Portfolio Management* 24(2), 11–26.
- Campbell, J. Y., and R. J. Shiller, 2001, "Valuation Ratios and the Long-Run Stock Market Outlook: An Update," NBER Working Paper 8221.
- Campbell, J. Y., and L. M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, New York.
- Campbell, J. Y., and M. Yogo, 2006, "Efficient Tests of Stock Return Predictability", *Journal of Financial Economics*, 81, 27–60.
- Cavanagh, C. L., G. Elliott, and J. H. Stock, 1995, "Inference in Models with Nearly Integrated Regressors," *Econometric Theory*, 11, 1131–1147.
- Clark, T. E., and K. D. West, 2005, "Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference," NBER Technical Paper 305.
- Cochrane, J. H., 2006, "The Dog That Did Not Bark: A Defense of Return Predictability," Unpublished paper, University of Chicago.
- Fama, E. F., and K. R. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.
- Fama, E. F., and K. R. French, 1989, "Business Conditions and Expected Returns on Stocks and Bonds," *Journal of Financial Economics*, 25, 23–49.
- Fama, E. F., and K. R. French, 2002, "The Equity Premium," *Journal of Finance*, 57, 637–659.
- Fama, E. F., and G. W. Schwert, 1977, "Asset Returns and Inflation," *Journal of Financial Economics*, 5, 115–146.
- Ferson, W. E., S. Sarkissian, and T. T. Simin, 2003, "Spurious Regressions in Financial Economics?," *Journal of Finance*, 58, 1393–1413.
- Foster, F. D., T. Smith, and R. E. Whaley, 1997, "Assessing Goodness-of-Fit of Asset Pricing Models: The Distribution of the Maximal R^2 ," *Journal of Finance*, 52, 591–607.
- Gordon, M., 1962, *The Investment, Financing, and Valuation of the Corporation*, Irwin, Homewood, IL.

- Goyal, A., and I. Welch, 2003, "Predicting the Equity Premium with Dividend Ratios," *Management Science*, 49, 639–654.
- Goyal, A., and I. Welch, 2006, "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," forthcoming *The Review of Financial Studies*.
- Graham, B., and D. L. Dodd, 1934, *Security Analysis*, First edition, McGraw Hill, New York, NY.
- Hodrick, R. J., 1992, "Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement," *The Review of Financial Studies*, 5, 257–286.
- Inoue, A., and L. Kilian, 2004, "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?," *Econometric Reviews*, 23, 371–402.
- Jansson, M., and M. J. Moreira, 2006, "Optimal Inference in Regression Models with Nearly Integrated Regressors," *Econometrica*, 74, 681–715.
- Keim, D. B., and R. F. Stambaugh, 1986, "Predicting Returns in the Stock and Bond Markets," *Journal of Financial Economics*, 17, 357–390.
- Kilian, L., 1999, "Exchange Rates and Monetary Fundamentals: What Do We Learn from Long-Horizon Regressions?," *Journal of Applied Econometrics*, 14, 491–510.
- Kothari, S.P., and J. Shanken, 1997, "Book-to-Market, Dividend Yield, and Expected Market Returns: A Time-Series Analysis," *Journal of Financial Economics*, 44, 169–203.
- Lamont, O., 1998, "Earnings and Expected Returns," *Journal of Finance*, 53, 1563–1587.
- Lettau, M., and S. Ludvigson, 2001, "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 56, 815–849.
- Lewellen, J., 2004, "Predicting Returns with Financial Ratios," *Journal of Financial Economics*, 74, 209–235.
- Litterman, R., 1986, "Forecasting with Bayesian Vector Autoregressions: Five Years of Experience," *Journal of Business and Economic Statistics*, 4, 25–38.

- Mark, N. C., 1995, "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability," *American Economic Review*, 85, 201–218.
- Merton, R. C., 1969, "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics*, 51, 247–257.
- Nelson, C., and M. Kim, 1993, "Predictable Stock Returns: The Role of Small Sample Bias," *Journal of Finance*, 48, 641–661.
- Polk, C., S. Thompson, and T. Vuolteenaho, 2006, "Cross-Sectional Forecasts of the Equity Premium," *Journal of Financial Economics*, 81, 101–141.
- Pontiff, J., and L. D. Schall, 1998, "Book-to-Market Ratios as Predictors of Market Returns," *Journal of Financial Economics*, 49, 141–160.
- Rozeff, M. S., 1984, "Dividend Yields are Equity Risk Premiums," *Journal of Portfolio Management*, 11(1), 68–75.
- Sentana, E., 2005, "Least Squares Predictions and Mean-Variance Analysis," *Journal of Financial Econometrics*, 3, 56–78.
- Siegel, J., 1994, *Stocks for the Long Run*, Norton, New York.
- Stambaugh, R. F., 1999, "Predictive Regressions," *Journal of Financial Economics*, 54, 375–421.
- Torous, W., R. Valkanov, and S. Yan, 2004, "On Predicting Stock Returns with Nearly Integrated Explanatory Variables," *Journal of Business*, 77, 937–966.
- Wachter, J. A., and M. Warusawitharana, 2006, "Predictable Returns and Asset Allocation: Should a Skeptical Investor Time the Market?," Unpublished paper, University of Pennsylvania.

Table 1: Excess return prediction with regression constraints

	Sample Begin	Forecast Begin	In-Sample t-statistic	In-Sample R-squared	Out-of-Sample R-squared with Different Constraints			
					Unconstrained	Positive Slope	Pos. Forecast	Both
Panel A: Monthly Returns								
Dividend-price ratio	1872m2	1927m1	1.25	1.13%	-0.65%	0.05%	0.07%	0.08%
Earnings-price ratio	1872m2	1927m1	2.29	0.71	0.12	0.18	0.14	0.18
Smooth earnings-price ratio	1881m2	1927m1	1.85	1.36	0.33	0.42	0.38	0.43
Book-to-Market	1926m6	1946m6	1.96	0.61	-0.43	-0.43	0.00	0.00
ROE	1936m6	1956m6	0.36	0.02	-0.93	-0.06	-0.93	-0.06
T-Bill rate	1920m1	1940m1	2.44	0.86	0.52	0.51	0.57	0.55
Long-term yield	1870m1	1927m1	1.46	0.19	-0.19	-0.19	0.20	0.20
Term spread	1920m1	1940m1	2.16	0.65	0.46	0.47	0.45	0.46
Default spread	1919m1	1939m1	0.74	0.10	-0.19	-0.19	-0.19	-0.19
Inflation	1871m5	1927m1	0.39	0.06	-0.22	-0.21	-0.18	-0.17
Net equity issuance	1927m12	1947m12	1.74	0.48	0.34	0.34	0.50	0.50
Consumption-wealth ratio	1951m12	1971m12	4.57	2.60	-1.36	-1.36	0.27	0.27
Panel B: Annual Returns								
Dividend-price ratio	1872m2	1927m1	2.69	10.89%	5.53%	5.53%	5.63%	5.63%
Earnings-price ratio	1872m2	1927m1	2.84	6.78	4.93	4.93	4.94	4.94
Smooth earnings-price ratio	1881m2	1927m1	3.01	13.57	7.89	7.89	7.85	7.85
Book-to-Market	1926m6	1946m6	1.98	8.26	-3.38	-3.38	1.39	1.39
ROE	1936m6	1956m6	0.35	0.32	-8.60	-0.03	-8.35	-0.03
T-Bill rate	1920m1	1940m1	1.77	4.26	5.54	5.54	7.47	7.47
Long-term yield	1870m1	1927m1	0.91	0.77	-0.15	-0.15	2.26	2.26
Term spread	1920m1	1940m1	1.72	3.10	4.79	4.79	4.74	4.74
Default spread	1919m1	1939m1	0.07	0.01	-3.81	-3.81	-3.81	-3.81
Inflation	1871m5	1927m1	0.17	0.07	-0.71	-0.71	-0.71	-0.71
Net equity issuance	1927m12	1947m12	0.54	0.35	-4.27	-4.27	-2.38	-2.38
Consumption-wealth ratio	1951m12	1971m12	3.76	19.87	-7.75	-7.75	-1.48	-1.48

This table presents statistics on forecast errors for stock returns. We use S&P500 total returns (dividend included) where data prior to January 1927 was obtained from Robert Shiller's website. "Sample Begin" denotes when the predictor was first available. All statistics are for the period that starts at "Forecast Begin" and ends on December 31, 2005 (for monthly forecasts) or December 31, 2004 (for annual forecasts). The "Unconstrained" out-of-sample R-squared compares the forecast error of the historical mean versus the forecast from unconstrained ordinary least squares. "Positive Slope" introduces the restriction that the coefficient on the predictor must be of the right sign, otherwise the historical mean is used as a predictor instead. "Pos. Forecast" requires that the prediction be positive, otherwise we use zero as the forecast. "Both" indicates that we impose both restrictions. The in-sample t-statistic and R-squared both come from the last regression (i.e., over the whole sample from the "Sample Begin" date to December 2005 for monthly data and December 2004 for annual data). The "In-Sample t statistic" is heteroskedasticity robust. The Annual forecasts in Panel B are based on 12 overlapping annual returns per year. The t-statistics in Panel B are corrected to take into account correlation induced by the overlapping nature of the dependent variable. For the consumption-wealth ratio, we report the robust F-statistic that all the slope coefficients are zero.

Table 2: Excess return prediction with valuation constraints

					Out-of-Sample R-squared with Different Constraints			
	Sample Begin	Forecast Begin	In-Sample t-statistic	In-Sample R-squared	Unconstrained	Positive Slope, Pos. Forecast	Pos. Intercept, Bounded Slope	Fixed Coefs
Panel A: Monthly Returns								
Dividend/price	1872m2	1927m1	1.25	1.12%	-0.66%	0.08%	0.19%	0.42%
Earnings/price	1872m2	1927m1	2.28	0.71	0.12	0.18	0.25	0.76
Smooth earnings/price	1881m2	1927m1	1.85	1.35	0.32	0.43	0.43	0.97
Dividend/price + growth	1891m2	1927m1	1.40	1.03	-0.05	0.20	0.17	0.63
Earnings/price + growth	1892m2	1927m1	1.82	0.49	-0.05	0.08	0.07	0.57
Smooth earnings/price + growth	1892m2	1927m1	2.00	1.10	0.11	0.25	0.21	0.72
Book-To-Market + growth	1936m6	1956m6	1.61	0.33	-0.35	-0.34	-0.34	0.33
Dividend/price + growth - real rate	1891m5	1927m1	1.47	0.86	-0.02	0.21	0.18	0.41
Earnings/price + growth - real rate	1892m2	1927m1	1.53	0.36	0.00	0.12	0.09	0.39
Smooth earnings/price + growth - real rate	1892m2	1927m1	1.97	0.84	0.15	0.26	0.23	0.52
Book-To-Market + growth - real rate	1936m6	1956m6	1.68	0.36	-0.45	-0.45	-0.42	0.24
Panel B: Annual Returns								
Dividend/price	1872m2	1927m1	2.69	10.89%	5.53%	5.63%	3.76%	2.20%
Earnings/price	1872m2	1927m1	2.84	6.78	4.93	4.94	4.34	5.87
Smooth earnings/price	1881m2	1927m1	3.01	13.57	7.89	7.85	6.44	7.99
Dividend/price + growth	1891m2	1927m1	1.77	9.30	2.49	2.99	2.67	4.35
Earnings/price + growth	1892m2	1927m1	1.42	4.44	1.69	2.11	1.80	3.89
Smooth earnings/price + growth	1892m2	1927m1	1.75	10.45	3.16	3.33	3.23	5.39
Book-To-Market + growth	1936m6	1956m6	1.97	5.45	-3.53	-0.64	-2.39	3.63
Dividend/price + growth - real rate	1891m5	1927m1	1.46	7.69	2.87	3.24	2.95	1.89
Earnings/price + growth - real rate	1892m2	1927m1	1.13	3.27	2.01	2.05	2.04	1.85
Smooth earnings/price + growth - real rate	1892m2	1927m1	1.53	7.90	3.35	3.35	3.38	3.22
Book-To-Market + growth - real rate	1936m6	1956m6	2.03	5.77	-1.73	-1.12	-1.82	2.33

The table presents forecast statistics for value predictors under various constraints. The predictor label "+ growth" indicates that we add an earnings growth forecast to the predictor. The predictor label "- real rate" indicates that we subtract a forecast of the real risk-free rate from the predictor. See the text for details. The "In-Sample" statistics are defined as in Table 1. The "Unconstrained" and "Positive Slope, Pos. Forecast" columns are described in Table 1. "Pos. Intercept, Bounded Slope" indicates that we constrain the intercept to be positive and the slope to be between zero and one. "Fixed Coefs" indicates that we fix the intercept at zero and the slope at one.

Table 3: Subsample stability

	Sample: 1927-1956			Sample: 1956-1980			Sample: 1980-2005		
	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs
Panel A: Monthly Returns									
Dividend/price	-0.86%	0.21%	0.63%	0.88%	0.57%	0.67%	-1.30%	-0.21%	-0.54%
Earnings/price	0.16	0.28	1.04	0.56	0.45	0.30	-0.53	-0.09	0.07
Smooth earnings/price	0.56	0.53	1.33	0.80	0.48	0.51	-1.06	-0.06	0.01
Dividend/price + growth	-0.15	0.18	0.78	0.18	0.18	0.59	0.11	0.11	0.14
Earnings/price + growth	-0.06	0.12	0.73	-0.12	-0.12	0.33	0.05	0.05	0.16
Smooth earnings/price + growth	0.09	0.25	0.93	0.19	0.19	0.47	0.06	0.06	0.16
Book-To-Market + growth				-0.62	-0.73	0.73	-0.12	-0.02	0.00
Dividend/price + growth - real rate	-0.01	0.30	0.45	-0.24	-0.24	0.76	0.11	0.11	-0.08
Earnings/price + growth - real rate	0.06	0.20	0.41	-0.34	-0.34	0.66	0.06	0.06	0.03
Smooth earnings/price + growth - real rate	0.27	0.39	0.60	-0.28	-0.28	0.74	0.04	0.04	0.02
Book-To-Market + growth - real rate				-0.82	-0.91	0.89	-0.14	-0.02	-0.27
Panel B: Annual Returns									
Dividend/price	9.95	4.53	3.67	9.46	5.99	6.88	-16.19	-1.38	-7.98
Earnings/price	7.45	5.34	7.58	5.08	3.25	2.56	-6.06	0.88	1.47
Smooth earnings/price	12.51	8.22	10.49	4.93	3.71	3.71	-8.86	1.33	1.33
Dividend/price + growth	2.77	3.05	4.83	1.76	1.74	6.61	1.87	1.82	0.28
Earnings/price + growth	2.21	2.38	4.37	-0.85	-0.85	3.97	1.63	1.63	1.60
Smooth earnings/price + growth	3.73	3.87	6.38	1.27	1.19	4.65	2.30	2.23	1.81
Book-To-Market + growth				-7.09	-5.16	10.34	-0.24	0.14	-2.43
Dividend/price + growth - real rate	4.40	4.51	1.67	-3.57	-3.56	8.28	2.19	2.19	-2.95
Earnings/price + growth - real rate	3.44	3.49	1.25	-4.68	-4.68	7.16	1.88	1.88	-0.64
Smooth earnings/price + growth - real rate	5.34	5.37	3.19	-4.91	-4.84	7.32	2.36	2.36	-0.47
Book-To-Market + growth - real rate				-3.36	-4.22	11.85	-0.25	0.35	-6.20

The table provides out-of-sample R-squared statistics from predicting the equity premium with a forecasting variable versus the historical mean. The subsamples roughly but not exactly divide the data into thirds. The column labels "Unconstrained," "Pos. Intercept, Bounded Slope," and "Fixed Coefs" are all defined in Tables 1 and 2. We do not provide results for the "Book-To-Market + Growth" predictor in the 1927-1956 subsample because we do not have 20 years of data until 1956.

Table 4: Portfolio choice

	Sample: 1927-1956			Sample: 1956-1980			Sample: 1980-2005		
	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs	Unconstrained	Pos. Intercept, Bounded Slope	Fixed Coefs
Panel A: Monthly Returns									
Dividend/price	-0.03%	0.43%	0.11%	1.93%	0.92%	1.46%	-3.69%	-0.22%	-1.32%
Earnings/price	0.86	1.42	1.80	0.28	0.12	-0.51	-0.80	0.07	0.74
Smooth earnings/price	0.53	1.14	1.89	0.75	0.05	-0.04	-2.73	0.13	0.46
Dividend/price + growth	-0.28	0.64	0.43	0.46	0.46	0.64	0.22	0.22	0.34
Earnings/price + growth	-1.05	0.36	1.60	0.08	0.08	-0.21	0.00	0.00	0.19
Smooth earnings/price + growth	-0.47	0.77	0.89	0.04	0.04	0.13	0.04	0.06	0.24
Book-To-Market + growth				0.43	0.15	1.38	0.08	-0.02	0.23
Dividend/price + growth - real rate	0.42	0.96	-0.74	0.01	0.01	1.49	0.04	0.04	0.18
Earnings/price + growth - real rate	-0.31	0.61	0.79	-0.17	-0.17	0.73	-0.09	-0.09	0.17
Smooth earnings/price + growth - real rate	0.36	1.14	-0.20	-0.38	-0.38	1.07	0.04	0.04	0.21
Book-To-Market + growth - real rate				0.38	0.05	1.94	-0.04	-0.10	-0.11
Panel B: Annual Returns									
Dividend/price	0.55	0.76	0.34	1.56	0.59	0.95	-4.74	0.34	-0.95
Earnings/price	1.29	1.30	0.78	1.41	0.41	0.40	-1.44	0.53	0.68
Smooth earnings/price	1.81	1.91	1.42	1.53	0.44	0.44	-3.06	0.81	0.81
Dividend/price + growth	-0.18	0.34	1.04	0.35	0.35	0.70	0.19	0.17	0.34
Earnings/price + growth	-0.14	0.28	1.52	0.15	0.15	0.49	0.04	0.04	0.24
Smooth earnings/price + growth	0.05	0.53	1.34	0.28	0.25	0.51	0.41	0.36	0.50
Book-To-Market + growth				-0.01	0.14	1.27	-0.10	-0.13	0.26
Dividend/price + growth - real rate	0.44	0.66	0.02	0.01	0.01	1.06	0.06	0.06	0.23
Earnings/price + growth - real rate	0.34	0.49	0.60	-0.09	-0.09	0.85	-0.02	-0.02	0.15
Smooth earnings/price + growth - real rate	0.68	0.87	0.39	-0.20	-0.18	0.94	0.24	0.24	0.36
Book-To-Market + growth - real rate				0.25	0.13	1.93	-0.17	-0.14	0.26

This table presents out-of-sample portfolio choice results. The numbers are the change in average utility from forecasting the market with the predictor instead of the historical mean. All numbers are annualized, so we multiply the monthly numbers by 12. "Unconstrained" indicates that we use the unconstrained OLS predictor of the equity premium. "Pos. Intercept, Bounded Slope" indicates that we use the forecast with the intercept bounded above zero and the slope bounded between zero and one. "Fixed Coefs" indicates that we use the forecast that sets the intercept to zero and the slope to one. The utility function is $E(R_p) - (\gamma/2)\text{Var}(R_p)$, where R_p is the portfolio return and $\gamma=3$. All utility changes are annualized, so we multiply monthly utility changes by 12.

Figure 1, Panel A: Forecasts Based on Dividend Yield
 ("oos ols" denotes out-of-sample ordinary least squares)

Figure 1, Panel B: Cumulative R-Squared Based on Dividend Yield
 (relative to historical mean)

Figure 2, Panel A: Forecasts Based on Smoothed Earnings Yield

**Figure 2, Panel B: Cumulative R-Squared Based on Smoothed Earnings Yield
(relative to historical mean)**

Figure legends

Figure 1: Forecasting excess returns with the dividend yield. The top panel shows the historical annualized excess return forecasts for the historical mean and three different regression models, as labeled in the figure, where “oos” refers to “out-of-sample” and “ols” refers to ordinary least squares. The bottom panel shows the cumulative out-of-sample R-squared up to each point in the historical sample, for the three regression models relative to the historical mean.

Figure 2: Forecasting excess returns with the smoothed earnings yield. The top panel shows the historical annualized excess return forecasts for the historical mean, an unrestricted regression on the smoothed earnings yield, and three theoretically restricted forecasts from the smoothed earnings yield, as labeled in the figure. The bottom panel shows the cumulative out-of-sample R-squared up to each point in the historical sample, for the four models that use the earnings yield relative to the historical mean.