Bond Risk Premia and the "Return Forecasting Factor"

Agustin Gutierrez
University of Chicago

Constantino Hevia
Universidad Torcuato Di Tella

Martin Sola
Universidad Torcuato Di Tella

October 4, 2018
(First version: April 3, 2017)

Abstract

The return forecasting factor is a linear combination of forward rates that seems to predict one-year excess bond returns of bond of all maturities better than traditional measures obtained from the yield curve. If this single factor actually captures all the relevant fluctuations in bond risk premia, then it should also summarize all the economically relevant variations in excess returns considering different holding periods. We find that it does not. We conclude that including the return forecasting factor as the main driver of risk premia in a term structure model, as has been suggested, is not supported by the data.

Keywords: Excess returns, bond risk premia, return forecasting factor, affine term structure models.

1 Introduction

For many years researchers agreed that the first three principal components (PC) of bond yields contain all the relevant information in bond prices. As a result, a large class of affine term structure models assumes that only those three factors determine the evolution of bond prices. Recent studies, however, suggest that other variables contain additional information that is useful for predicting bond returns. This evidence comes from predictive regressions of bond returns on variables other than (and possibly including) the first three PC.[1]

In an influential paper, Cochrane and Piazzesi (2005) (CP), argue that a linear combination of forward rates, labeled the "return forecasting factor" (RFF), predicts one-year excess bond returns of *each* maturity with an R^2 of about 0.35.[2] One implication of this observation is that a *single* factor may capture all the economically relevant variations in expected bond returns. As a result, Cochrane and Piazzesi (2008) propose to augment the three-factor affine term structure model by adding a fourth (RFF) factor as the main driver of risk premia.

The RFF is the linear combination of forwards rates that maximizes the predictive ability in *one-year* excess returns regressions. Yet, if the RFF is the only driver of bond risk premia, then it should also summarize all the relevant variations in excess returns considering holding periods different from one year. We find that it does not. The one-year RFF outperforms the first three PC only for one-year excess return regressions. In regressions of excess returns for holding periods longer than a year, the first three PC provide a better fit of excess returns. And for sufficiently long holding periods, the PC have more than three times the explanatory power of the RFF.

We also reject the null hypothesis that there is a single factor structure in expected excess returns over different holding horizons. Although CP also reject the single factor model, they claim that the RFF captures all the economically relevant variations in one-year excess returns. When considering different holding horizons, however, we find that it does not.

Finally, if the RFF is a relevant state variable in a term structure model, it should also be a

[1] Bauer and Hamilton (2018) define the "spanning hypothesis" as the claim that the first three PC contain all the relevant information for predicting future interest rates and bond returns. That paper also contains an extensive list of papers that claim to find violations of the spanning hypothesis. However, they have questioned the value of variables other than the first three PC as predictors of bond returns. The conventional tests that support including those variables in predictive regressions are subject to small-sample distortions that lead to oversized tests.

[2] The traditional Fama and Bliss (1987) regressions of excess returns on forwards rates attain R^2s of about 0.17, while regressions on the first three principal components deliver R^2s of about 0.27.

significant explanatory variable in unrestricted regressions of bond yields on the RFF and the first three PC. Using the robust t-test developed by Ibragimov and Müller (2010), we reject the significance of the RFF on such regressions, while the first three PC are highly significant.

Taken together, these findings contradict the single factor structure in bond returns and imply that including the RFF as the main driver of bond risk premia in a term structure model is not supported by the data.

2 The return forecasting factor and excess bond returns

We first define notation. Let $p_t^{(n)}$ denote the log price of an n-year zero coupon bond, where time and maturity are both expressed in years. The log yield is $y_t^{(n)} = -p_t^{(n)}/n$ and the log forward rate for loans starting at $t+n$ and maturing at $t+n+h$ is $f_t^{(n,n+h)} = p_t^{(n)} - p_t^{(n+h)}$. Let

$$r_{t\to t+h}^{(n)} = (p_{t+h}^{(n-h)} - p_t^{(n)})/h$$

denote the annualized h-year log holding return from buying an n-period zero coupon bond at time t and selling it at time $t+h$ as an $(n-h)$-year zero coupon bond. The excess return from holding an n-year zero coupon bond for h years is defined as

$$rx_{t\to t+h}^{(n)} = r_{t\to t+h}^{(n)} - y_t^{(h)}.$$

Finally, we denote the average h-year excess return across maturities by $\overline{rx}_{t\to t+h} = \frac{1}{H}\sum_{n>h} rx_{t\to t+h}^{(n)}$, where H is the number of bonds with maturities longer than h periods ahead.

2.1 Data

We use monthly data on U.S. bond prices for the period Jan-1964 through December-2003. We focus on the same sample period used by Cochrane and Piazzesi (2005) to make sure that our results do not change because the sample is different. We gather data from two sources. First, we use the Fama-Bliss data (Fama and Bliss, 1987) of one through five-year zero coupon bond prices. From these data, we construct five one-year forward rates $(y_t^{(1)}, f_t^{(1,2)}, f_t^{(2,3)}, f_t^{(3,4)}, f_t^{(4,5)})$ and one- through four-year excess holding returns. One limitation of the Fama-Bliss data is that

it is not possible to construct excess returns for holding periods different from those. For that reason, we also use the Gürkaynak, Sack and Wright (2007) dataset (GSW) which contains monthly records of zero coupon treasury yields with yearly maturities $(1/12, 2/12, ..., 120/12)$. These yields are generated using the extended Nelson-Siegel parametrization of the yield curve proposed by Svensson (1994), whose coefficients are available on a daily basis. From the fitted curves we recover bond prices, yields, and forward rates of any maturity ranging from one month to ten years.

2.2 The predictive one-year excess return regressions

Cochrane and Piazzesi (2005) run predictive regressions of one-year excess returns on government bonds on the five Fama-Bliss forward rates at the beginning of the period. Their main result is that a single linear combination of the five Fama-Bliss forward rates (the RFF) predicts one-year excess returns of bonds of each maturity with an R^2 of about 0.35, and up to 0.44 if one includes lagged forward rates as regressors. Furthermore, the RFF outperforms the first three PC extracted from the yield curve in terms of forecasting ability.

CP consider a single-factor predictive regression of excess returns of the form

$$rx_{t\to t+1}^{(n)} = b_n \left(\gamma' \mathbf{f}_t\right) + \epsilon_{t+1}^{(n)}, \tag{1}$$

where $n = 2, 3, 4, 5$ and $\mathbf{f}_t = \left[1,\, y_t^{(1)},\, f_t^{(1,2)},\, f_t^{(2,3)},\, f_t^{(3,4)},\, f_t^{(4,5)}\right]'$ is a vector with a constant and five Fama-Bliss forwards rates at time t. Equation (1) imposes a single factor structure on one-year excess returns.[3] The alternative, unrestricted, regressions are given by

$$rx_{t\to t+1}^{(n)} = \boldsymbol{\beta}_n' \mathbf{f}_t + \epsilon_{t+1}^{(n)}, \tag{2}$$

where the coefficients $\boldsymbol{\beta}_n$ are estimated independently for each maturity. CP show that estimating the one-factor model (1) leads to almost identical R^2s than estimating the unrestricted regressions (2). Moreover, the RFF has more predictive power than the first three PC from bond yields.

[3] Cochrane and Piazzesi (2005) construct the return forecasting factor $\gamma' \mathbf{f}_t$ by regressing the average one-year excess return $\overline{rx}_{t\to t+1}$ on the five forward rates. In a second step, they estimate the factor loadings b_n by using $\hat{\gamma}' \mathbf{f}_t$ as a generated regressor in equation (1).

Table 1 reports the R^2s of regressions of the average one-year excess returns (across maturities) on the return forecasting factor and the principal components. The left panel of the table, labeled *Fama-Bliss data*, uses the five Fama-Bliss discount prices to perform the regressions. The right panel, labeled *GSW data*, displays the same results but using ten one-year forward rates $f_t = \left[1, y_t^{(1)}, f_t^{(1,2)}, f_t^{(2,3)}, ..., f_t^{(9,10)}\right]'$ constructed using the Gürkaynak, Sack and Wright (2007) data. In both panels, the column labeled "F-test" reports the p-values of the null hypothesis that the coefficient δ_j on the jth principal component is equal to zero.[4]

Table 1: R^2s of one-year excess return forecast regressions

	Fama-Bliss data			*GSW data*		
Regressor	Individual	Joint $PC_1 : PC_i$	F-test	Individual	Joint $PC_1 : PC_i$	F-test
RFF	0.351	—	—	0.358	—	—
PC_1	0.028	0.028	—	0.033	0.033	—
PC_2	0.213	0.242	0.0000	0.169	0.202	0.0000
PC_3	0.024	0.266	0.0001	0.059	0.261	0.0000
PC_4	0.085	0.351	0.0000	0.009	0.269	0.0195
PC_5	0.000	0.351	0.9140	0.021	0.291	0.0002
PC_6				0.026	0.318	0.0000
PC_7				0.013	0.331	0.0028
PC_8				0.014	0.344	0.0021
PC_9				0.010	0.355	0.0066
PC_{10}				0.003	0.358	0.1561

The table report the R^2s of regressions of the average one-year excess returns $\overline{rx}_{t\to t+1}$ on the RFF and the PC of bond yields, and the p-values of F-tests of the null hypothesis that a model using with $j-1$ principal components is as good as a model with j principal components for $j = 2, 3, ...$ The columns labeled *Fama-Bliss data* are based on the five Fama-Bliss discount bond prices. The columns labeled *GSW data* uses the treasury yields constructed by Gürkaynak, Sack and Wright (2007). The sample period is that used by Cochrane and Piazzesi (2005), Jan-1964 to Dec-2003.

These results suggest that variations in one-year excess returns may be driven by a single factor since the R^2 of the regressions using only the RFF is at least as good as those using the principal components. Therefore, Cochrane and Piazzesi (2005) conclude that conventional three-factor models of the yield curve miss an important component of bond returns and that researchers

[4] In particular, consider the regressions $\overline{rx}_{t\to t+1} = \delta_0 + \delta_1 PC_{1t} + \delta_2 PC_{2t} + ... + \delta_j PC_{jt} + u_{t+1}$ for $j = 2, 3, ...$ We compare the sample fit of the model with all the regressors with that of a model with $\delta_j = 0$ using the test statistic $F = \frac{RSS_r - RSS_u}{RSS_u/(T-j)}$, where RSS_u is the residual sum of squares of the unrestricted model, RSS_r is the residual sum of squares of the restricted model that imposes $\delta_j = 0$, and T is the sample size.

should augment the model by adding a fourth (RFF) factor as the main driver of bond risk premia.

2.3 The predictive h-year excess return regressions

To be regarded as the main driver of risk premia, the one-year RFF should also summarize variations in expected bond returns considering holding periods different from one year. In this section we show that it does not. The one-year RFF outperforms the first three PC only in one-year excess return regressions such as those in Table 1. In regressions of excess returns for holding period longer than a year, the first three PC provide a better fit of excess returns. And for sufficiently long holding periods, the PC have three times the explanatory power of the RFF.

We first regress the average one-year excess return $\overline{rx}_{t\to t+1}$ on the forward rates f_t using both the FB and GSW datasets. We call this the "one-year-ahead" RFF because it maximizes the forecasting performance in a regression of the average one-year excess return. Likewise, we label the "m-year-ahead" RFF as the fitted value of a regression of the average m-year excess return $\overline{rx}_{t\to t+m}$ on the forward rates f_t. Let $RFF_t^{(m)}$ denote the m-year-ahead return forecasting factor computed using this procedure (below we consider $m = 1, 2, 3, 4$).

Next, we run predictive regressions of the expected h-year excess return averaged across available maturities $\overline{rx}_{t\to t+h}$, on the m-year return forecasting factor, $RFF_t^{(m)}$ and on the first three principal components PC_{1t}, PC_{2t}, and PC_{3t}, extracted from the same datasets,[5]

$$\overline{rx}_{t\to t+h} = \alpha_0^{(h)} + \alpha_1^{(h)} RFF_t^{(m)} + \epsilon_{t+h} \tag{3}$$

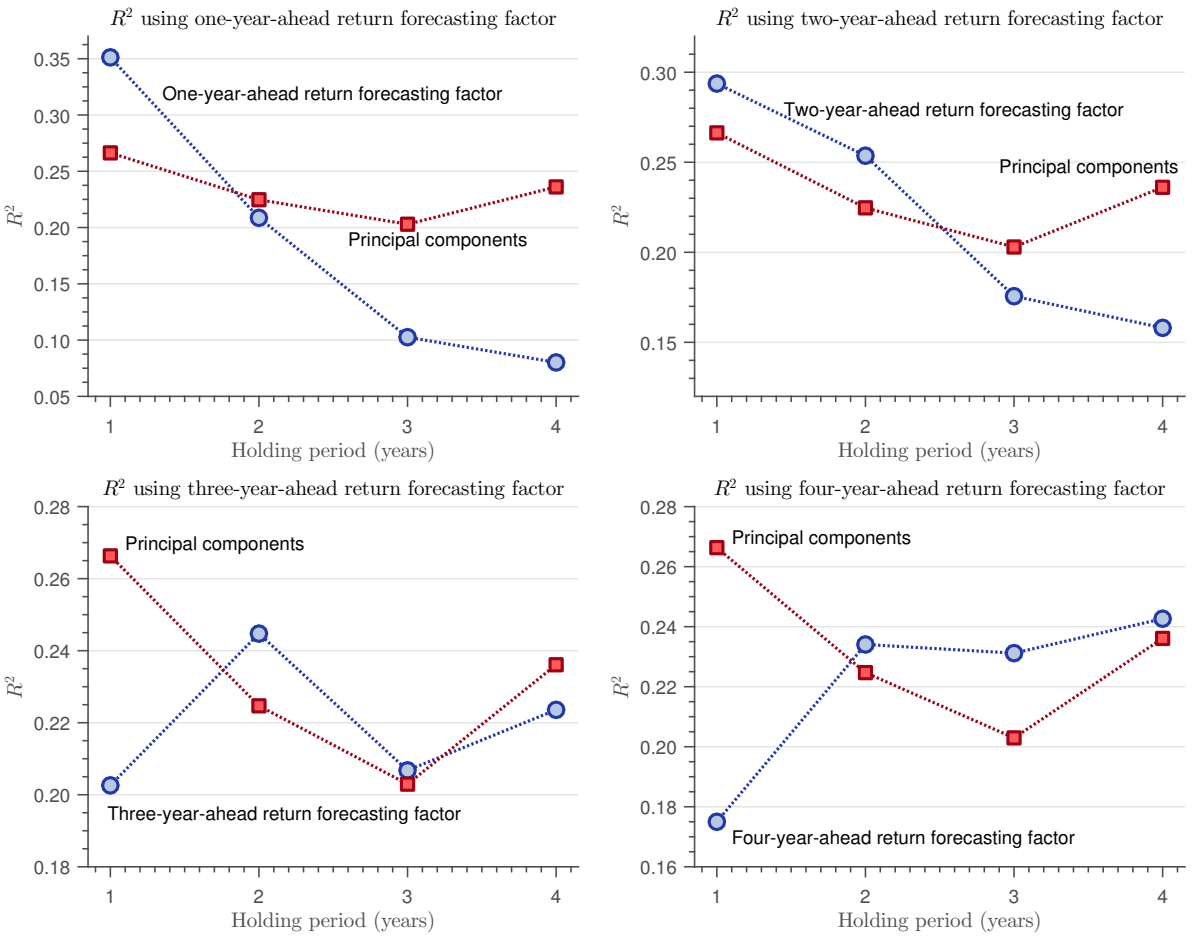
$$\overline{rx}_{t\to t+h} = \delta_0^{(h)} + \delta_1^{(h)} PC_{1t} + \delta_2^{(h)} PC_{2t} + \delta_3^{(h)} PC_{3t} + \eta_{t+h}. \tag{4}$$

Figure 1 displays the R^2 of the regressions of the excess holding returns on the return forecasting factors using the Fama-Bliss dataset. For example, the top left panel of the figure reports the R^2 of the regressing the one- through four-year ahead excess returns on the one-year-ahead return forecasting factor (dotted-circled line) and on the first three principal components (citted-squared line) as a function of the holding period $h = 1, 2, 3, 4$. While the RFF has more predictive ability than the PC at the one-year holding horizon, the relation is reversed at longer holding periods: the RFF explains less than 10 percent of the variability of the 4-year excess return

[5] Appendix A displays the estimated coefficients of the following regressions.

while the PC explain about 25 percent of the variance.

Figure 1: R^2 of the average h-year excess holding return regressions using Fama-Bliss yields



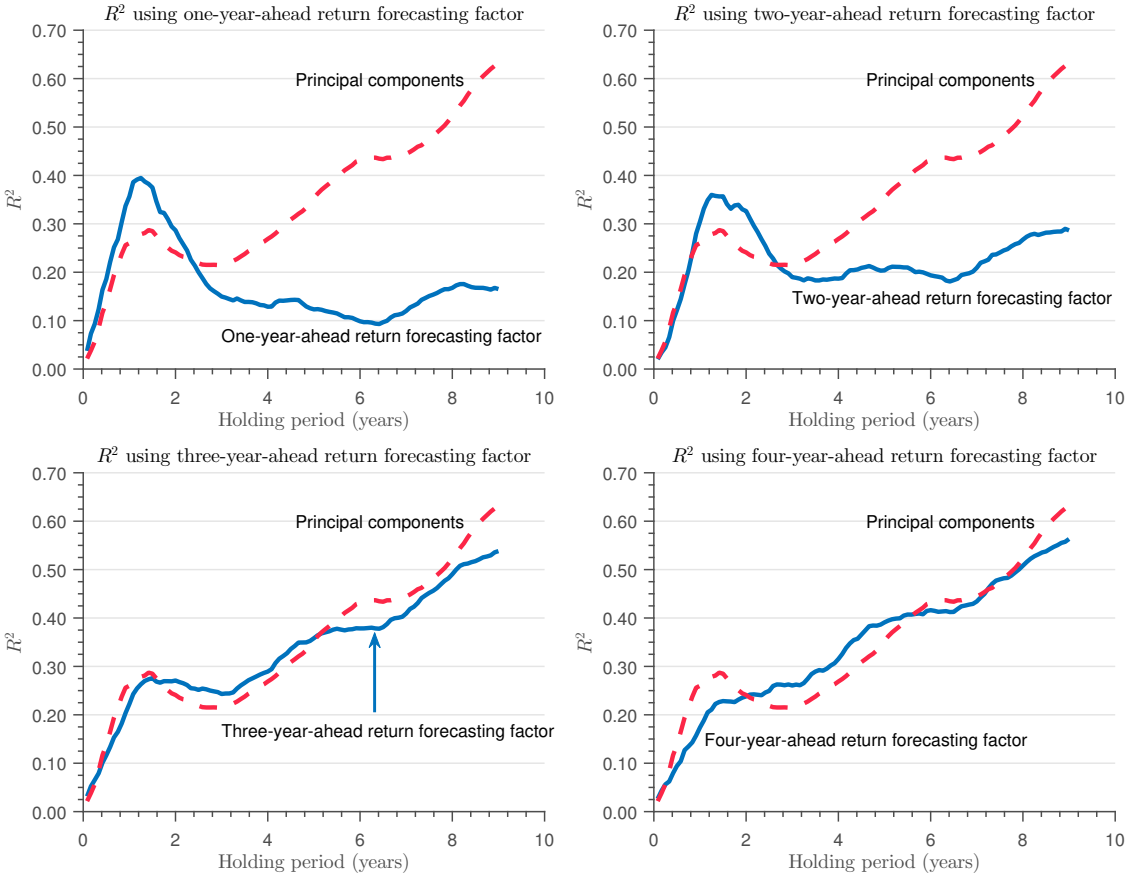
The figure shows the R^2s of a regression of the average h-period excess holding returns (h=1,2,3,4) on the return forecasting factor and on the first three principal components PC_{1t}, PC_{2t}, and PC_{3t} together. In the upper left panel the return forecasting factor is computed using a regression of one-year excess returns on the current forwards; in the upper right panel the return return forecasting factor is computed with a regression of two-year excess returns on the current forwards, and so forth.

In a similar fashion, the upper right panel of Figure 1 shows the R^2 of regressions of the average h-period excess holding returns ($h = 1, 2, 3, 4$) on the two-year ahead return forecasting factor, $RFF^{(2)}$, and the PC, and likewise for the lower two panels of the figure but using $RFF_t^{(3)}$ and $RFF_t^{(4)}$. Again, while the h-year-ahead RFF outperforms the PC at the h-year holding horizon, the PC tend to do better than the RFF at other holding horizons.

This pattern is reinforced when we consider the GSW database. Figure 2 is identical to Figure 1 but using a larger set of holding periods h. For example, the upper left panel of the figure shows

the R^2 of regressions of the average excess holding returns $\overline{rx}_{t\to t+h}$ for $h = 1/12, 2/12, ..., 108/12$ on the one-year-ahead return forecasting factor $RFF_t^{(1)}$ and the first three principal components, as a function of the holding period h. Note that for longer holding horizons, the predictive power of the principal components is more than three times larger than that of the one-year-ahead return forecasting factor $RFF_t^{(1)}$, and twice as large as that of the two-year-ahead return forecasting factor. Finally, the other two return forecasting factors, $RFF_t^{(3)}$ and $RFF_t^{(4)}$, do not provide any predictive improvement over the principal components.[6]

Figure 2: R^2 of h-year excess holding return regressions using GSW data



The figure shows R^2s of regressions of the average h-year excess holding returns on the one- though four-year ahead RFF and on the first three PC. In the upper upper left panel the RFF is computed using a regression of one-year excess returns on one- through ten-year forwards; in the upper right panel the RFF is computed with a regression of two-year excess returns on one- through ten-year forwards, and so forth.

[6]There are two differences between Figures 1 and 2. First, while the Fama-Bliss dataset only allows us to construct excess returns for holding periods of 1, 2, 3 and 4 years, the GSW data allows us to consider a richer set of holding periods, ranging from 1 month (1/12 years) through 108 months (9 years). Hence, relative to Figure 1, the average excess holding returns used in Figure 2 are constructed using a larger set of bond maturities and the figure displays excess returns for a finer and longer set of holding periods. The results using the GSW data restricted to holding periods of one through four years are very similar to those in Figure 1 (results available upon request).

The observation that the one-year-ahead return forecasting factor performs poorly when attempting to explain excess holding returns for other horizons can be rationalized in several ways. The one-year expected excess return obtained from the regressions can be thought as a model-free version of the expected excess return obtained from a structural model or an affine model of bond prices. For example, Bansal and Shaliastovich (2013) propose an equilibrium long-run risk model with Epstein-Zin preferences in which bond risk premia depend on expected economic growth and inflation. This model is able to reproduce one-year excess returns remarkably well. Yet, since in the model investors prefer early resolution of uncertainty, risks that affect returns in the short run can be very different from those that affect long-run returns.

Then we can understand the previous results in two alternative ways. The first is that the determinants of short and long run risks are different. Therefore, trying to fit short and long period excess returns using a single one-year return forecasting factor is most likely deemed to fail. Alternatively, since the Bansal and Shaliastovich (2013) model predicts a positive relation between bond risk premia and the preference for early resolution of uncertainty, estimating the model using data for different holding horizons will most likely generate different measures of risk premia. The empirical counterpart of this observation is that the h-year return forecasting factor will be a good predictor of the h-year excess return but not necessarily for other holding horizons. Trying to fit a four-year excess return using the one-year return forecasting factor is akin to wrongly calibrating the model.

A possible exception to the previous argument is when there exists a single factor structure in all excess returns, a property that is not supported by the data, as we show in Section 2.4. But before doing that we evaluate the out-of-sample forecast properties of the empirical models discussed above.

We now analyze the relative merits of the one-year-ahead RFF, the first three PCs, and a constant in terms of their ability to forecast excess returns.[7] We compute the out-of-sample forecasts based on a series of recursive forecasts beginning in 2003M12 and extending through 2014M12 for a forecast horizon of one month and for holding periods of 12, 24, 36, and 48 months.[8] Table 2 reports the results of this exercise. It is evident that the forecasting ability of the two

[7] We include a constant as a benchmark to capture no predictability of excess returns.
[8] Appendix B provides details of the forecasting algorithm.

empirical models (using the RFF and the PCs) are poor, and both are outperformed by a constant for all holding periods. Which of the two models (RFF or PCs) forecast returns better depend on which holding horizon is forecast. Perhaps surprisingly, the principal components outperform the one-year-ahead RFF only for a holding period of one year. Figure 3 displays the out-of-sample forecats of the excess returns for the four holding periods. Overall, these results show the poor forecasting performance of both models.

Table 2: Out-of-sample forecast accuracy

	Root-mean-squared-error		
Horizon	one-year RFF	Principal components	Constant
h = 12	3.305	3.065	1.163
h = 24	1.944	2.018	1.600
h = 36	1.136	1.315	1.007
h = 48	0.658	0.814	0.593

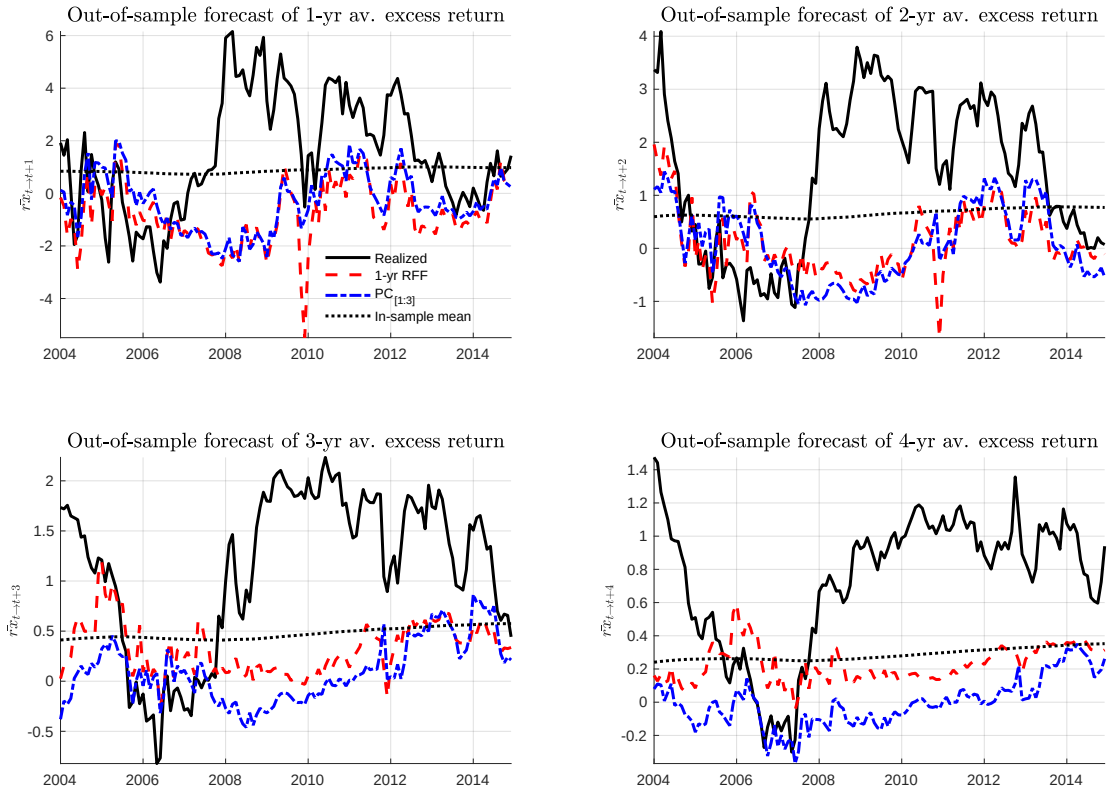
The table shows the root-mean-squared errors of one-month-ahead forecasts of excess bond returns for different holding periods $h = 12, 24, 36, 48$. The column one-year-ahead RFF computes the forecasts using Cochrane and Piazzesi (2005) RFF, the column labeled Principal components does so using the first three principal components of bond yields. In the last column the forecast of the excess return is the unconditional mean.

2.4 A single factor structure in excess returns?

Cochrane and Piazzesi (2008) propose to augment the traditional affine three-factor model by including an additional (return forecasting) factor that drives fluctuations in all excess returns. The principal components would then only serve to match the cross-section of yields but not to influence bond returns. The relevance of this result is that, to the extent that this dimensionality reduction of the data holds for all the possible manifestations of bond risk premia, models of the term structure of interest rates could be greatly simplified: while the traditional first three principal components may accurately explain the cross-section of bond yields, compensation for risk would depend only on a single return forecasting factor.

Although CP actually statistically reject the single factor model in one-year excess return, they claim that, nevertheless, the RFF captures all the economically relevant movements in one-year

Figure 3: Out-of-sample performance of the empirical models



One-month ahead out of sample forecasts of h-period excess returns of the one-year-ahead RFF, the principal components, and a constant.

excess returns.[9] Not surprisingly, we also reject a single factor model in excess returns. But we also find that the single factor model does not capture all the economically relevant variations in excess returns once we consider different holding periods.

Using average excess returns constructed from the Fama-Bliss data, we estimate by maximum

[9]From the unrestricted regressions (2), CP compute an eigenvalue decomposition of one-year expected excess returns $\hat{\beta}_n' f_t$ for $n = 2, 3, 4, 5$ and find that its first principal component, which is almost identical to one-year-ahead RFF, accounts for 99.5 percent of the variance of one-year excess returns.

likelihood an unrestricted system of equations

$$\begin{bmatrix} \overline{rx}_{t\to t+1} \\ \overline{rx}_{t\to t+2} \\ \overline{rx}_{t\to t+3} \\ \overline{rx}_{t\to t+4} \end{bmatrix} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} + \begin{bmatrix} \beta_{1,1} & \beta_{1,2} & \beta_{1,3} & \beta_{1,4} & \beta_{1,5} \\ \beta_{2,1} & \beta_{2,2} & \beta_{2,3} & \beta_{2,4} & \beta_{2,5} \\ \beta_{3,1} & \beta_{3,2} & \beta_{3,3} & \beta_{3,4} & \beta_{3,5} \\ \beta_{4,1} & \beta_{4,2} & \beta_{4,3} & \beta_{4,4} & \beta_{4,5} \end{bmatrix} \begin{bmatrix} y_t^{(1)} \\ f_t^{(1,2)} \\ f_t^{(2,3)} \\ f_t^{(3,4)} \\ f_t^{(4,5)} \end{bmatrix} + \begin{bmatrix} \epsilon_{t+1}^{(1)} \\ \epsilon_{t+2}^{(2)} \\ \epsilon_{t+3}^{(3)} \\ \epsilon_{t+4}^{(4)} \end{bmatrix}, \tag{5}$$

and a restricted, one-factor, system of equations given by

$$\begin{bmatrix} \overline{rx}_{t\to t+1} \\ \overline{rx}_{t\to t+2} \\ \overline{rx}_{t\to t+3} \\ \overline{rx}_{t\to t+4} \end{bmatrix} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \\ \alpha_4 \end{bmatrix} + \begin{bmatrix} \gamma_1 & \gamma_2 & \gamma_3 & \gamma_4 & \gamma_5 \\ b_2\gamma_1 & b_2\gamma_2 & b_2\gamma_3 & b_2\gamma_4 & b_2\gamma_5 \\ b_3\gamma_1 & b_3\gamma_2 & b_3\gamma_3 & b_3\gamma_4 & b_3\gamma_5 \\ b_4\gamma_1 & b_4\gamma_2 & b_4\gamma_3 & b_4\gamma_4 & b_4\gamma_5 \end{bmatrix} \begin{bmatrix} y_t^{(1)} \\ f_t^{(1,2)} \\ f_t^{(2,3)} \\ f_t^{(3,4)} \\ f_t^{(4,5)} \end{bmatrix} + \begin{bmatrix} \epsilon_{t+1}^{(1)} \\ \epsilon_{t+2}^{(2)} \\ \epsilon_{t+3}^{(3)} \\ \epsilon_{t+4}^{(4)} \end{bmatrix}, \tag{6}$$

assuming normal errors. A likelihood ratio test strongly rejects the single factor model with a likelihood ratio statistic of 133.2 (distribututed as chi-square with 12 degrees of freedom).[10]

Moreover, the statistical rejection is economically relevant. Performing an eigenvalue decomposition of the expected returns obtained from equation (5), we find that the first principal component explains less than 90 percent of the variance of expected returns, much smaller than the 99.5 percent reported in CP when considering only one-year excess returns. In other words, a single factor model fails to capture all the economically relevant variations in expected returns.

3 The return forecasting factor and bond yields

In this section we run regressions of bond yields on the RFF and the first three PC. If the RFF is a relevant state variable in a term structure model, it should also be a significant predictor of yields.

[10]Table A.1 in Appendix A displays the estimated coefficients of the system of equations (5) and (6). We also performed tests with individual returns and also reject the single factor model.

We estimate the regressions

$$y_t^{(n)} = \zeta_0^{(n)} + \lambda^{(n)}(\hat{\gamma}' f_t) + \zeta_1^{(n)} PC_{1t} + \zeta_2^{(n)} PC_{2t} + \zeta_3^{(n)} PC_{3t} + \eta_t^{(n)}, \tag{7}$$

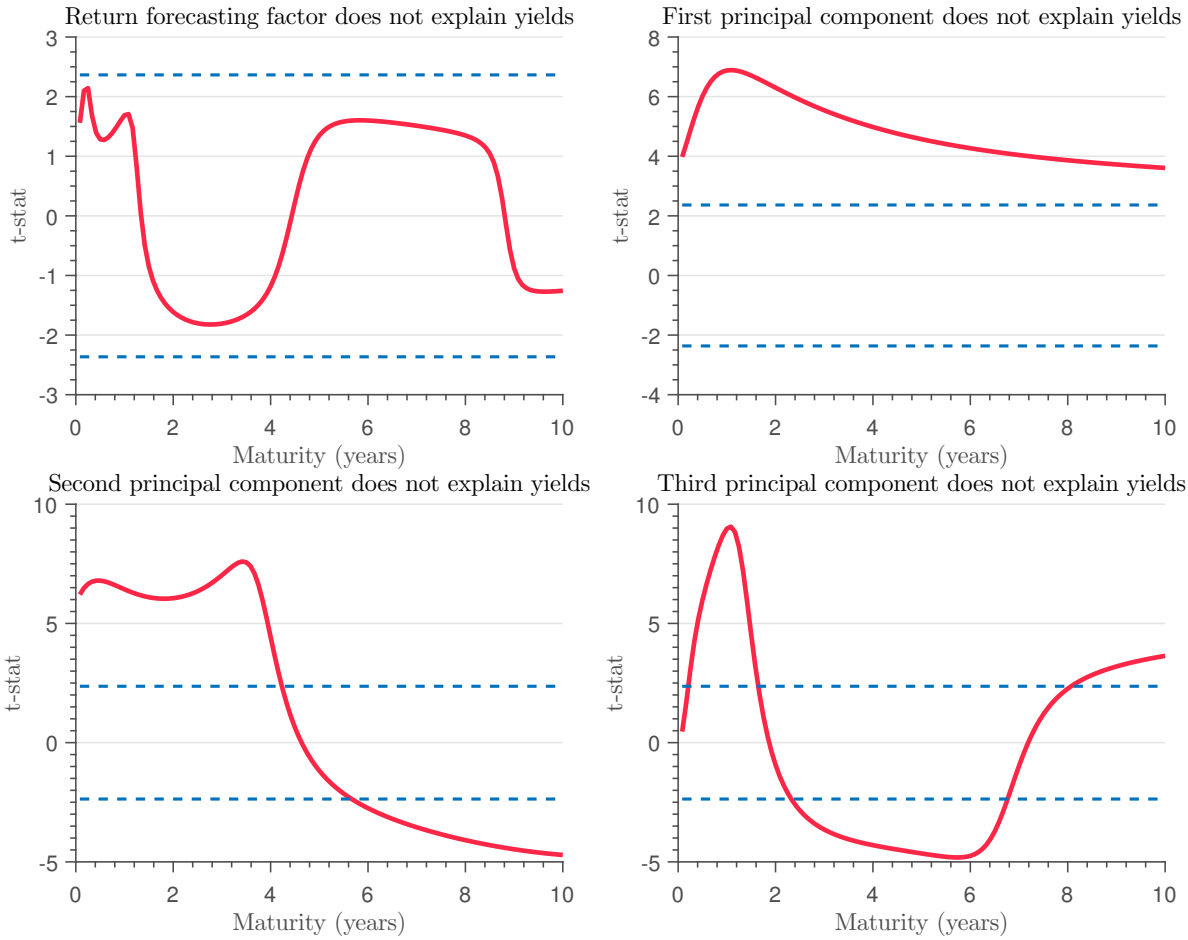
for $n = 1/12, 2/12..., 120/12$ and test the significance of the coefficient $\lambda^{(n)}$. To test hypotheses, it is usual to perform a Newey-West correction due to the serial correlation of the regression residuals. This type of correction, however, may be inappropriate for two reasons. First, the RFF is likely to be correlated with the residuals: equation (7) omits PC that are relevant to explain yields, perhaps with a small but still significant coefficient. Those missing PC are in the residual. This omission is not a problem in a regression of yields on the first three PC since all PC are orthogonal. But if we add the RFF to the regression, which is constructed using the data from which we extract the PC, it is likely that it will be correlated with the missing PC and, hence, with the residuals.[11] And second, the right hand variables of the regression are persistent and conventional HAC corrections work poorly with strongly dependent data (Müller, 2014).

To deal with this potential bias, we use the t-tests developed by Ibragimov and Müller (2010) and recommended by Bauer and Hamilton (2018) in the context of return regressions. We partition the sample into q subsamples and estimate regression (7) in each of them. If the estimates across subsamples are approximately independent, we can test significance of individual coefficients using a t-test with $q-1$ degrees of freedom. In our setting, however, this independence assumption is unlikely to hold because the PC and RFF are both constructed using the entire sample. Hence, we divide the sample into $q = 8$ equally sized subsamples but construct the PC and RFF separately for each subsample. We then estimate regression (7) for each subsample, and perform the t-test of the null hypothesis that the coefficient on the RFF is zero.

The upper left panel of Figure 4 plots the t-statistics of the null hypothesis that $\lambda^{(n)} = 0$ for each bond maturity. The tests do not reject the null hypothesis for any maturity. That is, the RFF is not a significant predictor of bond yields. In contrast, the first three PC are highly significant predictors of bond yields.

[11] In effect, using GSW data we find that the RFF is correlated with the missing PCs.

Figure 4: Testing if factors do not explain yields using Ibrahimov-Müller t-test



Robust t-statistics and $(0.025, 0.975)$ critical values (dashed lines) of the null hypothesis that the coefficients of the RFF and of the first three PC are zero in regression (7).

4 Concluding remarks

Cochrane and Piazzesi (2005) claim that variations in one-year excess bond returns of all maturities can be summarized by a single common factor. We reevaluate this dimensionality reduction in bond risk premia. What differentiates our analysis from most empirical work, is that we consider different holding periods. That is, we analyze different manifestations of bond risk premia besides one-year excess bond returns. We find that a given return forecasting factor only outperforms the first principal components for the horizon for which it has been constructed. We also show that the first three principal components predict bond returns better than the return forecasting factor, particularly for long holding horizons, and that there is no single factor model that captures

all the economically relevant variations in bond returns. Moreover, the coefficient of the return forecasting factor in unrestricted regressions of bond yields is insignificant for all bond maturities after controlling for the principal components. When considering out-of-sample forecasts of excess returns, however, we find that both, the return forecasting factor and the principal components, predict poorly and are dominated by a constant for all holding horizons.

Put differently, our results suggest that the restrictions necessary to summarize bond risk premia with a single factor are not supported by the data.

References

Bansal, Ravi and Ivan Shaliastovich, "A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets," *The Review of Financial Studies*, 2013, *26* (1), 1–33.

Bauer, Michael D. and James D. Hamilton, "Robust Bond Risk Premia," *The Review of Financial Studies*, 2018, *31* (2), 399–448.

Cochrane, John H. and Monika Piazzesi, "Bond Risk Premia," *American Economic Review*, March 2005, *95* (1), 138–160.

_ and _ , "Decomposing the Yield Curve," Technical Report 18, Society for Economic Dynamics 2008.

Fama, Eugene and Robert R. Bliss, "The Information in Long-Maturity Forward Rates," *American Economic Review*, 1987, *77* (4), 680–92.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury yield curve: 1961 to the present," *Journal of Monetary Economics*, November 2007, *54* (8), 2291–2304.

Ibragimov, Rustam and Ulrich K. Müller, "t-statistic Based Correlation and Heterogeneity Robust Inference," *Journal of Business & Economic Statistics*, 2010, *28* (4), 453–468.

Müller, Ulrich K., "HAC Corrections for Strongly Autocorrelated Time Series," *Journal of Business & Economic Statistics*, 2014, *32* (3), 311–322.

Svensson, Lars E.O., "Estimating and Interpreting Forward Interest Rates: Sweden 1992 - 1994," Working Paper 4871, National Bureau of Economic Research September 1994.

A Additional figures and tables

Figure A.1 displays the estimated coefficients $\hat{\alpha}_1^{(h)}$, $\hat{\delta}_j^{(h)}$, $\hat{\delta}_2^{(h)}$, and $\hat{\delta}_3^{(h)}$ of the regressions (3) and (4) using the Fama-Bliss dataset. Each line in the left panel represent the estimated coefficient as a function of the holding period h for a given return forecasting factor $RFF_t^{(m)}$ for $m = 1, 2, 3, 4$. For example, the blue circled-dotted line represent the estimated coefficients $\hat{\alpha}_1^{(h)}$ as a function of the holding period h when using the one-year-ahead return forecasting factor $RFF_t^{(1)}$ as the regressor in equation (3). The other lines display the same coefficients but using the return forecasting factor $RFF_t^{(2)}$, $RFF_t^{(3)}$, and $RFF_t^{(4)}$. Likewise, the right panel of the figure displays the estimated coefficients $\hat{\delta}_1^{(h)}$, $\hat{\delta}_2^{(h)}$, and $\hat{\delta}_3^{(h)}$ of the regression (3) on the principal components, for $h = 1, 2, 3, 4$.

Figure A.1: Estimated coefficients using the Fama-Bliss data set.

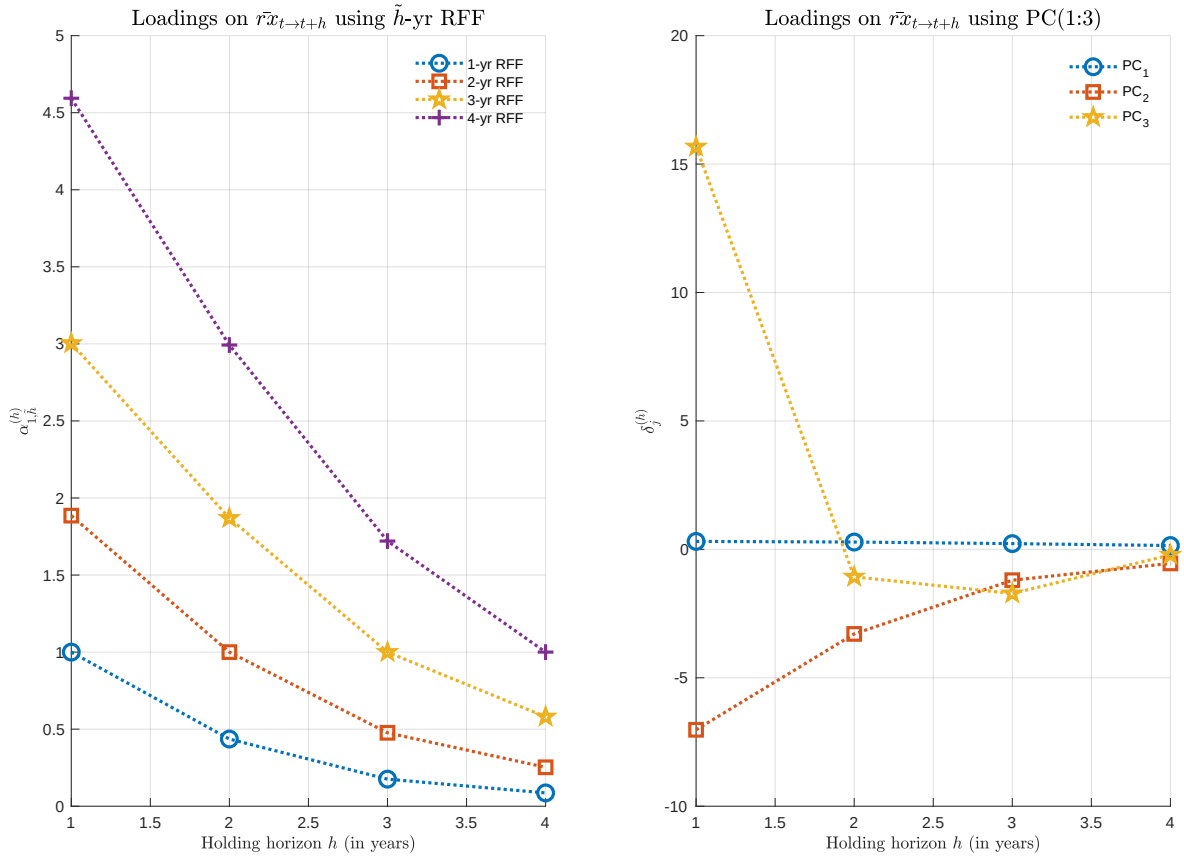


Figure A.2 is identical to Figure A.1 but using the GSW dataset and holding periods from 1

through 9 years ahead.

Figure A.2: Estimated coefficients using the GSW data set.

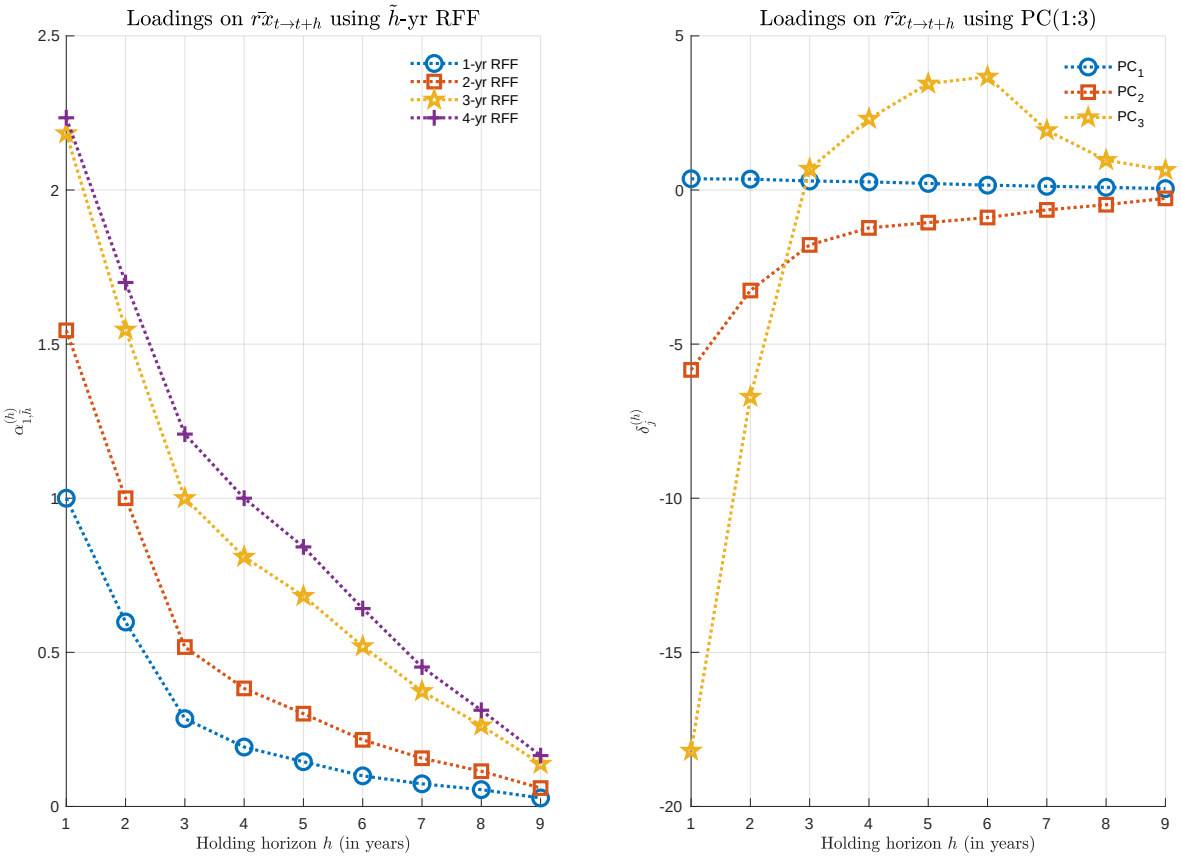


Table A.1 reports the estimates of the constrained and unconstrained system of equations (5) and (6) that we use to test the null hypothesis of a single factor structure on excess returns.

Table A.1: **A single factor structure in excess returns?**

Parameters of the unrestricted excess return regression

$$\alpha \qquad \beta$$

$$\begin{bmatrix} 4.499 \\ (0.747) \\ -3.336 \\ (0.517) \\ -3.046 \\ (0.459) \\ -0.977 \\ (0.183) \end{bmatrix} \qquad \begin{bmatrix} 0.531 & -3.597 & -7.306 & -1.416 & 10.415 \\ (0.079) & (0.815) & (1.078) & (1.286) & (1.361) \\ 0.431 & -0.714 & -2.744 & -2.829 & 1.583 \\ (0.062) & (0.518) & (0.752) & (1.026) & (1.048) \\ 0.397 & -0.564 & -2.164 & -2.747 & 0.897 \\ (0.056) & (0.479) & (0.694) & (0.930) & (0.943) \\ 0.155 & -0.178 & -0.633 & 0.058 & -0.316 \\ (0.025) & (0.232) & (0.386) & (0.418) & (0.417) \end{bmatrix}$$

Log-likelihood = 2208.896 Akaike = 10.337 Scharz=10.564

Parameters of the restricted single-factor excess return regression

$$\alpha \qquad \gamma \qquad b$$

$$\begin{bmatrix} -5.712 \\ (0.747) \\ -2.736 \\ (0.429) \\ -2.382 \\ (0.376) \\ -0.486 \\ (0.134) \end{bmatrix} \qquad \begin{bmatrix} 0.698 \\ (0.082) \\ -2.794 \\ (0.790) \\ -6.336 \\ (0.989) \\ -2.685 \\ (1.310) \\ 6.871 \\ (1.395) \end{bmatrix} \qquad \begin{bmatrix} 0.509 \\ (0.036) \\ 0.447 \\ (0.033) \\ 0.116 \\ (0.017) \end{bmatrix}$$

Log-likelihood = 2275.493 Akaike = 10.590 Scharz=10.703

LR statistic = 133.193 p-value = 0.0000

B Out-of-sample forecasts

This appendix describes the algorithm that we use to compute the recursive forecasts of excess returns. For the forecasting exercise we consider only bonds with integer yearly maturities $n = 1, 2, ..., N$. Let T be the total number of observations and $T_s < T$ the number of observations used for estimation Thus, $T - T_s$ observations are used for forecasting.

Let $\mathbf{y}_t = \left[y_t^{(1)}, y_t^{(2)}, \ldots, y_t^{(N)}\right]$ be a row vector of yields at time t, and Λ_s an $N \times N$ matrix with columns containing the eigenvectors covariance matrix of $\mathbf{y}$. Here, the subindex s is used to denote that the loadings are constructed using in-sample data. In addition, let $\mathbf{PC}_t$ be a $1 \times N$ row vector with the principal components of yields at time t, that is, $\mathbf{PC}_t = \mathbf{y}_t \Lambda_s$.

We consider two models of average excess returns (across bond maturities). The first uses on-year ahead return forecasting factor of Cochrane and Piazzesi (2005); the second, the first three principal components of bond yields,

$$\overline{rx}_{t\to t+h} = \alpha_0^{(h)} + \alpha_1^{(h} \left(\gamma_0 + \gamma_1 y_t^{(1)} + \gamma_2 f_t^{(1,2)} + \gamma_3 f_t^{(2,3)} + \cdots + \gamma_n f_t^{(n-1,n)}\right) + \overline{\epsilon}_{t+h} \tag{B.1}$$

$$\overline{rx}_{t\to t+h} = \delta_0^{(h)} + \delta_1^{(h)} PC_{1t} + \delta_2^{(h)} PC_{2t} + \delta_3^{(h)} PC_{3t} + \overline{\eta}_{t+h}. \tag{B.2}$$

Let $\mathbf{f}_t$ to denote the $1 \times N$ row vector of 1 year forwards rates. Since the forwards constitute N different linear combinations of the yields $\mathbf{y}_t$, the set of forwards and the set of yields span the same information. Moreover, since the N principal component span the same information of the set of yields, we can express model (B.1) in terms of the vector of principal components,

$$\overline{rx}_{t\to t+h} = \alpha_0^{(h)} + \alpha_1^{(h)} \left(\omega_0 + \omega_1 PC_{1t} + \omega_2 PC_{2t} + \cdots + \omega_N PC_{Nt}\right) + \overline{\epsilon}_{t+h} \tag{B.3}$$

We now describe how to produce recursive out-of-sample forecasts for the empirical models (B.2) and (B.3). We focus exclusively on the Fama-Bliss data set and consider four holding horizons, $h = \{12, 24, 36, 48\}$. For the in-sample regressions, use observations from 1964M01 through 2003M12; to perform out-of-sample forecasts, we use data from 2004M1 through 2014M12.

We first fit in-sample a vector autoregression of the principal components with two lags

(selected using the Schwarz information criterion),

$$\mathbf{PC}_t = \Phi_0 + \Phi_1 \mathbf{PC}_{t-1} + \Phi_2 \mathbf{PC}_{t-2} + \xi_t.$$

To produce out-of-sample forecasts of the principal components, we reestimate recursively the vector autoregression adding one observation at a time. Begin by setting $\hat{\Phi}_{0,0} = \Phi_0$, $\hat{\Phi}_{1,0} = \Phi_1$, and $\hat{\Phi}_{2,0} = \Phi_2$, for the regression of the principal components; $\widehat{\omega}_{j,0} = \omega_0$ for $j = 1, 2, ..., N$, for the equation relating the RFF with the principal components; $\widehat{\alpha}_{0,0}^{(h)} = \alpha_0^{(h)}$ and $\widehat{\alpha}_{1,0}^{(h)} = \alpha_1^{(h)}$, for the equation relating the excess returns with the RFF; and $\widehat{\delta}_{0,0}^{(h)} = \delta_0^{(h)}$, $\widehat{\delta}_{1,0}^{(h)} = \delta_1^{(h)}$, and $\widehat{\delta}_{2,0}^{(h)} = \delta_2^{(h)}$, for the equation relating the excess returns with the first three principal components. Then, set $k = 1$ and for holding period $h = \{12, 24, 36, 48\}$ do the following,

1. **Construct the forecast of the excess return $\overline{rx}_{T_s+k \to T_s+k+h}$ at $T_s + k$:**

   - Given data for $t = 1, ..., T_s + k - 1$, construct the forecast of the principal components,

   $$\widehat{\mathbf{PC}}_{T_s+k} = \hat{\Phi}_{0,k-1} + \hat{\Phi}_{1,k-1} \mathbf{PC}_{T_s+k-1} + \hat{\Phi}_{2,k-1} \mathbf{PC}_{T_s+k-2},$$

   - Using the forecast of the principal components, compute the forecast of the RFF,

   $$\widehat{\overline{RFF}}_{T_s+k} = \widehat{\omega}_{0,k-1} + \widehat{\omega}_{1,k-1} \widehat{PC}_{1,T_s+k} + \widehat{\omega}_{2,k-1} \widehat{PC}_{2,T_s+k} + \cdots + \widehat{\omega}_{n,k-1} \widehat{PC}_{n,T_s+k}$$

   - Compute the forecast of the excess returns using the RFF and the first three PCs,

   $$\widehat{\overline{rx}}_{T_s+k \to T_s+k+h} = \widehat{\alpha}_{0,k-1}^{(h)} + \widehat{\alpha}_{1,k-1}^{(h)} \widehat{\overline{RFF}}_{T_s+k}$$
   $$\widehat{\overline{rx}}_{T_s+k \to T_s+k+h} = \widehat{\delta}_{0,k-1}^{(h)} + \widehat{\delta}_{1,k-1}^{(h)} \widehat{PC}_{1,T_s+k} + \widehat{\delta}_{2,k-1}^{(h)} \widehat{PC}_{2,T_s+k} + \widehat{\delta}_{3,k-1}^{(h)} \widehat{PC}_{3,T_s+k}.$$

2. **Update sample and parameters:**

   Add one observation to the sample, so that the sample period is now $t = 1, 2, ...T_s + k$. Using the new sample, reestimate the principal components, the factor loadings Λ_s, k, and update the coefficients $\widehat{\alpha}_{0,k}^{(h)}$, $\widehat{\alpha}_{1,k}^{(h)}$, $\widehat{\delta}_{0,k}^{(h)}$, $\widehat{\delta}_{1,k}^{(h)}$, $\widehat{\delta}_{2,k}^{(h)}$, $\widehat{\delta}_{3,k}^{(h)}$, $\hat{\Phi}_{0,k}$, $\hat{\Phi}_{1,k}$, $\hat{\Phi}_{2,k}$, and $\widehat{\omega}_{j,k}^{(h)}$ for $j = 0, 1, ..., N$

estimating the regression equations

$$
\begin{aligned}
\overline{rx}_{t\to t+12} &= \widehat{\omega}_{0,k}^{(h)} + \widehat{\omega}_{1,k}^{(h)} PC_{1t} + ... + \widehat{\omega}_{N,k}^{(h)} PC_{Nt} + \overline{\epsilon}_{t\to t+12} \\
RFF_t &= \widehat{\omega}_{0,k}^{(h)} + \widehat{\omega}_{1,k}^{(h)} PC_{1t} + ... + \widehat{\omega}_{N,k}^{(h)} PC_{Nt} \\
\overline{rx}_{t\to t+h} &= \widehat{\alpha}_{0,k}^{(h)} + \widehat{\alpha}_{1,k}^{(h)} RFF_t + \overline{\epsilon}_{t+h} \\
\overline{rx}_{t\to t+h} &= \widehat{\delta}_{0,k}^{(h)} + \widehat{\delta}_{1,k}^{(h)} PC_{1t} + \widehat{\delta}_{2,k}^{(h)} PC_{2t} + \widehat{\delta}_{3,k}^{(h)} PC_{3t} + \overline{\eta}_{t+h} \\
\mathbf{PC}_t &= \widehat{\Phi}_{0,k} + \widehat{\Phi}_{1,k}\mathbf{PC}_{t-1} + \widehat{\Phi}_{2,k}\mathbf{PC}_{t-2} + \xi_t.
\end{aligned}
$$

3. **Set** $k \to k+1$ **and return to step 1.** The algorithm stops when $T_s + k + 12 = T$.