

The Equity Risk Premium: A Consensus of Models

By [Fernando Duarte](#) and [Carlo Rosa](#)

Abstract

We estimate the equity risk premium by combining information from twenty models. Our main finding is that there is broad agreement across models that the equity premium reached historical heights in July 2013 even when the models are substantially different from each other and use more than one hundred different economic variables. Our preferred estimator places the one-year-ahead equity premium in July 2013 at 14.5 percent, the highest level in fifty years and well above the 10.5 percent that was reached during the financial crisis in 2009. The models also show broad agreement that the term structure of equity risk premia is high and flat: expected excess returns at all foreseeable horizons are just as high as at the one-year horizon. A high equity premium that is not expected to mean-revert in the near future is an unprecedented phenomenon. Because expected dividend growth has not been above average in 2013, we conclude the high equity premium is mostly due to unusually low discount rates at all horizons.

1. Introduction

The equity risk premium—the expected return of stocks in excess of the risk-free rate—is a fundamental quantity in all of asset pricing, both for theoretical and practical reasons. It is a key measure of aggregate risk-aversion and an important determinant of the cost of capital for corporations, saving decisions of individuals and budgeting plans for governments. Recently, the equity risk premium (ERP) has also returned to the forefront of policymaking as a leading indicator of the evolution of the economy, a potential explanation for the jobless recovery and a gauge of financial stability. As an indicator of future activity, a high ERP at short horizons tends to be followed by higher GDP growth, higher inflation and lower unemployment, thus informing both fiscal and monetary decisions. Bloom (2009) and new research by Duarte, Kogan and Livdan (2013) point to large effects of the ERP on real aggregate investment, a component that has been lagging in the present recovery compared to policymakers' forecasts in the current cycle and actual performance in past cycles. As a potential explanation of the jobless recovery, Hall (2013) and Kuehn, Petrosky-Nadeau and Zhang (2012) have proposed that increased risk-aversion has prevented firms from hiring as much as would be expected in today's macroeconomic environment. From the perspective of financial stability, the so-called "great rotation" from bonds to stocks could be exacerbated in speed and magnitude if the ERP is persistently high. A sudden flow of money out of the bond market into stocks could spell large capital losses for fixed income investors, including the Federal Reserve. Low returns in other asset classes could provide incentives for investors to engage in potentially unsafe "reach for yield" either through excessive use of leverage or through other forms of risk-taking. The ERP is also important from the perspective of unconventional monetary policy: a high ERP may make the portfolio channel of Large Scale Asset Purchases more effective because it further increases the demand for risky assets.

In this article, we estimate the ERP by combining information from twenty models that are prominently used by practitioners and featured in the academic literature. Our main finding is that there is broad agreement across models that the ERP has reached historical heights even when the models are substantially different from each other and use more than one hundred different economic variables. Our preferred estimator places the one-year-ahead ERP in July 2013 at 14.5 percent, the highest level in fifty years and well above the 10.5 percent that was reached during the financial crisis in 2009. The models also show broad agreement that expected excess returns at all foreseeable horizons are just as high as at the one-year horizon. A high equity premium that is not expected to mean-revert in the near future is an unprecedented phenomenon.

In addition to estimating the level of the ERP, it is useful for policymakers and other economic agents to know why the ERP is high. We conclude the ERP is high at all foreseeable horizons because Treasury yields are unusually low at all maturities. In other words, the term structure of equity premia is high and flat because the term structure of interest rates is low and flat. Current and expected future dividend and earnings growth play only a minor role. A high ERP caused by low bond yields indicates that a stock market correction is likely to occur only when bond yields start to rise. Additionally, a bond-driven ERP makes it more unlikely that irrational exuberance can take hold in equity markets, especially at times of increasing expectations for a steepening of the yield curve. Another implication of a bond-driven ERP is that we should no longer rely on traditional indicators of the ERP like the price-dividend or price-earnings ratios, which all but ignore the term structure of risk-free rates.

As a second contribution, we evaluate the performance of different ERP models. Statements about the implications of a high ERP are valid only to the extent that expected returns predict future realized returns. For the models we consider, predictability is weak but present. We first categorize the twenty models we study into five groups: predictors that use historical mean returns, dividend-discount models, cross-sectional regressions, time-series regressions and surveys. To assess whether models can indeed predict returns, we regress realized excess returns on the corresponding ERP given by the models. We then use the out-of-sample R-squared for these regressions as a measure of success. We find that dividend-discount models perform best at short horizons, while cross-sectional regressions perform best at longer horizons. Combining all models into a single principal component — our preferred measure— reduces noise. A mean-variance investor with unit risk aversion using the principal component as an investment signal would have earned 15 percentage points more over the last fifty years (30 basis points per year) than if she had assumed expected returns are equal to past mean returns.

2. The Equity Risk Premium: Definition

Conceptually, the ERP is the compensation investors require to make them indifferent between holding the risky market portfolio and a risk-free bond. Because this compensation depends on the future performance of stocks, the ERP incorporates expectations of future stock market returns, which are not directly observable. At the end of the day, any model of the ERP is a model of investor expectations. Additionally, it is not clear what truly constitutes the market return and the risk-free rate in the real world. In practice, the most common measures of market returns are given by broad stock market indices, like the S&P 500 or the Dow Jones Industrial Average, but those indices do not include the whole universe of traded stocks and miss several other components of wealth. Even if we included all traded stocks, we still have several choices to make, such as whether to use value or equal-weighted indices, or whether to exclude penny stocks or rarely traded stocks. A similar problem arises with the risk-free rate. While we almost always use Treasury yields as measures of risk-free rates, they are not completely riskless since nominal Treasuries are exposed to inflation and liquidity risks. In this paper, we follow common practice and always use the S&P 500 as a measure of stock market prices and either nominal or real Treasury yields as risk-free rates. The models we consider differ only in how expectations are computed.

While implementing the concept of the ERP has pitfalls, we can precisely define the ERP mathematically. First, we decompose stock returns into an expected component and an unpredictable random component:

$$R_{t+k} = E_t[R_{t+k}] + error_{t+k} \quad (1)$$

In equation (1), R_{t+k} are net *realized* returns between t and $t+k$, $E_t[R_{t+k}]$ are the returns that were expected from t to $t+k$ using information available at time t and $error_{t+k}$ is a mean-zero random variable that is unknown at time t but is realized at $t+k$. The ERP at time t for horizon k is defined as

$$ERP_t(k) = E_t[R_{t+k}] - R_{t+k}^f \quad (2)$$

where R_{t+k}^f is the net risk-free rate for investing from t to $t+k$ (which, being risk-free, is known at time t).

This definition shows three important aspects of the ERP. First, because the unexpected component $error_{t+k}$ is stochastic and orthogonal to expected returns, the ERP is always less volatile than realized excess returns. Therefore, while realized stock returns are very volatile compared to bonds, we expect good ERP estimates to be somewhat smoother. Second, the ERP itself is a random variable, since expectations can change through time when new information arrives. Third, the ERP has an investment horizon k embedded in it, since we can consider expected excess returns over, say, one month, one year or five years from today. If we fix t , and let k vary, we trace the *term structure* of the equity risk premium.

3. Data

In constructing all estimates of the ERP we use over one hundred variables. The sources and definitions are standard. The nominal and real price, earnings and dividends for the S&P 500 are from Shiller. Inflation, the “cyclically adjusted price-earnings ratio” and the ten-year nominal treasury yield are also from Shiller. Expected earnings per share are mean analyst forecasts from Thomson Reuters I/B/E/S. Nominal bond yields for all maturities except 10-years¹ and all TIPS yields are from the Federal Reserve Board. Fama-French and momentum factors and portfolios are from Professor French’s website. Corporate bond spreads and the NBER recession indicator are from the St. Louis Federal Reserve (FRED). Book value per share for the S&P 500 is from Compustat. Debt issuance and equity issuance are from Jeffrey Wurgler’s website. Consumption to wealth ratio measured by *cay* is from Martin Lettau’s website (Ludvigson and Lettau, 2001). ERP estimates from CFOs are from the Duke CFO survey. The sentiment measure of Baker and Wurgler is from Jeffrey Wurgler’s website. Professor Damodaran’s estimates of the ERP are from his website. All variables are monthly from January 1960 to July 2013, except for *cay* and CFO surveys, which are quarterly, and book value per share and Damodaran’s ERP estimates, which are annual. Other variables are constructed using the variables mentioned before. A detailed description is in Appendix A.

4. Models of the Equity Risk Premium

We classify models of the ERP into five categories and discuss their advantages and disadvantages. We also describe in detail the models we use within each category and how to obtain a term-structure of the ERP for each one. Of course, there are many more models of the ERP than the ones we consider. We selected which models to include in our study based on the recent academic literature and widespread use by practitioners. All models are constructed in real time², so that an investor who lived through the

¹ Except for the 10-year yield, which, as described above, is from Shiller. We use Shiller’s 10-year yield for ease of comparability with the existing literature. Results are virtually unchanged if we use all yields, including the 10-year yield, from the Federal Reserve Board.

² The one exception is Adrian, Crump and Moench’s (2013) cross-sectional model, which is constructed using full-sample regression estimates. Our out-of-sample predictability results are essentially unchanged if we omit this model from the analysis.

sample would have been able to construct the measures at each point in time using available information only. This helps avoid look-ahead bias and makes the out-of-sample evaluation of the models meaningful.

4.1 Historical mean

The easiest approach to estimating the ERP is to assume it is equal to the historical mean of realized market returns in excess of the contemporaneous risk-free rate. The main choice is how far into the past to go when computing the historical mean. This model is very simple and, as we show in Section 8, quite difficult to improve upon when considering out-of-sample performance measures. The main drawbacks are that it is purely backward looking, and assumes that the future will behave like the past, i.e. it assumes the conditional mean of excess returns is not time-varying, giving very little time-variation in the ERP.

To trace the term structure of the ERP using the historical mean method, we simply use returns computed over different horizons and the corresponding maturity risk-free rate before taking the mean.

Model 1: We compute the historical mean going as far back into the past as the data allows.

Model 2: Same as Model 1 but we compute the mean using the previous 5-years of data only (i.e. we use a backward looking 5-year rolling window).

4.2 Dividend discount models (DDM)

All DDM start with the basic intuition that the value of a stock is determined by no more and no less than the cash flows it produces for its shareholders (Gordon 1962). Today's stock price should then be the sum of all expected future dividends, discounted at an appropriate rate to take into account their riskiness and the time value of money. The formula that reflects this intuition is

$$P_t = E_t \sum_{k=0}^{\infty} \frac{D_{t+k}}{\rho_{t+k}} \quad (3)$$

where E_t is the conditional expectations operator, P_t is the current price of the stock, D_t is the current level of dividends, D_{t+k} is the level of dividends k periods from now, and ρ_{t+k} is the discount rate for time $t+k$. The discount rate can be decomposed into

$$\rho_{t+k} = 1 + R_{t+k}^f + ERP_t(k) \quad (4)$$

When using a DDM, we refer to $ERP_t(k)$ as the *implied* ERP, since we plug in observed or estimated values for the price, dividends and the risk-free rates, and figure out what value of $ERP_t(k)$ makes the right-hand side equal to the left-hand side in equation (3). In this framework, the risk-free rate captures the discounting associated with the time value of money and the ERP captures the discounting associated with the riskiness of the dividends.

DDM are forward looking and are consistent with no arbitrage. In fact, equation (3) is an equilibrium condition that must hold in any bubble-free economy with no arbitrage. Another advantage of DDM is that they are easy to implement. A drawback of DDM is that the results are sensitive to how we measure expectations of future dividends. In addition, ignoring the bubble term, i.e. assuming that

$$\lim_{k \rightarrow \infty} E_t \left[\frac{D_{t+k}}{\rho_{t+k}} \right] = 0$$

may impute a higher ERP whenever a bubble is present but not considered in the model.

Even though DDM do not require the term structure of the ERP to be flat, in practice all DDM assume that $ERP_t(k) = ERP_t(j)$ for all k and j . With a single ERP measure for all horizons, equation (3) pins down the ERP completely, while if we had different ERP estimates for different horizons, equation (3) would become a single equation in several unknowns and the ERP would not be identified.

Model 3: The simplest DDM assumes a constant growth rate of dividends and a flat yield curve in addition to a flat term structure of the ERP (Gordon 1962). Under these assumptions, equation (3) becomes

$$P_t = \sum_{k=0}^{\infty} \frac{D_t(1+g)^k}{(1+R_t^f + ERP_t)^k} = \frac{D_t}{R_t^f + ERP_t - g}$$

Solving for the ERP gives

$$ERP_t = \frac{D_t}{P_t} - (R_t^f - g) \quad (5)$$

Note that even though the term-structure of the ERP is assumed to be flat, this model does not assume a constant ERP or a constant risk-free rate. In practice, there are several ways to operationalize equation (5). Model 3, called³ the “Fed Model”, uses the nominal ten-year Treasury yield as an estimate of $R_t^f - g$ and current earnings E_t as a proxy for current dividends D_t .

Model 4: The “Shiller model”. Same as Model 3 but uses Shiller’s cyclically adjusted price-earnings ratio (CAPE) as a proxy for the price-dividend ratio. CAPE is the current price of the S&P 500 divided by a trailing twelve month average of earnings.

Model 5: Same as Model 3, but uses the real ten-year Treasury yield as an estimate of $R_t^f - g$ (computed as the ten-year nominal Treasury rate minus the ten year breakeven inflation implied by TIPS). There are two typical justifications for this choice. First, in the long run, the growth rate of dividends g should be at least approximately equal to breakeven inflation. Second, the dividend-price ratio D_t/P_t , being the ratio of two nominal variables, is a real variable. Thus, it should be compared to the real risk-free rate and not to the nominal one as in Model 3.

Model 6: Same as Model 3, but uses one-year ahead expected earnings as a proxy for dividends. The usual justification is that including future expectations should better capture the forward-looking nature of the DDM.

Model 7: A variation in the assumptions in Models 3, 4 and 5: it uses the nominal ten-year Treasury yield and one-year ahead expected earnings.

³ The name “Fed Model” was coined by Ed Yardeni, at Deutsche Morgan Grenfell, in reference to a report issued in 1997 by the Federal Reserve that used the model. However, the Federal Reserve has never endorsed this model. See Asness (2003) for a critical view of the “Fed Model”.

Model 8: A two-stage DDM from Panigirtzoglou and Loeyes (2005) where the first stage corresponds to the first five years, and the second stage corresponds to years 6 and onwards. In this case, formula (3) becomes⁴

$$P_t = \frac{D_t[1 + g_{LR} + 5(g_{SR} - g_{LR})]}{R_t^f + ERP_t - g_{LR}}$$

where R_t^f is the ten year nominal Treasury yield; D_t is estimated by the current (observable) level of earnings-per share multiplied by a payout ratio assumed to be 50%; g_{LR} is the long-run estimate for earnings growth and assumed to be 2.2 percent; g_{SR} is the estimated growth rate of earnings over the first five years, which is estimated by using the fitted values in a regression of average realized earnings growth over the last five years on its lag and lagged earnings-price ratio. The main advantage of having two stages instead of a single one (as in Models 3 through 7) is that it allows for changes in the growth rate of dividends, a useful feature when growth rates are far away from their long-run level.

Model 9: A multi-stage DDM constructed by Damodaran (2012). We simplify equation (3) by assuming there are 6 stages. Each of the first five stages corresponds to each of the first five years, while the last stage corresponds to years six and onwards. Dividends are assumed to grow at a rate g_t for each of the first five stages, and then at a rate equal to the ten year nominal Treasury yield for the final stage. The discount rate is assumed to be constant over different horizons, so that $\rho_{t+k} = \rho_t^k$. With these assumptions, equation (3) becomes

$$P_t = \sum_{k=1}^5 \frac{D_t(1 + g_t)^k}{\rho_t^k} + \frac{D_{t+6}(1 + g_t)^6}{(\rho_t - R_t^f)\rho_t^5} \quad (6)$$

where $\rho_t = 1 + R_t^f + ERP_t$ and R_t^f is the ten year nominal Treasury yield. Given P_t, D_t, R_t^f and g_t , equation (6) determines a unique ERP_t .

Model 10: Is the same as Model 9 –and also proposed by Damodaran (2012)— but includes stock buybacks in cash flows. The idea is that investors care about total cash flows, not just dividends, and that buybacks are significant enough to affect measures of the ERP. In practice, we use free-cash-flow-to-equity as a proxy for dividends plus stock buybacks. Damodaran (2012) estimates that buybacks can increase the ERP by one to four percentage points per year.

4.3 Cross-sectional regressions

This method exploits the variation in returns and exposures to the S&P 500 of different assets to infer the ERP.⁵ Intuitively, this method finds the ERP by answering the following question: what is the level of the ERP that makes expected returns of a variety of stocks consistent with their exposure to the S&P 500? Because we need to explain the relationship between returns and exposures for multiple assets with a single value for the ERP (and perhaps a small number of other controls), this model imposes tight restrictions on the estimation of the ERP.

The first step is to find the exposures of assets to the S&P 500 by estimating an equation of the following form:

⁴ For a derivation, see Fuller and Hsia (1984).

⁵ See Polk, Thompson and Vuolteenaho (2006) and Adrian, Crump and Moench (2012) for a detailed description of this method.

$$R_{t+k}^i - R_{t+k}^f = \alpha_{t+k}^i \times \text{state variables}_{t+k} + \beta_{t+k}^i \times \text{risk factors}_{t+k} + \text{idiosyncratic risk}_{t+k}^i \quad (7)$$

In equation (7), R_{t+k}^i is the realized return on a stock or portfolio from time t to $t+k$. *State variables* $_{t+k}$ are any economic indicators that help identify changes in the investment opportunity set (possibly including a constant). *Risk factors* $_{t+k}$ are any measures of systematic contemporaneous co-variation in returns across all stocks or portfolios. Finally, *idiosyncratic risk* $_{t+k}^i$ is the component of returns that is particular to each individual stock or portfolio that is not explained by *State variables* $_{t+k}$ or *Risk factors* $_{t+k}$. Examples of state variables are inflation, unemployment, the term spread, the yield spread between Aaa and Baa bonds and the S&P 500's dividend-to-price ratio. It is crucial that we include the excess return on the S&P 500 as a risk-factor in the estimation so that we can infer the ERP. Other risk-factors usually used are the Fama-French (1992) factors and the momentum factor of Carhart (1997). The value of α_{t+k}^i gives the strength of asset-specific return predictability and β_{t+k}^i gives the asset-specific risk exposures we are trying to estimate. For the cross-section of assets, we can use the whole universe of traded stocks, a subset of them, or portfolios of stocks grouped, for example, by industry, size, book-to-market or recent performance.

The second step is to find the ERP associated with the S&P 500 by estimating the cross-sectional equation

$$R_{t+k}^i - R_{t+k}^f = ERP_t(k) \times \hat{\beta}_{t+k}^i \quad (8)$$

where $\hat{\beta}_{t+k}^i$ are the values found when estimating equation (7). Equation (8) attempts to find the single number $ERP_t(k)$ (or vector of numbers, if we have more than one risk factor) that makes exposures $\hat{\beta}_{t+k}^i$ consistent with realized excess returns of all stocks or portfolios considered. The term structure of the ERP is obtained by computing returns over different horizons on the left hand side of equations (7) and (8).

One advantage of the cross-sectional regression method is that it uses more asset prices than other models, which provide more independent information about the ERP. Cross-sectional regressions also have sound theoretical foundations, since they are one way to implement Merton's (1973) Intertemporal Capital Asset Pricing Model (ICAPM). Finally, this method nests many of the other models considered. The two main drawbacks of this method are that results are dependent on what portfolios, state variables and risk factors are used and that it is not easy to implement.

Model 11: The most widely used cross-sectional model is the Fama-French model (Fama and French 1992). The only state variable is a constant, and there are three risk factors: the returns on the market portfolio, a size portfolio and a book-to-market portfolio. Equation (7) is estimated by running rolling OLS regressions over the previous five years, and equation (8) is estimated by OLS without a constant⁶.

Model 12: Same as Model 11, but includes momentum as an additional risk factor (Carhart 1997).

Model 13: Same as Model 12, but also includes inflation as a risk factor, which has been shown to account for a substantial part of the equity premium beyond the four factors of Carhart's model (Duarte 2013). Additionally, the time-varying coefficients α_{t+k}^i and β_{t+k}^i are estimated with the non-parametric kernel estimator of Ang and Kristensen (2012).

Model 14: This model is from Adrian, Crump, and Moench (2012). The state variables are the dividend yield, the default spread, and the risk free rate, which are commonly thought to capture changes in the

⁶ Using OLS *with* a constant is an equally valid procedure; whether to include a constant depends on the familiar tradeoff between efficiency and robustness (Cochrane 2001).

investment opportunity set. The inclusion of these state variables allows the model to capture dynamics of the pricing kernel not captured by Models 11 through 13. The risk free rate is the one-month Treasury bill rate; the dividend yield is for the S&P500; and the default spread is calculated as the difference between Moody’s seasoned Baa corporate bond yield and the 20-year Treasury bond yield at constant maturity. The market is the single risk factor. The model is estimated using a three step regression approach. First, the market return is orthogonalized with respect to the state variables and the residual of that regression is the considered the risk factor. Then each stock or portfolio’s excess return is regressed on the lagged state variables and the risk factor to obtain the coefficients α_{t+k}^i and β_{t+k}^i . Finally, the ERP is obtained by estimating equation (8) using OLS.

4.4 Time-series regressions

This method uses the relationship between economic variables and stock returns to estimate the ERP. The idea is to run a linear regression of realized excess returns on lagged “fundamentals”:

$$R_{t+k} - R_{t+k}^f = a + b \times \text{Fundamental}_t + \text{error}_t \quad (9)$$

Once estimates \hat{a} and \hat{b} for a and b are obtained, the ERP is obtained by ignoring the error term:

$$\text{ERP}_t(k) = \hat{a} + \hat{b} \times \text{Fundamental}_t \quad (10)$$

In other words, we estimate only the forecastable or expected component of excess returns. This method attempts to implement equations (1) and (2) as directly as possible in equations (9) and (10), with the assumption that “fundamentals” are the right sources of information to look at when computing expected returns and that the conditional expectation is a linear function.

The use of time-series regression requires minimal assumptions; there is no concept of equilibrium and no absence of arbitrage necessary for the method to be valid. In addition, implementation is quite simple, since it involves running univariate OLS regressions. The challenge of this method is to select the variables to include in the right-hand side of equation (9), since results can change substantially depending on what fundamental variables are used. In addition, including more than a single variable gives poor out-of-sample predictions even if economic theory may suggest a role for many variables to be used as predictors⁷. Finally, time-series regressions ignore information in the cross-section of stock returns.

The term structure of the ERP, as equations (9) and (10) suggest, is easily obtained in this method by simply running the predictive regressions with excess returns computed over different horizons.

Model 15: This model uses the dividend-price ratio as the only predictive variable. The key rationale is that the dividend-price ratio is first-order stationary so that it should eventually return to its long-run mean⁸: Values of the dividend-price ratio above its mean should forecast either low returns or high dividends going forward (and vice-versa for low values). Empirically, a high dividend-price ratio forecasts higher returns, not lower dividends, so the price-dividend ratio contains information about the ERP (Cochrane 2011).

Model 16: Same as Model 15, but uses the twelve predictive variables proposed by Goyal and Welch (2008). We use each variable independently and all of them together. At each point in time, we select

⁷ Goyal and Welch (2008).

⁸ See Lettau and Van Nieuwerburgh (2008) for an argument against first-order stationarity and its implications for predictability of returns.

the specification that performs the best out-of-sample (see Section 7 for a detailed description of how we do this) and use that specification for the next period. In the following period, we repeat the procedure; it is possible that this method uses different predictors depending on which one is performing best at each point in time.

Model 17: Same as Model 16, but imposes two restrictions on the estimation. First, the coefficient b in equation (9) is replaced by zero if it has the “wrong” theoretical sign. For example, if the price-dividend ratio has a negative coefficient, then we replace b by zero. Second, we replace the estimate of the ERP by zero if the estimation otherwise finds a negative ERP. These two restrictions are imposed one at a time and then together, and considered for the same twelve predictive variables considered in Model 16. The best specification at time t is used for prediction of $t+k$ returns, so specifications can be changing over time. This model is advocated by Campbell and Thompson (2008), who argue that the restrictions, being based on theory, should improve estimation efficiency compared to unrestricted estimation.

Model 18: Uses as predictors the price-dividend ratio adjusted by the growth rate of earnings RE_t , dividends RD_t or stock prices RP_t :

$$RE_t = \frac{D_t}{P_{t-1}} + \left(\frac{E_t}{D_{t-1}} \right) \frac{CPI_{t-1}}{CPI_t} - 1 \quad (11)$$

$$RD_t = \frac{D_t}{P_{t-1}} + \left(\frac{D_t}{D_{t-1}} \right) \frac{CPI_{t-1}}{CPI_t} - 1 \quad (12)$$

$$RP_t = \frac{D_t}{P_{t-1}} + \left(\frac{P_t}{P_{t-1}} \right) \frac{CPI_{t-1}}{CPI_t} - 1 \quad (13)$$

where D_t are dividends, E_t are earnings, P_{t-1} is the lagged price of the S&P 500 and CPI_t is the consumer price index. We also consider the three measures constructed above, but subtracting the ten year nominal Treasury yield from each of them. The idea behind these measures is to impose—rather than assume—that stationary variables must eventually return to their long-run mean. As in models 16 and 17, at time t we use the predictor that has the best out-of-sample performance until $t - 1$, which leads to different measures being used at different points in time. This model was proposed by Fama and French (2002) who argue that stock returns have been too high compared to dividend or earnings growth, and therefore must have been in part due to luck (positive shocks). A way to account for this sample-specific realization is to “correct” the dividend-price ratio as in equations (11), (12) and (13).

Model 19: The predictor is Baker and Wurgler’s (2007) sentiment measure. The measure is constructed by finding the most predictive linear combination of five variables: the closed-end fund discount, NYSE share turnover, the number and average first-day returns on IPOs, the equity share in new issues, and the dividend premium. Baker and Wurgler (2007) have a more detailed explanation.

4.5 Surveys

The survey approach consists in asking economic agents what they think the ERP is. Surveys incorporate the views of many people, some of which are very sophisticated and/or make real investment decisions based on the level of the ERP. Surveys should also be good forecasters of the ERP because in principle stock prices are determined by supply and demand of investors such as the ones taking the surveys. On the other hand, Greenwood and Shleifer (2012) document that investor

expectations of future stock market returns are positively correlated with past stock returns and with the current level of the stock market, but strongly *negatively* correlated with model-based expected returns and future realized stock market returns. Other studies such as Easton and Sommers (2007) also argue that survey measures of the ERP can be systematically biased.

The term structure of the ERP can only be obtained from surveys to the extent that questions are asked about the ERP at different horizons. To the authors' knowledge, the only consistent survey with a long enough time-series for analysis that asks about point estimates of the ERP at different horizons is the Duke CFO survey by Graham and Harvey (2012), which we use for the next model.

Model 20: Chief financial officers (CFOs) are asked about the one and ten-year-ahead ERP. A typical question in the survey is the following:

*On November 19, 2007 the annual yield on 10-yr treasury bonds was 4.1 percent.
Please complete the following: Over the next year, I expect the average annual S&P
500 return will be:*

The survey has grown over time and now has around 600 respondents. We take the mean of all responses as our measure of the ERP⁹. We construct the term-structure of the ERP by linearly interpolating the one and ten-year-ahead ERP estimates given by respondents. For this model, we do not construct ERP measures for horizons shorter than a year.

5. The Equity Risk Premium has Reached a Historic High

We summarize the behavior of the twenty models we consider by their first principal component. Let X be the matrix containing the demeaned ERP estimates at a monthly horizon from the different models we consider, with columns corresponding to models and rows corresponding to observation periods. The matrix is a 643-by-20 matrix, since we have 643 monthly observations and 20 models. The first principal component is the eigenvector of the variance-covariance matrix of X associated with the largest eigenvalue. Because X was demeaned, this principal component has mean zero. We take as our preferred ERP estimate the sum of the first principal component and the unconditional mean of ERP estimates across all models (i.e. the average of all elements of X). We repeat this process using ERP estimates at different horizons to obtain a single ERP time-series for each horizon. We call these estimates our *preferred measures*. The share of the variance explained by these measures ranges between 81% and 94%, suggesting that they are good summary statistics for the behavior of the models.

One challenge that arises in computing the principal component is that the matrix X has missing observations, either because some models can only be obtained at frequencies lower than monthly or because the necessary data is not available for all time periods (Appendix A contains a detailed description of when this happens). To overcome this challenge, we use an iterative linear projection method¹⁰. On the first iteration, we make a guess for the principal component and regress the non-missing elements of each row of X on the guess and a constant. We then find the first principal component of the variance-covariance matrix of the fitted values of these regressions, and use it as the guess for the next iteration. The process ends when the norm of the difference between consecutive estimates is small enough.

Figure 1 displays our preferred measures for the one-month and one-year-ahead ERP in blue and red, respectively. Recessions are indicated by shaded bars. The correlation between the two measures is

⁹ Taking the median does not substantially alter results.

¹⁰ We thank Richard Crump for suggesting this method and providing code for its implementation.

86%, but we do see that the one-month ahead ERP is sometimes above and sometimes below the one-year-ahead ERP, indicating that the slope of the term structure of the ERP is time-varying. As expected, the ERP measures tend to peak during financial turmoil, recessions and periods of low real GDP growth or high inflation. The ERP tends to bottom out after periods of sustained bullish stock markets and high real GDP growth.

The one year ahead ERP is at 14.5 percent in July of 2013, the highest it has even been. The one month ahead ERP is at 11.5 percent in July 2013, nearing the record levels obtained in February 2009, July 2012 and the early 1980s, but still below the peak of 15 percent in September of 1974.

The current high levels of the ERP are unusual in that we are not currently in a recession and we have just experienced an extended period of high stock returns, with 60 percent returns since July 2010 and almost 20 percent since the beginning of 2013. During previous periods, the ERP has always decreased during periods of sustained high realized returns. This is also the only period in which the ERP is elevated and the one-year ahead ERP is significantly higher than the one month ahead ERP.

Figure 2 displays in red the standard deviation of one-year-ahead ERP estimates across models for each time period. The standard deviation has been steadily decreasing since 2000 except for a few months during the financial crisis and has reached an all-time low in the last three months. A low standard deviation can be interpreted as models displaying a high degree of agreement – in this case, agreement that the ERP is high. Figure 2 also shows the reason for the recent increase in agreement and in the ERP by plotting the 25th and 75th percentiles of the distribution of models in blue and green, respectively. The interquartile range –the difference between the 25th and 75th percentiles— has compressed, mostly because the models in the bottom of the distribution have had higher ERP estimates since 2010. It is also interesting to note that the 75th percentile has remained fairly constant over the last 10 years, and is actually somewhat below its long-run mean.

6. The Term Structure of Equity Risk Premia

In Section 4, we described how each of the different models can trace out a term structure of the ERP – what expected excess returns are over different time horizons¹¹. Figure 3 plots our preferred ERP measures as a function of investment horizon (rather than time) for some selected dates. The black line shows the average of the term structure across all periods. It is slightly upward sloping, with a short-term ERP at just over 6% and a three-year ERP at almost 7%. We selected the other dates because they are typical dates for when the ERP was unusually high or unusually low at the one-month horizon. We see that the ERP is strongly mean-reverting, with the term structure sloping downward for high one-month ERP periods, and sloping upward for low one-month ERP periods. In contrast, the ERP in July 2013 is upward sloping, something that has never happened in periods of elevated ERP.

7. Why is the Equity Risk Premium High?

The last two sections showed evidence that the ERP is high at all horizons, and that this is an unusual occurrence given the current economic and financial environment. There are two reasons why the ERP can be high: low discount rates and high current or expected future cash flows.

¹¹ For other ways to estimate the term structure of the ERP using equilibrium models or derivatives, see Ait-Sahalia, Karaman and Mancini (2012), Ang and Ulrich (2012), Berg (2013), Boguth, Carlson, Fisher and Simutin (2012), Croce, Lettau and Ludvigson (2012), Lemke and Werner (2009), Lettau and Wachter (2011), Muir (2013), among others.

Figure 4 shows that earnings are likely not the reason why the ERP is high. The blue line shows the realized monthly growth rates of real earnings for the S&P500 expressed in annualized percentage points. Since 2010, earnings growth has been declining, hovering around zero for the last few months of the sample. It currently stands at 2.5%, which is near its long-run average. Perhaps more importantly for the equity premium, the expectations of future earnings growth since 2010 have also been moderate to low. The red line in Figure 4 shows the year-on-year change in the mean expectation of one-year-ahead earnings per share for the S&P500. Similarly to realized earnings growth, earnings per share have been declining over the last three years, making expected earnings growth an unlikely reason for why the ERP is near its all-time high.

Nominal and real bond yields, on the other hand, have been exceptionally low since the end of the financial crisis. Figure 5 displays the term structure of the ERP under two counterfactual scenarios, in addition to the mean and current term structures already displayed in Figure 4. In the first counterfactual scenario, we leave expected stock returns unmodified but change the risk-free rates from the current values of nominal bond yields to the average nominal bond yields over 1960-2013. In other words, we replace R_{t+k}^f in equation (2) by the mean of R_{t+k}^f over t . The result of this counterfactual is shown in Figure 5 in orange. Using average levels of bond yields brings the whole term structure of the ERP much closer to its mean level (the black line), especially at short horizons. This shows that a “normalization” of bond yields, everything else being equal, would bring the ERP down substantially. In our second counterfactual exercise, we do not keep expected stock returns unchanged, but instead estimate the following regression:

$$\Delta ERP_t(k) = a(k) + b(k) \times \Delta y_t^1 + c(k) \times \Delta y_t^5 + d(k) \times \Delta y_t^{10} + error_t(k) \quad (14)$$

where y_t^1 , y_t^5 and y_t^{10} are nominal yields for one, five and ten-year constant maturity bonds, and Δ is the first-difference operator, i.e. $\Delta x_t = x_t - x_{t-1}$ for any variable x . Equation (14) can be thought of as regressing $\Delta ERP_t(k)$ on basic level, slope and curvature factors of the nominal yield curve since these factors are linear combinations of the three bond yields y_t^1 , y_t^5 and y_t^{10} . We chose to run regression (14) in differences to avoid spurious regression bias, since bond yields and the ERP are persistent variables¹². We then add to the current term structure of the ERP the fitted values of regression (14) that result from plugging in the values of Δy_t^1 , Δy_t^5 and Δy_t^{10} that would bring bond yields from their current levels to their historical levels:

$$\Delta y_t^j = \bar{y}^j - y_{July\ 2013}^j$$

for $j = 1, 5$ and 10 years, where \bar{y}^j is the time-average of y_t^j . The resulting counterfactual term structure of the ERP is shown in green in Figure 5. Unlike the case in which expected returns were held constant, this counterfactual assumes that expected returns respond to changes in yields in the same way that they have responded in the past. The resulting counterfactual term structure of the ERP is now flat and substantially below its average value. This means that if yields increased to their average levels and expected returns reacted to this increase as they have in the past, the ERP would decrease below its average levels at all horizons. This exercise shows that the current environment of exceptionally low bond yields is capable, quantitatively speaking, of causing an ERP as high as we are currently observing.

¹² An augmented Dickey-Fuller test fails to reject a unit root in the EPR at the 5% level using any number of lags between 5 and 18 (the maximum lag chosen by the Schwert criterion). The tau test statistic using 15 lags (the optimal number of lags obtained by the Ng-Perron procedure) is -1.96, which is smaller than the critical value of -2.83. A similar analysis for bond yields also fails to reject the null hypothesis of a unit root.

8. Are excess returns predictable?

In this section, we analyze how ERP models perform when trying to predict the realized equity premium. There is substantial debate in the academic literature¹³ on whether any model can explain or predict the ERP better than the historical mean of realized excess returns. For this reason, we choose performance metrics that are relative to the historical mean. The historical mean itself is only a weak explanatory and predictive variable. Its correlation with the realized ERP is about 6 percent, and the R^2 in one to sixty month predictive regressions is less than 1 percent.

The measure we use for how well a model predicts the ERP is the out-of-sample R^2 , popularized by Campbell and Thompson (2008):

$$R_{OOS}^2 = 100 \times \left(1 - \frac{\sum_{t=1}^T (R_t^e - Model\ ERP_t)^2}{\sum_{t=1}^T (R_t^e - Historical\ mean_t)^2} \right) \quad (15)$$

Here R_t^e are the realized excess returns at time t , $Model\ ERP_t$ is the time t real-time estimate of the ERP given by some model and $Historical\ mean_t$ is the real-time mean of the ERP since the beginning of the sample. Because $Model\ ERP_t$ is computed using information available at $t - 1$ but R_t^e is only realized at time t , we interpret R_{OOS}^2 as a measure of how well the model predicts the ERP compared to the historical mean.

The out-of sample R^2 in equation (15) ranges from minus infinity to +100 percent. If the R^2 is 0, the historical mean and the model in question perform equally well, whereas a positive R^2 implies that the model outperforms the historical mean. Note that this measure is not a traditional R^2 , so its units cannot be interpreted as the percentage of the variance of the dependent variable explained by the model. However, the R^2 numbers have an intuitive economic interpretation. A mean-variance investor with coefficient of relative risk aversion equal to γ , when using the $Model\ ERP_t$ as the measure of expected excess returns, will earn returns over the whole sample in excess of those predicted by the historical mean equal to R_{OOS}^2/γ .¹⁴ For example, if $\gamma = 2$, and $R_{OOS}^2 = 10$ percent, then the investor can, though better predictability, earn extra returns of 5 percent over the T periods considered. Of course, if R_{OOS}^2 is negative, using the model would lead to returns that are lower than those obtained by using the historical mean as the sole estimate for the ERP.

Figures 6 and 7 show the results. In Figure 6, we show the R_{OOS}^2 for the models that perform the best within each category. For DDM, the Shiller model (model 4) outperformed all other DDM at all horizons. For cross-sectional regressions, the model by Adrian, Crump and Moench (model 14) did best at the one month to two year horizons, but was outperformed by Fama-French and momentum (model 12) at the three and four year horizons, and by the cross-sectional regression with inflation (model 13) at the five year horizon. For time-series regressions, the results were mixed: The Goyal and Welch predictors (model 16) were the best at the one month and three year horizons; The predictors in Fama and French (model 18) were the best at the four and five year horizons; The dividend-price ratio (model 15) was the best at the other horizons. For surveys, all results correspond to CFO surveys (model 20), which is the only survey we analyze.

The main conclusion is that the R-squares are small, which means that none of the models drastically outperform the historical mean out of sample. For example, cross-sectional regressions (the green line with crosses) starts at almost 12%, which means that the mean-variance investor with coefficient of

¹³ See, for example, Goyal and Welch (2008), Lettau and Van Nieuwerburgh (2008), Campbell and Thompson (2008), Cochrane (2011), Fama and French (2002) and references therein.

¹⁴ For a derivation of this fact, see Campbell and Thompson (2008).

relative risk aversion equal to one would have made an extra 12% over the last 53 years compared to just using the historical mean as a measure of the ERP. This amounts to about 23 basis points per year. In addition, although surveys are clearly inferior predictors, all other predictors are comparable at all horizons.

Figure 7 displays the same analysis but using principal components instead of the individual models for each category. The principal components were computed in real time following the procedure explained in Section 4. The predictability of the principal component need not be better or worse than the best model in each category. On the one hand, the principal component may reduce noise and aggregate useful information from the many models. On the other hand, it puts some weight on models that have worse predictability than the best models. Figure 7 shows that, as a group, dividend discount models perform substantially better than other models at short horizons but are worst at long horizons, while cross-sectional regressions perform best at long horizons but are worst at short horizons. The principal component of all models, in the solid purple line, performs well across all horizons and is always close to the model with the best predictability. The good performance of the principal component reinforces its usefulness as a summary statistic.

9. Conclusion

Estimates for the ERP as high as we have found should give policymakers pause. We have argued that it is unusual for the ERP to be at its present level in the current stage of the business cycle, especially when expectations are that it will continue to rise over the next three years. Because the ERP is a key input in many important decisions of economic agents, an unusually high ERP can herald unusual behavior. Our analysis provides evidence that is consistent with a bond-driven ERP: expected excess stock returns are high not because stocks are expected to have high returns, but because bond yields are exceptionally low. In such an environment, we should expect monetary policy—both conventional and unconventional—to have a large impact on asset prices and hence the real economy.

Our study of the ERP has many limitations. The main one is that stocks returns are very difficult to predict, if they are predictable at all. We have shown how to improve upon current estimates by using principal components yet still found weak evidence in favor of predictability, at least at horizons shorter than five years. Any conclusions that rely on ERP estimates must be weighted by how strongly it predicts future returns. Another limitation is that even though we have conducted all of our out-of-sample tests in real time (using information available at time t for $t + 1$ estimates), some of the models we use had not been yet proposed for many periods of our sample, so there is some selection and forward looking biases. Finally, we have not focused on the possibility that a bubble—rational or irrational—could be a further driver of the recent high realized and expected returns, a topic we consider outside of the scope of the broadly used models we consider.

References

- Adrian, T., Crump, R., Moench E., 2012. Efficient regression-based estimation of dynamic asset pricing models. Federal Reserve Bank of New York Staff Reports, Number 439, May 2011.
- Ait-Sahalia, Y., Karaman, M., Mancini, L., 2012. The term structure of variance swaps, risk premia and the expectation hypothesis. Risk Premia and the Expectation Hypothesis (August 24, 2012).
- Ang, A., Kristensen, D., 2012. Testing conditional factor models. NBER Working Paper 17561.
- Ang, A., Ulrich, M., 2012. Nominal Bonds, Real Bonds, and Equity.
- Asness, C., 2003. Fight the Fed model. Journal of Portfolio Management (Fall).11-24.
- Berg, T., 2013. The Term Structure of Risk Premia: Evidence from CDS Spreads.
- Bloom, N., 2009. The impact of uncertainty shocks. Econometrica, . 77 (3), 623-685.
- Boguth, O., Carlson, M., Fisher, A. J., Simutin, M., 2012. Leverage and the limits of arbitrage pricing: Implications for dividend strips and the term structure of equity risk premia. manuscript, Arizona State University.
- Campbell, J.Y., Thompson, S., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies 21 (4), 1509-1531.
- Carhart, M.M., 1997. On persistence in mutual fund performance. Journal of Finance 52, 57–82.
- Cochrane, J. H., 2001. Asset pricing. Princeton University Press, New Jersey.
- Cochrane, J.H., 2011. Presidential address: discount rates. Journal of Finance 66 (4), 1047-1108.
- Croce, M., Lettau, M., Ludvigson, S., 2012. Investor information, long-run risk, and the term structure of equity. Unpublished paper, New York University.[40, 49].
- Duarte, F., 2013. Inflation and the cross-section of stock returns. Federal Reserve Bank of New York Staff Reports, Number 621.
- Duarte, F., Kogan, L., Livdan, D., 2013. Aggregate Investment and Stock Returns. Unpublished paper, M.I.T.
- Damodaran, A., 2012. Equity Risk Premiums (ERP): Determinants, Estimation and Implications – The 2012 Edition. Working paper New York University.
- Easton, P.D., Sommers, G.A., 2007. Effect of analysts' optimism on estimates of the expected return implied by earnings forecasts. Journal of Accounting Research 45, 983-1016.
- Fama, E.F., French, K.R., 1992. The cross-section of expected stock returns. Journal of Finance 47 (2), 427-465.
- Fama, E.F., French, K.R., 2002. The equity premium, Journal of Finance, 57 (2), 637-659

- Fuller, R., Hsia, C., 1984. A Simplified Common Stock Valuation Model, *Financial Analysts Journal*, September-October, 1984.
- Gordon, M. J., 1962. *The investment, financing, and valuation of the corporation*. Greenwood Press.
- Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium predictions. *Review of Financial Studies* 21 (4), 1455-1508.
- Graham, J., Harvey, C., 2012. *The Equity Risk Premium in 2012*. Working paper, Duke University.
- Greenwood, R., Shleifer, A., 2012. *Expectations of returns and expected returns*. Working paper, Harvard University.
- Gürkaynak, R. S., Sack, B., & Wright, J. H., 2007. The US Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8), 2291-2304.
- Hall, R. E., 2013. *The Routes into and out of the Zero Lower Bound*. Working paper, Stanford University.
- Kuehn, L. A., Petrosky-Nadeau, N., & Zhang, L., 2012. *An equilibrium asset pricing model with labor market search* (No. w17742). National Bureau of Economic Research
- Lettau, M., Van Nieuwerburgh, S., 2008. Reconciling the return predictability evidence. *The Review of Financial Studies* 21 (4), 1607-1652.
- Lettau, M., Wachter, J. A., 2011. The term structures of equity and interest rates. *Journal of Financial Economics*, 101(1), 90-113.
- Ludvigson, S.C., Lettau, M., 2001. Consumption, aggregate wealth, and expected stock returns. *Journal of Finance* 56 (3), 815-849.
- Merton, R.C., 1973. An intertemporal capital asset pricing model. *Econometrica* 41 (5), 867-887.
- Muir, T., 2013. *Financial Crises, Risk Premia, and the Term Structure of Risky Assets*.
- Panigirtzoglou, N., Loeys, J., 2005. *A fair value model for US bonds, credit and equities*. JP Morgan, June 2005.
- Polk, C., Thompson, S., Vuolteenaho, T., 2006. Cross-sectional forecasts of the equity premium. *Journal of Financial Economics* 81, 101-141.

Appendix A: Data Variables

Variable	Source	Original Frequency	First period	Last Period
Sentiment measure from "Investor Sentiment in the Stock Market," Journal of Economic Perspectives, 2007	Jeffrey Wurgler	monthly	07/01/65	12/01/10
Equity issuance	Jeffrey Wurgler	monthly	01/01/58	12/01/10
Debt issuance	Jeffrey Wurgler	monthly	01/01/58	12/01/10
CAY	Martin Lettau	quarterly	03/01/52	09/01/12
One-year ahead ERP from CFO survey	Duke CFO Survey	quarterly	06/06/00	06/05/13
Ten-year ahead ERP from CFO survey	Duke CFO Survey	quarterly	06/06/00	06/05/13
Book value per share	Compustat	annual	12/31/77	12/31/12
S&P 500 closing price	Compustat	monthly	02/28/62	06/30/13
ERP from a dividend discount model	Damodaran	annual	12/01/61	12/01/12
ERP from a dividend discount model using free cash flow	Damodaran	annual	12/01/61	12/01/12
Size and book-to-market sorted portfolios	Fama-French	monthly	07/01/26	06/01/13
Realized excess returns for the market	Fama-French	monthly	07/01/26	06/01/13
Size factor	Fama-French	monthly	07/01/26	06/01/13
Book-to-market factor	Fama-French	monthly	07/01/26	06/01/13
Risk free rate	Fama-French	monthly	07/01/26	06/01/13
Baa minus Aaa bond yield spread	FRED	monthly	01/01/19	07/01/13
NBER recession indicator	FRED	monthly	01/02/00	06/01/13
Momentum portfolios	Fama-French	monthly	01/01/27	12/01/12
Momentum factor	Fama-French	monthly	01/01/27	06/01/13
ERP as constructed in Adrian, Crump and Moench (2013)	NY Fed	monthly	01/01/63	07/01/13
Nominal price for the S&P 500	Shiller	monthly	01/02/00	07/01/13
Nominal dividends for the S&P 500	Shiller	monthly	01/02/00	06/01/13
Nominal earnings for the S&P 500	Shiller	monthly	01/02/00	03/01/13
Consumer Price Index	Shiller	monthly	01/02/00	07/01/13
10 year nominal treasury yield	Shiller	monthly	01/02/00	07/01/13
Real price for the S&P 500	Shiller	monthly	01/02/00	07/01/13
Real dividends for the S&P 500	Shiller	monthly	01/02/00	06/01/13
Real earnings for the S&P 500	Shiller	monthly	01/02/00	03/01/13
Cyclically Adjusted Price-Earnings ratio	Shiller	monthly	01/02/00	07/01/13
Realized Earnings per Share for the S&P 500	Thomson Reuters I/B/E/S	annual	12/31/81	12/31/12
Mean analyst forecast for Earnings per share for the S&P500	Thomson Reuters I/B/E/S	monthly	01/14/82	04/18/13
Yield on TIPS	Fed Board	monthly	01/01/03	07/01/13
Nominal yields	Fed Board (Gurkaynak, Sack and Wright)	daily	06/14/61	08/12/13

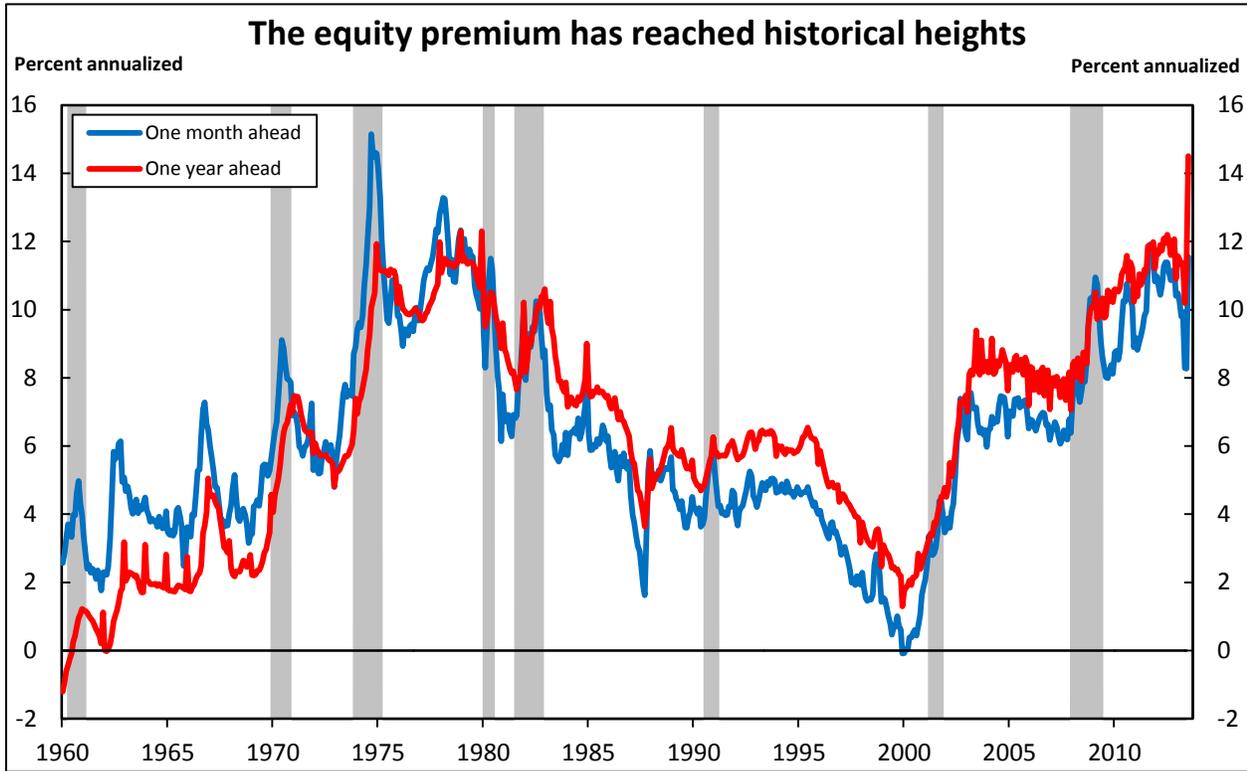


Figure 1 The equity risk premium (expected excess returns) over a one year ahead and one month ahead horizons are the first principal components of 20 models of the equity premium. The models include time-series and cross-sectional regressions, dividend discount models and surveys. Shaded bars are NBER recessions.

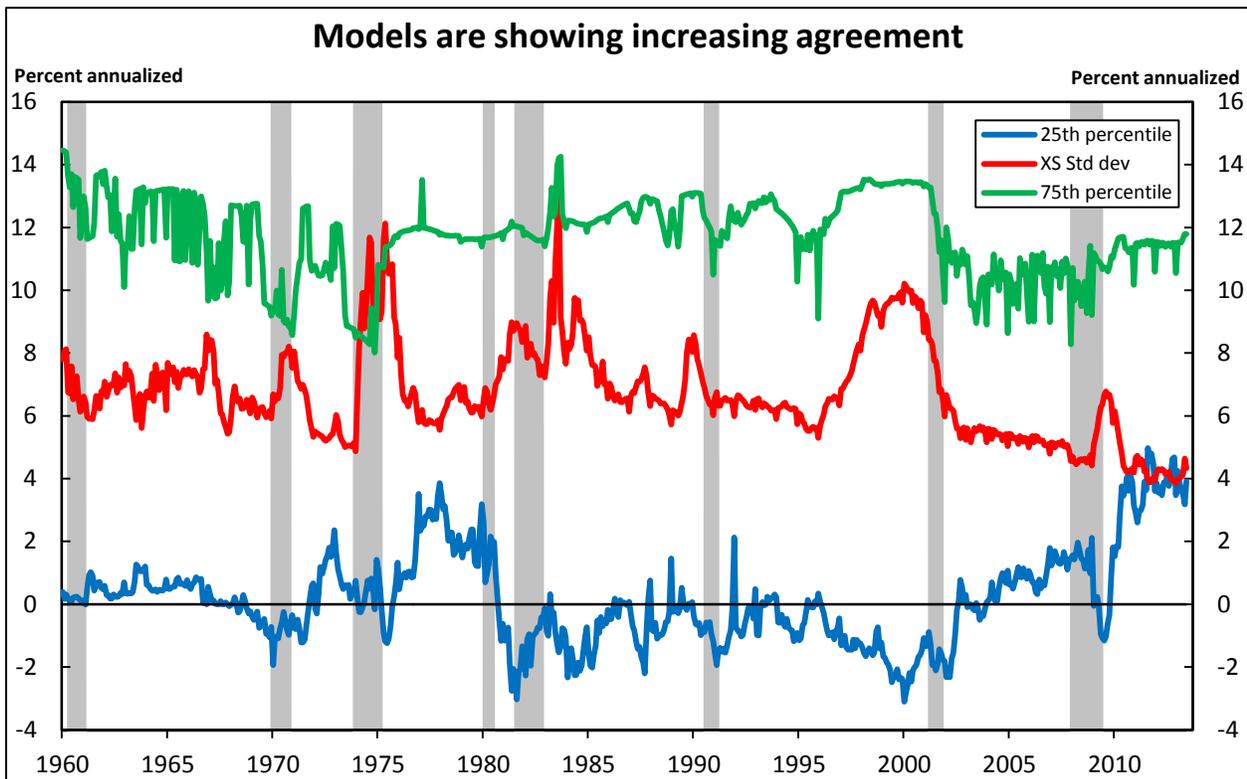


Figure 2 The cross-sectional standard deviation (labeled “XS Std dev”, in red) computes, at each time period, the standard deviation of the 20 equity risk premium estimates given by the different models. The 25th and 75th percentiles (in blue and green, respectively) give the corresponding quartile of the 20 estimates for each time period. Shaded bars are NBER recessions.

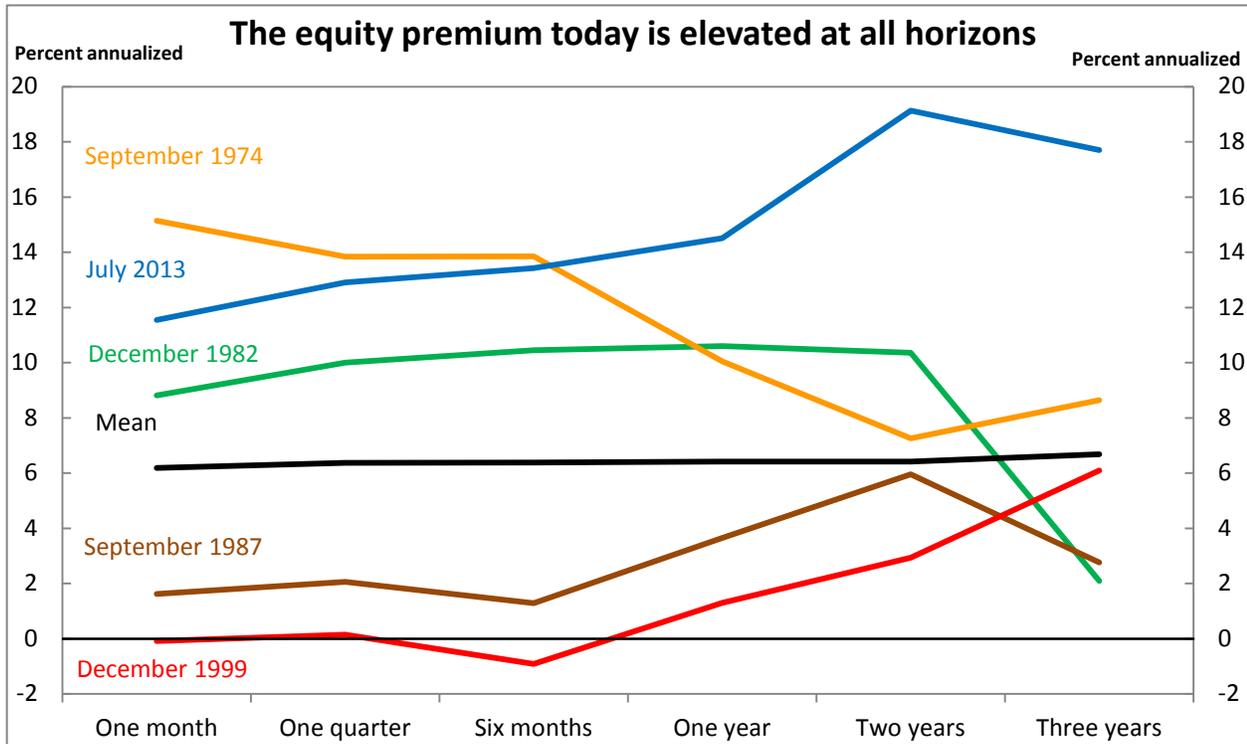


Figure 3 The equity risk premium at different horizons are the first principal component of 20 estimates of expected excess returns at different horizons. The estimates are obtained from cross-sectional and time-series regressions, dividend discount models and surveys. The black line (labeled “Mean”) shows the mean of expected excess returns at different horizons over the sample 1960-2013. The most recent estimates of the term structure of the equity risk premium (labeled “July 2013” in blue), does not show mean reversion, unlike other periods when the equity risk premium was substantially above or below its mean at the one-month horizon.

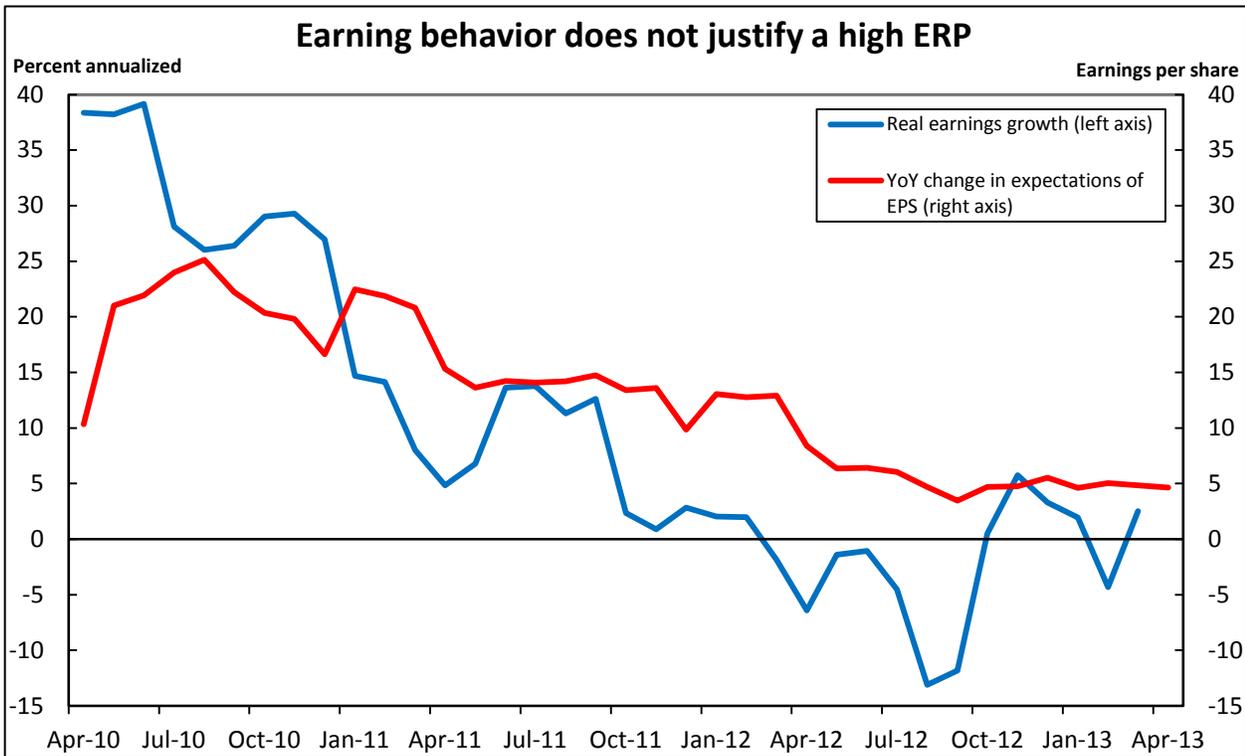


Figure 4 The blue line shows the monthly growth rate of real S&P 500 earnings, annualized and in percentage points. The red line shows the year-on-year change in the mean expectation of one-year ahead earnings per share for the S&P 500 from analysts from Thomson Reuters I/B/E/S.

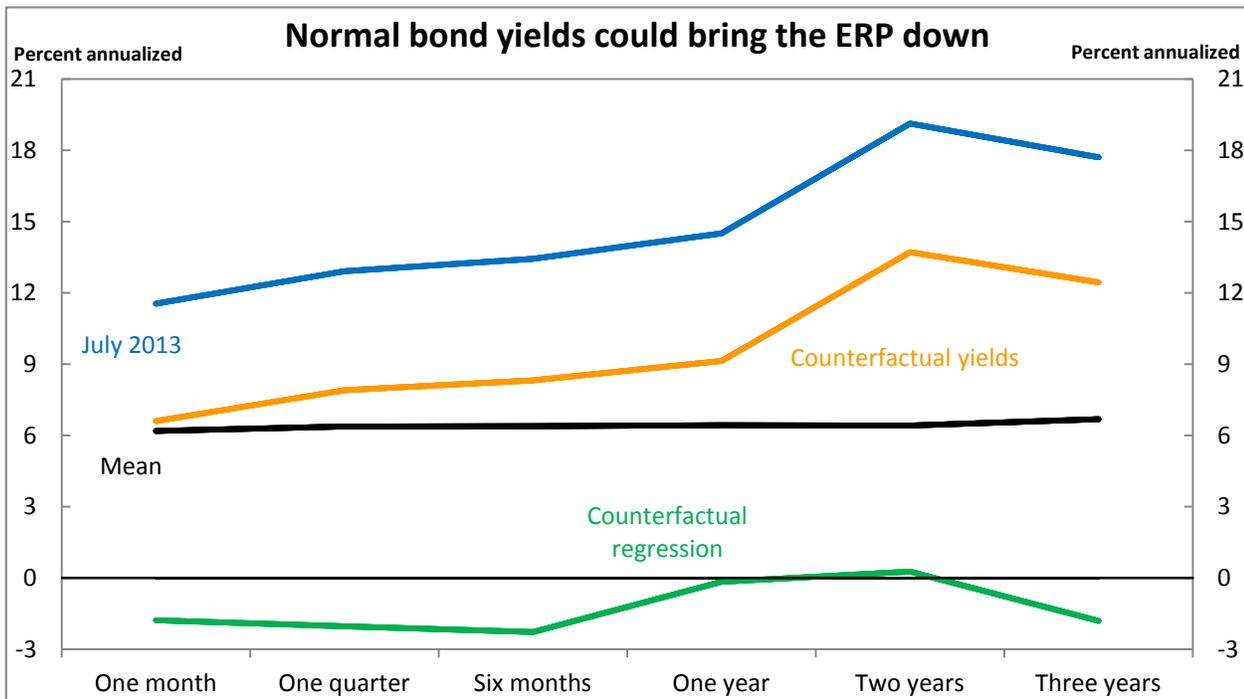


Figure 5 The black line (labeled “Mean”) shows the mean term structure of the equity risk premium over the sample 1960-2013. The blue line (labeled “July 2013”) shows the most recent estimates. The orange line (labeled “Counterfactual yields”) shows what the term structure of the equity risk premium would be in July 2013 if instead of subtracting today’s yield curve from expected returns we subtracted the average yield curve for the period 1960-2013. The green line shows an estimate of what the term structure of the equity premium would be if yields rose to their average historical levels and expected stock returns co-moved with yields with the same correlation as during 1960-2013.

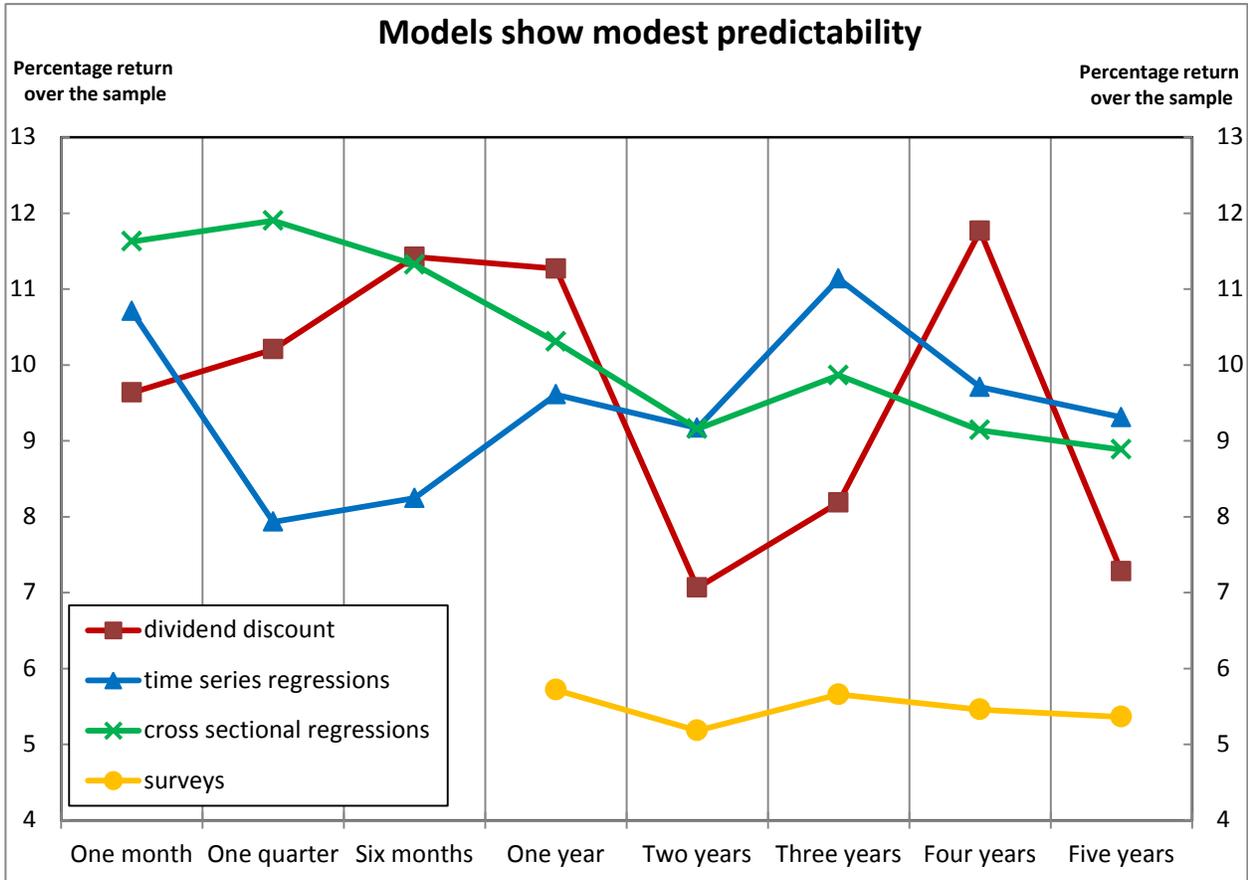


Figure 6 Each data point corresponds to the returns that a mean-variance investor with unit coefficient of relative risk aversion would have earned over the period 1960-2013 if she had used one of the equity risk premium models over and above the returns she would have made if she had assumed that expected excess returns are equal to their historical mean at each point in time. The x-axis shows the investment horizon of the investor (how often the portfolio is rebalanced and hence how far ahead excess returns must be forecast).

For each class of model (dividend discount, time-series regressions, cross-sectional regressions, surveys) we report the model that had the best predictability. For dividend discount models, the Shiller model (model 4) outperformed all other discount models at all horizons. For cross-sectional regressions, the model by Adrian, Crump and Moench (model 14) did best at the one month to two year horizon, but was outperformed by Fama-French and momentum (model 12) at the three and four year horizons, and by the cross-section with inflation (model 13) at the five year horizon. For time-series regressions, the results were mixed: The Goyal and Welch predictors (model 16) were the best at a one month and three year horizons; The predictors in Fama and French (model 18) were the best at the four and five year horizons; The dividend-price ratio (model 15) was the best at the other horizons. For surveys, all results correspond to CFO surveys (model 20), which is the only survey we analyze.

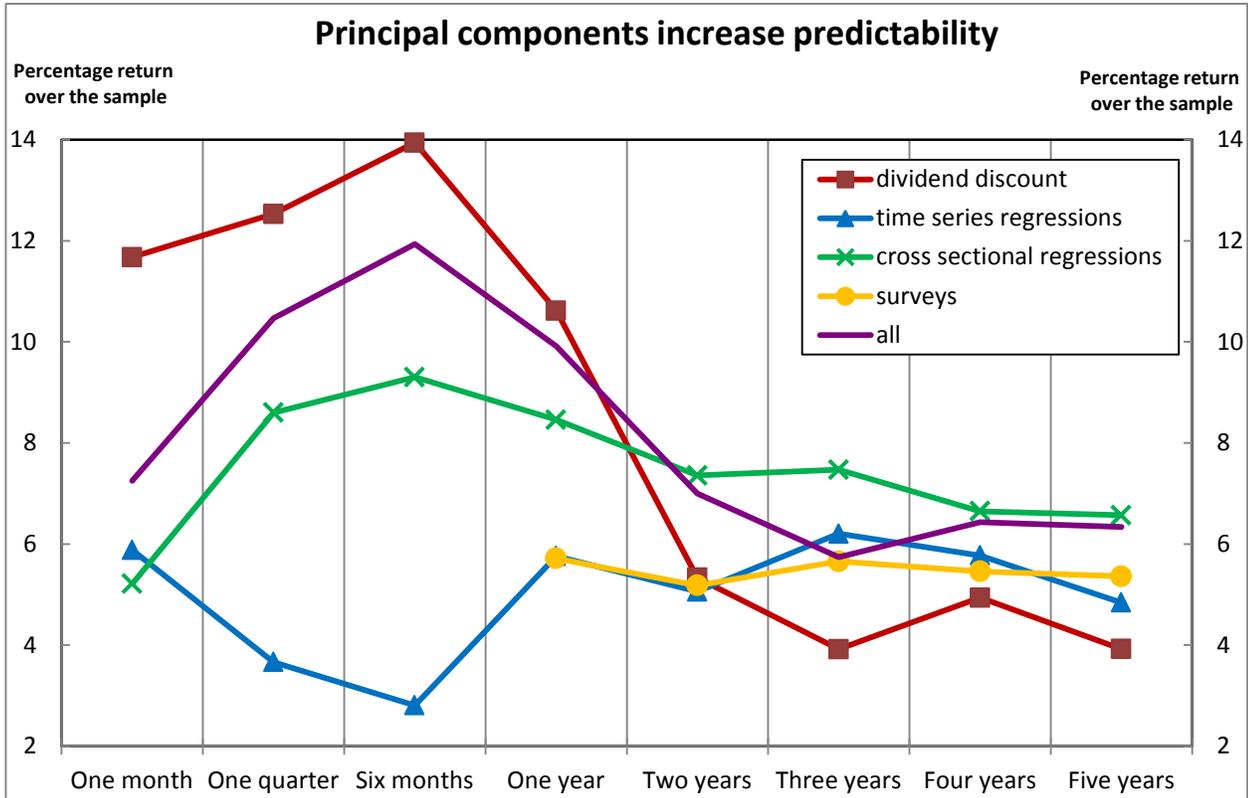


Figure 7 Each data point corresponds to the returns that a mean-variance investor with unit coefficient of relative risk aversion would have earned over the period 1960-2013 if she had used the first principal component of all models within a certain class (dividend discount, time-series regressions, cross-sectional regressions, surveys) over and above the returns she would have made if she had assumed that expected excess returns are equal to their historical mean at each point in time. The x-axis shows the investment horizon of the investor (how often the portfolio is rebalanced and hence how far ahead excess returns must be forecast). The line labeled “all” corresponds to the principal component of all models (our preferred measure).