

Duarte, Fernando; Rosa, Carlo

Working Paper

The equity risk premium: A review of models

Staff Report, No. 714

Provided in Cooperation with:
Federal Reserve Bank of New York

Suggested Citation: Duarte, Fernando; Rosa, Carlo (2015) : The equity risk premium: A review of models, Staff Report, No. 714, Federal Reserve Bank of New York, New York, NY

This Version is available at:
<https://hdl.handle.net/10419/120825>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

The Equity Risk Premium: A Review of Models

Fernando Duarte
Carlo Rosa

Staff Report No. 714
February 2015



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Equity Risk Premium: A Review of Models

Fernando Duarte and Carlo Rosa

Federal Reserve Bank of New York Staff Reports, no. 714

February 2015

JEL classification: C58, G00, G12, G17

Abstract

We estimate the equity risk premium (ERP) by combining information from twenty models. The ERP in 2012 and 2013 reached heightened levels—of around 12 percent—not seen since the 1970s. We conclude that the high ERP was caused by unusually low Treasury yields.

Key words: equity premium, stock returns

Duarte, Rosa: Federal Reserve Bank of New York (e-mail: fernando.duarte@ny.frb.org, carlo.rosa@ny.frb.org). The authors thank Tobias Adrian and James Egelhof for helpful comments on earlier drafts. This article is an update of, and a more comprehensive and rigorous treatment than, our blog post in *Liberty Street Economics* (<http://libertystreeteconomics.newyorkfed.org/2013/05/are-stocks-cheap-a-review-of-the-evidence.html>). The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1. Introduction

The equity risk premium —the expected return on 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, savings decisions of individuals and budgeting plans for governments. Recently, the equity risk premium (ERP) has also returned to the forefront as a leading indicator of the evolution of the economy, a potential explanation for jobless recoveries and a gauge of financial stability³.

In this article, we estimate the ERP by combining information from twenty prominent models used by practitioners and featured in the academic literature. Our main finding is that the ERP has reached heightened levels. The first principal component of all models —a linear combination that explains as much of the variance of the underlying data as possible— places the one-year-ahead ERP in June 2012 at 12.2 percent, above the 10.5 percent that was reached during the financial crisis in 2009 and at levels similar to those in the mid and late 1970s. Since June 2012 and until the end of our sample in June 2013, the ERP has remained little changed, despite substantial positive realized returns. It is worth keeping in mind, however, that there is considerable uncertainty around these estimates. In fact, the issue of whether stock returns are predictable is still an active area of research.⁴ Nevertheless, we find that the dispersion in estimates across models, while quite large, has been shrinking, potentially signaling increased agreement

³ 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. See, for example, Piazzesi and Schneider (2007), Stock and Watson (2003), and Damodaran (2012). Bloom (2009) and Duarte, Kogan and Livdan (2013) study connections between the ERP and real aggregate investment. As a potential explanation of the jobless recovery, Hall (2014) and Kuehn, Petrosky-Nadeau and Zhang (2012) propose that increased risk-aversion has prevented firms from hiring as much as would be expected in the post-crisis macroeconomic environment. Among many others, Adrian, Covitz and Liang (2013) analyze the role of equity and other asset prices in monitoring financial stability.

⁴ A few important references among a vast literature are Ang and Bekaert (2007), Goyal and Welch (2008), Campbell and Thompson (2008), Kelly and Pruitt (2013), Chen, Da and Zhao (2013), Neely, Rapach, Tu and Zhou (2014).

even when the models are substantially different from each other and use more than one hundred different economic variables.

In addition to estimating the level of the ERP, we investigate the reasons behind its recent behavior. Because the ERP is the difference between expected stock returns and the risk-free rate, a high estimate can be due to expected stock returns being high or risk-free rates being low. We conclude the ERP is high because Treasury yields are unusually low. Current and expected future dividend and earnings growth play a smaller role. In fact, expected stock returns are close to their long-run mean. One implication of a bond-yield-driven ERP is that traditional indicators of the ERP like the price-dividend or price-earnings ratios, which do not use data from the term structure of risk-free rates, may not be as good a guide to future excess returns as they have been in the past.

As a second contribution, we present a concise and coherent taxonomy of ERP models. We categorize the twenty models into five groups: predictors that use historical mean returns only, dividend-discount models, cross-sectional regressions, time-series regressions and surveys. We explain the methodological and practical differences among these classes of models, including the assumptions and data sources that each require.

2. The Equity Risk Premium: Definition

Conceptually, the ERP is the compensation investors require to make them indifferent at the margin 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. One challenge in estimating the ERP is that 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 total market returns are based on broad stock market indices, such as 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 such as housing, private equity and non-tradable human capital. Even if we restricted ourselves to all traded stocks, we still have several choices to make, such as whether to use value or equal-weighted indices, and whether to exclude penny or infrequently 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 inflation⁵ and liquidity risks even if we were to assume there is no prospect of outright default. In this paper, we want to focus on how expectations are estimated in different models, and not on measurement issues regarding market returns and the risk-free rate. Thus, 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 so that our models are comparable with each other and with most of the literature.

While implementing the concept of the ERP in practice has its challenges, we can precisely define the ERP mathematically. First, we decompose stock returns⁶ into an expected component and a random component:

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

In equation (1), R_{t+k} are *realized* returns between t and $t+k$, and $E_t[R_{t+k}]$ are the returns that were expected from t to $t+k$ using information available at time t . The variable $error_{t+k}$ is a random variable that is unknown at time t and realized at $t+k$. Under rational expectations, $error_{t+k}$ has a mean of zero and is orthogonal to $E_t[R_{t+k}]$. We keep the discussion as general as possible and do not assume rational

⁵ Note that inflation risk in an otherwise risk-free nominal asset does not invalidate its usefulness to compute the ERP. If stock returns and the risk-free rate are expressed in nominal terms, their difference has little or no inflation risk. This follows from the following formula, which holds exactly in continuous time and to a first order approximation in discrete time: real stock returns – real risk-free rate = (nominal stock returns – expected inflation) – (nominal risk-free rate – expected inflation) = nominal stock returns – nominal risk-free rate. Hence, there is no distinction between a nominal and a real ERP.

⁶ Throughout this article, all returns are *net* returns. For example, a five percent return corresponds to a net return of 0.05 as opposed to a *gross* return of 1.05.

expectations at this stage, although it will be a feature of many of the models we consider. 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 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, future expected returns and the future ERP are stochastic, since expectations depend on the arrival of new information that has a random component not known in advance⁷. Second, 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. Third, if expectations are rational, because the unexpected component $error_{t+k}$ is stochastic and orthogonal to expected returns, the ERP is always less volatile than realized excess returns. In this case, we expect ERP estimates to be smoother than realized excess returns.

3. Models of the Equity Risk Premium

We describe twenty models of the equity risk premium, comparing their advantages, disadvantages and ease of implementation. Of course, there are many more models of the ERP than the ones we consider. We selected the models in our study based on the recent academic literature, their widespread use by practitioners and data availability. Table I describes the data we use and their sources, all of which are either readily available or standard in the literature⁸. With a few exceptions, all data is monthly from January 1960 to June 2013. Appendix A provides more details.

[Insert Table I here]

⁷ More precisely, $E_t[R_{t+k}]$ and $ERP_t(k)$ are known at time t but random from the perspective of all earlier periods.

⁸ In fact, except for data from I/B/E/S and Compustat, all sources are public.

We classify the twenty models into five categories based on their underlying assumptions; models in the same category tend to give similar estimates for the ERP. The five categories are: models based on the historical mean of realized returns, dividend discount models, cross-sectional regressions, time-series regressions and surveys.

All but one of the estimates of the ERP are constructed in real time, so that an investor who lived through the sample would have been able to construct the measures at each point in time using available information only⁹. This helps minimize look-ahead bias and makes any out-of-sample evaluation of the models more meaningful. Clearly, most of the models themselves were designed only recently and were not available to investors in real time, potentially introducing another source of forward-looking and selection biases that are much more difficult to quantify and eliminate.

3.1 Historical mean of realized returns

The easiest approach to estimating the ERP is to use the historical mean of realized market returns in excess of the contemporaneous risk-free rate. This model is very simple and, as shown in Goyal and Welch (2008), quite difficult to improve upon when considering out-of-sample predictability 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 mean of excess returns is either constant or very slow moving over time, giving very little time-variation in the ERP. The main choice is how far back into the past we should go when computing the historical mean. Table II shows the two versions of historical mean models that we use.

[Insert Table II here]

⁹ The one exception is Adrian, Crump and Moench's (2014) cross-sectional model, which is constructed using full-sample regression estimates.

3.2 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, as in Gordon (1962). Today's stock price should then be the sum of all expected future cash flows, discounted at an appropriate rate to take into account their riskiness and the time value of money. The formula that reflects this intuition is

$$P_t = \frac{D_t}{\rho_t} + \frac{E_t[D_{t+1}]}{\rho_{t+1}} + \frac{E_t[D_{t+2}]}{\rho_{t+2}} + \frac{E_t[D_{t+3}]}{\rho_{t+3}} + \dots, \quad (3)$$

where P_t is the current price of the stock, D_t are current cash flows, $E_t[D_{t+k}]$ are the cash flows k periods from now expected as of time t , and ρ_{t+k} is the discount rate for time $t+k$ from the perspective of time t . Cash flows to stockholders certainly include dividends, but can also arise from spin-offs, buy-outs, mergers, buy-backs, etc. In general, the literature focuses on dividend distributions because they are readily available data-wise and account for the vast majority of cash flows. The discount rate can be decomposed into

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

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 dividends. When using a DDM, we refer to $ERP_t(k)$ as the *implied* ERP. The reason is that we plug in prices, risk-free rates and estimated expected future dividends into equation (3), and then derive what value of $ERP_t(k)$ makes the right-hand side equal to the left-hand side in the equation, i.e. what ERP value is *implied* by equation (3).

DDM are forward looking and are consistent with no arbitrage. In fact, equation (3) must hold in any economy with no arbitrage¹⁰. Another advantage of DDM is that they are easy to implement. A drawback of DDM is that the results are sensitive to how we compute expectations of future dividends. Table III displays the DDM we consider and a brief description of their different assumptions.

[Insert Table III here]

3.3 Cross-sectional regressions

This method exploits the variation in returns and exposures to the S&P 500 of different assets to infer the ERP¹¹. Intuitively, cross-sectional regressions find the ERP by answering the following question: what is the level of the ERP that makes expected returns on 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 stocks with a single value for the ERP (and perhaps a small number of other variables), this model imposes tight restrictions on estimates 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:

$$R_{t+k}^i - R_{t+k}^f = \alpha^i \times \text{state variables}_{t+k} + \beta^i \times \text{risk factors}_{t+k} + \text{idiosyncratic risk}_{t+k}^i. \quad (5)$$

In equation (5), R_{t+k}^i is the realized return on a stock or portfolio i from time t to $t+k$.

State variables $_{t+k}$ are any economic indicators that help identify the state of the economy and its likely future path. *Risk factors* $_{t+k}$ are any measures of systematic contemporaneous co-variation in returns across all stocks or portfolios. Of course, some economic indicators can be both state variables and risk

¹⁰ Note that when performing the infinite summation in equation (3) we have not assumed the n^{th} term goes to zero as n tends to infinity, which allows for rational bubbles. In this sense, DDM do allow for a specific kind of bubble.

¹¹ See Polk, Thompson and Vuolteenaho (2006) and Adrian, Crump and Moench (2014) for a detailed description of this method.

factors at the same time. Finally, *idiosyncratic risk* r_{t+k}^i is the component of returns that is particular to each individual stock or portfolio that is not explained by *state variables* s_{t+k} or *risk factors* f_{t+k} (both of which, importantly, are common to all stocks and hence not indexed by i). Examples of state variables are inflation, unemployment, the yield spread between Aaa and Baa bonds, the yield spread between short and long term Treasuries, and the S&P 500's dividend-to-price ratio. The most important risk factor is the excess return on the S&P 500, which we must include if we want to infer the ERP consistent with the cross-section of stock returns. Other risk-factors usually used are the Fama-French (1992) factors and the momentum factor of Carhart (1997). The values in the vector α^i give the strength of asset-specific return predictability and the values in the vector β^i give the asset-specific exposures to risk factors¹². For the cross-section of assets indexed by i , 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. It is important to point out that equation (5) is not a predictive regression; the left and right-hand side variables are both associated with time $t + k$.

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

$$R_{t+k}^i - R_{t+k}^f = \lambda_t(k) \times \hat{\beta}^i, \quad (6)$$

where $\hat{\beta}^i$ are the values found when estimating equation (5). Equation (6) attempts to find, at each point in time, the vector of numbers $\lambda_t(k)$ that makes exposures β^i as consistent as possible with realized excess returns of all stocks or portfolios considered. The element in the vector $\hat{\lambda}_t(k)$ that is multiplied by

¹² The vectors α^i and β^i could also be time-varying, reflecting a more dynamic relation between returns and their explanatory variables. In this case, the estimation of equation (5) is more complicated and requires making further assumptions. The model by Adrian, Crump and Moench (2014) is the only cross-sectional model we examine that uses time-varying α^i and β^i .

the element in the $\hat{\beta}^i$ vector corresponding to the S&P 500 is $ERP_t(k)$, the equity risk premium we are seeking.

One advantage of cross-sectional regressions is that they use information from more asset prices than other models. Cross-sectional regressions also have sound theoretical foundations, since they provide one way to implement Merton's (1973) Intertemporal Capital Asset Pricing Model. 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 (Harvey, Liu and Zhu (2014)), and that it is not as easy to implement as most of the other options. Table IV displays the cross-sectional models in our study, together with the state variables and risk factors they use.

[Insert Table IV here]

3.4 Time-series regressions

Time-series regressions use the relationship between economic variables and stock returns to estimate the ERP. The idea is to run a predictive linear regression of realized excess returns on lagged "fundamentals":

$$R_{t+k} - R_{t+k}^f = a + b \times Fundamental_t + error_t. \quad (7)$$

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

$$ERP_t(k) = \hat{a} + \hat{b} \times Fundamental_t. \quad (8)$$

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 (7) and (8), with the assumption that "fundamentals" are the right sources of information to look at when computing expected returns, and that a linear equation is the correct functional specification.

The use of time-series regressions requires minimal assumptions; there is no concept of equilibrium and no absence of arbitrage necessary for the method to be valid¹³. In addition, implementation is quite simple, since it only involves running ordinary least-square regressions. The challenge is to select what variables to include on the right-hand side of equation (7), since results can change substantially depending on what variables are used to take the role of “fundamentals”. In addition, including more than one predictor gives poor out-of-sample predictions even if economic theory may suggest a role for many variables to be used simultaneously (Goyal and Welch (2008)). Finally, time-series regressions ignore information in the cross-section of stock returns. Table V shows the time-series regression models that we study.

[Insert Table V here]

3.5 Surveys

The survey approach consists of asking economic agents about the current level of the ERP. 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 predictors of excess returns 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 (2014) 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. In this paper, we use the survey of CFOs by Graham and Harvey (2012), which to our knowledge is the only large-scale ERP survey that has more than just a few years of data (see Table VI).

[Insert Table VI here]

¹³ However, the Arbitrage Pricing Theory of Ross (1976) provides a strong theoretical underpinning for time-series regressions by using no-arbitrage conditions.

4. Estimation of the Equity Risk Premium

We now study the behavior of the twenty models we consider by conducting principal component analysis. Since forecast accuracy can be substantially improved through the combination of multiple forecasts¹⁴, the optimal strategy to forecast excess stock returns may consist of combining together all these models. The first principal component of the twenty models that we use is the linear combination of ERP estimates that captures as much of the variation in the data as possible. The second, third, and successive principal components are the linear combinations of the twenty models that explain as much of the variation of the data as possible and are also uncorrelated to all the preceding principal components. If the first few principal components —say one or two— account for most of the variation of the data, then we can use them as a good summary for the variation in all the measures over time, reducing the dimensionality from twenty to one or two. In addition, in the presence of classical measurement error, the first few principal components can achieve a higher signal-to-noise ratio than other summary measures like the cross-sectional mean of all models (Geiger and Kubin (2013)).

To compute the first principal component, we proceed in three steps. We first de-mean all ERP estimates and find their variance-covariance matrix. In the second step, we find the linear combination that explains as much of the variance of the de-meaned models as possible. The weights in the linear combination are the elements of the eigenvector associated with the largest eigenvalue of the variance-covariance matrix found in the first step. In the third step, we add to the linear combination just obtained, which has mean zero, the average of ERP estimates across all models and all time periods. Under the assumption that each of the models is an unbiased and consistent estimator of the ERP, the average across all models and all time periods is an unbiased and consistent estimator of the unconditional mean of the ERP. The time

¹⁴ See, *inter alia*, Clemen (1989), Diebold and Lopez (1996) and Timmermann (2006).

variation in the first principal component then provides an estimate of the conditional ERP¹⁵. The share of the variance of the underlying models explained by this principal component is 76 percent, suggesting that there is not too much to gain from examining principal components beyond the first¹⁶.

We now focus on the one-year-ahead ERP estimates and study other horizons in the next section.

The first two columns in Table VII show the mean and standard deviation of each model's estimates. The unconditional mean of the ERP across all models is 5.7 percent, with an average standard deviation of 3.2 percent. DDM give the lowest mean ERP estimates and have moderate standard deviations. In contrast, cross-sectional models tend to have mean ERP estimates on the high end of the distribution and very smooth time-series. Mean ERP estimates for time-series regressions are mixed, with high and low values depending on the predictors used, but uniformly large variances. The survey of CFOs has a mean and standard deviation that are both about half as large as in the overall population of models. The picture that emerges from Table VII is that there is considerable heterogeneity across model types, and even sometimes within model types, thereby underscoring the difficulty inherent in finding precise estimates of the ERP.

¹⁵ As is customary in the literature, we perform the analysis using ERP estimates in levels, even though they are quite persistent. Results in first-differences do not give economically reasonable estimates since they feature a pro-cyclical ERP and unreasonable magnitudes.

One challenge that arises in computing the principal component is when we have 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 method, which conceptually preserves the idea behind principal components. Let X be the matrix that has observations for different models in its columns and for different time periods in its rows. 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. We thank Richard Crump for suggesting this method and providing the code for its implementation.

¹⁶ The second and third principal components account for 13 and 8 percent of the variance, respectively.

[Insert Table VII here]

Figure 1 shows the time-series for all one-year-ahead ERP model estimates, with each class of models in a different panel. The green lines are the ERP estimates from the twenty underlying models. The black line, reproduced in each of the panels, is the principal component of all twenty models. The shaded areas are NBER recessions. The figure gives a sense of how the time-series move together, and how much they co-vary with the first principal component. Table VIII shows the correlations among models. Figure 1 and Table VIII give the same message: despite some outliers, there is a fairly strong correlation within each of the five classes of models. Across classes, however, correlations are small and even negative. Interestingly, the correlation between some DDM and cross-sectional models is as low as -91 percent. This negative correlation, however, disappears if we look at lower frequencies. When aggregated to quarterly frequency, the smallest correlation between DDM and cross-sectional models is -22 percent, while at the annual frequency it is 12 percent.

[Insert Figure 1 here]

[Insert Table VIII here]

Figure 1 also shows that the first principal component co-varies negatively with historical mean models, but positively with DDM and cross-sectional regression models. Time-series regression models are also positively correlated with the first principal component, although this is not so clearly seen in Panel 4 of Figure 1 because of the high volatility of time-series ERP estimates. The last panel shows that the survey of CFOs does track the first principal component quite well at low frequencies (e.g. annual), although any conclusions about survey estimates should be interpreted with caution given the short length of the sample.

As explained earlier, the first principal component is a linear combination of the twenty underlying ERP models:

$$PC_t^{(1)} = \sum_{m=1}^{20} w^{(m)} ERP_t^{(m)}. \quad (9)$$

In the above equation, m indexes the different models, $PC_t^{(1)}$ is the first principal component, $ERP_t^{(m)}$ is the estimate from model m and $w^{(m)}$ is the weight that the principal component places on model m . The third column in Table VII, labeled “PC coefficients”, shows the weights $w^{(m)}$ normalized to sum up to one to facilitate comparison, i.e. the table reports the weights $\hat{w}^{(m)}$ where

$$\hat{w}^{(m)} = \frac{w^{(m)}}{\sum_{m=1}^{20} w^{(m)}}. \quad (10)$$

The first principal component puts positive weight on models based on the historical mean, cross-sectional regressions and the survey of CFOs. It weights DDM and time-series regressions mostly negatively. The absolute values of the weights are very similar for many of the models, and there is no single model or class of models that dominates. This means that the first principal component uses information from many of the models.

The last column in Table VII, labeled “Exposure to PC”, shows the extent to which models *load* on the first principal component. By construction, each of the twenty ERP models can be written as a linear combination of twenty principal components:

$$ERP_t^{(m)} = \sum_{i=1}^{20} load_i^{(m)} PC_t^{(i)}, \quad (11)$$

where m indexes the model and i indexes the principal components. The values in the last column of Table VII are the loadings on the first principal component ($i = 1$) for each model ($m = 1, 2, \dots, 20$), again normalized to one for ease of comparability:

$$\widehat{load}_1^{(m)} = \frac{load_1^{(m)}}{\sum_{m=1}^{20} load_i^{(m)}}. \quad (12)$$

Most models have a positive loading on the first principal component; whenever the loading is negative, it tends to be relatively small. This means the first principal component, as expected, is a good explanatory variable for most models. Looking at the third and fourth columns of Table VII together, we can obtain additional information. For example, a model with a very high loading (fourth column) accompanied by a very small PC coefficient (third column) is likely to mean that the model is almost redundant, in the sense that it is close to being a linear combination of all other models and does not provide much independent information to the principal component. On the other hand, if the PC coefficient and loading are both high, the corresponding model is likely providing information not contained in other measures.

Figure 2 shows the first principal component of all twenty models in black, with recessions indicated by shaded bars (the black line is the same principal component shown in black in each of the panels of Figure 1). As expected, the principal component tends to peak during financial turmoil, recessions and periods of low real GDP growth or high inflation. It tends to bottom out after periods of sustained bullish stock markets and high real GDP growth. Evaluated by the first principal component, the one-year-ahead ERP reaches a local peak in June of 2012 at 12.2 percent. The surrounding months have ERP estimates of similar magnitude, with the most recent estimate in June 2013 at 11.2 percent. This behavior is not so clearly seen by simply looking at the collection of individual models in Figure 1, highlighting the usefulness of principal components analysis. Similarly high levels were seen in the mid and late 1970s, during a period of stagflation, while the recent financial crisis had slightly lower ERP estimates closer to 10 percent.

[Insert Figure 2 here]

Figure 2 also displays the 10th, 25th, 75th and 90th percentiles of the cross-sectional distribution of models. These bands can be interpreted as confidence intervals, since they give the range of the distribution of ERP estimates at each point in time. However, they do not incorporate other relevant sources of uncertainty, such as the errors that occur during the estimation of each individual model, the degree of doubt in the correctness of each model, and the correlation structure between these and all other kinds of errors. Standard error bands that capture all sources of uncertainty are therefore likely to be wider.

The difference in high and low percentiles can also be interpreted as measures of agreement across models. 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 at a level somewhat below its long-run mean. The cross-sectional standard deviation in ERP estimates (not shown in the graph) also decreased from 10.2% in January of 2000 to 4.3% in June of 2013, confirming that the disagreement among models has decreased.

Another *a priori* reasonable summary statistic for the ERP is the cross-sectional mean of estimates across models. In Figure 3, we can see that by this measure the ERP has also been increasing since the crisis. However, unlike the principal component, it has not reached elevated levels compared to past values. The cross-sectional mean can be useful, but it has a few undesirable features as an overall measure of the ERP compared to the first principal component. First, it is procyclical, which contradicts the economic intuition that expected returns are highest in recessions, when risk aversion is high and future prospects look brighter than current ones. Second, it overloads on DDM simply because there is a higher number of DDM models in our sample. Lastly, it has a smaller correlation with the realized returns it is supposed to predict.

[Insert Figure 3 here]

5. The Term Structure of Equity Risk Premia

In Section 2, we described the term structure of the ERP – what expected excess returns are over different investment horizons. In practical terms, we estimate the ERP at different horizons by using the inputs for all the models at the corresponding horizons¹⁷. For example, if we want to take the historical mean of returns as our estimate, we can take the mean of returns over one month, six months, or a one-year period. In cross-sectional and time-series regressions, we can predict monthly, quarterly or annual returns using monthly, quarterly or annual right-hand side variables. DDM, on the other hand, have little variation across horizons. In fact, all the DDM we consider have a constant term structure of expected stock returns, and the only term structure variation in ERP estimates comes from risk-free rates¹⁸.

Figure 4 plots the first principal components of the ERP as a function of investment horizon for some selected dates. We picked the dates because they are typical dates for when the ERP was unusually high or unusually low at the one-month horizon. As was the case for one-year-ahead ERP estimates, we can capture the majority of the variance of the underlying models at all horizons by a single principal component. The shares of the variance explained by the first principal components at horizons of one month to three years range between 68 and 94 percent. The grey line in Figure 4 shows the average of the term structure across all periods. It is slightly upward sloping, with a short-term ERP at just over 6 percent and a three-year ERP at almost 7 percent.

[Insert Figure 4 here]

¹⁷ For other ways to estimate the term structure of the ERP using equilibrium models or derivatives, see Ait-Sahalia, Karaman and Mancini (2014), Ang and Ulrich (2012), van Binsbergen, Hueskes, Koijen and Vrugt (2014), Boguth, Carlson, Fisher and Simutin (2012), Durham (2013), Croce, Lettau and Ludvigson (2014), Lemke and Werner (2009), Lettau and Wachter (2011), Muir (2013), among others.

¹⁸ In equation (3), ρ_{t+k} is assumed to be the same for all k , while risk-free rates are allowed to vary over the investment horizon k in equation (4). Of course, with additional assumptions, it is possible to have DDM with a non-constant term structure of expected excess returns.

The first observation is that the term structure of the ERP has significant time variation and can be flat, upward or downward sloping. Figure 4 also shows some examples that hint at lower future expected excess returns when the one-month-ahead ERP is elevated and the term structure is downward sloping, and higher future expected excess returns when the one-month-ahead ERP is low and the term structure is upward sloping. In fact, this is generally true: There is a strong negative correlation between the level and the slope of the ERP term structure of -71 percent. Figure 5 plots monthly observations of the one-month-ahead ERP against the slope of the ERP term structure (the three-year-ahead minus the one-month-ahead ERP) together with the corresponding ordinary least squares regression line in black. Of course, this is only a statistical pattern and should not be interpreted as a causal relation.

[Insert Figure 5 here]

6. Why is the Equity Risk Premium High?

There are two reasons why the ERP can be high: low discount rates and high current or expected future cash flows.

Figure 6 shows that earnings are unlikely to be the reason why the ERP is high. The green line shows the year-on-year change in the mean expectation of one-year-ahead earnings per share for the S&P 500. These expectations are obtained from surveys conducted by the Institutional Brokers' Estimate System (I/B/E/S) and available from Thomson Reuters. Expected earnings per share have been declining from 2010 to 2013, making earnings growth an unlikely reason for why the ERP was high in the corresponding period. The black line shows the realized monthly growth rates of real earnings for the S&P 500 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 percent, which is near its long-run average.

[Insert Figure 6 here]

Another way to examine whether a high ERP is due to discount rates or cash flows is shown in Figure 7. The black line is the same one-year-ahead ERP estimate shown in Figure 2. The green line simply adds the realized one-year Treasury yield to obtain expected stock returns. The figure shows expected stock returns have increased since 2000, similarly to the ERP. However, unlike the ERP, expected stock returns are close to their long-run mean, and nowhere near their highest levels, achieved in 1980. The discrepancies between the two lines are due to exceptionally low bond yields since the end of the financial crisis.

[Insert Figure 7 here]

Figure 8 displays the term structure of the ERP under a simple counterfactual scenario, in addition to the mean and current term structures already displayed in Figure 4. In this scenario, we leave expected stock returns unmodified but change the risk-free rates in June 2012 from their actual values 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 8 in green. Using average levels of bond yields brings the whole term structure of the ERP much closer to its mean level (the grey line), especially at intermediate horizons. This shows that a “normalization” of bond yields, everything else being equal, would bring the ERP close to its historical norm. This exercise shows that the current environment of low bond yields is capable, quantitatively speaking, of significantly contributing to an ERP as high as was observed in 2012-2013.

[Insert Figure 8 here]

7. Conclusion

We have analyzed twenty different models of the ERP by considering the assumptions and data required to implement them, and how they relate to each other. When it comes to the ERP, we find that there is substantial heterogeneity in estimation methodology and final estimates. We then extract the first

principal component of the twenty models, which signals that the ERP in 2012 and 2013 is at heightened levels compared to previous periods. Our analysis provides evidence that the current level of the ERP is consistent with a bond-driven ERP: expected excess stock returns are elevated not because stocks are expected to have high returns, but because bond yields are exceptionally low. The models we consider suggest that expected stock returns, on their own, are close to average levels.

References

Adrian, T., D. M. Covitz, and N. Liang. 2013. "Financial stability monitoring." Federal Reserve Bank of New York Staff Reports, no. 601 (February).

Adrian, T., R. Crump, and E. Moench. 2014. "Regression-based estimation of dynamic asset pricing models." Federal Reserve Bank of New York Staff Reports, no. 493 (June).

Ait-Sahalia, Y., M. Karaman, and L. Mancini. 2014. "The term structure of variance swaps, risk premia and the expectation hypothesis. Risk Premia and the Expectation Hypothesis." (May 10, 2014). Available at SSRN: <http://ssrn.com/abstract=2136820>

Ang, A., and G. Bekaert. 2007. "Stock return predictability: Is it there?" *Review of Financial Studies* 20, no. 3 (May): 651-707.

Ang, A., and M. Ulrich. 2012. "Nominal Bonds, Real Bonds, and Equity." Unpublished paper, Columbia University.

Baker, M., and J. Wurgler. 2007. "Investor sentiment in the stock market." *The Journal of Economic Perspectives*: 129-51.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt. 2013. "Equity yields." *Journal of Financial Economics* 110, no. 3 (December): 503-19.

Bloom, N. 2009. "The impact of uncertainty shocks." *Econometrica* 77, no. 3 (May): 623-85.

Boguth, O., M. Carlson, A. J. Fisher, and M. Simutin. 2012. "Leverage and the limits of arbitrage pricing: Implications for dividend strips and the term structure of equity risk premia." Unpublished paper, Arizona State University.

Campbell, J.Y., and S. Thompson. 2008. "Predicting excess stock returns out of sample: Can anything beat the historical average?" *Review of Financial Studies* 21, no. 4 (July): 1509-31.

Carhart, M. M. 1997. "On persistence in mutual fund performance." *Journal of Finance* 52, no 1 (March): 57-82.

Chen, L., Z. Da, and X. Zhao. 2013. "What drives stock price movements?" *Review of Financial Studies* 26, no. 4 (April): 841-76.

Clemen, R.T., 1989. "Combining forecasts: a review and annotated bibliography". *International Journal of Forecasting* 5 (4), 559-83.

Croce, M., M. Lettau, and S. Ludvigson. Forthcoming. 2014. "Investor information, long-run risk, and the term structure of equity." *Review of Financial Studies*.

- Damodaran, A.* 2012. "Equity Risk Premiums (ERP): Determinants, Estimation and Implications – The 2012 Edition." Unpublished paper, New York University.
- Diebold, F.X., and J. Lopez,* 1996. "Forecast Evaluation and Combination". In G.S. Maddala and C.R. Rao (eds.), *Handbook of Statistics*, 241-268, 1996. Amsterdam: North-Holland.
- Duarte, F.* 2013. "Inflation and the cross-section of stock returns." Federal Reserve Bank of New York Staff Reports, no. 621 (May).
- Durham, J. B.* 2013. "Arbitrage-free models of stocks and bonds." Federal Reserve Bank of New York Staff Reports, no. 656 (December).
- Duarte, F., L. Kogan, and D. Livdan.* 2013. "Aggregate Investment and Stock Returns." Unpublished paper, M.I.T.
- Easton, P.D., and G. A. Sommers.* 2007. "Effect of analysts' optimism on estimates of the expected return implied by earnings forecasts." *Journal of Accounting Research* 45, no. 5 (December): 983-1016.
- Fama, E.F., and K. R. French.* 1988. "Dividend yields and expected stock returns." *Journal of Financial Economics* 22, no. 1 (October), 3-25.
- Fama, E.F., and K. R. French.* 1992. "The cross-section of expected stock returns." *Journal of Finance* 47, no. 2 (June): 427-65.
- Fama, E.F., and K. R. French.* 2002. "The equity premium." *Journal of Finance* 57, no. 2 (April): 637-59.
- Geiger, B., and G. Kubin.* 2013. "Signal Enhancement as Minimization of Relevant Information Loss". Proceedings of ITG Conference on Systems, Communication and Coding
- Gordon, M. J.* 1962. "The investment, financing, and valuation of the corporation." Greenwood Press.
- Goyal, A., and I. Welch.* 2008. "A comprehensive look at the empirical performance of equity premium predictions." *Review of Financial Studies* 21, no. 4 (July): 1455-1508.
- Graham, J., and C. Harvey.* 2012. "The Equity Risk Premium in 2012." Unpublished paper, Duke University.
- Greenwood, R., and A. Shleifer.* 2014. "Expectations of returns and expected returns." *Review of Financial Studies* 27, no. 3 (March): 714-46
- Gurkaynak, R. S., B. Sack, and J. H. Wright.* 2007. "The US Treasury yield curve: 1961 to the present." *Journal of Monetary Economics* 54, no. 8 (November): 2291-2304.
- Gurkaynak, R. S., S. Refet, B. Sack, and J. H. Wright.* 2010. "The TIPS yield curve and inflation compensation." *American Economic Journal: Macroeconomics* 2, no. 1 (January): 70-92.

- Hall, R. E.* Forthcoming. 2014. "The Routes into and out of the Zero Lower Bound." Federal Reserve Bank of Kansas City Symposium (Jackson Hole).
- Harvey, C. R., and Y. Liu, and H. Zhu.* 2014. "...and the Cross-Section of Expected Returns." (October 7, 2014). Available at SSRN: <http://ssrn.com/abstract=2249314>
- Kelly, B., and S. Pruitt.* 2013. "Market Expectations in the Cross-Section of Present Values." *The Journal of Finance* 68, no. 5 (October): 1721-56.
- Kuehn, L. A., N. Petrosky-Nadeau, and L. Zhang.* 2012. "An equilibrium asset pricing model with labor market search." NBER Working Paper no. 17742.
- Lemke, W., and T. Werner.* 2009. "The term structure of equity premia in an affine arbitrage-free model of bond and stock market dynamics." European Central Bank, Working Paper No.1045.
- Lettau, M., and J. A. Wachter.* 2011. "The term structures of equity and interest rates." *Journal of Financial Economics* 101, no. 1 (July): 90-113.
- Merton, R.C.* 1973. "An intertemporal capital asset pricing model." *Econometrica* 41, no. 5 (September): 867-87.
- Muir, T.* 2013. "Financial crises, risk premia, and the term structure of risky assets." Unpublished paper, Yale School of Management.
- Neely, C. J., D.E. Rapach, J. Tu, and G. Zhou.* 2014. "Forecasting the equity risk premium: the role of technical indicators." *Management Science* 60, no. 7 (July): 1772-91.
- Panigirtzoglou, N., and J. Loeys.* 2005. "A fair value model for US bonds, credit and equities." JP Morgan (January).
- Piazzesi, M., and M. Schneider.* 2007. "Equilibrium yield curves." In *NBER Macroeconomics Annual 2006* 21: 389-472. MIT Press, 2007.
- Polk, C., S. Thompson, and T. Vuolteenaho.* 2006. "Cross-sectional forecasts of the equity premium." *Journal of Financial Economics* 81, no. 1 (July): 101-41.
- Ross, S.* 1976. "The arbitrage theory of capital asset pricing." *Journal of Economic Theory* 13, no. 3 (December): 341-60.
- Shiller, R. J.* 2005. "Irrational Exuberance". Princeton University Press 2000, Broadway Books 2001, 2nd ed.
- Stock, J.H., and M. W. Watson.* 2003. "Forecasting output and inflation: The role of asset prices." *Journal of Economic Literature* 41, no. 3 (September): 788-829.

Timmermann, A. 2006. "Forecast combinations". Handbook of Economic Forecasting, Volume 1 (chapter 4), 135-96.

Appendix A: Data Variables

Fama and French (1992)	http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html Monthly frequency; 1/1/1960 to 6/30/2013. We use 25 portfolios sorted on size and book to market, 10 portfolios sorted on momentum, realized excess market returns, HML, SMB, and the momentum factor.
Shiller (2005)	http://www.econ.yale.edu/~shiller/data.htm Monthly frequency; 1/1/1960 to 6/30/2013. We use the nominal and real price, nominal and real dividends and nominal and real earnings for the S&P 500, CPI, and 10 year nominal treasury yield.
Baker and Wurgler (2007)	http://people.stern.nyu.edu/jwurgler/data/Investor_Sentiment_Data_v23_POST.xlsx Monthly frequency; 7/1/1965 to 12/1/2010. We use the “sentiment measure”.
Graham and Harvey (2012)	http://www.cfosurvey.org/index.htm Quarterly frequency; 6/6/2000 to 6/5/2013. We use the answer to the question “Over the next 10 years, I expect the average annual S&P 500 return will be: Expected return:” and the analogous one that asks about the next year.
Damodaran (2012)	http://www.stern.nyu.edu/~adamodar/pc/datasets/histimpl.xls Annual frequency; 1/1/1960 to 12/1/2012. We use the ERP estimates from his dividend discount models (one uses free-cash flow, the other one doesn’t).
Gurkaynak, Sack and Wright (2007)	http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html Daily frequency; starting on 6/14/61 for one- to seven-year yields; 8/16/71 for nine- and ten-year yields; 11/15/71 for eleven- to fifteen-year yields; 7/2/81 for sixteen- to twenty-year yields; 11/25/85 for twenty-one- to thirty-year yields. We use all series until 6/30/2013.
Gurkaynak, Refet, Sack and Wright (2010)	http://www.federalreserve.gov/econresdata/researchdata.htm Monthly frequency; 1/1/2003 to 7/1/2013. We use yields on TIPS of all maturities available.
Compustat	Variable BKVLPS Annual frequency; 12/31/1977 to 12/31/2012.
Thomson Reuters I/B/E/S	Variables EPS 1 2 3 4 5 Monthly frequency; 1/14/1982 to 4/18/2013 for current and next year forecasts; 9/20/84 to 4/18/2013 for two-year-ahead forecasts; 9/19/85 to 3/15/2012 for three-year-ahead forecasts; 2/18/88 to 3/15/07 for four-year-ahead forecasts.
FRED (St. Louis Federal Reserve)	http://research.stlouisfed.org/fred2/graph/?g=D9J and http://research.stlouisfed.org/fred2/graph/?g=KKk Monthly frequency. 1/1/1960 to 7/1/2013 for Baa minus Aaa bond yield spread and recession indicator.

Tables and Figures

Table I: Data sources	
Fama and French (1992)	Fama-French factors, momentum factor, twenty-five portfolios sorted on size and book-to-market
Shiller (2005)	Inflation and ten-year nominal treasury yield. Nominal price, real price, earnings, dividends and cyclically adjusted price-earnings ratio for the S&P 500
Baker and Wurgler (2007)	Debt issuance, equity issuance, sentiment measure
Graham and Harvey (2012)	ERP estimates from the Duke CFO survey
Damodaran (2012)	ERP estimates
Gurkaynak, Sack and Wright (2007)	Zero coupon nominal bond yields for all maturities ¹⁹
Gurkaynak, Refet, Sack and Wright (2010)	Zero coupon TIPS yields for all maturities
Compustat	Book value per share for the S&P 500
Thomson Reuters I/B/E/S	Mean analyst forecast of expected earnings per share
FRED (St. Louis Federal Reserve)	Corporate bond Baa-Aaa spread and the NBER recession indicator

Note: All variables start in January 1960 (or later, if unavailable for early periods) and end in June 2013 (or until no longer available). CFO surveys are quarterly; book value per share and ERP estimates by Damodaran (2012) are annual; all other variables are monthly. Appendix A provides more details.

¹⁹ Except for the 10-year yield, which is from Shiller (2005). We use the 10-year yield from Shiller (2005) for ease of comparability with the existing literature. Results are virtually unchanged if we use all yields, including the 10-year yield, from Gurkaynak, Sack and Wright (2007).

Table II: Models based on the historical mean of realized returns

Long-run mean	Average of realized S&P 500 returns minus the risk-free rate using all available historical data
Mean of the previous five years	Average of realized S&P 500 returns minus the risk-free rate using only data for the previous five years

Table III: Dividend Discount Models

Gordon (1962) with nominal yields	S&P 500 dividend-to-price ratio minus the ten-year nominal Treasury yield
Shiller (2005)	Cyclically adjusted price-earnings ratio (CAPE) minus the ten-year nominal Treasury yield
Gordon (1962) with real yields	S&P 500 dividend-to-price ratio minus the ten year real Treasury yield (computed as the ten-year nominal Treasury rate minus the ten year breakeven inflation implied by TIPS)
Gordon (1962) with earnings forecasts	S&P 500 expected earnings-to-price ratio minus the ten-year nominal Treasury yield
Gordon (1962) with real yields and earnings forecasts	S&P 500 expected earnings-to-price ratio minus the ten-year real Treasury yield (computed as the ten-year nominal Treasury rate minus the ten-year breakeven inflation implied by TIPS)
Panigirtzoglou and Loeys (2005)	Two-stage DDM. The growth rate of earnings over the first five years 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 growth rate of earnings from years six and onwards is 2.2 percent
Damodaran (2012)	A six-stage DDM. Dividend growth the first five stages are estimated from analyst's earnings forecasts. Dividend growth in the sixth stage is the ten-year nominal Treasury yield
Damodaran (2012) free cash flow	Same as Damodaran (2012), but uses free-cash-flow-to-equity as a proxy for dividends plus stock buybacks

Table IV: Models with cross-sectional regressions

Fama and French (1992)	Uses the excess returns on the market portfolio, a size portfolio and a book-to-market portfolio as risk factors
Carhart (1997)	Identical to Fama and French (1992) but adds the momentum measure of Carhart (1997) as an additional risk factor
Duarte (2013)	Identical to Carhart (1997) but adds an inflation risk factor
Adrian, Crump and Moench (2014)	Uses the excess returns on the market portfolio as the single risk factor. The state variables are the dividend yield, the default spread, and the risk free rate

Table V: Models with time-series regressions

Fama and French (1988)	Only predictor is the dividend-price ratio of the S&P 500
Goyal and Welch (2008)	Uses, at each point in time, the best out-of-sample predictor out of twelve predictive variables proposed by Goyal and Welch (2008)
Campbell and Thompson (2008)	Same as Goyal and Welch (2008), 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. Second, we replace the estimate of the ERP by zero if the estimation otherwise finds a negative ERP
Fama and French (2002)	Uses, at each point in time, the best out-of-sample predictor out of three variables: the price-dividend ratio adjusted by the growth rate of earnings, dividends or stock prices
Baker and Wurgler (2007)	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

Table VI: Surveys

Graham and Harvey (2012)	Chief financial officers (CFOs) are asked since 1996 about the one and ten-year-ahead ERP. We take the mean of all responses
---------------------------------	--

Table VII: ERP models

		Mean	Std. dev.	PC coefficients $\widehat{w}^{(m)}$	Exposure to PC $\widehat{load}_1^{(m)}$
Based on historical mean	Long-run mean	9.3	1.3	0.78	-0.065
	Mean of previous five years	5.7	5.8	0.42	-0.160
DDM	Gordon (1926): E/P minus nominal 10yr yield	-0.1	2.1	-0.01	0.001
	Shiller (2005): 1/CAPE minus nominal 10yr yield	-0.4	1.8	-0.10	0.011
	Gordon (1962): E/P minus real 10yr yield	3.5	2.1	0.69	-0.077
	Gordon (1962): Expected E/P minus real 10yr yield	5.3	1.7	-0.78	0.208
	Gordon (1962): Expected E/P minus nominal 10yr yield	0.4	2.3	-0.79	0.077
	Panigirtzoglou and Loeys (2005): Two-stage DDM	-1.0	2.3	0.07	-0.011
	Damodaran (2012): Six-stage DDM	3.4	1.3	-0.26	0.032
	Damodaran (2012): Six-stage free cash flow DDM	4.0	1.1	-0.62	0.053
Cross-sectional regressions	Fama and French (1992)	12.6	0.7	0.80	-0.040
	Carhart (1997): Fama-French and momentum	13.1	0.8	0.81	-0.042
	Duarte (2013): Fama-French, momentum and inflation	13.1	0.8	0.82	-0.044
	Adrian, Crump and Moench (2014)	6.5	6.9	-0.05	0.114
Time-series regressions	Fama and French (1988): D/P	2.4	4.0	-0.27	0.069
	Best predictor in Goyal and Welch (2008)	14.5	5.2	-0.07	0.023
	Best predictor in Campbell and Thompson (2008)	3.1	9.8	-0.12	0.081
	Best predictor in Fama French (2002)	11.9	6.8	-0.72	0.321
	Baker and Wurgler (2007) sentiment measure	3.0	4.7	-0.32	0.184
Surveys	Graham and Harvey (2012) survey of CFOs	3.6	1.8	0.72	0.264
All models		5.7	3.2	0.78	-0.065

For each of the twenty models of the equity risk premium, we show four statistics. The first two are the time-series means and standard deviations for monthly observations from January 1960 to June 2013 (except for surveys, which are quarterly). The units are annualized percentage points. The third statistic, “PC coefficients $\widehat{w}^{(m)}$ ”, is the weight that the first principal component places on each model (normalized to sum to one). The fourth is the “Exposure to PC $\widehat{load}_1^{(m)}$ ”, the weight on the first principal component when each model is written as a weighted sum of all principal components (also normalized to sum to one).

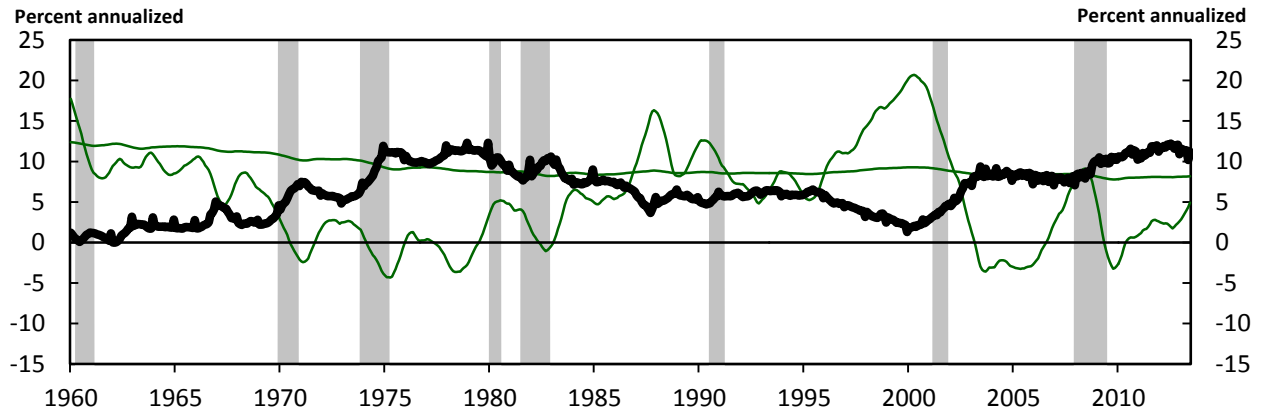
Table VIII: Correlation of ERP models

	LR mean	Mean past 5yr	E/P - 10yr	1/CAPE-10yr	E/P-real 10yr	Exp E/P-real 10yr	Exp E/P- 10yr	Two-stage DDM	Six-stage DDM	Free cash flow	FF	Carhart	Duarte	ACM	D/P	G and W	C and T	FF	Sentiment	CFO Survey
LR mean	100																			
Mean past 5yr	32	100																		
E/P - 10yr	8	15	100																	
1/CAPE-10yr	-9	0	78	100																
E/P-real 10yr	-11	25	98	23	100															
Exp E/P-real 10yr	-58	42	70	84	60	100														
Exp E/P- 10yr	-83	-61	84	95	46	98	100													
Two-stage DDM	17	27	88	54	89	66	79	100												
Six-stage DDM	3	-38	26	39	-30	32	52	-31	100											
Free cash flow	-43	-55	59	70	35	80	94	27	62	100										
FF	69	29	-8	-36	-21	-69	-91	9	-29	-77	100									
Carhart	71	30	-5	-31	-24	-71	-91	10	-25	-75	99	100								
Duarte	71	30	-3	-29	-22	-70	-91	11	-28	-74	99	100	100							
ACM	-1	-52	36	62	6	54	63	27	23	33	-28	-28	-25	100						
D/P	49	12	27	12	27	42	54	24	74	42	44	54	55	21	100					
G and W	25	12	25	21	-7	-36	-60	20	29	-9	7	13	14	-24	61	100				
C and T	27	31	14	-7	81	49	-60	28	-51	-40	60	57	58	-33	54	50	100			
FF	1	-30	-24	-29	37	-27	-37	-18	22	38	36	38	37	-9	40	23	43	100		
Sentiment	-10	33	-4	-20	68	-23	-29	27	-38	-20	18	17	18	-12	-38	-8	21	6	100	
CFO survey	-43	-33	12	30	1	1	13	16	5	-3	-36	-37	-39	60	14	-21	-32	-3	-36	100

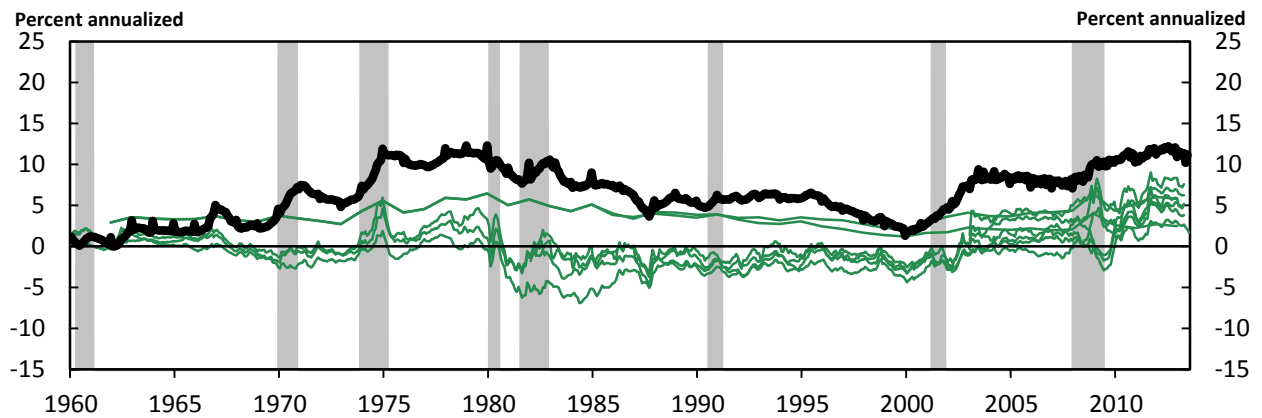
This table shows the correlation matrix of the twenty equity risk premium models we consider. Numbers are rounded to the nearest integer. Thick lines group models by their type (see Tables II to VI). Except for the CFO survey, the observations used to compute correlations are monthly for January 1960 to June 2013. For the CFO survey, correlations are computed by taking the last observation in the quarter for monthly series and then computing quarterly correlations.

Figure 1: ERP estimates for all models

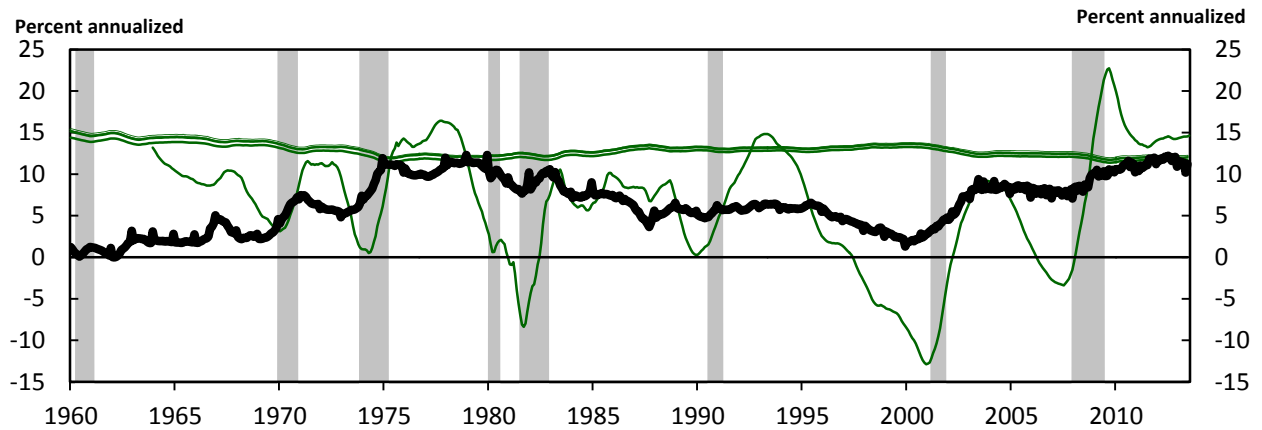
Panel 1: ERP models based on the historical mean of excess returns



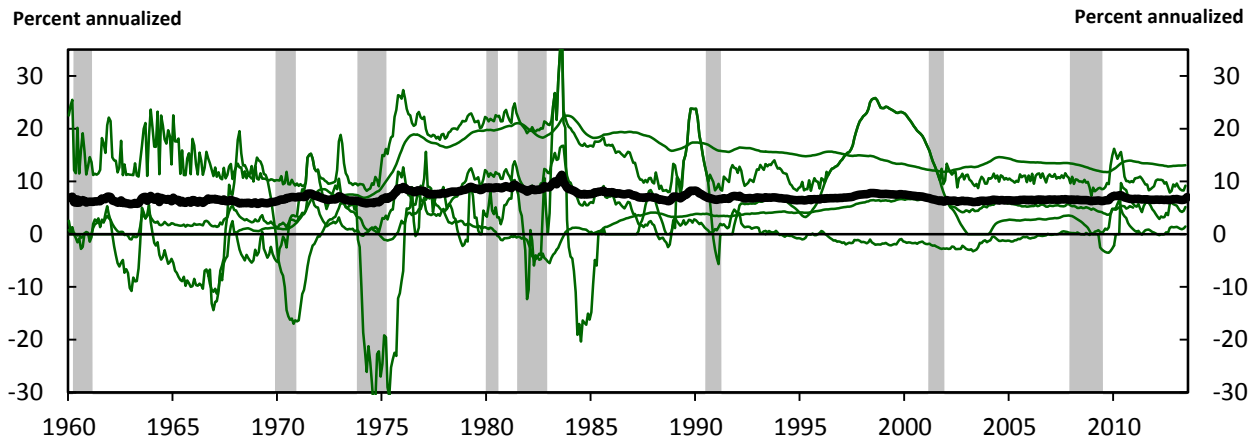
Panel 2: ERP dividend discount models (DDM)



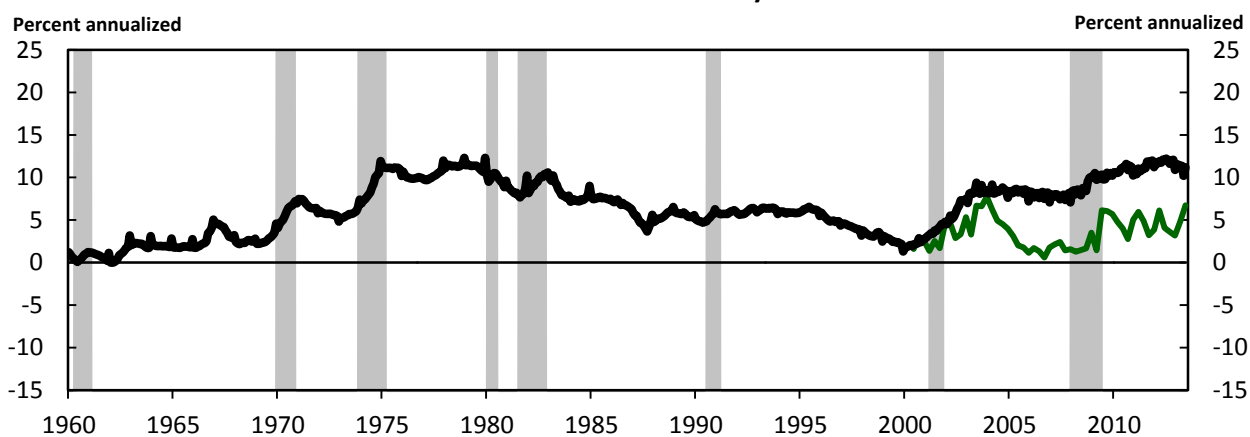
Panel 3: ERP cross sectional models



Panel 4: ERP time series models



Panel 5: ERP surveys

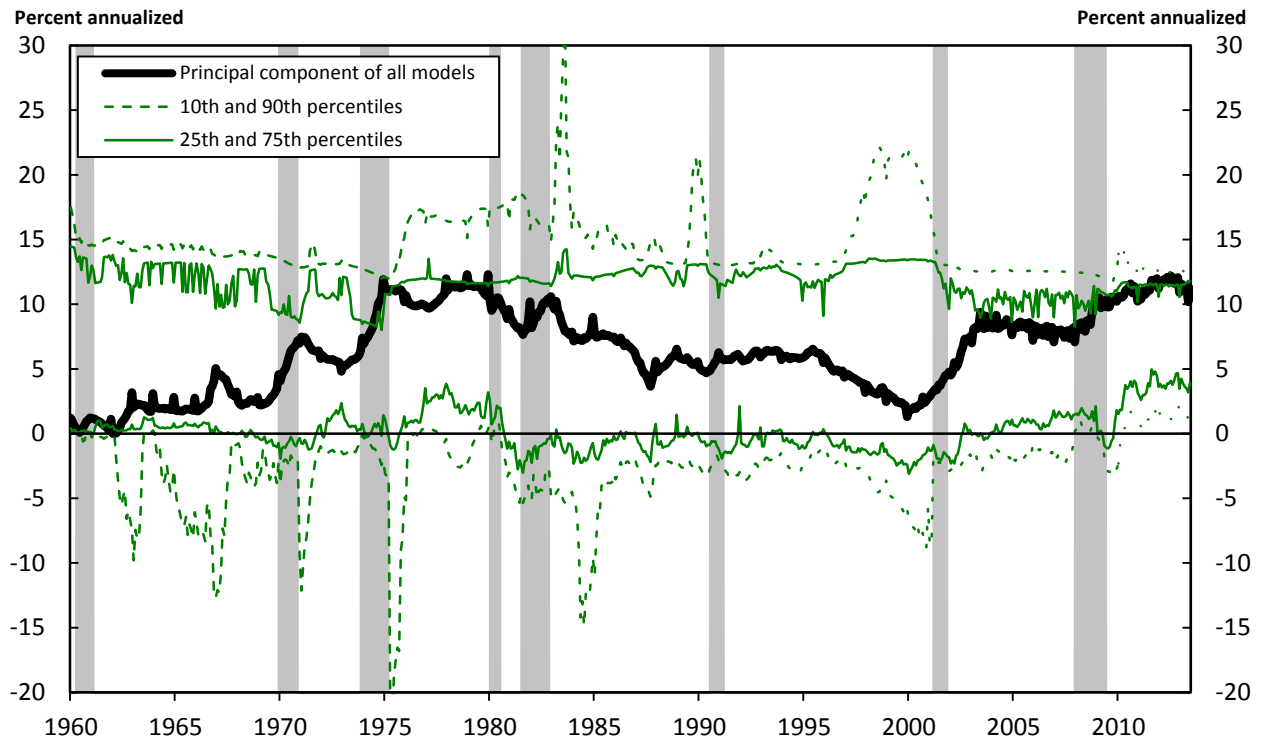


Each green line gives the one-year-ahead equity risk premium from each of the models listed in Tables II to VI. All numbers are in annualized percentage points.

Panel 1 shows the estimates for models based on the historical mean of excess returns, which are listed in Table II. Panel 2 shows estimates computed by the dividend discount models in Table III. Panel 3 uses the cross-sectional regression models from Table IV. Panel 4 shows the equity risk premium computed by the time-series regression models in Table V. Panel 5 gives the estimate obtained from the survey cited in Table VI.

In all panels, the black line is the first principal component of all twenty models (it can look different across panels due to different scales in the y-axis).

Figure 2: One-year-ahead ERP

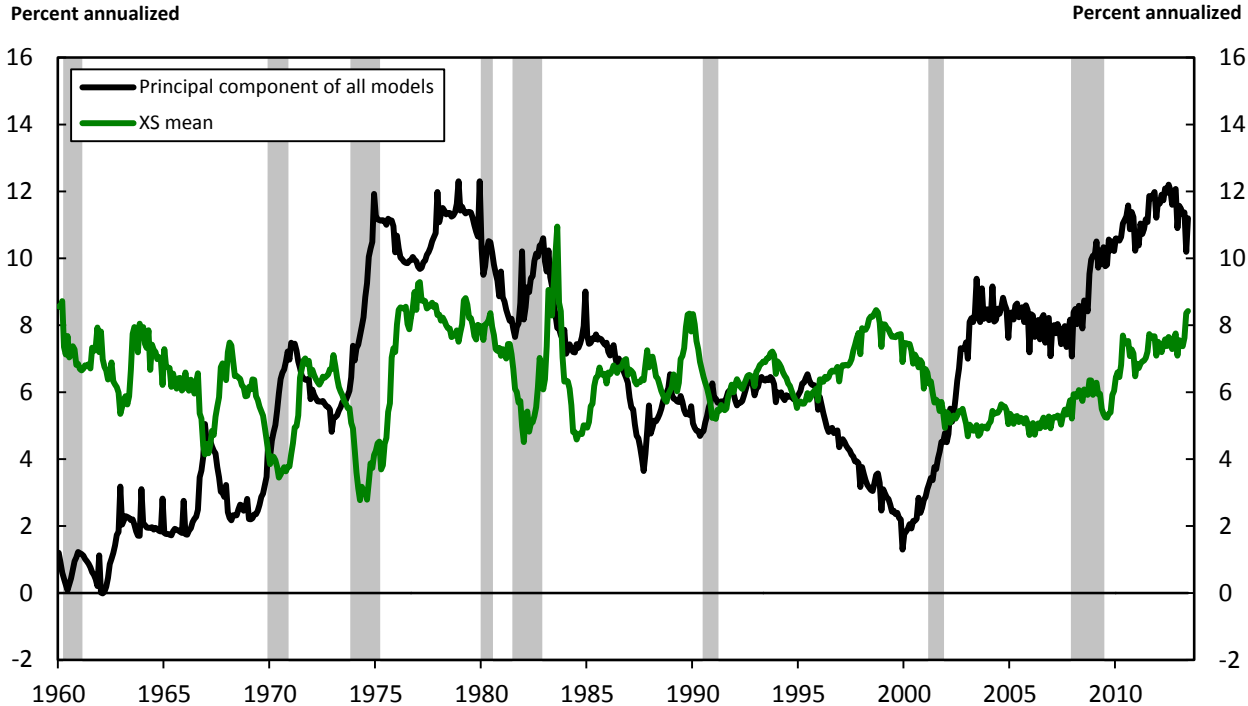


The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (this is the same principal component shown in black in all panels of Figure 1). The models are listed in Tables II to VI.

The 25th and 75th percentiles (solid green lines) give the corresponding quartile of the 20 estimates for each time period, and similarly for the 10th and 90th percentiles (dashed green line).

Shaded bars indicate NBER recessions.

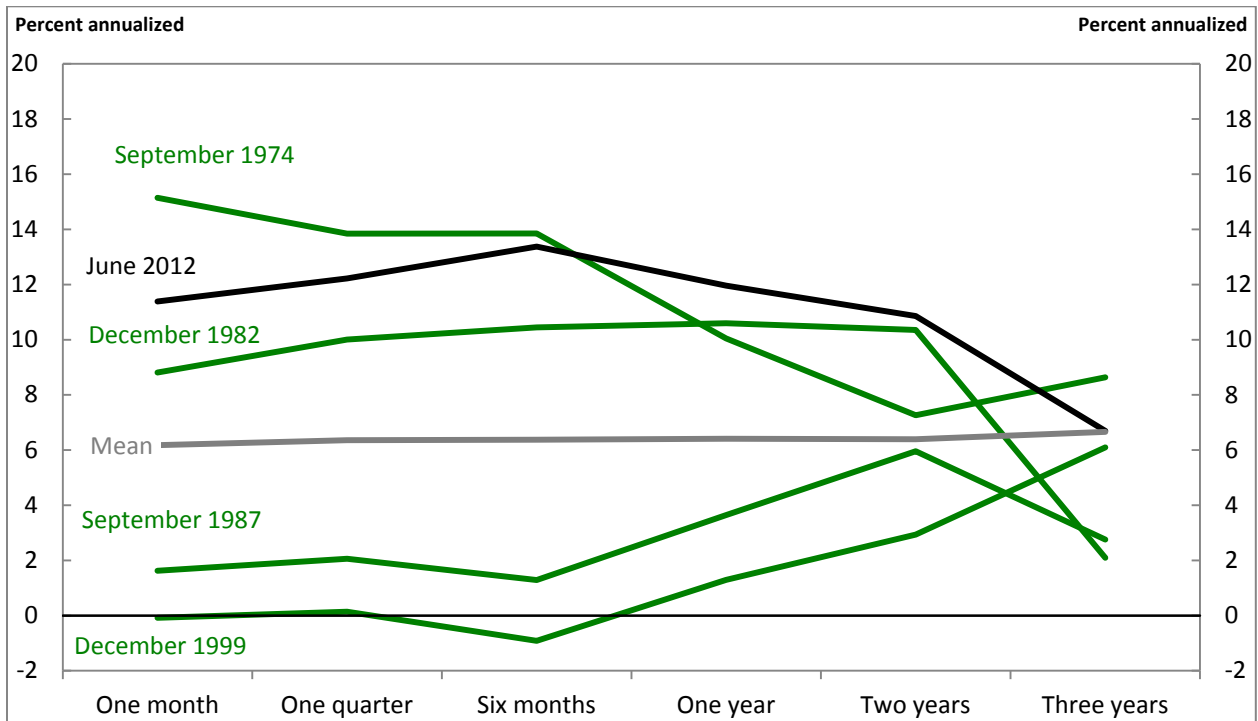
Figure 3: One-year-ahead ERP and cross-sectional mean of models



The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Figures 1 and 2). The green line is the cross-sectional average of models for each time period.

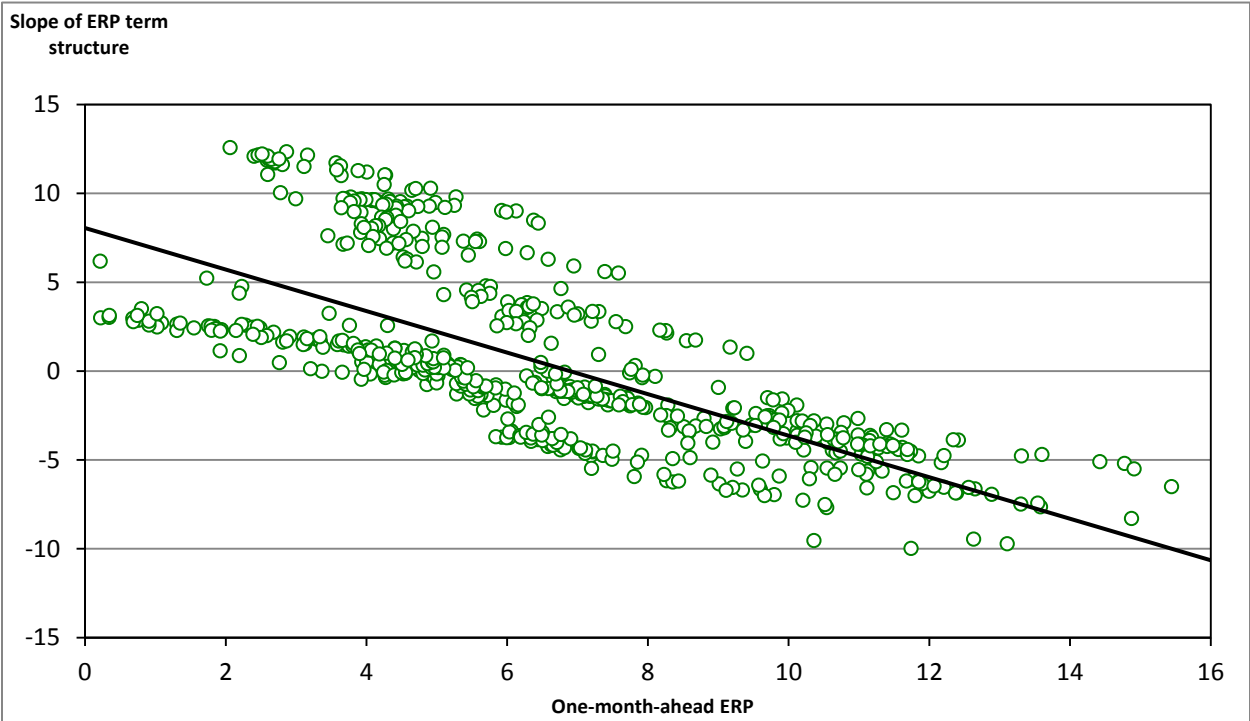
Shaded bars are NBER recessions.

Figure 4: Term structure of the ERP



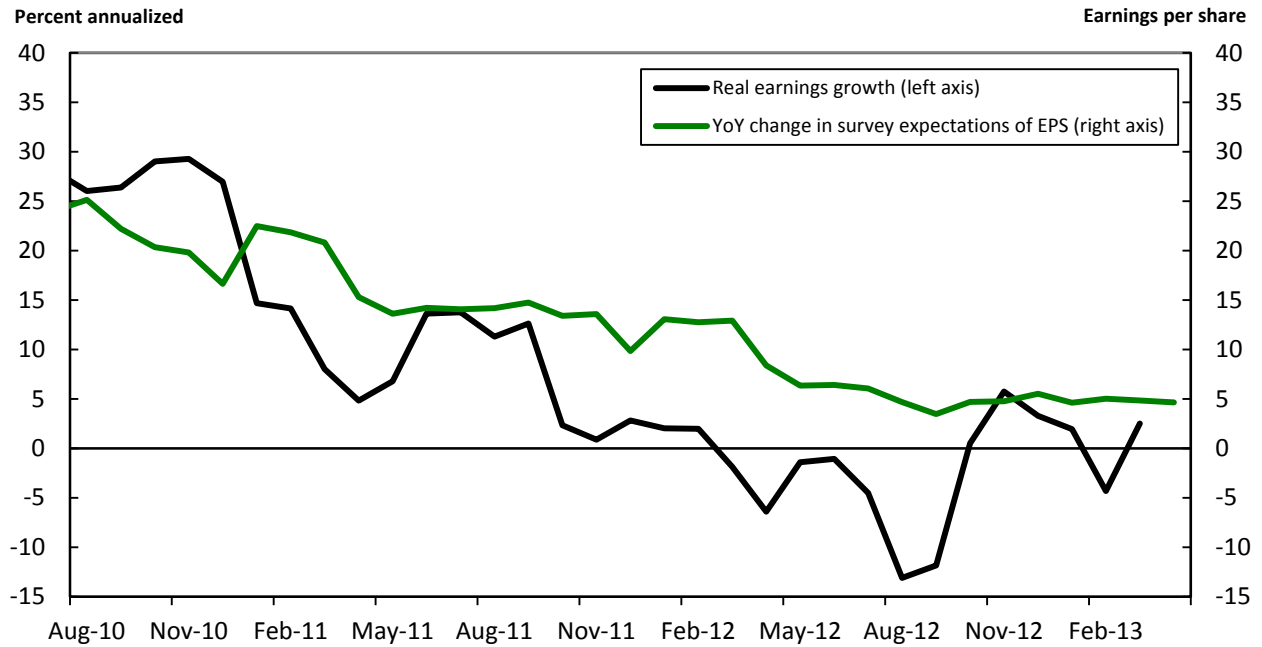
Each line, except for the grey one, shows equity risk premia as a function of investment horizon for some specific months in our sample. We consider horizons of one month, one quarter, six months, one year, two years and three years. The grey line (labeled “Mean”) shows the average risk premium at different horizons over the whole sample January 1960 to June 2013. September 1987 and December 1999 were low points in one-month-ahead equity premia. In contrast, September 1974, December 1982 and June 2012 were peaks in the one-month-ahead equity premium.

Figure 5: Regression of the slope of the ERP term structure on one-month-ahead ERP



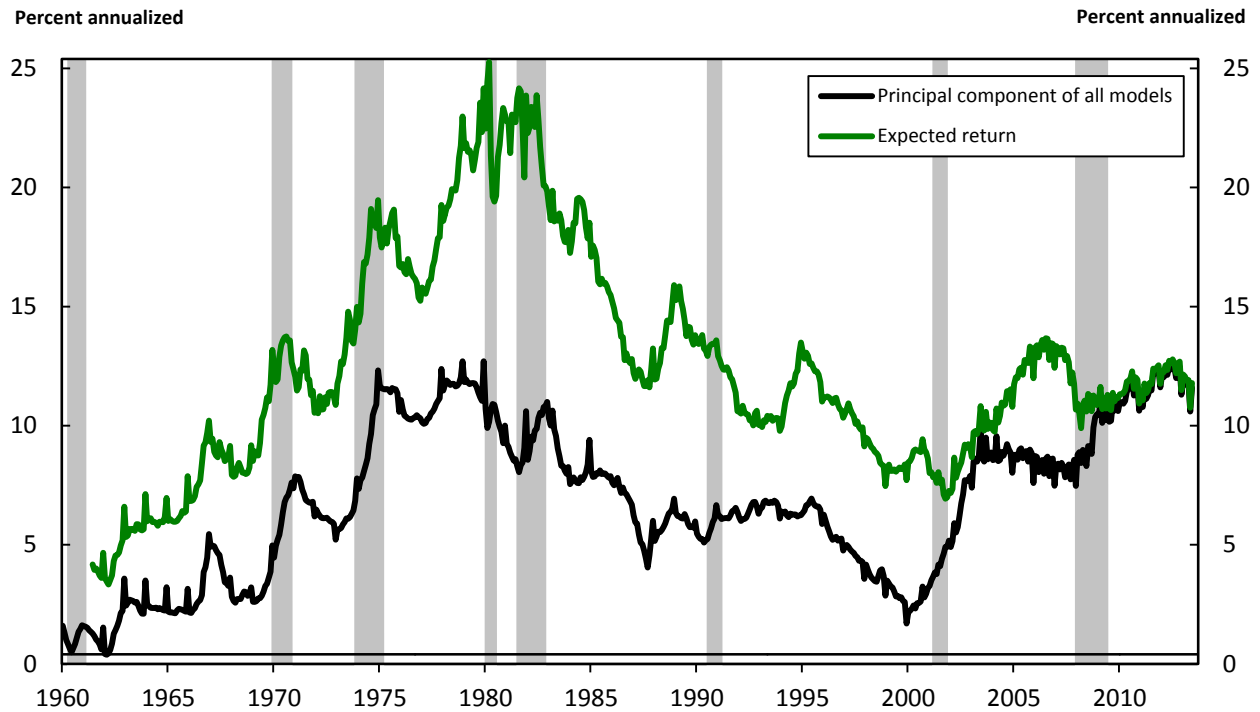
The figure shows monthly observations and the corresponding OLS regression for of the one-month-ahead ERP plotted against the slope of the ERP term structure for the period January 1960 to June 2013. The slope of the ERP term structure is the difference between the three-year-ahead ERP and the one-month-ahead ERP. All units are in annualized percentage points. The one-month-ahead and three-year-ahead ERP estimates used are the first principal components of twenty one-month-ahead or three-year-ahead ERP estimates from models described in Tables II-VI. The OLS regression slope is -1.17 (significant at the 99 percent level) and the R^2 is 50.1 percent.

Figure 6: Earnings behavior



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

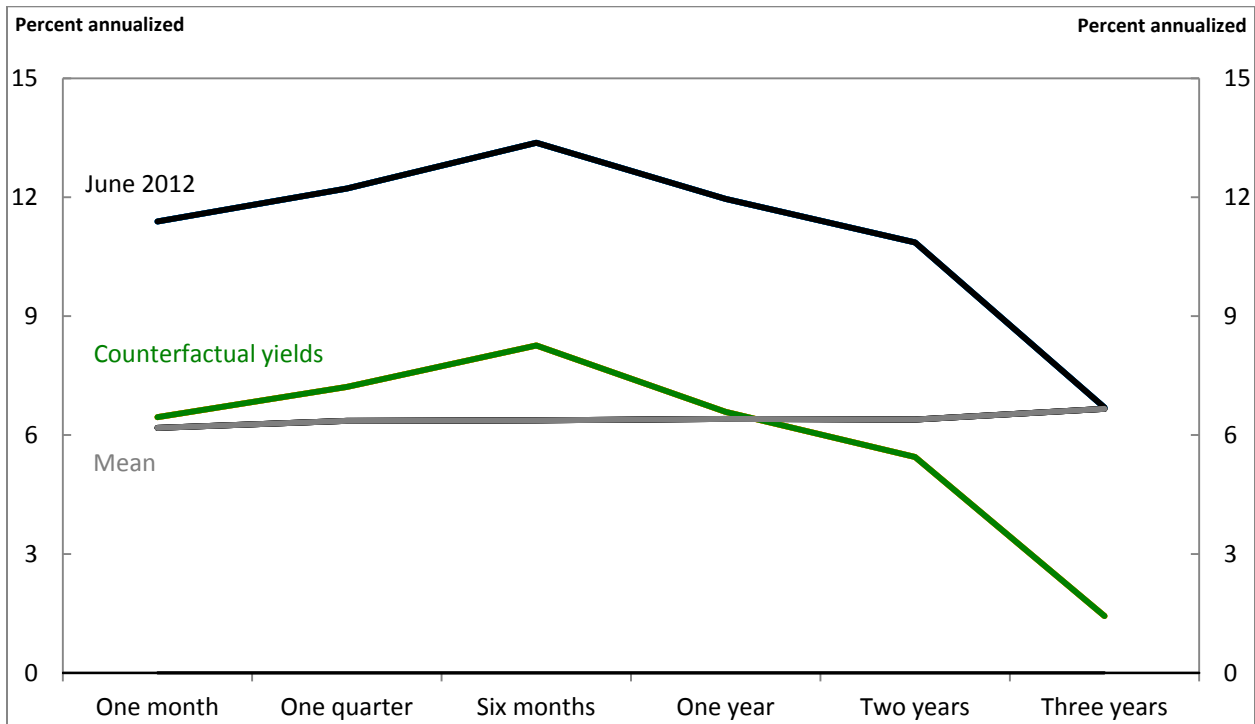
Figure 7: One-year-ahead ERP and expected returns



The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Figures 1, 2 and 3). The green line is the one-year-ahead expected return on the S&P 500, obtained by adding the realized one-year maturity Treasury yield from the principal component (the black line).

Shaded bars are NBER recessions.

Figure 8: Term structure of ERP using counterfactual bond yields



The grey line, labeled “Mean”, shows the mean term structure of the equity risk premium over the sample January 1960 to June 2013. The black line, labeled “June 2012”, shows the term structure for the most recent peak in the one-month-ahead ERP. These two lines are the same as in Figure 4. The green line, labeled “Counterfactual yields”, shows what the term structure of equity risk premia would be in June 2012 if instead of subtracting June 2012’s yield curve from expected returns we subtracted the average yield curve for January 1960 to June 2013.