Fernando Duarte and Carlo Rosa

The Equity Risk Premium: A Review of Models

- The equity risk premium (ERP) is a key quantity in asset pricing that serves as an indicator of economic activity and a determinant of firms' cost of capital, individuals' savings decisions, and government budgeting plans.
- This study estimates the ERP by combining data from twenty models. It finds that the ERP in 2012 and 2013 reached heightened levels—of about 12 percent—not seen since the 1970s.
- The authors attribute the high ERP to unusually low Treasury yields rather than to expectations that stocks would have high returns.
- One implication of the ERP being driven by bond yields rather than expected stock returns is that traditional indicators of the ERP, such as simple valuation ratios, may not be as good a guide to future excess returns as they have been in the past.

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, the savings decisions of individuals, and budgeting plans for governments. Recently, the equity risk premium (ERP) has also moved to the forefront as a leading indicator of the evolution of the economy, a potential explanation for jobless recoveries, and a gauge of financial stability.[1]

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

[1] 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. Offering 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 readily 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.

Fernando Duarte is an economist and Carlo Rosa a senior economist at the Federal Reserve Bank of New York.

fernando.duarte@ny.frb.org; carlo.rosa@ny.frb.org

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/author_disclosure/ad_epr_2015_equity-risk-premium.html.

The authors thank Tobias Adrian and James Egelhof for helpful comments on earlier drafts. All remaining errors are the authors'. This article updates the authors' May 2013 blog post in *Liberty Street Economics,* available at http://libertystreeteconomics.newyorkfed .org/2013/05/are-stocks-cheap -a-review-of-the-evidence.html, and offers a more rigorous and comprehensive treatment than the post. The views expressed in this article 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.

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 reached during the financial crisis in 2009 and at levels similar to those in the mid- and late 1970s. From June 2012 to the end of our sample in June 2013, the ERP has changed little, 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.[2] Nevertheless, we find that the dispersion in estimates across models, while quite large, has been shrinking, potentially signaling increased agreement even when the models differ substantially from one another 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 the result of expected stock returns being high or risk-free rates being low. We conclude that 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 diverse assumptions and data sources that they require.

2. The Equity Risk Premium: Definition

Conceptually, the ERP is the compensation that 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 return are based on broad stock market indexes, such as the S&P 500 or the Dow Jones Industrial Average, which do not include the whole universe of traded stocks and miss several other components of wealth such as housing, private equity, and nontradable human capital. Even if we restricted ourselves to all traded stocks, we would still have several choices to make, such as whether to use value or equal-weighted indexes, 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, nominal Treasury securities are not completely riskless since they are exposed to inflation[3] and liquidity risks, even if we were to assume that there is no prospect of outright default. In this article, we 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 one another and with most of the literature.

While putting the concept of the ERP into practice has its challenges, we can precisely define the ERP mathematically. First, we decompose stock returns[4] 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 expectations at this stage, although it will be a feature

[2] 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), and Neely et al. (2014).

[3] 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.

[4] Throughout this article, all returns are net returns. For example, a 5 percent return corresponds to a net return of 0.05 as opposed to a gross return of 1.05.

Table 1
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 University/CFO Magazine Global Business Outlook Survey
Damodaran (2012)	ERP estimates
Gurkaynak, Sack, and Wright (2007)	Zero-coupon nominal bond yields for all maturities[a]
Gurkaynak, Refet, Sack, and Wright (2010)	Zero-coupon TIPS (Treasury Inflation-Protected Securities) 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 (Federal Reserve Bank of St. Louis)	Corporate bond Baa-Aaa spread and the National Bureau of Economic Research recession indicator

Notes: 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.

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

of many of the models we consider. The ERP at time t for horizon k is defined as

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

where R^f_{t+k} 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 advance.[5] 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 $error_{t+k}$ has mean zero and is 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.

[5] 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.

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 those we consider. We selected the models in our study based on three criteria: the recent academic literature, widespread use of the models by practitioners, and data availability. Table 1 describes the data we use and their sources, all of which are either readily available or standard in the literature.[6] With a few exceptions, all data are monthly from January 1960 to June 2013. Appendix A provides further detail.

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 only.[7] This helps minimize look-ahead bias and makes any out-of-sample evaluation of

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

[7] The one exception is the cross-sectional model of Adrian, Crump, and Moench (2014), which is constructed using full-sample regression estimates.

Table 2
Models Based on the Historical Mean of Excess 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

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 that it assumes the future will behave like the past—in other words, that 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 2 shows the two versions of historical mean models that we use.

3.2 Dividend Discount Models

All dividend discount models (DDMs) 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

$$P_t = \frac{D_t}{\rho_t} + \frac{E_t\,[D_{t+1}]}{\rho_{t+1}} + \frac{E_t\,[D_{t+2}]}{\rho_{t+2}} + \frac{E_t[D_{t+3}]}{\rho_{t+3}} + \ldots \,, \tag{3}$$

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 they can also arise from spinoffs, buyouts, mergers, and buybacks. 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

$$\rho_{t+k} = 1 + R^f_{t+k} + ERP_t(k). \tag{4}$$

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 for this is that we plug 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—in other words, what ERP value is *implied* by equation (3).

DDMs are forward-looking and are consistent with no arbitrage. In fact, equation (3) must hold in any economy with no arbitrage.[8] Another advantage of DDMs is that they are easy to implement. A drawback of DDMs is that the results are sensitive to how we compute expectations of future dividends. Table 3 displays the DDMs that we consider and a brief description of their different assumptions.

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 ERP.[9] 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.

[8] 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, DDMs do allow for a specific kind of bubble.

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

Table 3
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 [Treasury Inflation-Protected Securities])
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 dividend discount model. 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 year six and onward is 2.2 percent
Damodaran (2012)	Six-stage dividend discount model. Dividend growth in the first five stages is estimated from analysts' 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

Sources: See Appendix A and Table 1 for full source details.

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

$$(5) \qquad R^i_{t+k} - R^f_{t+k} = \alpha^i \times \textit{state variables}_{t+k} + \beta^i \times \textit{risk factors}_{t+k} + \textit{idiosyncratic risk}^i_{t+k}.$$

In equation (5), R^i_{t+k} is the realized return on a stock or portfolio i from time t to $t+k$. $\textit{State variables}_{t+k}$ are any economic indicators that help identify the state of the economy and its likely future path. $\textit{Risk factors}_{t+k}$ are any measures of systematic contemporaneous covariation in returns across all stocks or portfolios. Of course, some economic indicators can be both state variables and risk factors at the same time. Finally, $\textit{idiosyncratic risk}^i_{t+k}$ is the component of returns that is particular to each individual stock or portfolio that is not explained by $\textit{state variables}_{t+k}$ or $\textit{risk factors}_{t+k}$ (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 Treasury securities, 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 factors.[10] 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

$$(6) \qquad R^i_{t+k} - R^f_{t+k} = \lambda_t(k) \times \hat{\beta}^i,$$

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 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 the portfolios, state variables, and risk factors that are used (Harvey, Liu, and Zhu 2014), and that it is not as easy to implement as most of the other options. Table 4 displays the

[10] 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.

Table 4
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

Sources: See Appendix A and Table 1 for full source details.

Table 5
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, the estimate of the ERP is replaced 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 six variables: the closed-end fund discount, New York Stock Exchange share turnover, the number of initial public offerings, the average first-day returns on initial public offerings, the equity share in new issues, and the dividend premium

Sources: See Appendix A and Table 1 for full source details.

cross-sectional models in our study, together with the state variables and risk factors they use.

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":

$$(7) \qquad R_{t+k} - R^f_{t+k} = a + b \times Fundamental_t + error_t \, .$$

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

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

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 a minimal number of assumptions; there is no concept of equilibrium and no absence of arbitrage necessary for the method to be valid.[11] In addition, implementation is quite simple, since it only involves running ordinary least-square regressions. The challenge is to select the variables to include on the right-hand side of equation (7), since results can change substantially depending on the variables that are used to take the role of "fundamentals." Including more than one predictor gives poor out-of-sample performance 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 5 shows the time-series regression models that we study.

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

Table 6
Surveys

Graham and Harvey (2012)	Since 1996, the Duke University/CFO Magazine Global Business Outlook Survey has asked chief financial officers about the one- and ten-year-ahead ERP. We take the mean of all responses

Sources: See Appendix A and Table 1 for full source details.

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 whom 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 the supply and demand of investors such as the ones taking the surveys. However, 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 article, we use the Duke University/CFO Magazine Global Business Outlook Survey of chief financial officers 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 6).

4. Estimation of the Equity Risk Premium

We now study the behavior of the twenty models under consideration by conducting principal component analysis. Since forecast accuracy can be substantially improved through the combination of multiple forecasts,[12] the optimal strategy to forecast excess stock returns may consist of combining all of 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 of 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. First, we de-mean all ERP estimates and find their variance-covariance matrix. Second, 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. Third, we add to the linear combination just obtained, which has a mean of 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 variation in the first principal component then provides an estimate of the conditional ERP.[13] The share of the variance of the underlying models explained by this principal component is 76 percent, suggesting that little would be gained from examining principal components beyond the first.[14]

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

The first two columns in Table 7 show the mean and standard deviation of each model's estimates. The unconditional

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

[13] 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 observations are missing, either because some models can only be obtained at frequencies lower than monthly or because the necessary data are 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 nonmissing 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.

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

Table 7

ERP Models

		Mean	Standard Deviation	PC Coefficients $\hat{w}^{(m)}$	Exposure to PC $load_1^{(m)}$
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
Dividend Discount Models (DDM)	Gordon (1962): E/P minus nominal ten-year yield	-0.1	2.1	-0.01	0.001
	Shiller (2005): 1/CAPE minus nominal ten-year yield	-0.4	1.8	-0.10	0.011
	Gordon (1962): E/P minus real ten-year yield	3.5	2.1	0.69	-0.077
	Gordon (1962): Expected E/P minus real ten-year yield	5.3	1.7	-0.78	0.208
	Gordon (1962): Expected E/P minus nominal ten-year 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 and 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) Duke University/ CFO Magazine Global Business Outlook Survey	3.6	1.8	0.72	0.264
	All models	5.7	3.2	0.78	-0.065

Sources: See Appendix A and Table 1 for full source details.

Notes: 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 percentages. The third statistic, "PC Coefficients $\hat{w}^{(m)}$", is the weight that the first principal component places on each model (normalized to sum to one). The fourth is the "Exposure to PC $load_1^{(m)}$", 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). E/P is earnings-to-price. CAPE is cyclically adjusted price-to-earnings. D/P is dividend-to-price.

mean of the ERP across all models is 5.7 percent, with an average standard deviation of 3.2 percent. DDMs 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 7 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.

Chart 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 chart gives a sense of how the time series move together and how much they covary with the first principal component. Table 8 shows the correlations among models. Chart 1 and Table 8 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 DDMs 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.

Chart 1 also shows that the first principal component covaries negatively with historical mean models but positively with DDMs and cross-sectional regression models. Time-series regression models are also positively correlated with the first principal component, although this is not so

Chart 1
ERP Estimates for All Models

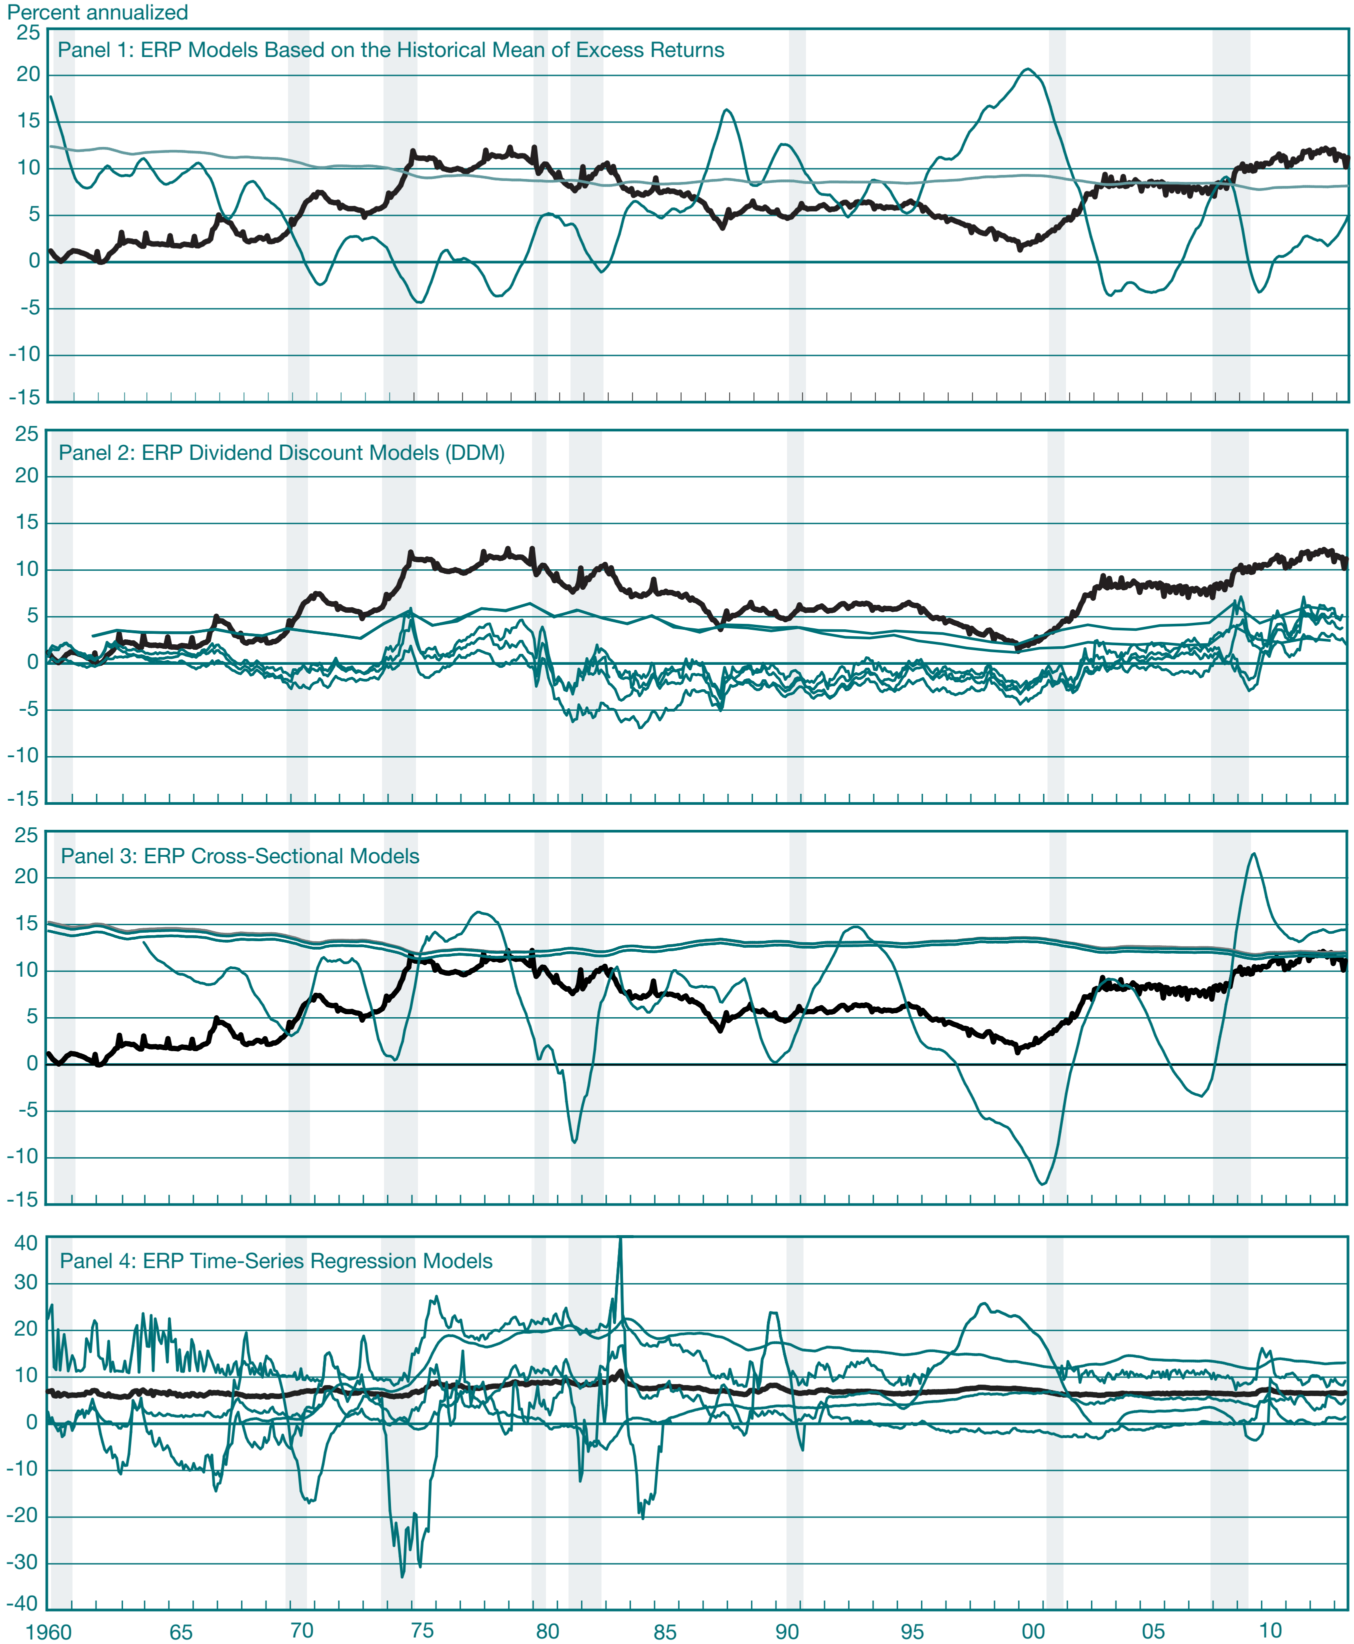

Chart 1 (Continued)

ERP Estimates for All Models

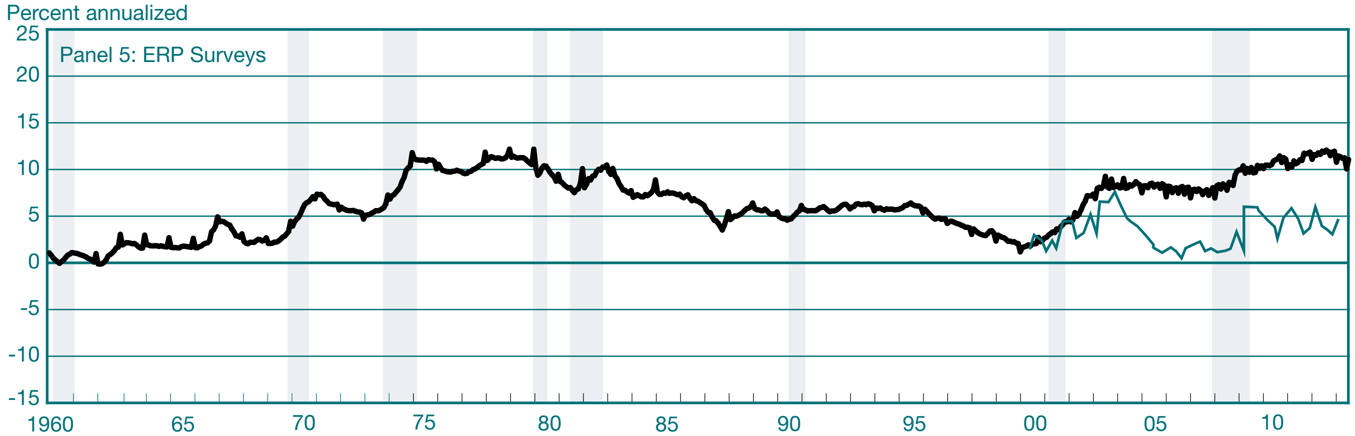

Sources: See Appendix A and Table 1 for detailed source information.

Notes: Each green line gives the one-year-ahead equity risk premium from one of the models listed in Tables 2 through 6. Panel 1 shows the estimates for models based on the historical mean of excess returns; these models are listed in Table 2. Panel 2 shows estimates computed by the dividend discount models listed in Table 3. Panel 3 uses the cross-sectional regression models listed in Table 4. Panel 4 uses the time-series regression models listed in Table 5. Panel 5 shows the estimate obtained from the survey cited in Table 6. In all panels, the black line is the first principal component of all twenty models (it can look different across panels because of different scales used in the y-axis.) The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

clearly seen in Panel 4 of Chart 1 because of the high volatility of time-series regression ERP estimates. The last panel shows that the survey of CFOs does track the first principal component quite well at low frequencies (for example, 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:

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

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 7, labeled "PC Coefficients," shows the weights $w^{(m)}$ normalized to sum up to one to facilitate comparison; in other words, the table reports the weights $\widehat{w}^{(m)}$, where

(10) $$\widehat{w}^{(m)} = \frac{w^{(m)}}{\Sigma_{m=1}^{20} w^{(m)}}.$$

The first principal component puts positive weight on models based on the historical mean, cross-sectional regressions, and the survey of CFOs. It weights DDMs 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 7, 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 the twenty principal components:

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

where m indexes the model and i indexes the principal components. The values in the last column of Table 7 are the loadings on the first principal component ($i = 1$) for each model ($m = 1, 2, \ldots, 20$), again normalized to one for ease of comparability:

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

Most models have a positive loading on the first principal component; whenever the loading is negative, it tends to be relatively small. This means that the first principal component, as expected, is a good explanatory variable for most models. Looking at the third and fourth columns of Table 7 together,

Table 8
Correlation of ERP Models

	Long-run mean	Mean past five years	E/P-ten year	1/CAPE-ten year	E/P-real ten year	Exp E/P-real ten year	Exp E/P-ten year	Two-stage DDM	Six-stage DDM	Free cash flow	Fama and French	Carhart	Duarte	Adrian, Crump, and Moench	D/P	Goyal and Welch	Campbell and Thompson	Fama and French	Sentiment	CFO survey
Long-run mean	100																			
Mean past five years	32	100																		
E/P-ten year	8	15	100																	
1/CAPE-ten year	-9	0	78	100																
E/P-real ten year	-11	25	98	23	100															
Exp E/P-real ten year	-58	42	70	84	60	100														
Exp E/P-ten year	-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										
Fama and French	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							
Adrian, Crump, and Moench	-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					
Goyal and Welch	25	12	25	21	-7	-36	-60	20	29	-9	7	13	14	-24	61	100				
Campbell and Thompson	27	31	14	-7	81	49	-60	28	-51	-40	60	57	58	-33	54	50	100			
Fama and French	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

Sources: See Appendix A and Table 1 for additional source details.

Notes: 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 2-6). Except for the chief financial officer (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 the monthly series and then computing quarterly correlations. E/P is earnings-to-price. CAPE is cyclically adjusted price-to-earnings. DDM is dividend discount model. D/P is dividend-to-price.

we can obtain additional information. For example, when a model has a very high loading (fourth column) accompanied by a very small PC coefficient (third column), it likely means 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. However, if the PC coefficient and loading are both high, the corresponding model is likely providing information not contained in other measures.

CHART 2

One-Year-Ahead ERP

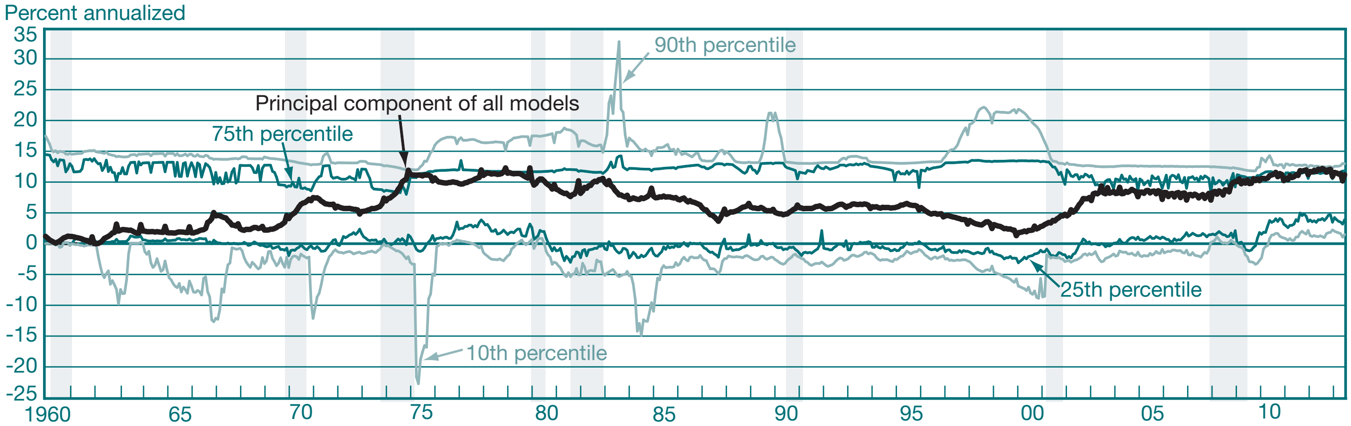

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (this is the same principal component shown in black in all panels of Chart 1). The models are listed in Tables 2 through 6. The green lines give the corresponding percentiles of the twenty estimates for each time period. The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

Chart 2 shows the first principal component of all twenty models in black (the black line is the same principal component shown in black in each of the panels of Chart 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 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 Chart 1, a finding that highlights the usefulness of principal component analysis. Similarly high levels were observed 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.

Chart 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—is 11.6 percent on average. It has recently compressed, mostly because the models in the bottom of the distribution have had higher ERP estimates since 2010 while the 75th percentile has remained fairly constant. The lowest value for the interquartile range, 6.8 percent, was reached in 2012. The cross-sectional standard deviation in ERP estimates (not shown in the chart) also decreased from 10.2 percent in January 2000 to 4.3 percent in June 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 Chart 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 with past values. The cross-sectional mean can be useful, but compared with the first principal component, it has a few undesirable features as an overall measure of the ERP. 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. And last, it has a smaller correlation with the realized returns it is supposed to predict.

Chart 3

One-Year-Ahead ERP and Cross-Sectional Mean of Models

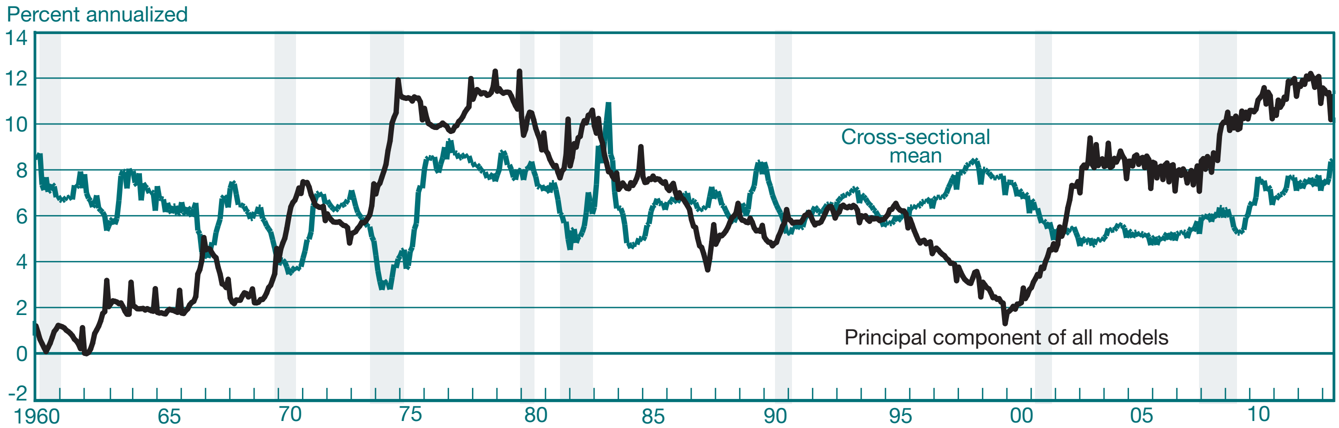

Sources: See Appendix A and Table 1 for detailed source information.

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Charts 1 and 2). The green line is the cross-sectional average of models for each time period. The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

5. The Term Structure of Equity Risk Premia

In Section 2, we described the term structure of the ERP—the expected excess returns 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 horizons.[15] For example, if we want to take the historical mean of returns as our estimate, we can take the mean of returns over a one-month, six-month, or 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. DDMs, on the other hand, have little variation across horizons. In fact, all the DDMs we consider have a constant term structure of expected stock returns, and the only term structure variation in ERP estimates comes from risk-free rates.[16]

Chart 4 plots the first principal components of the ERP as a function of investment horizon for some dates when the ERP

Chart 4

Term Structure of the ERP

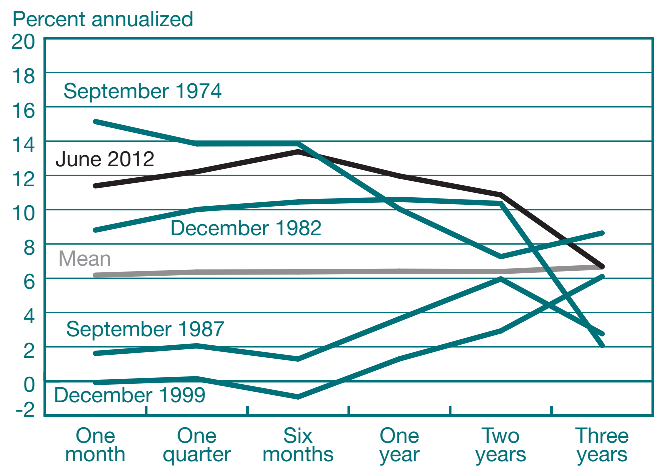

Notes: Each line, except for the gray one, shows equity risk premia as a function of investment horizon for some specific months in our sample. We consider horizons of one month, one quarter, six months, one year, two years, and three years. The gray line (labeled "Mean") shows the average risk premium at different horizons over the full 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.

[15] For other ways to estimate the term structure of the ERP using equilibrium models or derivatives, see Äit-Sahalia, Karaman, and Mancini (2014), Ang and Ulrich (2012), van Binsbergen et al. (2014), Boguth et al. (2012), Durham (2013), Croce, Lettau, and Ludvigson (2015), Lemke and Werner (2009), Lettau and Wachter (2011), and Muir (2013), among others.

[16] 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 DDMs with a nonconstant term structure of expected excess returns.

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 from 68 to 94 percent. The gray line in Chart 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.

The first observation is that the term structure of the ERP has significant time variation and can be flat, upward sloping, or downward sloping. Chart 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 true more generally: there is a strong negative correlation between the level and the slope of the ERP term structure of -71 percent. Chart 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.

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.

Chart 6 shows that earnings are unlikely to be the reason 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 declined from 2010 to 2013, making earnings growth an unlikely reason for the high ERP in the corresponding period. The black line shows the realized monthly growth rates of real earnings for the S&P 500 expressed in annualized percentages. Since 2010, earnings growth has been declining, hovering around zero for the last few months of the sample. At the end of the sample, it stands at 2.5 percent, which is near its long-run average.

Another way to examine whether a high ERP is caused by discount rates or cash flows is shown in Chart 7. The

Chart 5

Regression of the Slope of the ERP Term Structure on One-Month-Ahead ERP

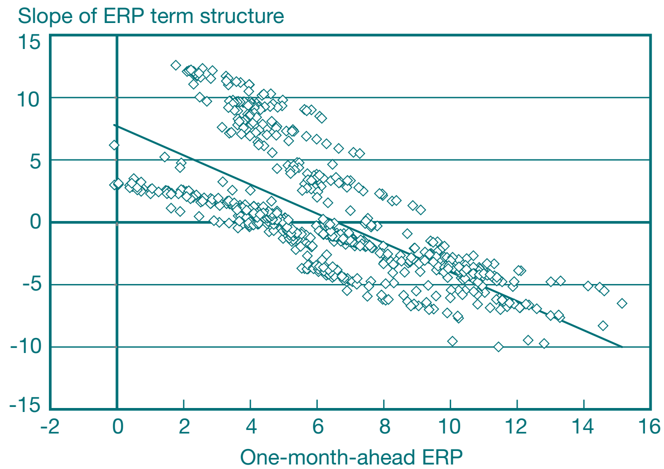

Notes: The chart shows monthly observations and the corresponding OLS regression for the one-month-ahead ERP plotted against the slope of the ERP term structure for the period January 1960 to June 2013. The slope of the ERP term structure is the difference between the three-year-ahead ERP and the one-month-ahead ERP. All units are in annualized percentages. 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 2-6. The OLS regression slope is -1.17 (significant at the 99 percent level) and the R^2 is 50.1 percent.

Chart 6

Earnings Behavior

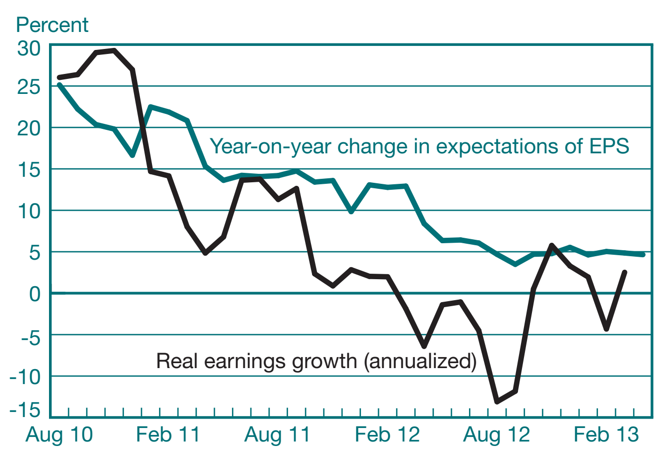

Notes: The black line shows the monthly growth rate of real S&P 500 earnings, annualized and in percentages. The green line shows the year-on-year change in the mean expectation of one-year-ahead earnings per share (EPS) for the S&P 500 from a survey of analysts provided by Thomson Reuters I/B/E/S.

Chart 7

One-Year-Ahead ERP and Expected Returns

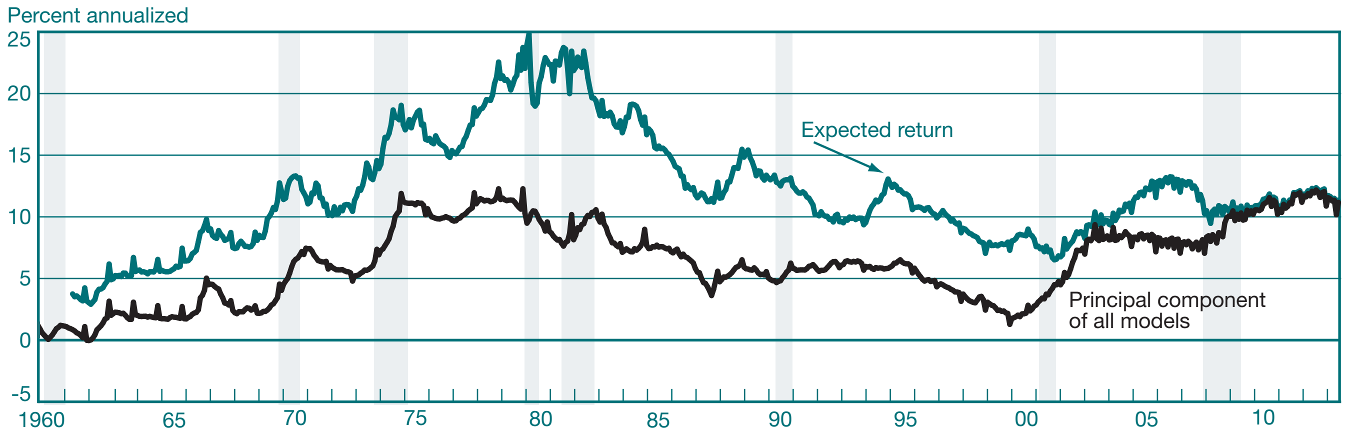

Notes: See Appendix A and Table 1 for detailed source information.

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Charts 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). The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

black line is the same one-year-ahead ERP estimate shown in Chart 2. The green line simply adds the realized one-year Treasury yield to obtain expected stock returns. The chart shows that expected stock returns have increased since 2000, similar 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 the result of exceptionally low bond yields since the end of the financial crisis.

Chart 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 Chart 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 the period 1960-2013. In other words, we replace R^f_{t+k} in equation (2) by the mean of R^f_{t+k} over t. The result of this counterfactual is shown in Chart 8 in green. Using average levels of bond yields brings the whole term structure of the ERP much closer to its mean level (the gray 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 2012-13.

Chart 8

Term Structure of the ERP Using Counterfactual Bond Yield

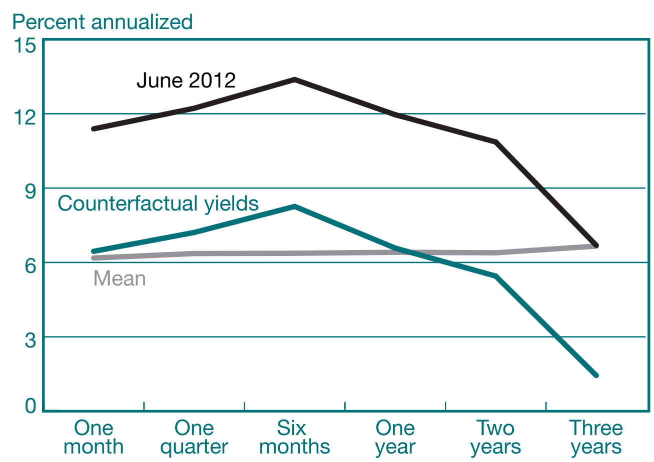

Notes: The gray line (labeled "Mean") shows the mean term structure of the equity risk premium over the full sample, January 1960 to June 2013. The black line (labeled "June 2012") shows the term structure for the most recent peak in the one-month-ahead ERP. These two lines are the same as in Chart 4. 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.

7. Conclusion

In this article, we analyze twenty different models of the equity risk premium by considering the assumptions and data required to implement them, and how the models relate to one another. 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 with 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.

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 twenty-five portfolios sorted on size and book-to-market, ten 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 ten-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.html Quarterly frequency; 6/6/2000 to 6/5/2013. We use the answer to the question "Over the next ten years, I expect the average annual S&P 500 return will be: expected return:" and the analogous question 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 does not).
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; 1/1/1960 to 7/1/2013 for Baa minus Aaa bond yield spread and recession indicator.
Compustat	Book value per share (variable BKVLPS) Annual frequency; 12/31/1977 to 12/31/2012.
Thomson Reuters I/B/E/S	Earnings per share (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 (Federal Reserve Bank of St. Louis)	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.

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." 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." Journal of Economic Perspectives 21, no. 2 (Spring): 129-51.

Binsbergen, J. V., 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, no. 4: 559-83.

Croce, M., M. Lettau, and S. Ludvigson. 2015. "Investor Information, Long-Run Risk, and the Term Structure of Equity." Review of Financial Studies 28, no. 3: 706-74.

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, 241-68. Amsterdam: North-Holland.

Duarte, F. 2013. "Inflation Risk and the Cross Section of Stock Returns." Federal Reserve Bank of New York Staff Reports, no. 621, May.

Duarte, F., L. Kogan, and D. Livdan. 2013. "Aggregate Investment and Stock Returns." Unpublished paper, M.I.T.

Durham, J. B. 2013. "Arbitrage-Free Models of Stocks and Bonds." Federal Reserve Bank of New York Staff Reports, no. 656, December.

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.

———. 1992. "The Cross Section of Expected Stock Returns." Journal of Finance 47, no. 2 (June): 427-65.

———. 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.

References (Continued)

Gordon, M. J. 1962. The Investment, Financing, and Valuation of the Corporation. Homewood, Illinois: 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): 1455-508.

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., 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.

Gurkaynak, R. S., B. Sack, and J. H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." Journal of Monetary Economics 54, no. 8 (November): 2291-304.

Hall, R. E., 2014. "The Routes Into and Out of the Zero Lower Bound." Proceedings of the 2013 Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City.

Harvey, C. R., Y. Liu, and H. Zhu. 2014. ". . . and the Cross Section of Expected Returns." NBER Working Paper no. 20592, October. Available at SSRN: http://ssrn.com/abstract=2249314

Kelly, B., and S. Pruitt. 2013. "Market Expectations in the Cross Section of Present Values." 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, January.

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, April.

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 U.S. Bonds, Credit and Equities." JP Morgan Investment Strategies Series, no. 11 (January).

Piazzesi, M., and M. Schneider. 2007. "Equilibrium Yield Curves." In NBER Macroeconomics Annual 2006 21: 389-472. Cambridge, Mass.: MIT Press.

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.* 2nd ed. Princeton: Princeton University Press.

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." In G. Elliott, C. W. J. Granger, and A. Timmermann, eds., Handbook of Economic Forecasting, vol. 1, 135-96. Amsterdam: Elsevier.