Understanding Bond Risk Premia Uncovered by the Term Structure

Juliusz F. Radwański*

Vienna Graduate School of Finance (VGSF)

November 1, 2010

This paper makes a step towards understanding the term-structure forecasts of bond risk premia. Two economically interpretable variables, the level of nominal forward rates, and one-year-ahead expected inflation extracted from the forwards (IE factor) are enough to summarize virtually all of predictive power for excess bond returns contained in the factor of Cochrane and Piazzesi [2005] (CP). The IE is constructed by regressing realized inflation on five forward rates, which improves substantially upon predictive regressions for inflation that use the term spread and the spot rate only. The intuition can be well explained in terms of natural cointegration between nominal level and the IE factor. Deviations from this long-term relation explain 95 percent of the variation in the CP factor, and have clear business cycle properties. Apart from that, they are large around inflationary 70's and 80's, and more moderate in the 50's, 60's, and in the most recent period which includes the financial crisis. This explains weaker predictability of excess bond returns in the subsamples that cover the most recent years.

Keywords: Term structure, expectations hypothesis, predictability, bond risk premia, inflation expectations, business cycle, crisis.

JEL Classification Numbers: E31, E43, E44, G12.

---

*Address: Vienna Graduate School of Finance (VGSF), c/o WU (Vienna University of Economics and Business), Heiligenstädter Strasse 46-48, DG 1.21, 1190 Vienna, Austria, telephone: +43-1-31336-6322, e-mail: *juliusz.radwanski@vgsf.ac.at*. The author is grateful to Alois Geyer and Josef Zechner for advice and suggestions, and to John Cochrane and Monika Piazzesi for helpful comments.

1 Introduction

How to uncover the risk premium in the bond market in the best possible way using the information included in the term structure? What is the relation between the risks uncovered by the yields or forward rates, and macroeconomic variables? And finally, can statistical yield-curve factors be supported by meaningful economic interpretations?

The answer to the first question can be found in the seminal paper of Cochrane and Piazzesi [2005], who run predictive regressions of realized excess bond returns on five forward rates with maturities between one and five years. Their tent-shaped factor contains significant information on future bond returns, predicting them with high R^2 coefficients.

This paper provides the answers to the remaining two questions. The information about excess bond returns contained in forward rates can be given a relatively simple economic interpretation, which explains why statistical factor of Cochrane and Piazzesi [2005] (CP) predicts excess bond returns.[1] I show, that two economically interpretable variables, the level of nominal forward rates and the one-year-ahead inflation expectation contained in the term structure (IE factor) are enough to summarize virtually all of their predictability. This is remarkable, given that neither of my factors is designed *a priori* for this task.

The level of the term structure is obtained by means of principal components analysis. I extract the information about expected inflation directly from forward rates, running predictive regressions with realized CPI growth on the left hand side on five forward rates with maturity between one and five years, replicating the logic in Cochrane and Piazzesi [2005], but in a different context. This improves significantly upon more traditional predictive regressions for inflation, which usually use the term spread and the spot rate as the only explanatory variables. Interestingly, the coefficients of the linear combination of forward rates that form the IE factor look like an inverted tent, very similar to the one from the regressions of Cochrane and Piazzesi [2005], which is the first hint that it may be helpful in predicting bond returns.

[1]In a way, this lends additional support to their findings.

Neither of my two variables predict excess bond returns in univariate regressions. It is the distance between the two that matters, which points to an interpretation in terms of cointegration relation. In the long run, the nominal level of the term structure is naturally co-integrated with the IE. Since expected inflation is persistent, IE summarizes well expected inflation for some years in the future. Sometimes, expected inflation is further below the level than implied by the long-run relation. In other words, market prices of bonds are seemingly below the fundamentals, which results in a correction, during which high excess bond returns are realized. The correction comes from both level and IE, and lasts about two years. Importantly, the deviations are very highly correlated with the CP factor (0.97), which allows to think of both as approximately the same thing. Since both level and expected inflation are formed using market prices, it is very unlikely that irrationality or inefficiency played a role in observed predictability patterns, lending support to risk-based explanations. The deviations from cointegration are found to be relatively short lived on average, with a half-life of exactly one year. After two years, only 24 percent of the initial CP factor remains. At longer horizons, bond risk premium is virtually zero in expectations.

The interpretation of the CP factor in terms of cointegration leads to important conclusions about testing the expectations hypothesis. It is known, that in the presence of cointegration, one should estimate it using the whole sample, exploiting the super-consistency property of the estimators.[2] In other words, out-of-sample tests of the expectations hypothesis may have limited power to reject the incorrect null of no predictability. I also show, that it is important to include the inflationary 80's in the sample, because the deviations were quite large at that time, which allows to estimate cointegrating relation efficiently.[3] In the more recent times including the crisis, and in the relatively calm 50's and 60's, the level and expected inflation were quite highly correlated, which resulted in exactly offsetting effect from both of them on expected returns.

In this study, I concentrate on predictability of the single factor driving excess bond

[2] See Lettau and Ludvigson [2010].

[3] This is conditional on assuming no structural breaks in the whole sample.

returns of all maturities in the same direction, which was found by Cochrane and Piazzesi [2005] to contribute 99.5 percent to the total predictable variance in bond risk premia. By disaggregating realized excess returns into more primitive realized forward premia, I find the single factor less important but still dominant, with a contribution of about 90 percent to the total predictable variance.

1.1 Related literature

There are several papers closely related to the current study. I start the overview with Cochrane and Piazzesi [2005], and Ludvigson and Ng [2009]. Cochrane and Piazzesi extend the study of Fama and Bliss [1987], who test the ability of the forward-spot spread to predict excess returns of bonds with different maturities. The extension of Cochrane and Piazzesi adds all available forward rates[4] to predictive regressions. Both studies strongly reject the expectations hypothesis. Cochrane and Piazzesi find additionally, that a single factor predicts returns of bonds with all maturities (between two and five years), and that intermediate forward rates add predictive power to that obtained by using the term spread only. For interested readers, I provide an updated evidence on the results of both papers in Appendix A. Here, it is only important to note that neither of them provides economic intuition behind the statistical factors they use. The present paper can be thus considered complementary to their work.

Ludvigson and Ng [2009] study the ability of macroeconomic fundamentals to predict excess bond returns. They use a large set of 132 measures of economic activity to estimate common factors in the fundamentals, using dynamic factor analysis. The results suggest, that macro factors explain 26% of variance in bond risk premia. Interestingly, the predictive power of their factors is largely independent of that contained in the CP factor, which remains an important predictor in multivariate regressions. Together with the CP, macro factors explain up to 44% of the variation in excess bond returns, which means that about

[4]They use the so-called Fama-Bliss dataset, described in more detail in the next section of this paper.

18% must be attributed to the former. In contrast, the current paper shows that economically interpretable factors can account for *all* variation in excess bond returns, once the information set is measured carefully. It seems, that macroeconomic time series do not capture well the expectations of market participants, and that one needs to use market-based information to draw inference about risk premia.

In an independent study, Cieslak and Povala [2010] also stress the importance of equilibrium relationships between yields and slow-moving components of macroeconomic variables. They exploit information in a moving average of inflation and savings to construct a cyclical variable that predicts excess bond returns with impressive R^2 coefficient of up to 60 percent. The difference between their study and mine is that I focus on the information content of the yield curve, and provide interpretation of yield-curve based forecasting factors, which is complementary to their study.

The results in my paper are also loosely related to Campbell and Ammer [1993], which is a classical paper on the properties of realized excess bond returns. They decompose them into components that come from revised expectations of future excess returns, future interest rates, and future inflation.[5] Their VAR specification includes the term spread as one of predictors of excess returns and inflation. In the early sub-periods of their data between 1952 and 1979, only the latter of the three components contributed significantly to the total variance of the excess returns. For the sub-periods that included the 80's, innovations to future excess returns correlated negatively with innovations to expected inflation, which tended to have offsetting effect on bond prices. The current paper can be considered an extension of their study, by providing some insights into their results.

This paper also contributes to the literature on forecasting macroeconomic variables using the term structure. One of the very first papers that pointed to predictive content of the nominal interest rate on inflation was Fama [1975], who found that the short-term nominal rate contained all information about expected inflation available to market participants in

[5] They provide a similar decomposition for stocks as well.

the data prior to the 70's. Fama [1990] uses five-year term spread to predict inflation one year ahead, and the present study can be considered as an extension of this approach. Ang, Bekaert and Wei [2007] point out, that inflation forecasts extracted from the term structure are less precise than survey-based forecasts, or even purely statistical ARMA models. However, they use the information in the term spread only, which I show to be inferior to the approach employed here. Moreover, it is worth emphasis that the results of my study hold irrespectively of whether the term-structure inflation forecasts are the most efficient ones, because I explicitly consider only the information in the yield curve.

Several recent studies use survey forecasts in an attempt to predict bond risk premia, with a hope that *forward looking* variables may add significant power to predictive specifications. Piazzesi and Schneider [2010], Wright [2009], and Söderlind [2009] are examples. Dick, Schmeling and Schrimpf [2010] construct a real-time proxy for expected term premium changes for long-term bonds, and find significant, albeit not very strong predictive power of their variable (R^2 of about 13 percent at one-year horizons). I argue here, that forecasts based on the yield curve have an important advantage of being derived from information compiled by one of the most liquid markets, where investors back their bets with real money. In contrast to this, even professional forecasters may sometimes have weak incentives to formulate their forecasts accurately, or may sometimes be willing to conceal their true opinions.

My paper is only loosely related to the literature on theoretical modeling of the term structure. For example, Ang and Piazzesi [2003], and Diebold, Rudebusch and Arouba [2006] consider VAR specifications with both term structure factors and measures of inflation and economic activity. Since macroeconomic variables have been shown to give a poor fit when regressed on the term structure factors,[6] some models explicitly model the information in the macroeconomic variables as orthogonal to the term structure. This approach is pursued in Rudebusch and Wu [2008], Kim [2009], Orphanides and Wei [2010] or Joslin, Priebsch and Singleton [2009]. Finally, Duffee [2009] estimates a latent factor orthogonal to the term

[6]See Gürkaynak and Wright [2010].

structure, which he interprets as containing market expectations about economic activity, that have offsetting effect on expected yield changes and term premia.

Yet another strand of the literature, including D'Amico, Kim and Wei [2010], Christensen, Lopez and Rudebusch [2010] or Joyce, Lildholdt and Sorensen [2010], considers models for nominal and real yields separately, which allows for explicit modeling of expected inflation and inflation risk premium. The focus in these papers is on the *break-even* inflation, which is defined as the variable that makes the real and nominal yield curves distinct. These models attempt to match the term structure of inflation-protected securities (TIPS). As noted by the first authors, the TIPS market faces liquidity problems, and it is unclear, to what extent it is informative about expected inflation and the real yields.

Finally, there are papers that extend the results of Cochrane and Piazzesi [2005] and employ their factor in other contexts. Koijen, Lustig and van Nieuwerburgh [2010] find that value stocks are more sensitive to the CP factor, and they document predictive power of the CP factor for macroeconomic fundamentals at business cycle frequencies. Taken together, these findings rationalize the value premium puzzle. Krishnan, Ritchken and Thompson [2010] provide evidence, that the CP factor adds significant predictive power to the regressions of credit spread changes at the firm-level. Dahlquist and Hasseltoft [2010] extend the evidence in Cochrane and Piazzesi [2005] to international bond markets. They find significant improvements from the CP specification over Fama and Bliss [1987], and show that the CP factor in the US helps to predict excess bond returns in other countries. The study of Kessler and Scherer [2009] is similar in focus.

2 Econometric Analysis

This section starts by defining the variables used in the analysis.

2.1 Definitions of Variables

Let p_t^n denote the log-price of n-year zero-coupon bond at time t. The continuously compounded forward rates[7] can be defined by

$$f_t^n \equiv p_t^{n-1} - p_t^n. \tag{1}$$

The yield of an n-maturity bond is

$$y_t^n \equiv -\frac{1}{n} p_t^n. \tag{2}$$

These definitions imply $f_t^1 = y_t^1 = -p_t^1$, which is the return on a bond with maturity one year. I will refer to this return as the risk-free return (rate), or just the spot rate. Throughout the paper, I adopt the convention of denoting it by f_t^1.

Consider a strategy of buying an n-year zero coupon bond today, and reselling it after one year. The realized return is

$$r_{t+1}^n \equiv p_{t+1}^{n-1} - p_t^n,$$

which is risky for $n \geq 2$, since tomorrow's prices are unknown (with the only exception of one-year bonds, which have known notional amount of 1).

The excess return of an n-year zero coupon bond over the risk free rate is

$$rx_{t+1}^n \equiv p_{t+1}^{n-1} - p_t^n - f_t^1. \tag{3}$$

Equation (1) can be thought of as a difference equation that can be solved for p_t^n, with the boundary condition $p_1^0 = 0$. The solution is

$$p_t^n = -(f_t^1 + f_t^2 + \cdots + f_t^n),$$

[7]Forward rates are the interest rates for future loans, between $t+n$ and $t+n+k$, where $k > 0$. By convention, if $n = 0$, the name "forward" is replaced by "spot". All forward rates are known today (at t), and are determined by no-arbitrage conditions. Throughout the paper, all forward rates are for one year loans, that is, k=1.

which can be used to substitute prices away from (3). In effect, one can write the excess return

$$rx^n_{t+1} = \sum_{i=2}^{n} f^i_t - \sum_{i=2}^{n} f^{i-1}_{t+1} = \sum_{i=2}^{n}(f^i_t - f^{i-1}_{t+1}), \tag{4}$$

which illustrates how excess bond returns depend on the set of present forward rates, and the corresponding set of next year forward rates with maturities decreased by one year.

It is clear from (4), that realized excess return on a bond with maturity two is also present in realized excess returns of all other bonds. Similarly, the excess return realized on the three-year bond is contained in all returns of bonds with maturities larger than three. It is therefore useful to decompose the realized excess returns into more basic elements, which I call realized forward premia

$$fpr^n_{t+1} \equiv f^n_t - f^{n-1}_{t+1}. \tag{5}$$

Similarly, we can define expected forward premia, by taking rational expectations of (5),

$$epr^n_{t+1} \equiv E^P_t[f^n_t - f^{n-1}_{t+1}], \tag{6}$$

where the expectation is taken under the real-world probability measure P, conditional on all information available at t.

2.2 The Data

I use the so-called Fama-Bliss dataset on the end-of-month synthetic constant-maturity bond prices, which covers full-year maturities between one and five years. The data are available from the Center for Research in Security Prices (CRSP), and cover the period between June 1952 and December 2009. The description of the methodology used to construct the data can be found in Appendix A of Fama and Bliss [1987].[8]

[8]Some studies using the Fama-Bliss dataset, including Fama and Bliss [1987] and Cochrane and Piazzesi [2005], discard the observations prior to 1964, arguing that only after that year there was at least one bond traded in the market with maturity falling between every pair of the constant maturities, which could be used to bootstrap the discount rate function. See Fama and Bliss [1987], Appendix A, for more detailed

I compute continuously compounded forward rates according to equation (1). Returns are computed as continuously compounded over twelve months ahead, so that effectively the sample ends at December 2008. Table 1 shows the most important descriptive statistics: the means, cross-correlations, and yearly autocorrelations of the forward rates (panel A), and of yearly excess bond returns (panel B).

Realized inflation is calculated as continuously compounded yearly growth rate of the Consumer Price Index for All Urban Consumers (CPIAUCNS, All Items), obtained from the FRED database of the Federal Reserve Bank of St. Louis.[9] Figure 1 plots the time series of realized inflation, together with one-year and five-year forward rates from the full Fama-Bliss dataset. Panel C of Table 1 shows the descriptive statistics of inflation. It is correlated unconditionally at 0.61 with the one-year nominal (spot) rate, and at 0.49 with five-year maturity forward rate. It is negatively correlated with realized excess bond returns.

For reasons of robustness and completeness, I often consider subsamples of the data. The "older" subsample is defined as the period between June 1952 and May 1997, and the "newer" subsample is between January 1964 and December 2008. Both have 540 monthly observations, and the latter is the same as in Cochrane and Piazzesi [2005], with the exception of adding more recent data points. To provide more detailed evidence on the validity of results, I also employ the rolling-window strategy, by repeating the analysis with different samples of the data. I set the window width to 25 years, which seems to trade-off relatively well the need for a relatively large number of independent windows, and the need for reliable in-sample estimates within each window.

2.3 Predicting Excess Bond Returns

The focus of this paper is is on the common factor that drives excess bond returns of all maturities in the same direction. Cochrane and Piazzesi [2005] argue, that this factor accounts

explanation of this issue. Other studies, including Fama [2006], or Koijen, Lustig and van Nieuwerburgh [2010] use the whole range of available data. In the appendix A of this paper, I follow the former authors to provide direct comparison with their results. In the main part, I use all the data available.

[9] See http://research.stlouisfed.org/fred2/series/CPIAUCNS.

for 99.5% of variance in expected excess bond returns.[10]. However, by using the principal component analysis directly on *excess returns*, their approach misses the information about potentially important predictable variation in individual forward premia with maturities other than two years: as already argued, excess returns on all bonds contain by construction realized excess return on the two-year bond, which is enough to introduce common variation artificially. The first question I ask is therefore about the importance of the single factor.

2.3.1 How important is the single factor?

If there is a factor which moves the forward premia of all maturities defined in (5) in the same direction, it must also move all realized excess bond returns in one direction, as seen in (4). I first construct the time series of realized forward premia, and apply the principal component decomposition to their unconditional covariance matrix. Since the data allow to construct four time series for maturities between two and four years, there is four eigenvectors, associated with four eigenvalues. I order the eigenvectors according to the magnitudes of the eigenvalues. I refer to the factor scores associated with the eigenvectors simply as the "principal components" or "factors" (with ordering as above), hoping that no confusion can arise.

Figure 2 plots two eigenvectors associated with the largest eigenvalues. The first (solid line) has the "level" interpretation. It's slightly negative inclination means that the excess bond returns (cumulative sums of the values at each maturity) increase less than proportionately with maturity, when all forward premia move along this line. The other eigenvector (dashed line) has positive slope, and crosses zero at maturity of three years. It makes the excess returns of two- and three-year bonds go up, and those of five maturity bonds go down. The remaining two factors have very irregular shapes, and are not plotted in the graph.

The first factor accounts for 90 percent of unconditional variance in the forward premia, and the second for only 6 percent, implying that the other two contribute 4 percent together.

[10]See the on-line Appendix to their paper: http://www.stanford.edu/ piazzesi/cpapp.pdf.

To compute the fractions of each factor in *predictable* variance, I first regress the four factors separately on five forward rates with a constant, in the spirit of Cochrane and Piazzesi [2005]. Then, I multiply the unconditional total variances of each factor by corresponding R^2, to obtain the predictable fractions. Finally, I compute the shares that each factor contributes to the total predictable variance of the forward premia.

The results seem to generally support the findings in Cochrane and Piazzesi [2005], but are much less remarkable. The first factor explains 84% of predictable variance, and the second one just 4%, leaving 12% for the two less important factors. I argue however, that the economic reliability of the latter number is doubtful, since the remaining two factors have very small unconditional variances, so they are very prone to over-fitting by five forward rates. In their case, predictability may also come from measurement errors on individual forward rates, which disappear much more quickly than within one year, resulting in spurious predictability patterns. These arguments suggest, that the value of 84% should be thought of as the lower bound on the importance of predictability contained in the first principal component. And, the portion of predictability due to the second one seems negligible, at least from the point of view of this study.

Because of the problems with Fama-Bliss data for the earliest years of the sample prior to 1964 (see the section on the data), it is very likely that the single factor hypothesis may not hold very well in samples that include those points. To test this suspicion, I repeat the steps above using the rolling window of length 25 years, moving from the first month in the sample, until December 2008 is included. I compute cumulative contributions of all factors to the predictable variance in each window, and plot the results in figure 3. The single-factor hypothesis gets a much better support for most of the subsamples starting at least 100 months from June 1952, with the share of the first principal component of about 90%. It seems, that the number 84% is heavily influenced by extremely low values at the beginning of the available data.

2.3.2 Properties of the CP factor

I define the Cochrane-Piazzesi factor (CP) as the linear combination of forward rates (including the constant), that predicts the first principal component of forward premia in the best possible way, in the sense of the OLS regression. In other words, it is the fitted value from

$$F^1_{t+1} = \alpha_0 + \alpha_1 f^1_t + \alpha_2 f^2_t + \alpha_3 f^3_t + \alpha_4 f^4_t + \alpha_5 f^5_t + \epsilon_{t+1}, \tag{7}$$

where F^1_{t+1} is the factor score of the first principal component of forward premia.

What is the relation between the CP factor defined here, and its original counterpart defined in Cochrane and Piazzesi [2005] (see (22) in Appendix A)? It turns out, that the correlation between the two is 0.9984, which means that one is just the linear transformation of the other, so that they carry exactly the same information.

Figure 4 presents the CP factor, estimated with all available data. It has clear business cycle properties, low at the beginning of recessions, and rising sharply during their continuance. It also shows interesting low-frequency patterns, it is elevated throughout the 80's and part of the 90's, which seems to occur independently from what its business-cycle patterns. These patterns are described for example in Koijen, Lustig and van Nieuwerburgh [2010].

Table 2 and figure 5 report coefficient estimates for (7), for three subsamples defined in the section on the data. The characteristic tent is a robust feature of the data, despite the R^2 coefficients of between 19 and 26 percents are significantly lower than those reported in Cochrane and Piazzesi [2005] (see Appendix A).

Also, the significance of some coefficients is rather low. This is because there is high collinearity between the forward rates (which is impossible to avoid), and because the Newey-West procedure adjusts the standard errors to account for the autocorrelation in the residuals in the overlapping data.[11] Since I define the CP factor to have the maximum in-sample

[11] Yearly returns on monthly data generate an MA(12) term in the residuals. In all regressions, I use the automatic lag truncation choice at $q = floor(4(T/100)^{2/9})$, suggested in Newey and West [1987]. Cochrane and Piazzesi [2005] use the a lag truncation of 18 in their restricted specification, but a replication of their study with automatic choice of $q = 5$ yields virtually the same standard errors.

predictive power, I use all five forward rates in its construction.

2.4 Predicting Inflation

This section is concerned with the information in the yield curve about future inflation, which is seemingly unrelated to bond risk premia, but since expected inflation measures the expectations of real cash flows to the bondholders, it can be a priori expected to contain useful information about bond excess returns, if one controls for bond prices (forward rates). Here, I concentrate on how to extract inflation expectations from the forward curve in the most efficient way.[12]

I first regress the inflation rate realized between today and the next year on five forward rates (with a constant), in a way analogous to how Cochrane and Piazzesi [2005] proceed with excess bond returns,

$$\pi_{t+1} = \delta_0 + \delta_1 f_t^1 + \delta_2 f_t^2 + \delta_3 f_t^3 + \delta_4 f_t^4 + \delta_5 f_t^5 + \varepsilon_{t+1}. \tag{8}$$

Table 3 presents the estimation results for three subsamples: the full sample, the older sample from June 1952 to May 1997, and the newer sample from January 1964 to December 2008. The R^2 coefficients are very high, if compared to the ones from the forward premia regressions in the previous subsection, between 39 and 48 percent. Figure 6 plots the coefficients. Interestingly, they form inverted tents, similar to the ones in figure 5.

Which coefficients are really significant in these regressions? Looking at the table and at the graph, those related to forward rates of maturity two and four years are not significantly different from zero, thus removing them should increase the efficiency by reducing multicollinearity problem. When instead of estimating (8), I run a regression of the form

$$\pi_{t+1} = \delta_0 + \delta_1 f_t^1 + \delta_3 f_t^3 + \delta_5 f_t^5 + \varepsilon_{t+1}, \tag{9}$$

[12] Campbell and Ammer [1993] show that revisions in inflation expectations are correlated with revisions of excess bond returns in their subsamples that cover the 80's, which points to a potential role of expected inflation to uncover bond risk premia.

the R^2 coefficients remain unchanged, and the coefficients on three remaining forward rates are indeed estimated with much higher precision (unreported). However, since the focus of this study is on *in-sample* relations, there is no loss of correctness associated with using (8) to construct the inflation-forecasting factor. Also, using all five forward rates makes it easier to compare the inflation-forecasting factor to the CP factor.[13]

How large is the difference between the precision of forecasts based on the whole term structure of five forward rates, versus the more traditional approach of using the short rate and the term spread only? To provide the answer, I run regressions of the form

$$\pi_{t+1} = \beta_0 + \beta_1 f_t^1 + \beta_2(f_t^5 - f_t^1) + \varepsilon_{t+1}, \tag{10}$$

trying to predict changes in the log price level by means of the spot rate and a long-maturity forward-spot spread.[14] I analyze the three subsamples already employed, and compare the R^2 coefficients from specifications (9) and (10). The losses in the adjusted R^2 in the full sample, older sample, and the newer sample are 6, 7, and 9 percentage points, which are relatively large values.

The main message from this section is therefore, that the overall shape of the forward curve can be used to extract the information about future inflation. It seems reasonable to conclude, that simple predictive regressions that use only the term spread and the risk-free rate may be mis-specified in the sense of omitting important information. This is analogous to the conclusion of Cochrane and Piazzesi [2005] in the context of predictability of excess bond returns by their factor versus the term spread only.

[13] In an unreported exercise, I verify that using (2.3.2), or its restricted version is of zero relevance for the conclusions of this paper, because the fitted values from both specifications have correlation of one. However, using the restricted specification has clear advantage, if one is willing to use the results of this section to predict inflation out of sample.

[14] This formulation is equivalent to removing three intermediate forward rates from the specification (8).

2.4.1 The IE factor

I define the IE factor (a shortcut for inflation expectations) as the linear combination of forward rates, which contains all information about one-year-ahead inflation revealed by the term structure, and which can be extracted by means of a linear regression. In other words, it is the fitted value from equation (8).

Does the IE factor contain all information about future inflation rate, that is available to the investors?[15] Here, I test whether one-year lagged inflation rate improves inflation predictions based on IE only, by running regressions

$$\pi_{t+1} = a_0 + a_1\pi_t + a_2 IE_t + \varepsilon_{t+1}. \tag{11}$$

It turns out, that the AR term has significant explanatory power on inflation over and above the one contained in the IE factor only. Including AR boosts the R^2 coefficients by 16-18 percentage points, if the three already analyzed subsamples are used.

How should one interpret this result? To further test the need for the AR term in (11), I re-do this regression on rolling windows of 300 months (25 years), starting from June 1953, and moving by one month at each repetition, until the last time-point of the data is reached. I collect the R^2 coefficients for two specifications, one with, and one without the AR term. As seen in figure 6, they virtually coincide for the subsamples that are contained within the first 330, or the last 350 months of the dataset. These are the periods, when both inflation and the overall yield curve were increasing (the former), or decreasing (the latter), and the regressions find it easier to fit realized inflation only to the forward curves. On the other hand, if the subsamples contain the turning point of the 80's, when inflation started to drop but interest rates did not, the AR term becomes significant.[16]

Since the additional predictability of inflation by the AR term comes from a factor or-

[15] This is a question about bond market efficiency, as defined in Fama [1975].

[16] Another important observation from the figure is that after the inflation dropped in the first half of the 80's, it became much harder to predict, by either specification.

thogonal to the term structure of forward rates, a question arises whether it belongs to the agents' information set. If yes, we should expect, that it also predicts excess bond returns marginally better, than the term-structure variables alone. I show later on, that this is indeed the case, so that the results of this paper are further supported. At the same time, the forward-rate factor IE remains a valid variable to predict excess bond returns, because it is essentially a projection of the agents' information set on the interest rates. Obviously, it will predict the returns with less precision, but analogous conclusion is equally true for the CP factor as well, since it also uses only the information in the forward rates.

2.5 Inflation Forecasting Factor and Bond Return Predictability

In figure 7, I plot the coefficients that form the IE factor estimated in the full sample (solid line), the older sample (dashed line), and the newer sample (dotted line), where the samples are defined as before. A comparison with Figure 5 reveals striking similarities between the CP and the IE factors. The shape of the latter looks like an inverted version of the former. This raises a question about the relation between inflation forecasts contained in the forward curve and expected excess bond returns, which is the main topic of this section.

The time series of factor scores corresponding to the two factors are shown in figure 8. There is striking similarity between the two series at business-cycle frequency. A rise in IE is almost always associated with a fall in CP. However, the correlation between the two is only -0.19, which seems to be due to important differences in their low-frequency components. While the CP factor is clearly stationary,[17] the IE is initially increasing, and then decreasing, pointing to a behavior close to unit-root. This is no coincidence, since the same is true for the nominal level of the term structure, which should contain expected inflation one-to-one, if the Fisher hypothesis holds in the long run. As argued in this section, there exists a close empirical relation between the level of the nominal term structure and expected inflation contained in the forward rates, that almost perfectly uncovers bond risk

[17]The augmented Dickey-Fuller test rejects unit root in CP at 1% significance level.

premia. This section documents this relationship, and the next re-interprets the findings in terms of cointegration analysis.

The level and expected inflation can be a priori suspected to contain important information on bond risk premia, because while the latter variable summarizes the expectations of real cash flows to the bondholders, the former controls for the general level of bond prices. If prices are "too high" with respect to the fundamentals, either the fundamentals should be expected to improve, or prices to fall.[18] I show here, that expected inflation extracted from the yield curve does a great job indeed, which means that it is close to sufficient in describing expected real cash flows. Controlling for the level (prices), it summarizes the risk premia in a way that comes almost indistinguishably close to what the CP factor achieves.

I run direct regressions of the first principal component of forward premia on the level and IE factors, and compare the in-sample predictive ability to the one from the CP regressions. I estimate

$$F^1_{t+1} = \kappa_0 + \kappa_1 L_t + \kappa_2 IE_t + \eta_{t+1}, \tag{12}$$

where F^1_{t+1} denotes the factor score of the first principal component of forward premia one year ahead, and L_t is the level. For comparison, I also run regressions equivalent to the restricted specification in Cochrane-Piazzesi [2005],

$$F^1_{t+1} = \kappa_0 + \kappa_1 CP_t + \eta_{t+1}. \tag{13}$$

The results are summarized in Table 4. Panels A and B show that both sets of factors (estimated using the whole sample) have roughly the same predictive power for predicting the first principle component of forward premia (the common component of excess bond returns of all maturities), across all three subsamples. The differences in the R^2 are between zero and two percentage points only.

Panels C and D are allow for the re-estimation of all factors used in the regressions

[18] This is the same type of argument, as with predictability of stock returns by the P/D ratio. See Cochrane [2005].

(12) and (13) on both sides of these equations, thus making it potentially more difficult for the level and IE to match the degree of in-sample predictability of the CP. However, the differences in the R^2 coefficients are again very small, between 1 and 2 percentage points, despite the fact, that now the CP factor is free to fit the data in the best possible way within each subsample.

To further test the hypothesis that the two sets of factors contain very similar, if not the same information, I re-estimate both (12) and (13) on 25-year rolling subsamples of the available data, allowing for the re-estimation of all factors within each window. Figure 9 compares the determination coefficients. It is striking, that the in-sample fit that uses the level and IE often performs almost as well as the CP factor, which by definition is the best in-sample predictor of the forward premia. In other words, the solid line in the figure can never be crossed. The combination of level and IE comes very close to this best outcome, despite neither of this factors was created with the purpose of predicting bond risk premia.

We have thus seen that the predictive ability of one set of regressors (CP) is almost the same as that of the other set (level plus IE). Still, it may be not clear, whether the two sets contain the same information. It could happen, for example, that the information corresponding to both sets can be, to some extent, independent. If this was the case, then merging the two should increase our ability to predict the forward premia. This cannot be true for a simple reason, the CP factor is constructed to have maximum in-sample predictive power, so no other linear combination of forward rates (level plus IE are such combinations) can increase it.

What would happen, if one used a better predictor of inflation to forecast excess bond returns? After all, investors may possess information about future inflation, that is not included in the yield curve. To answer this question, I first construct an alternative inflation-forecasting factor using equation (11), which uses inflation lagged one year in addition to the information already included in the IE. Figure 10 is constructed in an analogous way to figure 9, and provides some insights on the issue. In most of the subsamples, better prediction of inflation really leads to a better prediction of excess bond returns. In some early and later subsamples, which contain the inflationary turning point of the early 80's, the resulting predictability is even higher than given by the CP factor alone. Also, in most of the intermediate samples, the R^2 is higher than that obtained using the former measure of inflation expectations. Overall, this piece of evidence provides an extra support to the idea, that the IE factor (together with level) capture almost all information about the risk premia, that one is able to extract from the forward curve.

And what happens, if one uses the measure of inflation expectations based only on the term spread? I test this by first fitting realized inflation to the spot rate and 5-year term spread, and then use this variable together with the level. Estimated on the whole sample, the predictive ability of the two is reflected in an R^2 of 13 percent, which is significantly lower, than with my baseline specification. The results for the other subsamples yield similar results.

Why do we obtain a deterioration of fit for some of the latest subsamples in figure 10? Remind from figure 6, that the additional AR term does not help to forecast inflation at all in this region of the data, and that the overall ability of either specification to predict inflation is much lower. Including the additional AR term most likely leads to over-fitting the true inflation expectations in these particular subsamples.

The level coefficient in table 4 is positive, varying between 0.21 and 0.36 across samples. The IE coefficient is negative, between -0.85 and -1.11. Figure 11 plots the whole range of estimated coefficients in the rolling windows described above. The values of $\kappa 1$ and $\kappa 2$ seem to be negatively correlated, which points at potential collinearity problems in specification (12). This is confirmed in figure 12, which shows correlation coefficients between the level and the IE, for each 25-year rolling window. High correlations towards the beginning and the end of the sample explain the relative imprecision of the estimated coefficients in these subsamples. In other words, one has to be prepared for losses of efficiency if restricting the full sample available.Note however, that in the context of this paper, analyzing subsamples is correct, because the focus is on in-sample comparison of two different sets of predictors, not on out-of sample forecasting.

It can be observed from figure 9, that the overall predictive ability of either CP and IE (combined with level) has changed a lot over the years. It was especially high in the subsamples that included the inflationary period of the turning point of the 70's and 80's. I return to this issue in the next section, where I provide more economic intuition on the results.

Finally, it is interesting to point out, that univariate regressions of the excess bond returns (or forward premia) on either level alone or IE alone do very poorly with uncovering the predictability patterns presented above. In case of the level, the R^2 is only 3 percent in the full sample, 7 percent in the older sample, and 3 percent in th newer sample, with coefficients significant (at 5% level) only in the second case. Similarly, the R^2's are just 1, 1 and 6 percent, when only the IE is used. What causes the increase in predictive power when both are included into the regression? This can be understood by considering the effects of omitted variable bias. The correlation between the level and the IE is 0.73 (0.66 and 0.60 in the older and newer samples, respectively), and they both predict expected returns in opposite directions. Removing one of them from the regression must lead to a severe omitted variable bias, that pulls down the coefficient on the other towards zero. This also results in a decreased R^2, as explained in Appendix B.

3 Interpretations

This section discusses economic interpretation of the evidence presented. What is the reason for the combination of level and IE to predict excess bond returns? Given the results from the previous section, the question is really about the reasons of why does the CP factor work.[19] I argue here, that a cointegration relation exists between the level and the IE, and

[19]The economic meaning of the CP factor has remained mysterious. Koijen, Lustig and van Nieuwerburgh [2010] interpret this factor as the proportion of transitory component in total variance of the economy-wide

that deviations from this relation have the ability of predict excess bond returns with high power.

3.1 Cointegration

I first check the relation between the CP, and the other two factors. I run the following regression using the whole sample,

$$CP_t = c_0 + c_l L_t - c_e IE_t + u_t$$

The resulting coefficients are $c_0 = -0.002$, $c_l = 0.309$, and $c_e = 0.923$, and the R^2 coefficient is 0.95.

The equation above can be written in an approximate form (omitting the error term)

$$CP_t = c_0 + c_l L_t - c_e IE_t. \tag{14}$$

The right-hand side is proportional to the distance between (scaled) level and IE factors. Since the CP factor on the left-hand side is stationary,[20] the right-hand side must also be. If both level and the inflation-forecasting factors are non-stationary in the data,[21] then the right hand side can be interpreted by means of a cointegrating relation between the level and expected inflation (projected on forward rates). The long-run mean of CP_t defines the long-run relation, in which L_t and IE_t are, and the deviation of CP from its mean can be interpreted as a deviation from cointegration.

Figure 13 presents the evolution of the two components of this relation over time (I have

stochastic discount factor. Their results are based on the insights of Alvarez and Jermann [2005], who prove that if expected long-maturity bond return is high relative to the equity premium, the transitory component in the SDF must be large. Lettau and Wachter [2010] hypothesize, that the CP factor predicts excess returns, because it extracts the market-wide risk aversion factor from forward rates in the best possible way, which has to be true in general.

[20]See footnote 17.

[21]See Fama and Gibbons [1982], and Mankiw and Miron [1986], or Cochrane and Piazzesi [2008]. In the full sample between June 1952 and December 2008, the ADF test does not reject the existence of the unit root in the level.

subtracted the constant c_0 from the scaled level). One is able to observe interesting patterns, the IE factor usually leads the business cycle, in the sense that it starts to fall prior to the starting points of NBER recessions. At the same time, the level of forward rates tends to stay more or less unchanged, despite the fall in inflation expectations. This discrepancy is subsequently corrected by upward movements of the IE, downward movements in the level, or both. The latter are equivalent to high realized forward premia (excess bond returns).

This pattern is clear during the recessions of 69-70, 81-82, 90-91, and 2001. In the three recessions at the beginning of the sample (53-54, 57-58, and 60-61) it is not there, or even quite the opposite seems to be true (53-54). This pattern is also not present in the most recent recession, that started in December 2007. This provides an intuitive explanation of the findings in figure 9, that the predictive power of the CP factor was lower in the earliest and latest subsamples.

We can also see in figure 13, that after the recession of 81-82, expected inflation was very low while nominal interest rates remained high for an extended period of time. This was followed by a prolonged decline in the level of nominal rates during the 80's, which contributed to high predictive power of the CP factor in subsamples containing this period.

3.2 Tests of the expectations hypothesis

Since the relation between the CP factor and the deviations from cointegrating relation defined by (14) is robust across subsamples, the precision with which the CP factor is estimated in a given sample is closely related to the precision, with which cointegration is estimated. It is very likely, that using only subsamples of the data (for example, to perform out-of-sample tests) can lead to large efficiency losses. This is because one must have enough deviations from cointegration it in the data to estimate it. As already shown in figure 12, correlations between the level and the IE are very large for those subsamples, which do not include the inflationary 80's. Estimating cointegration in these subsamples must result in high imprecision of estimated parameters. The conclusion is that the period of the great inflation provides an opportunity to identify the parameters in (14). Another conclusion is that out-of-sample tests of expectations hypothesis that rely on small subsamples to re-estimate the forecasting relations are very likely to have limited power. Similar point is emphasized more formally by Inoue and Kilian [2004].[22]

3.3 Irrationality or risk premia?

Predictability patterns presented above can be reconciled with either behavioral or risk-based stories. For example, bond traders may react too slowly to changing inflation environment, so that the prices of bonds adjust with a lag, generating predictability. This would have to be interpreted as market inefficiency, since inflation forecasts contained in the yield curve must be known to bond traders. Because of the latter fact, I endorse the other explanation, based on time-varying risk premia required by bond market participants. Consider a hypothetical situation, in which the yield curve contains the information about expected inflation efficiently, and the premium is zero. Suppose now that a recession is coming. The usual pattern would be that inflation expected for one-year ahead begins to fall, while the level of interest rates initially stays unchanged. Since decreased inflation expectations are in fact positive cash-flow news, the unchanged level of interest rates (and therefore prices) implies that the rate at which investors discount real cash flows must have increased. When the recession comes to an end, the gap between relatively high level and low expected inflation gradually closes and the premium disappears. Part of this adjustment comes from the level, which is precisely when high excess bond returns materialize.

A similar pattern occurred in the 80's, when nominal rates stayed elevated, and then fell gradually towards lower values, consistent with long-run inflation expectations. The adjustment took much longer than a business cycle, which signals that the predictability of excess return in the data is not just a business-cycle phenomenon. In theoretical work it

[22] As argued in Lettau and Ludvigson [2010], when the sample is long enough, the super-consistency of the estimators of cointegration relation can be exploited. It is then correct to estimate cointegration in the first step, and then treat it as known in the out-of sample tests. Equivalently, when the existence of cointegration relation can be postulated, it is inefficient to re-estimate it in the subsamples.

may be necessary to account for two components of time-varying price of risk, one operating at business cycle frequencies, and one at lower frequency.[23]

How to interpret the risk premium, which is captured by the CP factor? Although the series of expected inflation can be efficiently used to uncover it, it would be generally wrong to tie this risk to the level of inflation expectations, or to the nominal level directly. It would be better to speak of generic business cycle risk, that prevents bond market participants from exploiting better investment opportunities, which arise with falling inflation expectations. Similarly, the CP factor will capture any other kind of risk, that makes bond investors miss (apparently) profitable investments.

3.4 Dynamic adjustments

What are the time series properties of the CP factor? Now, this question can be boiled down to the question about the properties of the right-hand side in (14). It will be of particular interest to see, how the deviations of the CP factor from its long-run mean are corrected over time. Start by demeaning the fitted relation (14)

$$CP_t - \bar{CP} = c_0 - \bar{CP} + c_l L_t - c_e IE_t.$$

Letting $\tilde{CP}_t = CP_t - \bar{CP}$ and $d_0 = c_0 - \bar{CP}$, re write this as

$$\tilde{CP}_t = d_0 + c_l L_t - c_e IE_t. \tag{15}$$

If $\tilde{CP}_t = 0$, the right-hand side defines the equilibrium relation between L_t and IE_t. Any deviation from this relation must be corrected in the long-run either by downward level

[23]Theoretical links between inflation and the yield curve are explored in Piazzesi and Schneider [2006], who argue that inflation is a leading business cycle indicator, and that inflation innovations carry bad news about future consumption growth, so that the yield curve is on average upward sloping. Piazzesi and Schneider [2010] introduce adaptive learning about joint distribution of inflation and consumption growth, to argue that investors perceived nominal bonds to be unusually risky after the period of the great inflation, which is consistent with Cochrane and Piazzesi [2005] and this paper.

movement, upward expected inflation movement, or both. In the short-run, part of the initial deviation may remain. The question is about the persistence of the deviations, or in other words how long do exceptionally high risk premia survive?

Consider the following identity, derived from (15)

$$\tilde{CP}_t = \tilde{CP}_{t+h} - c_l(L_{t+h} - L_t) + c_e(IE_{t+h} - IE_t), \tag{16}$$

where h is the time-shift in years ($h \geq 1$). The logic in this equation can be used to write three complementary regressions, for each of the three components of the right hand side,

$$\begin{aligned} \tilde{CP}_{t+h} &= \alpha_0^h + \alpha_1^h \tilde{CP}_t + \epsilon_{t+h}^1, \\ c_l(L_{t+h} - L_t) &= \beta_0^h + \beta_1^h \tilde{CP}_t + \epsilon_{t+h}^2, \\ c_e(IE_{t+h} - IE_t) &= \gamma_0^h + \gamma_1^h \tilde{CP}_t + \epsilon_{t+h}^3. \end{aligned} \tag{17}$$

The equations are complementary in the sense that the restrictions $\alpha_0 - \beta_0 + \gamma_0 = 0$, $\alpha_1 - \beta_1 + \gamma_1 = 1$, and $\epsilon_{t+h}^1 - \epsilon_{t+h}^2 + \epsilon_{t+h}^3 = 0$ hold identically (ex post) for all h. The estimation results are in table 5. The first row denoted α_1^h presents the slope coefficients from the first equation. One observes, that a deviation from cointegrating relation (15) has the half-life of exactly one year, which is consistent with the hypothesis that the business cycle is the primary force behind the CP factor. After two years, only a quarter of initial deviation remains, and after three years the gap virtually closes. Rows two and three of table 5 help us to understand the reasons, why bond risk premia disappear so quickly. The scaled level and expected inflation play important roles in closing the deviation at one year horizon, with coefficients of -0.24 and 0.26, respectively, although the R^2 statistics of 10 and 5 tell us, that the reaction from level is more certain to happen, while the response from the IE is more volatile. For horizons between two and five years, expected inflation plays a more important role in correcting the initial deviation, but the level is responsible for the correction of two thirds of the initial deviation after eight years. In other words, expected inflation seems to provide a trend to which nominal level slowly reverts. This feature of the data is at least partly a consequence of the events in the 80's.

4 Concluding remarks

I provide evidence, that the predictability of bond excess returns uncovered by the term structure of interest rates has a neat economic interpretation in terms of a natural cointegration between prices and fundamentals. I measure prices as the level of nominal forward rates, and the fundamentals by inflation expectations uncovered by the same forwards.

To obtain a measure of inflation expectations, I use regression of realized inflation on the term structure of forward rates with maturities between one and five years. The linear combination of forwards that forms my inflation-forecasting factor has an inverted-tent shape, and predicts subsequent price growth with significantly greater power than traditional term-structure based inflation predictors, like the term spread or the spot rate.

When bonds are relatively cheap, as measured by elevated level of the nominal term structure with respect to the measure of inflation expectations, there is a strong tendency in the data for the level to fall and expected inflation to rise. This situation is equivalent to high bond risk premia, which result in high realized excess returns. Indeed, the deviations from cointegration between my two factors are almost perfectly correlated (at 97 percent) with the single statistical factor of Cochrane and Piazzesi [2005], which gives a nice economic interpretation of their classic result. The deviations disappear after about two years, and have clear business-cycle pattern.

Predictability of excess bond returns is concentrated around the inflationary 70's and 80's, which is the period of the largest discrepancy between my two factors. Prior to that time, the two are highly correlated, which is also the case for the most recent period, which includes the financial crisis. This explains relatively weaker predictability in the older and the most recent subsamples.

I only concentrate on the informational content of the term structure, which is enough for my purpose. It has been shown that bond risk premia can be very well predicted by the information from outside the yield curve. Therefore it seems plausible, that if one finds a better measures of market expectations using the information that does not rely on forward rates, this will lead to improved expected return forecasts. To test for this possibility, I include an autoregressive term to the inflation-forecasting specification, and show that this improves excess bond return forecasts, to the extent to which it improves the measure of inflation expectations.

This study raises further questions, and can be extended along several dimensions. One idea is to construct an affine term structure model that incorporates the findings of this paper. In principle, three factors should be enough to capture the dynamics: level, slope and the inflation-forecasting factor. The model could then be used to study the dynamic behavior of the real term structure and the risk-free rate. Another interesting task would be to explain the nature of predictive power, that the CP factor has on equity returns, value premium and credit spreads.

Appendix

A Updated Evidence

The focus in Fama and Bliss [1987] is on testing the expectations hypothesis of the term structure, not directly on excess return predictability. They start with identities, which hold for all maturities n,[24]

$$f_t^n - f_t^1 = rx_{t+1}^n + (n-1)(y_{t+1}^{n-1} - y_t^{n-1}). \tag{18}$$

[24]To derive them, add and subtract from the left hand side all forward rates with maturities between 2 and $n-1$ years dated at time t, as well as all forward rates with maturities between 1 and $n-1$ years dated at $t+1$. Rearange to get the right-hand side.

The content of this equation is best seen for $n = 2$,

$$f_t^2 - f_t^1 = rx_{t+1}^2 + f_{t+1}^1 - f_t^1.$$

The sum of 2-year bond's excess return and the change in the risk-free rate between today and tomorrow is deterministic and known at t, although both of its components are random. This leads to the conclusion, that the slope forecasts excess returns only to the extent to which it *does not* forecast the risk-free rate. In what follows, I call variables $f_t^n - f_t^1$ the Fama-Bliss slopes.

Because the identities always hold, one can write two complementary regressions for every n

$$rx_{t+1}^n = a + b(f_t^n - f_t^1) + u_t^n, \tag{19}$$

$$(n-1)(y_{t+1}^{n-1} - y_t^{n-1}) = -a + (1-b)(f_t^n - f_t^1) - u_t^n. \tag{20}$$

The intercepts and error terms across the two equations sum to zero, and the slopes sum to one (so that (18) is always satisfied). The condition $b = 0$ is necessary for the expectations hypothesis to hold. In other words, the excess returns should not be predictable by anything, including the corresponding Fama-Bliss slope.

Table A1 provides some updated evidence on the results. Panels A and B use the sample between January 1964 and December 2008, while Panels C and D restrict the data span to coincide with the study of Cochrane and Piazzesi [2005] (for later reference). The coefficients b are positive and significant, while $1 - b$ are insignificantly different from zero, suggesting the failure of the expectations hypothesis at all maturities.

Cochrane and Piazzesi [2005] go a step further, and run predictive regressions for excess bond returns of all maturities, using the same set of five forward rates from the Fama-Bliss dataset (including the spot rate). Now, the focus is on finding a factor that predicts excess bond returns with maximum power. Their empirical (unrestricted) specification is

$$rx^n_{t+1} = \gamma^n_0 + \gamma^n_1 f^1_t + \gamma^n_2 f^2_t + \gamma^n_3 f^3_t + \gamma^n_4 f^4_t + \gamma^n_5 f^5_t + \epsilon^n_{t+1}. \tag{21}$$

Table A2 displays the regression coefficients for bond maturities between 2 and 5 years (standard errors beneath each row). Panel A uses the sample between January 1964 and December 2008, while Panel B restricts the data to coincide with their original study. Figures A1 and A2 show the results graphically.

Several conclusions are in order. First, the coefficients on individual forward rates form the famous tents. Second, the predictive ability of forward rates seems to have deteriorated since the Cochrane-Piazzesi study, the R^2 coefficients are now around 0.24-0.27, below the values of 0.31-0.36 originally reported.[25] Third, although the symmetry of the tent has changed slightly, the tent-like patterns of coefficients are proportional across maturities, which led Cochrane and Piazzesi [2005] to hypothesize a single factor behind expected excess returns of all maturity bonds. To uncover it, they run a regression

$$\frac{1}{4}\sum_{i=2}^{n} rx^n_{t+1} = \gamma_0 + \gamma_1 f^1_t + \gamma_2 f^2_t + \gamma_3 f^3_t + \gamma_4 f^4_t + \gamma_5 f^5_t, \tag{22}$$

and then proceed with regressions

$$rx^n_{t+1} = b_n(\gamma^T \mathbf{f}_t) + \varepsilon^n_t,$$

where $\gamma^T \mathbf{f}_t$ is the linear combination of forward rates from (22), including the constant.[26] The R^2 coefficients for the restricted regressions (22) are slightly below their unrestricted counterparts. However, they are still high, which suggests that the single-factor restriction

[25]The online Appendix to the paper of Cochrane and Piazzesi [2005] provides confidence intervals for the R^2 statistics computed by Monte Carlo experiments (see http://www.stanford.edu/ piazzesi/cpapp.pdf). The updated magnitudes fall well within the bounds reported therein.

[26]The coefficients b^n for the updated sample are 0.45, 0.85, 1.25 and 1.46 (in the order of increasing maturity).

is a good first-order approximation to the data.

A comparison of Tables A1 and A2 makes it clear, that the restricted linear combination of forward rates predicts excess bond returns better than individual Fama-Bliss slopes. Why does this happen, and can it be due to over-fitting? It turns out, that the restricted factor has to predict excess returns, because it also predicts yield changes in each of the complementary Fama-Bliss regressions, controlling for the Fama-Bliss slopes. Consider again the pairs of complementary Fama-Bliss regressions with the CP factor added as an additional regressor (with an opposite sign, so that after summing the two equations the identity (18) is preserved)

$$rx_{t+1}^n = a + b(f_t^n - f_t^1) + c(cp_t) + u_t^n, \tag{23}$$

$$(n-1)(y_{t+1}^{n-1} - y_t^{n-1}) = -a + (1-b)(f_t^n - f_t^1) - c(cp_t) - u_t^n. \tag{24}$$

Table A3 provides the estimates. The R^2 statistics are now relatively large in both panels. The Cochrane-Piazzesi factor has strong predictive power to explain future movements in *both* yields and excess returns, with opposite signs. This occurs despite the Fama-Bliss slopes are controlled for, so the new variable has important predictive content. Since the linear combination of all forward rates is essentially a linear combination of Fama-Bliss slopes (plus the spot rate), it seems that all of the slopes taken together have much greater ability to predict bond returns than individual slopes.[27]

B Omitted Variable Bias

Consider a simple two-variable setting, where the true specification is

$$y_{t+1} = \alpha_0 + \alpha_1 x_{1,t} + \alpha_2 x_{2,t} + u_{t+1},$$

with the error term orthogonal to both regressors.

[27] This argument is closely related to the predictability of the spot rate by the Cochrane-Piazzesi factor, explained in the appendix to their paper, cited in a footnote above.

Omitting $x_{2,t}$ is equivalent to moving its part orthogonal to $x_{1,t}$ from the regression to the error term. First, consider contemporaneous regression of $x_{2,t}$ on $x_{1,t}$,

$$x_{2,t} = \beta_0 + \beta_1 x_{1,t} + w_t,$$

where by construction w_t is orthogonal in sample to $x_{1,t}$.

Now, substitute the second equation into the first, to get

$$y_{t+1} = (\alpha_0 + \alpha_2\beta_0) + (\alpha_1 + \alpha_2\beta_1)x_{1,t} + (u_{t+1} + \alpha_2 w_t).$$

The error term is orthogonal to $x_{1,t}$, because both of its components are. The coefficient at $x_{1,t}$ is now biased with respect to the true specification. Similarly, the variance of the error term will be different than in the true equation. In particular, if $\alpha_2 > 0$, the R^2 will be biased downwards. If $\alpha_2\beta_1 < 0$, the coefficient will be biased towards zero.

References

Alvarez F., Jermann U., *Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth*, Econometrica 73, (2005), 1977-2016.

Ang A., Bekaert G., Wei M., *Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?*, Journal of Monetary Economics 54, (2007), 1163-1212.

Ang A., Piazzesi M., *A No-Arbitrage Vector Autoregression of Term Structure Dynamics With Macroeconomic and Latent Variables*, Journal of Monetary Economics 50, (2003), 475-787.

Campbell J. Y., Ammer J., *What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns*, Journal of Finance 48, (1993), 3-37.

CAMPBELL J. Y., SHILLER R. J., *Yield Spreads and Interest Rate Movements: A Bird's Eye View*, Review of Economic Studies 58, (1991), 495-514.

CHRISTENSEN J. H. E., LOPEZ J. A., RUDEBUSCH G. D., *Inflation Expectations and Risk Premiums in an Arbitrage-Free Model of Nominal and Real Bond Yields*, Working paper, (2010).

CIESLAK A., POVALA P., *Understanding Bond Term Premia*, Working paper, (2010).

COCHRANE J. H., *Asset Pricing*, Princeton, (2005).

COCHRANE J. H., PIAZZESI M., *Bond Risk Premia*, American Economic Review 95, (2005), 138-160.

COCHRANE J. H., PIAZZESI M., *Decomposing the Yield Curve*, Working paper, (2008).

DAHLQUIST M., HASSELTOFT H., *International Bond Risk Premia*, Working paper, (2010).

D'AMICO S., KIM D. H., WEI M., *Tips From the TIPS: the Informational Content of Treasury Inflation-Protected Security Prices*, Working paper, (2010).

DIEBOLD F. X., RUDEBUSCH G. D., AROUBA S. B., *The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach*, Journal of Econometrics 131, (2006), 309-338.

DICK C. D., SCHMELING M., SCHRIMPF A., *Macro Expectations, Aggregate Uncertainty, and Expected Term Premia*, Working Paper, (2010).

DUFFEE G. R., *Information in (and Not in) the Term Structure*, Working paper, (2009).

FAMA E. F., *Short-Term Interest Rates as Predictors of Inflation*, American Economic Review 65, (1975), 269-282.

FAMA E. F., *Term-Structure Forecasts of Interest Rates, Inflation and Real Returns*, Journal of Monetary Economics 25, (1990), 59-76.

Fama E., *The Behavior of Interest Rates*, Review of Financial Studies 19, (2006), 359-379.

Fama E. F., Gibbons M. R., *Inflation, Real Returns and Capital Investment*, Journal of Monetary Economics 9, (1982), 297-324.

Fama E. F., Bliss R. R., *The Information in Long Maturity Forward Rates*, American Economic Review 77, (1987), 680-692.

Guykarnak R. S., Wright J. H., *Macroeconomics an the Term Structure*, Working paper, (2010).

Inoue A., Kilian L., *In-Sample or Out-of-Sample Tests of Predictability: Which Ones Should We Use?*, Econometric Reviews 23, (2004), 371-402.

Joslin S., Priebsch M., Singleton K. J., *Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks*, Working paper, (2010).

Joyce M. A. S., Lidholdt P., Sorensen S., *Extracting Inflation Expectations and Inflation Risk Premia From the Term Structure: A Joint Model of the UK Nominal and Real Yield Curves*, Journal of Banking and Finance 34, (2010), 281-294.

Kessler S., Scherer B., *Varying Risk Premia in International Markets*, Journal of Banking and Finance 33, (2009), 1361-1375.

Kim D. H., *Challenges in Macro-Finance Modeling*, Federal Reserve Bank of Saint Louis Review 91, (2009), 519-544.

Koijen R., Lustig H., van Nieuwerburgh, S., *The Cross-Section and Time Series of Stock and Bond Returns*, Working paper, (2010).

Krishnan C. N. V., Ritchken P. H., Thompson J. B., *Predicting Credit Spreads*, Journal of Financial Intermediation 19, (2010), 529-563.

LETTAU M., LUDVIGSON S. C., *Measuring and Modeling Variation in the Risk-Return Tradeoff*, in Yacine Ait-Sahalia and Lars Peter Hansen, eds. *Handbook of Financial Econometrics*, Elsevier, North Holland (2010).

LETTAU M., WACHTER J. A., *The Term Structure of Equity and Interest Rates*, Working paper, (2010).

LUDVIGSON S. C., NG S., *Macro Factors in Bond Risk Premia*, Review of Financial Studies 22, (2009), 5027-5067.

MANKIW N. G., MIRON J. A., *The Changing Behavior of the Term Structure of Interest Rates*, Quarterly Journal of Economics C1-2, (1986), 211-228.

NEWEY W. K., WEST K. D., *A Simple, Positive Definite, Heteroskedasticity, and Autocorrelation Consistent Covariance Matrix*, Econometrica 55, (1987), 702-708.

ORPHANIDES A., WEI M., *Evolving Macroeconomic Perceptions and the term Structure of Interest Rates*, Working paper (2010).

PIAZZESI M., SCHNEIDER M., *Equilibrium Yield Curves*, NBER Working Paper 12609, (2006).

PIAZZESI M., SCHNEIDER M., *Trend and Cycle in Bond Premia*, Working paper, (2010).

RUDEBUSCH G. D., WU T., *A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy*, Economic Journal 118, (2008), 906-926.

SÖDERLIND P., *Inflation Risk Premia And Survey Evidence on Macroeconomic Uncertainty*, Working paper, St. Gallen, (2009).

WRIGHT J. H., *Term Premia and Inflation Uncertainty: International Evidence From an International Panel Dataset*, American Economic Review, forthcoming.

Panel A: forward rates (691 obs.)					
	1yr	2yr	3yr	4yr	5yr
means	5.37	5.77	6.09	6.28	6.33
st. dev.	2.93	2.87	2.72	2.72	2.61
cond. st. dev.	1.64	1.41	1.23	1.26	1.15
autocorrelations	0.77	0.81	0.84	0.84	0.87
correlations	1.00	0.97	0.94	0.91	0.89
		1.00	0.98	0.97	0.95
			1.00	0.98	0.97
				1.00	0.97
					1.00
Panel B: realized excess returns (679 obs.)					
means		0.42	0.74	0.92	0.94
st. dev.		1.71	3.12	4.32	5.29
autocorrelations		0.19	0.13	0.10	0.06
correlations		1.00	0.98	0.96	0.95
			1.00	0.99	0.98
				1.00	0.99
					1.00
Panel C: inflation rate (679 obs.)					
mean	4.42				
st. dev.	2.94				
cond. st. dev.	2.13				
autocorrelation	0.73				
correlations w. fwd	0.61	0.55	0.50	0.48	0.49
correlations w. rx		-0.35	-0.39	-0.38	-0.39

Table 1. Panel A presents the descriptive statistics of 5 forward rates from the Fama-Bliss dataset, from June 1952 to December 2009. Conditional standard deviations are standard deviations of yearly changes. The autocorrelations are of the order of 12 months. Panel B contains analogous statistics for one-year realized excess bond returns. Panel C presents similar statistics for the inflation rate, together with its contemporaneous correlations with five forward rates and four excess return series.

	α_0	α_1	α_2	α_3	α_4	α_5	R^2
full sample	-0.012	-1.05	0.25	0.94	0.68	-0.63	0.19
	0.006	0.29	0.45	0.48	0.32	0.32	
old sample	-0.017	-1.38	0.89	1.08	0.49	-0.87	0.26
	0.006	0.25	0.44	0.50	0.32	0.36	
new sample	-0.016	-1.15	-0.09	1.71	0.84	-1.12	0.25
	0.009	0.30	0.54	0.48	0.38	0.36	

Table 2. Estimation results for equation (7). The first row contains results for the full sample from June 1952 to December 2008, with standard errors beneath the coefficients. The second row corresponds to the older sample, from January 1964 to May 1997, and the third row to the newer sample, from January 1964 to December 2008. The standard errors use the Newey-West correction for autocorrelation and heteroskedasticity, with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987].

	δ_0	δ_1	δ_2	δ_3	δ_4	δ_5	R^2
full sample	0.010	1.48	-0.46	-1.17	0.04	0.68	0.45
	0.005	0.36	0.51	0.44	0.29	0.22	
old sample	0.013	1.75	-0.59	-1.25	0.07	0.57	0.48
	0.005	0.37	0.56	0.43	0.30	0.26	
new sample	0.023	1.41	-0.04	-1.68	-0.07	0.79	0.39
	0.006	0.37	0.59	0.48	0.34	0.25	

Table 3. Estimation results for equation (8). The first row contains results for the full sample from June 1952 to December 2008, with standard errors beneath the coefficients. The second row corresponds to the older sample, from January 1964 to May 1997, and the third row to the newer sample, from January 1964 to December 2008. The standard errors use the Newey-West correction for autocorrelation and heteroskedasticity, with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987].

Panel A: estimation results for (12)				
	κ_0	κ_1	κ_2	R^2
full sample	-0.002	0.31	-0.92	0.18
	0.006	0.07	0.19	
old sample	-0.009	0.36	-1.02	0.25
	0.006	0.07	0.18	
new sample	-0.003	0.36	-1.11	0.23
	0.007	0.07	0.18	
Panel B: estimation results for (13)				
	κ_0	κ_1	κ_2	R^2
full sample	0.000	1.00		0.19
	0.002	0.19		
old sample	-0.004	1.12		0.25
	0.002	0.19		
new sample	-0.002	1.20		0.25
	0.002	0.19		
Panel C: estimation results for (12), reestimated factors				
	κ_0	κ_1	κ_2	R^2
full sample	-0.002	0.31	-0.92	0.18
	0.006	0.07	0.19	
old sample	-0.009	0.32	-0.85	0.24
	0.006	0.06	0.15	
new sample	0.008	0.28	-1.02	0.24
	0.007	0.06	0.16	
Panel D: estimation results for (13), reestimated factors				
	κ_0	κ_1		R^2
full sample	0.000	1.00		0.19
	0.002	0.19		
old sample	0.000	1.00		0.26
	0.002	0.17		
new sample	0.000	1.00		0.25
	0.002	0.15		

Table 4. Forecasting the first principal component in forward premia using the level of the term structure of forward rates plus IE (panels A and C), and the CP factor (panels B and D). The first row in each panel corresponds to the full sample, the second to the "older", and the last to the "newer" subsamples (for the definitions, see the main text). Regressions in panels C and D use re-estimated factors. The standard errors are beneath the coefficients, and are autocorrelation- and heteroskedasticity-adjusted, using the Newey-West procedure, with lag truncation q set automatically to $q = floor(4(T/100)^{2/9})$, as suggested in Newey-West [1987].

h	1	2	3	4	5	6	7	8
α_1^h	0.50	0.24	0.08	0.05	0.13	0.13	0.17	0.12
s.e.	0.07	0.09	0.09	0.08	0.08	0.08	0.07	0.08
R^2	0.25	6	1	0	2	2	3	1
β_1^h	-0.24	-0.29	-0.31	-0.37	-0.42	-0.49	-0.57	-0.66
s.e.	0.07	0.10	0.10	0.10	0.12	0.11	0.14	0.16
R^2	10	8	7	8	9	10	12	13
γ_1^h	0.26	0.47	0.61	0.58	0.45	0.37	0.25	0.22
s.e.	0.10	0.13	0.16	0.16	0.15	0.15	0.17	0.19
R^2	5	8	11	9	6	4	2	1

Table 5. Estimation results for regressions (17). The table presents the slope coefficients, their standard errors and the R^2 statistics, computed using the whole sample from June 1952 to December 2008. The standard errors use the Newey-West correction for autocorrelation and heteroskedasticity, with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987].

Panel A: Fama-Bliss regressions, equation (19), full sample			
maturity	b	$s(b)$	R^2
2	0.83	0.10	0.11
3	1.13	0.13	0.13
4	1.37	0.15	0.14
5	1.03	0.17	0.06
Panel B: Fama-Bliss regressions, equation (20), full sample			
maturity	$(1-b)$	$s(b)$	R^2
2	0.17	0.10	0.01
3	-0.13	0.13	0.00
4	-0.37	0.15	0.01
5	-0.03	0.17	0.00
Panel C: Fama-Bliss regressions, equation (19), CP sample			
maturity	b	$s(b)$	R^2
2	0.99	0.19	0.16
3	1.35	0.24	0.17
4	1.61	0.30	0.18
5	1.27	0.42	0.08
Panel D: Fama-Bliss regressions, equation (20), CP sample			
maturity	$(1-b)$	$s(b)$	R^2
2	0.01	0.19	0.01
3	-0.35	0.24	0.01
4	-0.61	0.30	0.03
5	-0.27	0.42	0.00

Table A1. Panels A and B show estimation results for the pairs of complementary Fama-Bliss regressions as in equations (19) and (20), for the sample from January 1964 to December 2008. Panels C and D show the same results for the sample from January 1964 to December 2002 (used in Cochrane and Piazzesi [2005]). The standard errors use the Newey-West correction for autocorrelation and heteroskedasticity, with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987]. The R^2 statistics are adjusted for the degrees of freedom.

Panel A: Updated sample							
maturity	γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	R^2
2	-0.01	-0.69	0.13	0.96	0.42	-0.65	0.23
	0.01	0.19	0.36	0.32	0.23	0.22	
3	-0.02	-1.23	-0.30	2.58	0.63	-1.42	0.24
	0.01	0.35	0.64	0.57	0.43	0.41	
4	-0.02	-1.82	-0.25	2.91	1.62	-2.14	0.27
	0.01	0.48	0.85	0.77	0.59	0.57	
5	-0.03	-2.29	-0.14	3.20	1.69	-2.08	0.24
	0.02	0.60	1.06	0.95	0.75	0.71	
restricted	-0.02	-1.51	-0.14	2.41	1.09	-1.57	0.25
	0.01	0.40	0.73	0.65	0.50	0.48	
fprpc1	-0.02	-1.15	-0.09	1.71	0.84	-1.12	0.25
	0.01	0.30	0.54	0.48	0.38	0.36	
Panel B: Cochrane-Piazzesi sample							
maturity	γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	R^2
2	-0.02	-0.98	0.59	1.21	0.29	-0.89	0.31
	0.01	0.18	0.36	0.30	0.23	0.21	
3	-0.03	-1.78	0.54	3.07	0.38	-1.86	0.33
	0.01	0.31	0.64	0.54	0.42	0.40	
4	-0.04	-2.57	0.87	3.60	1.28	-2.73	0.36
	0.01	0.42	0.85	0.73	0.58	0.55	
5	-0.05	-3.21	1.25	4.10	1.25	-2.83	0.34
	0.02	0.53	1.05	0.92	0.73	0.69	
restricted	-0.03	-2.14	0.81	3.00	0.80	-2.08	0.34
	0.01	0.36	0.72	0.62	0.49	0.46	
fprpc1	-0.03	-1.62	0.62	2.16	0.62	-1.49	0.34
	0.01	0.27	0.53	0.47	0.37	0.35	

Table A2. The Cochrane-Piazzesi [2005] predictive regressions. In both panels, rows referring to maturities 2-5 show coefficient estimates in the unrestricted specifications, defined in (21). The row denoted "restricted" refers to the specification in equation (22), and the one with "fprpc1" presents results related to methodology described in subsection (2.3.2). Beneath each row the standard errors are presented, computed using the Newey-West correction for autocorrelation and heteroskedasticity with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987]. The R^2 statistics are adjusted for the degrees of freedom. Panel A shows estimation results for the sample from January 1964 to December 2008, while the results in Panel B are obtained using the sample between January 1964 and December 2002, the same as in Cochrane and Piazzesi [2005].

Panel A: Fama-Bliss regressions with CP, equation (23)					
maturity	b	$s(b)$	c	$s(c)$	R^2
2	-0.10	0.25	0.63	0.12	0.22
3	-0.16	0.34	1.22	0.26	0.24
4	-0.12	0.39	1.75	0.36	0.27
5	0.08	0.31	1.92	0.32	0.25
Panel B: Fama-Bliss regressions with CP, equation (24)					
maturity	$(1-b)$	$s(b)$	-c	$s(c)$	R^2
2	1.10	0.25	-0.63	0.12	0.13
3	1.16	0.34	-1.22	0.26	0.13
4	1.12	0.39	-1.75	0.36	0.16
5	0.92	0.31	-1.92	0.32	0.20

Table A3. Estimation results for the pairs of complementary Fama-Bliss regressions, defined in equations (23) and (24), for the sample from January 1964 to December 2008. The standard errors use the Newey-West correction for autocorrelation and heteroskedasticity, with lag truncation q set automatically to the value $q = floor(4(T/100)^{2/9})$, suggested in Newey-West [1987]. The R^2 statistics are adjusted for the degrees of freedom.

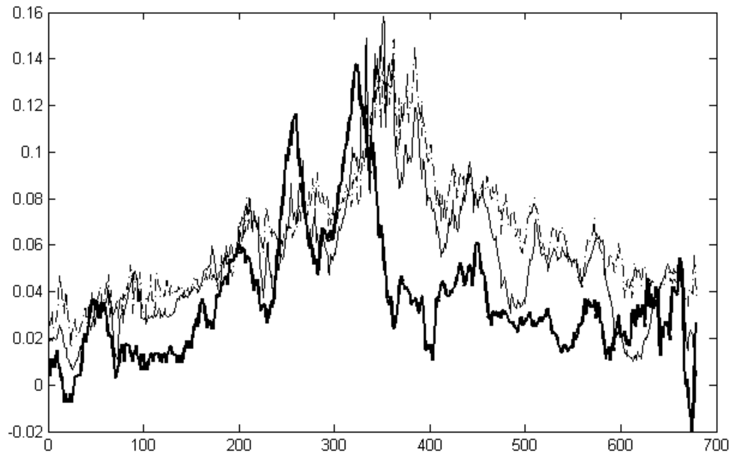

Figure 1. Yearly inflation rate realized over subsequent 12 months (solid line, thick), one-year spot rate (solid line, thin), and five-year forward rate (for contracts entered now and working between $t+4$ and $t+5$ years). The time-span is June 1952 - December 2008. The spot and forward rates are from the Fama-Bliss (FB) dataset.

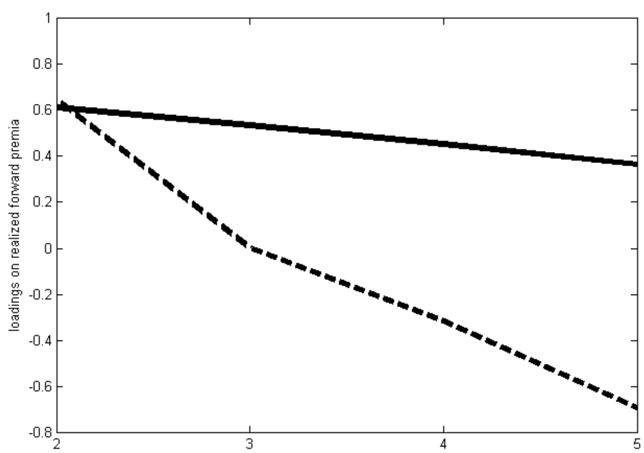

Figure 2. The solid and dotted lines are the eigenvectors of unconditional variance of the forward premia derived from the Fama-Bliss forward rates, corresponding to two largest eigenvalues. See equation 5.

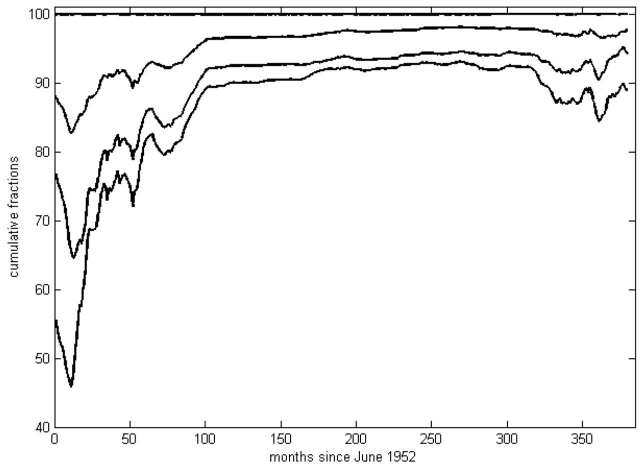

Figure 3. Fractions of predictable variation in the forward premia due to the four principal components (see section 2.3.1). The fractions that correspond to the most important components (as measured by the size of the eigenvalues) are drawn from the bottom to the top.

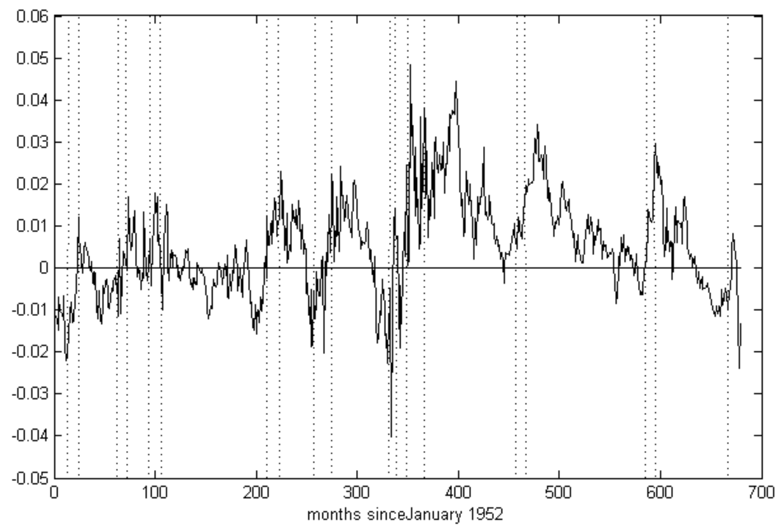

Figure 4. The CP factor and NBER recession dates.

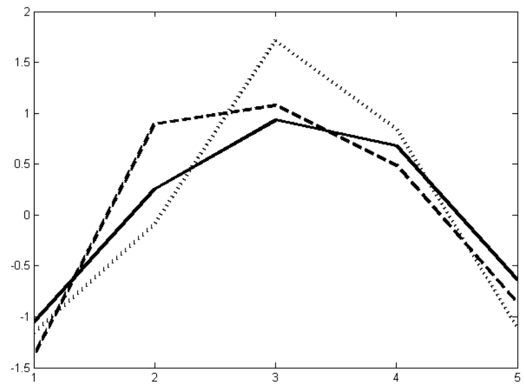

Figure 5. The tent-shaped pattern of coefficients in the linear combination of the forward rates defining the CP factor in equation (7). The solid line corresponds to the results in the full sample from June 1952 to December 2008. The dashed line corresponds to the older sample from January 1964 to May 1997, and the dotted line to the newer sample from January 1964 to December 2008.

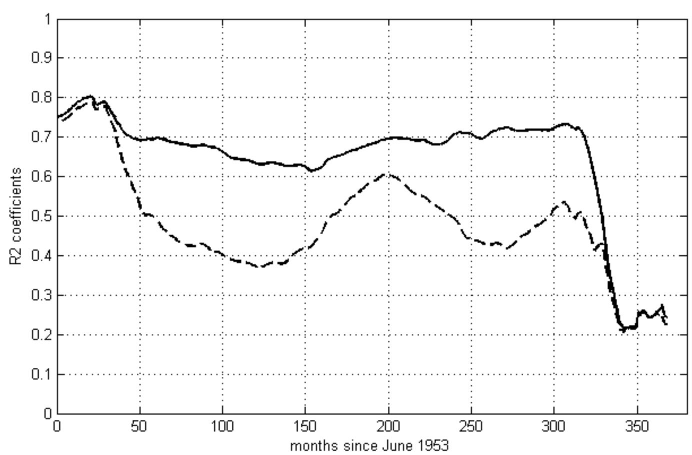

Figure 6. The R^2 coefficients from regression (11), with and without the AR term (solid and dotted lines, respectively). The x axis positions the beginning months of each of the 25-year wide rolling window, relative to June 1953.

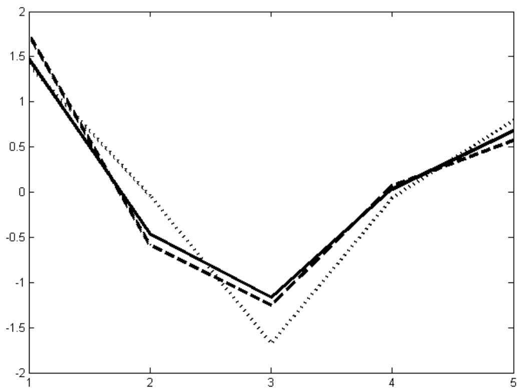

Figure 7. The tent-shaped pattern of coefficients in the linear combination of the forward rates defining the IE factor in equation (8). The solid line corresponds to the results in the full sample from June 1952 to December 2008. The dashed line corresponds to the older sample from January 1964 to May 1997, and the dotted line to the newer sample from January 1964 to December 2008.

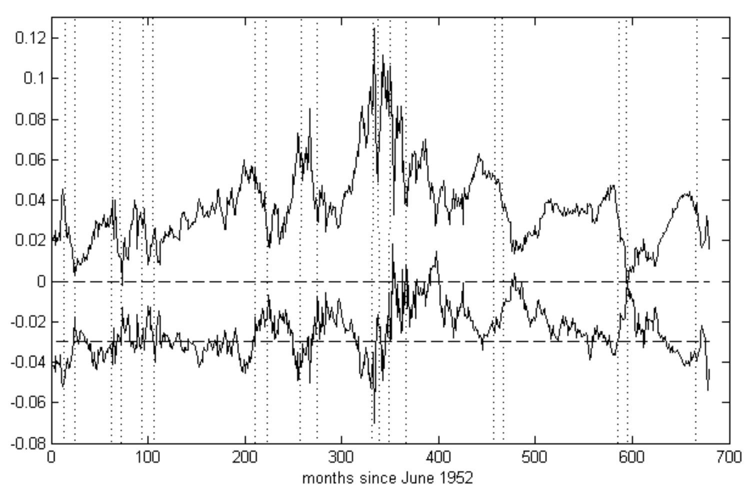

Figure 8. The CP factor (bottom) and the IE factor (top). The former is shifted down by 0.03, for better readability. The vertical lines are the beginning and ending dates of NBER recessions.

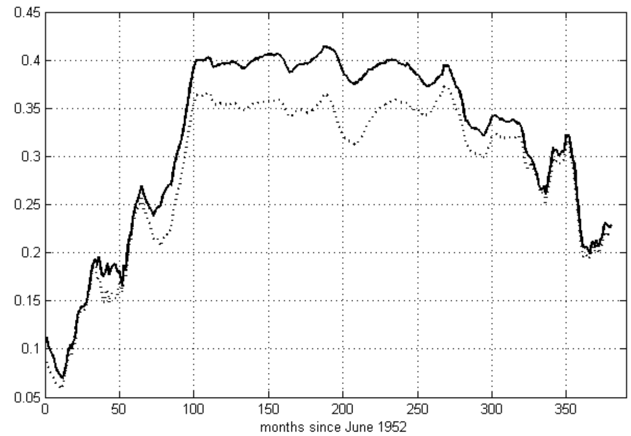

Figure 9. The R^2 coefficients from regressions (13) and (12). The solid line corresponds to the R^2 coefficients obtained by using the CP factor as the single regressor. The dashed line uses the level and the IE factors instead. The x axis positions initial months of each of the 25-year wide rolling windows, relative to June 1952. All factors on both sides of the regressions are re-estimated every time.

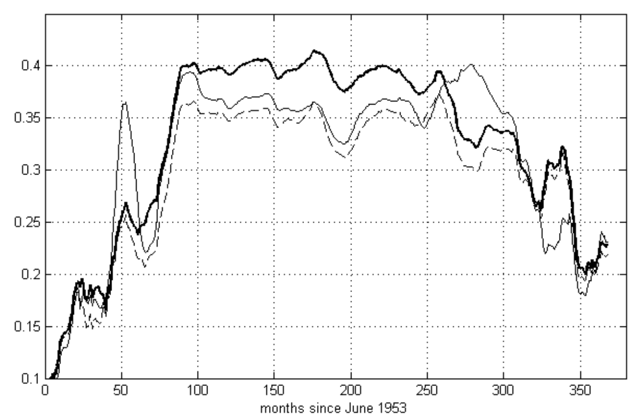

Figure 10. The R^2 coefficients from regressions (13) (thick solid line), and (12). The thin dashed line is the same as the dashed line in figure 9. The thin solid line is generated using an alternative specification of the IE factor (see the main text).

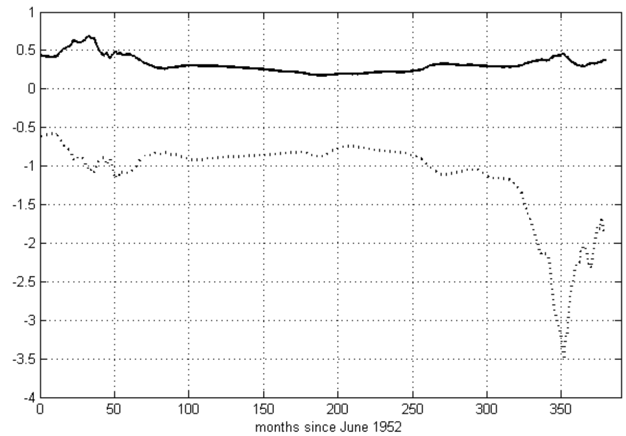

Figure 11. The slope coefficients in (12), estimated in rolling windows. The x axis positions initial months of each of the 25-year wide rolling windows, relative to June 1952. All factors on both sides of the regression are re-estimated.

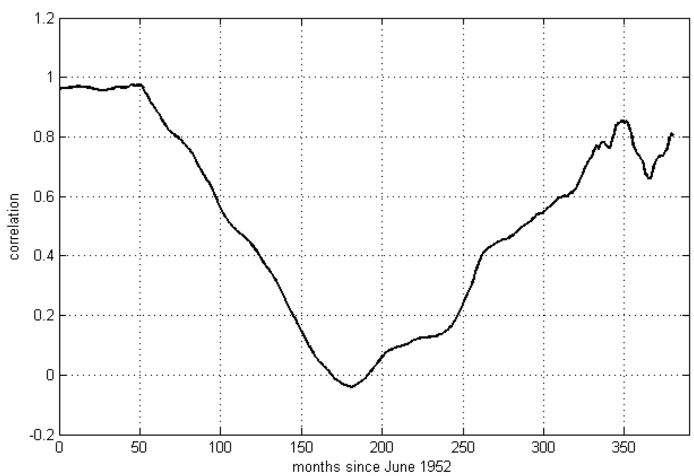

Figure 12. The correlation coefficients between the level and the IE factor, estimated in rolling windows. The x axis positions initial months of each of the 25-year wide windows, relative to June 1952. The factors are re-estimated every time the window changes position.

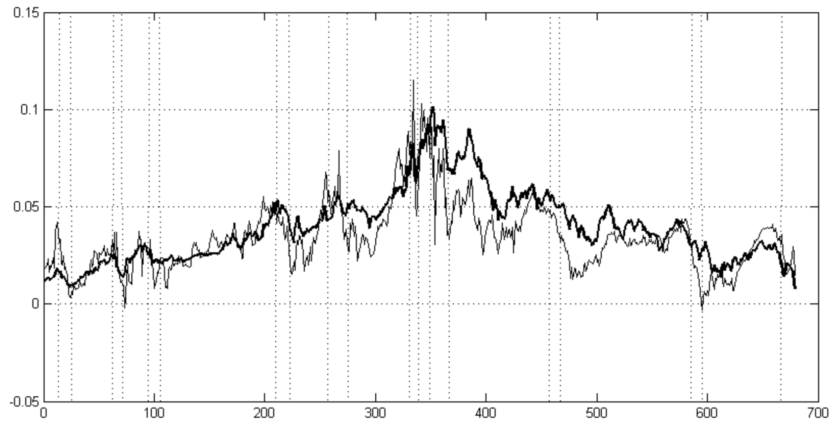

Figure 13. The components of cointegrating relation in equation (14). The thick solid line is the scaled level of forward rates. The thin line is the IE factor. Vertical dotted lines mark NBER recessions.

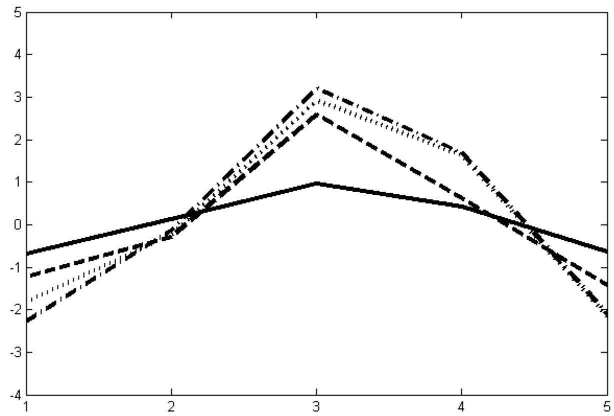

Figure A1. Cochrane-Piazzesi [2005] regressions. Graphical illustration of the estimation results of equation (21) for excess returns of bonds with maturities of two, three, four and five years (solid, dashed, dotted and dash-dotted lines, respectively), on forward rates corresponding to maturities between one and five years (the constant omitted from the graph). The sample is from January 1964 to December 2008.

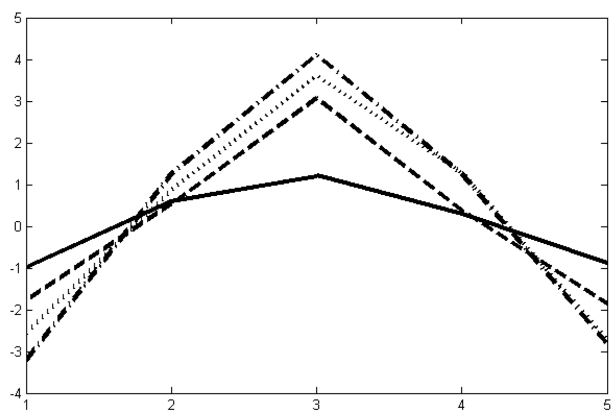

Figure A2. Cochrane-Piazzesi [2005] regressions. Graphical illustration of the estimation results of equation (21) for excess returns of bonds with maturities of two, three, four and five years (solid, dashed, dotted and dash-dotted lines, respectively), on forward rates corresponding to maturities between one and five years (the constant omitted from the graph). The sample is from January 1964 to December 2002.