

Understanding bond risk premia

ANNA CIESLAK and PAVOL POVALA*

We decompose yields into long-horizon expected inflation and maturity-related cycles to study the predictability of bond excess returns. Cycles capture the risk premium and the business cycle variation of short rate expectations. From cycles, we construct a forecasting factor that explains up to above 50% (30%) of in-sample (out-of-sample) variation of annual bond returns. The factor varies at a frequency higher than the business cycle, and predicts real activity at long horizons. It also aggregates information from different macro-finance predictors of bond returns. Our decomposition reveals why bond returns are predictable by a linear combination of forward rates or the term spread.

First version: March 16, 2010

This version: December 24, 2011

JEL classification: E32, E44, G12

Key words: term premia, bond return forecasting factor, macro factors

*Cieslak is at the Northwestern University, Kellogg School of Management. Povala at the University of Lugano, Switzerland. Cieslak: a-cieslak@kellogg.northwestern.edu, Department of Finance, Kellogg School of Management, Northwestern University, 2001 Sheridan Road Evanston, IL 60208. Povala: pavol.povala@usi.ch, University of Lugano, Institute of Finance, Via Buffi 13a, 6900 Lugano, Switzerland. We thank Torben Andersen, Ravi Bansal, Jules van Binsbergen, Greg Duffee, Jean-Sébastien Fontaine, Ralph Koijen, Arvind Krishnamurthy, Robert McDonald, Kenneth Singleton, Ivan Shaliastovich, Fabio Trojani, Pietro Veronesi, Annette Vissing-Jorgensen, Liuren Wu, and seminar participants at the NBER Asset Pricing Meetings, WFA Meetings, SoFiE Meetings, Stanford GSB, Columbia Business School, Duke Fuqua, Berkeley Haas, Northwestern Kellogg, Toronto Rotman, NY Fed, Fed Board, Blackrock, University of Texas at Austin McCombs, Dartmouth Tuck, Boston University, Wisconsin Business School, Economic Dynamics Working Group at the University of Chicago, University of Lugano, Bank of Canada, University of Geneva, HEC Lausanne SFI, and NCCR Finrisk Review Panel Zurich for comments. Cieslak gratefully acknowledges two grants of the Swiss National Science Foundation (SNSF).

UNDERSTANDING THE VARIATION of expected excess returns of Treasury bonds and their relation with the economy is one of the central questions in economics and finance, and an active area of ongoing research. The literature suggests that factors driving bond risk premia can take at least two forms. First, the yield curve itself seems to contain a component that, being hard to detect in the cross-section of yields, has a strong forecasting power for future bond returns. This important variable reveals itself through a linear combination of forward rates. Second, macroeconomic variables such as real activity, unemployment or inflation appear to contribute to the predictability of bond returns beyond what is explained by factors in the yield curve. Combining these two domains into a joint view of term premia and yields continues to present an open question. This is the question we address with the current paper.

Inflation expectations play a major role in determining the level of nominal yields in a world characterized by fiat money. We start with this notion to analyze the linkages between factors in yields and those determining expected bond returns. A key observation underlying our approach is that interest rates move on at least two different economic frequencies. Specifically, we decompose the yield curve into a persistent component and shorter-lived fluctuations particular to each bond maturity, which we term as cycles. The persistent component captures smooth adjustments in the short rate expectations that may take decades to unfold, and are related both economically and statistically with the shifting long-run inflation expectations. The cycles across different maturities combine the term structure of transitory short rate expectations with the term structure of bond term premia.

We proxy for the persistent inflation factor using the discounted moving average of past core inflation data. This single variable explains 87% of variation in the ten-year yield. Cycles, as we show, represent stationary deviations from the long-term relationship between yields and that slow-moving factor.

Working from the basic notion of a n -period yield $y^{(n)}$ as the sum of the short rate $y^{(1)}$ expectations and the risk premium $rpy^{(n)}$:

$$y_t^{(n)} = \frac{1}{n} E_t \sum_{i=0}^{n-1} y_{t+i}^{(1)} + rpy_t^{(n)}, \quad (1)$$

we exploit the cross-sectional composition of the cycles to construct a powerful predictor of excess bond returns. The underlying economic intuition is as follows: Being derived from a one-period risk-free bond, the cycle with the shortest maturity inherits stationary variation in short rate expectations but not in risk premia. As maturity increases, however, the transitory short rate expectations subside, and the variation in premia becomes stronger. In combination, we are able to trace out a term structure pattern of risk compensation throughout the yield curve.

We separate these effects into three economic frequencies: generational frequency related to persistent inflation expectations, business cycle frequency related to transitory short rate expectations, and the term premium frequency. This decomposition delivers new findings about the degree of bond return predictability, and about the link between term premia, yield curve factors and macroeconomic risks.

We start by revisiting the empirical predictability of bond excess returns. From cycles, we construct a factor that forecasts bond returns across all bond maturities. We label this factor \widehat{cf} . Predictive regressions of one-year excess bond returns on \widehat{cf} give R^2 's above 50% in the period 1971–2009 and at long maturities. The sources of this high predictability lie in the additional information about long-run inflation expectations that we provide to an otherwise standard predictive regression. This additional information alleviates the error-in-variables problem that yield-only predictors entail. As one example, the level factor (first principal component of yields) combines short rate expectations and term premia into one variable, and being dominated by the expectations component, shows little predictive ability for bonds returns. Our decomposition frees up an otherwise ignored piece of term premium information that the level embeds.

As a consequence of this view, we are able to explore the mechanism that makes forward rates a successful predictor of bond excess returns. We show that the forward rate factor proposed by Cochrane and Piazzesi (2005) can be interpreted as a specific linear combi-

nation of interest rate cycles, whose predictive power is constrained by the persistence of yields. If the information set of market participants contains only past history of forward rates, then the forward rate factor is the best measure of the term premia that both econometricians and investors could obtain. However, because it is plausible that investors change their long-horizon expectations about inflation, \widehat{cf} provides a viable benchmark for the degree of term premia variation.

One is ultimately interested in understanding the link between the term premia and macro-finance conditions. Taking \widehat{cf} as a benchmark, we can assess the marginal predictive content of macroeconomic fundamentals for bond returns. The presence of \widehat{cf} in the predictive regression renders most macro-finance variables insignificant, suggesting that our factor successfully aggregates different economic risks into a single quantity. This evidence points to a heterogeneity of economic factors driving term premia. Moreover, the half-life of \widehat{cf} of about ten months suggests that term premia vary at a frequency higher than the business cycle. Many of the large moves in bond excess returns and in \widehat{cf} appear in otherwise normal times, giving rise to an interest rate-specific cycle.

Our results have also implications for forecasting economic activity. An important literature documents that the positive slope of the term structure is associated with an increased economic activity in the future (Harvey, 1989; Estrella and Hardouvelis, 1991). Trying to understand the economic sources of this finding, we further decompose the slope into short rate expectations and term premia. Lower short rate expectations forecast stronger growth up to two years ahead, and account almost entirely for the predictive ability of the slope for the real activity. The effect of term premia on growth is different: While at horizons below one year, a higher term premium predicts lower growth, this effect reverts completely at longer forecast horizons so that a higher term premium today suggests stronger economic growth two and a half years ahead.

One may wonder whether our findings present a robust feature of Treasury bonds. We address this concern in several ways. First, we show that the bond return predictability is not an artifact of the discounted moving average by which we measure the persistent component of inflation expectations. Using long-horizon inflation surveys, we report very

similar predictability results. The expectations' horizon matters: Short-horizon measures of inflation expectations fail to deliver additional information about the sources of persistence in yields and, consequently, underestimate the predictable variation in the term premia. Second, we verify that the predictive power of the cycles for bond returns is not affected by the choice of the data set or by the splining of the zero curve. We support our results using actual bond portfolio returns, analyze the predictability at different holding periods, and for non-overlapping returns. Third, we show that our forecasting factor provides stable and positive out-of-sample performance. In sum, these robustness checks confirm that our results describe a reliable characteristic of the bond market.

Related literature

Our paper is related to the recent literature on predictability of bond returns. Cochrane and Piazzesi (2005, CP) show that a single linear combination of forward rates, the CP factor, predicts bond excess returns across a range of maturities, but has a low correlation with the principal components of yields. Subsequently, a large literature studies the link between term premia, the CP factor and macro-finance variables that are outside of the yield curve (Ludvigson and Ng, 2009; Cooper and Priestley, 2009; Huang and Shi, 2010; Fontaine and Garcia, 2011). We focus on explaining the variation of bond excess returns using a predictor, \widehat{cf} , that we form from the basic zero yield curve and a variable that captures long-run inflation expectations. We show that the \widehat{cf} factor encompasses many commonly used predictors of bond returns, and we reconcile this result with the predictive power of forward rates.

Our work also relates to Duffee (2011) and Joslin, Priebsch, and Singleton (2010, JPS) who study hidden factors in bond risk premia within affine term structure models. The theoretical implication of those models (absent hidden factors) is that variables other than yields add no content from a forecasting perspective. In an important modification of the standard setting, the authors allow for the possibility that yields have non-Markov dynamics under the historical measure. Our results are consistent with this notion. Specifically, we contribute to the discussion by showing that long-horizon inflation expectations are an

important element of investors' information set beyond yields themselves. We document that yield-only predictors of bond excess returns suffer from an error-in-variables problem. Interestingly, we find that long-horizon inflation expectations have a major role in explaining the level factor of yields, and thus the additional information that inflation expectations provide impacts significantly both the cross-section and the time-series dynamics of yields. A separate strand of research deals with the persistence of yields. Kozicki and Tinsley (1998, 2001a,b) point to sluggish changes in the market perceptions of the long-run monetary policy target for inflation. In a related fashion, Fama (2006) shows that the predictability of the short rate for horizons beyond one year comes from its reversion toward a time-varying rather than constant long-term mean, which he proxies with a moving average of past one-year yield. Following similar intuition, several authors adopt slow-moving means of variables to generate persistent long-term yields (Rudebusch and Wu, 2008; Dewachter and Lyrio, 2006; Orphanides and Wei, 2010; Dewachter and Iania, 2010; Piazzesi and Schneider, 2011; Koijen, Van Hemert, and Van Nieuwerburgh, 2009). To the best of our knowledge, our study is the first to establish the link between long-horizon inflation expectations, persistent and transitory short rate expectations, and the predictability of bond excess returns.

I. Data sources

We use end-of-month yield data obtained from the H.15 statistical release of the Fed. Since we want to cover a broad spectrum of maturities over a possibly long sample, we consider constant maturity Treasury (CMT) yields. The available maturities comprise six months and one, two, three, five, seven, ten and 20 years in the post Bretton Woods period from November 1971 through December 2009. We bootstrap the zero coupon curve by treating the CMTs as par yields.¹

The inflation data is from the FRED database. To construct the persistent inflation variable, we use core CPI which excludes volatile food and energy prices. There are two

¹In the online appendix, we provide a comparison of our zero curve and realized excess bond returns with other data sets (Fama-Bliss and Gürkaynak, Sack, and Wright (2006)). The comparison confirms a very close match between the different data sets at maturities that overlap.

reasons for using core CPI rather than all-items CPI . First, core CPI has been at the center of attention of the monetary policy makers.² Second, it is more suitable to compute the long-run expectations of inflation by excluding volatile components of prices. Nevertheless, we verify that our results remain robust to both core and all-items CPI. An important advantage of CPI over other measures of inflation (e.g. PCE inflation) is that CPI is not subject to revisions (Croushore and Stark, 1999; Croushore, 2008), and thus can be treated as real time data.

Inflation data for a given month are released in the middle of the next month. To account for the publication lag when constructing the persistent inflation component we use the CPI data that are available as of month end. For example, the estimate of the persistent component for April 2000 uses inflation data until March 2000 only. We also check that our results are not sensitive to whether or not we allow for the publication lag.

For robustness in Section III.C, we use inflation expectations compiled by different surveys. Specifically, we consider ten-year CPI inflation forecasts from the Livingston survey and from the Blue Chip Financial Indicators survey, one-year forecast of the GDP deflator from the Survey of Professional Forecasters, and the four-quarter ahead inflation forecast from the Blue Chip Financial Forecasts.³ The online appendix provides additional details about the data we use in the subsequent analysis.

II. Components in the yield curve

II.A. Basic intuition

The yield of an n -period bond can be expressed as the average future short rate $y_t^{(1)}$ expected over the life of the bond and the term premium, $rpy_t^{(n)}$ (assuming log normality). Reiterating equation (1):

²Fed officials rely on core inflation to gauge price trends. As one recent example, this view has been expressed by the Fed chairman Ben Bernanke in his semiannual testimony before the Senate Banking Committee on March 2, 2011: “Inflation can vary considerably in the short run. [...] Our objective is to hit low and stable inflation in the medium term.” Core inflation is a good predictor of the overall inflation over the next several years, which is the horizon of focus for the monetary policy makers.

³Philadelphia Fed web page provides a useful collection of different inflation surveys.

$$y_t^{(n)} = \frac{1}{n} E_t \sum_{i=0}^{n-1} y_{t+i}^{(1)} + rpy_t^{(n)}. \quad (2)$$

Since inflation is a key determinant of real cash flows from nominal government bonds, inflation expectations have a first order effect on nominal bond prices. To convey our intuition, we start from the simple Fisher equation, in which the one-period nominal interest rate $y_t^{(1)}$ is:

$$y_t^{(1)} = r_t + E_t \pi_{t+1}, \quad (3)$$

where r_t is a one-period real rate and $E_t \pi_{t+1}$ is the expected inflation over one period. Let the real rate r_t follow a stationary process, and let inflation be composed of a persistent (τ_t) and a transitory (π_t^c) component (e.g., Stock and Watson, 2007):

$$\pi_t = \tau_t + \pi_t^c \quad (4)$$

$$\tau_t = \tau_{t-1} + \varepsilon_t^\tau, \quad (5)$$

where ε_t^τ is uncorrelated with π_t^c , and π_t^c follows an AR(1). One can think of τ_t in equation (4) as a time varying long-run inflation mean: $\lim_{s \rightarrow \infty} E_t(\pi_{t+s}) = \tau_t$ (Kozicki and Tinsley, 2001a, 2006). Then, substituting for the short rate $y_t^{(1)}$ in equation (2), the yield with n -years to maturity $y_t^{(n)}$ can be written as:

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)} \tau_t + b_r^{(n)} r_t + b_\pi^{(n)} \pi_t^c + rpy_t^{(n)}, \quad (6)$$

where $b_0^{(n)}$ is a maturity dependent constant, $b_\tau^{(n)} = 1$ and $b_r^{(n)}$, $b_\pi^{(n)}$ decay with maturity n .⁴ Equation (6) captures the long-run Fisher effect, i.e. permanent changes in inflation are reflected approximately one-to-one in nominal yields.⁵

⁴Assuming that π_t^c and r_t follow an AR(1), these coefficients are of the form $b_x^{(n)} = \frac{1}{n} (\phi_x^n - 1) (\phi_x - 1)^{-1}$, $x = \{\pi^c, r\}$, where ϕ_x is the respective AR(1) coefficient.

⁵Mishkin (1992) provides a detailed discussion of the long-run Fisher and empirical evidence supporting its validity. If the Fisher effect does not hold at least in the long run, one has to resort to non-rational explanations of interest rates such as money illusion (e.g. Summers, 1982) or to believe in the long term impact of demand and supply effects on interest rates induced by inflation (Mundell, 1963). In the presence of taxes, the Fisher effect is higher than one implying that yields change by more than one-to-one with expected inflation due to taxation (Crowder and Wohar, 1999).

We will refer to the sum of the transitory components of short rate expectations, π_t^c and r_t , and the risk premium $ropy_t^{(n)}$ in equation (6) as “the cycle” defined as:

$$\tilde{c}_t^{(n)} = b_r^{(n)} r_t + b_{\pi^c}^{(n)} \pi_t^c + rpy_t^{(n)}. \quad (7)$$

The composition of $\tilde{c}_t^{(n)}$ changes with the maturity of the bond. For $n = 1$, $\tilde{c}_t^{(1)}$ captures variation in the real rate r_t and the cyclical part of inflation expectations π_t^c , but not in premia because $ropy_t^{(1)}$ is zero in nominal terms. We term $\tilde{c}_t^{(1)}$ as the transitory component of short rate expectations. As the maturity n increases, the transitory part of short rate expectations decays due to the mean reversion. Thus, cycles extracted from the long end of the yield curve should provide the most valuable information about expected excess returns. This intuition is strongly supported by the data and underlies the predictability of bond returns that we document below.

II.B. Measuring the persistent component of inflation expectations

Borrowing from the literature on adaptive learning (e.g., Evans and Honkapohja, 2009), we construct a proxy for the long-run mean of inflation τ_t in equation (6) as a discounted moving average of past core CPI inflation:

$$\tau_t^{CPI} = \frac{\sum_{i=0}^{t-1} v^i CPI_{t-i}}{\sum_{i=0}^{t-1} v^i}, \quad (8)$$

where $(1 - v)$ is the constant gain learning parameter. The discounted moving average captures the idea that economic agents update their long-run inflation expectations by a fraction of the deviation of current inflation from the previous long-run mean.⁶ Using various inflation surveys, we estimate the gain parameter at $v = 0.9868$ (standard error 0.0025), and truncate the sums in equation (8) at $N = 120$ months. The online appendix

⁶See e.g. Carceles-Poveda and Giannitsarou (2007) for a derivation of learning recursions.

provides details of the estimation of v .⁷ With those parameters, an observation from ten years ago still receives a weight of approximately 0.2.

As proxied by τ_t^{CPI} , inflation expectations explain a major part of yield variation with an R^2 of 86% on average across maturities from one to 20 years. Figure 1, panel *a*, superimposes a one-year and ten-year yield with τ_t^{CPI} showing that τ_t^{CPI} coincides with the low-frequency variation in interest rates. For comparison, in panel *b*, we plot the median inflation forecast from the Livingston survey one year ahead, collected in June and December each year. The limited forecast horizon drives shorter-lived variation in the survey-based measure, especially in the volatile periods; still, τ_t^{CPI} and surveys display a similar behavior over time.

[Figure 1 here.]

Our approach to constructing τ_t^{CPI} is deliberately simple, but it reflects the important feature that economic agents adapt their long-run inflation expectations with a lag (Barsky and DeLong, 1991; DeLong, 1997). In Section III.C, we confirm our results with long-horizon inflation surveys. However, due to the limitations on the sample length and the frequency of observations that long-horizon inflation surveys impose, in subsequent sections we rely on τ_t^{CPI} .

II.C. Cycles as deviations from the long-run relationship between yields and expectations

Cointegration between yields and long-run inflation expectations, as predicted by the long-run Fisher effect, provides an econometric argument for the intuition that cycles should

⁷A number of papers argue for a similar gain parameter for inflation: Kozicki and Tinsley (2001a) use $v = 0.985$ for monthly data, Piazzesi and Schneider (2011) and Orphanides and Wei (2010) use $v = 0.95$ and $v = 0.98$ for quarterly data, respectively. Kozicki and Tinsley (2005) estimate $v = 0.96$ and find that discounting past data at about 4% per quarter gives inflation forecasts that closely track the long-run inflation expectations from the Survey of Professional Forecasters. The truncation parameter $N = 120$ months is motivated by the recent research of Malmendier and Nagel (2009) who argue that individuals form their inflation expectations using an adaptive rule and learn from the data experienced over their lifetimes rather than from all the available history. We stress that the parameters v and N are not a knife edge choice that would determine our subsequent findings. A sensitivity analysis shows that varying N between 100 and 150 months and v between 0.975 and 0.995 leads to negligible quantitative differences in results and does not change our interpretation. These results are available in the online appendix.

predict bond excess returns. In our sample, yields and τ_t^{CPI} both feature nonstationary dynamics, as indicated by unit root tests.⁸ Following the standard approach (Engle and Granger, 1987), we regress yields on a contemporaneous value of τ_t^{CPI} :

$$y_t^{(n)} = b_0^{(n)} + b_\tau^{(n)} \tau_t^{CPI} + \epsilon_t^{(n)}, \quad (9)$$

and test for stationarity of the fitted residual. We denote the fitted residual of (9) as $c_t^{(n)}$ for individual yields, and \bar{c}_t for the average yield across maturities, i.e. $\bar{y}_t = \frac{1}{20} \sum_{i=1}^{20} y_t^{(i)}$. As an example, we provide the point estimates of equation (9) for the average yield \bar{y}_t :

$$\bar{c}_t = \bar{y}_t - \underbrace{\hat{b}_0}_{0.02 [4.7]} - \underbrace{\hat{b}_\tau}_{1.24 [14.2]} \tau_t^{CPI}, \quad R^2 = 0.86. \quad (10)$$

The numbers in brackets are the Newey-West corrected t -statistics. We report detailed results of stationarity tests in Appendix A, and here just state the main conclusions. We consistently reject the null hypothesis that $c_t^{(n)}$ contains a unit root for maturities from one to 20 years at the 1% level. Thus, the data strongly supports the cointegrating relation.

Note that $c_t^{(n)}$ gives an empirical content to the notion of cycles we have introduced in equation (7). By cointegration, cycles represent stationary deviations from the long-run relationship between yields and the slow moving component of inflation expectations. Therefore, by the Granger representation theorem, they should forecast either Δy_t or $\Delta \tau_t$, or both. To verify this prediction, we estimate the error correction representation for yield changes. We allow one lag of variable changes to account for short-run deviations from (9):

$$\Delta y_t^{(n)} = a_c c_{t-\Delta t}^{(n)} + a_y \Delta y_{t-\Delta t}^{(n)} + a_\tau \Delta \tau_{t-\Delta t} + a_0 + \varepsilon_t, \quad \Delta t = 1 \text{ month}. \quad (11)$$

⁸Even if the assumption of nonstationary interest rates may raise objections, the results of Campbell and Perron (1991) suggest that a near-integrated stationary variables are, in a finite sample, better modeled as containing a unit root, despite having an asymptotically stationary distribution. Indeed, many studies fail to reject the null hypothesis of a unit root in the US data (e.g. Jardet, Monfort, and Pegoraro, 2010; Joslin, Priebsch, and Singleton, 2010).

We focus on $\Delta y_t^{(n)}$ because we are interested in transitory adjustments of asset prices. Indeed, the error correction term, $c_{t-\Delta t}^{(n)}$, turns out significant precisely for this part of the system.

Table II presents the estimates of equation (11) for monthly data. The essence of the results is that cycles are highly significant predictors of monthly yield changes. The negative sign of a_c coefficients for all maturities suggests that a higher value of the cycle today predicts lower yields and thus higher excess bond returns in the future. This conforms with the intuition of equation (7) that cycles and term premia should be positively related. In the next section, we build on this observation to explore the predictability of excess bond returns by the cycles.

[Table II here.]

III. The predictability of bond excess returns

III.A. First look at predictive regressions

Following the literature, we focus on one-year holding period bond excess returns, and defer the analysis of other holding periods to Section VII.A. To fix notation, a one-year holding period excess log return on a bond with n years to maturity is defined as: $rx_{t+1}^{(n)} = p_{t+1}^{(n-1)} - p_t^{(n)} - y_t^{(1)}$, where $p_t^{(n)}$ is the log price of a zero bond, $p_t^{(n)} = -ny_t^{(n)}$, and $y_t^{(1)}$ is the one-year continuously compounded rate. The one-year forward rate locked in for the time between $t+n-1$ and $t+n$ is given by: $f_t^{(n)} = p_t^{(n-1)} - p_t^{(n)}$. In Table III, we report the descriptive statistics for bond excess returns.

[Table III here.]

We obtain cycles as fitted residuals from the regressions of yields on the persistent inflation factor in equation (9), i.e.:

$$c_t^{(n)}(\tau_t^{CPI}) = y_t^{(n)} - \widehat{b}_0^{(n)} - \widehat{b}_\tau^{(n)}\tau_t^{CPI}, \quad (12)$$

and estimate the predictive regression:

$$rx_{t+1}^{(n)} = \delta_0 + \sum_i \delta_i c_t^{(i)} (\tau_t^{CPI}) + \varepsilon_{t+1}^{(n)}, \quad (13)$$

where $i = \{1, 2, 5, 7, 10, 20\}$ years.⁹ In equation (12) and (13), we make explicit the dependence of $c_t^{(i)}$ on the cointegrating variable τ_t^{CPI} ; subsequently, we simply use $c_t^{(i)}$ as shorthand notation for $c_t^{(i)}(\tau_t^{CPI})$. To provide a benchmark for our results, we also estimate an analogous equation using forward rates instead of cycles:

$$rx_{t+1}^{(n)} = d_0 + \sum_i d_i f_t^{(i)} + \varepsilon_{t+1}^{(n)}. \quad (14)$$

For excess returns rx , we single out points along the yield curve with maturities of two, five, seven, ten, 15 and 20 years.¹⁰

Table IV summarizes the estimation results. We report the adjusted R^2 values and the Wald test statistics for the null hypothesis that all coefficients in (13) are jointly zero. The individual coefficient loadings are not reported, as by themselves they do not yield an economic interpretation (Section IV.A explains why). The cycles c_t 's forecast a remarkable portion of variation in excess bond returns. R^2 's increase from 42% up to 57% across maturities. On average, these numbers more than double the predictability achieved with forward rates. All coefficient on c_t 's are statistically different from zero. We confirm the statistical robustness of these conclusions later in Section VII.B.

[Table IV here.]

One may be worried about the small-sample reliability of our findings. For this reason, Table IV provides small sample (SS) bootstrapped confidence bounds on R^2 's. Even though $c_t^{(n)}$ is estimated with a high precision, the bootstrap procedure automatically accounts for its uncertainty (see Appendix E for details). Importantly, the lower 5% confidence bound

⁹This choice of maturities summarizes the relevant information in c_t 's.

¹⁰We note that in terms of its predictive power, regression (13) is equivalent to using a set of yields and τ_t^{CPI} as the explanatory variables. We follow the representation in terms of cycles because it offers a convenient interpretation of factors underlying the yield curve which we exploit below.

for the R^2 's obtained with the cycles consistently exceeds the large-sample R^2 obtained with forward rates. A similar observation holds true for the reported values of the Wald test.

III.B. The single return forecasting factor

The intuition following from equation (7) suggests that we can construct a single forecasting factor that predicts bond returns with different maturities. We proceed in two steps:

Step 1. Obtain the cycles $c_t^{(n)}$ as residuals from regressing yields of different maturities on τ_t^{CPI} as in equation (9) and (12).

Step 2. Project the average excess return across maturities $\overline{r\bar{x}}$ onto the transitory short rate expectations factor $c_t^{(1)}$ and the average cycle \bar{c}_t . The fitted linear combination of $c_t^{(1)}$ and \bar{c}_t gives the return forecasting factor which we term \widehat{cf}_t :

$$\overline{r\bar{x}}_{t+1} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \varepsilon_{t+1}, \quad \text{where} \quad \bar{c}_t = \frac{1}{m-1} \sum_{i=2}^m c_t^{(i)} \quad (15)$$

$$\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t. \quad (16)$$

Step 1 removes from the yield curve the slow moving component of inflation expectations; Step 2 removes the business cycle variation of short rate expectations. Empirically, \widehat{cf}_t has a faster mean reversion than the factors related to the short rate expectations, τ_t^{CPI} and $c_t^{(1)}$: Its half-life is slightly below ten months, giving rise to a distinctive term premium frequency. Figure 2 displays the evolution of the forecasting factor over time.

[Figure 2 here.]

In Table V, panel A, we report the estimates of equation (15). Low standard errors on the estimated coefficients indicate that we are able to identify a robust feature of the data. The negative sign of γ_1 is consistent with the decomposition of cycles into the premium and expectations components. We find that a 100 basis points change in the transitory short

rate expectations factor $c_t^{(1)}$ generates a 42 basis points reaction in longer-maturity cycles, on average.

[Table V here.]

With \widehat{cf}_t , we forecast individual excess bond returns:

$$rx_{t+1}^{(i)} = \beta_0^{(i)} + \beta_1^{(i)}\widehat{cf}_t + \varepsilon_{t+1}^{(i)}. \quad (17)$$

Panel B of Table V reports the predictability of individual bond returns achieved with the single factor. On average, during the 1971–2009 period, \widehat{cf}_t explains around 54% of variation in excess returns. The results are not significantly worse than those of the unrestricted regression in equation (13): That comparison is reflected in the row “ ΔR^2 .”

III.C. Alternative measures of long-horizon inflation expectations

The results we have reported so far, rely on discounted moving average of past core CPI inflation, τ_t^{CPI} , as a simple proxy for the long-run inflation expectations. Here, we discuss alternative measures of inflation expectations based on surveys and realized CPI. Specifically, we consider:

Long-horizon inflation expectations from surveys: combined median ten-year CPI inflation expectations from the Livingston survey and from the Blue Chip Economic Indicators survey (BCEI), 1979:12–2009:12, as provided by Philadelphia Fed;

Short-horizon inflation expectations from surveys: (i) median forecast of GDP deflator one-year ahead from the Survey of Professional Forecasters (SPF), 1971:11–2009:12; (ii) median four-quarters ahead inflation forecast from the Blue Chip Financial Forecasts (BCFF), 1984:07–2009:12;

Realized inflation measures: (i) current core CPI inflation; (ii) 12-month moving average of past core CPI inflation.

The sample lengths are dictated by the availability of the survey data.¹¹ We reconstruct the return forecasting factor by replacing τ_t^{CPI} with each of the inflation expectations proxies listed above, and re-estimate predictive regressions. Panels A through C in Table VI compare the regression R^2 's.

There are three main observations. First, using long-horizon surveys (panel A), we obtain R^2 's that are very similar to those based on τ_t^{CPI} , and significantly higher than those implied by the CP factor. Second, survey-based measures (panels A and B) outperform, while the realized inflation proxies (panel C) are approximately on par with the CP factor. Short-horizon surveys deliver a weaker predictability relative to long-horizon inflation forecasts, suggesting that the forecast horizon plays an important role in capturing the persistent component in yields. One-year surveys still contain a substantial portion of transitory deviations of inflation from its long-run mean. Third, and analogously, realized inflation measures (Panel C) do not capture well long-run inflation expectations and their R^2 's are considerably lower than the R^2 's obtained using surveys. In summary, survey-based measures of long-horizon inflation expectations capture the same economic effect as the adaptive learning proxy, τ_t^{CPI} .

Alternatively, one could measure the persistent component in yields by a moving average of past interest rates themselves. In an extreme case when the real rate is constant, the short nominal rate should perfectly reveal inflation expectations. We find, however, that bond return predictability weakens significantly when using the moving average of past yields instead of τ_t^{CPI} (see the online appendix for detailed results). The predictive R^2 's are close to those obtained with the CP factor, indicating that the moving average of past yields adds little content to the regression. We explain this finding with the behavior of real rates, which are not constant in real world. When using a moving average of past interest rates, one faces a tradeoff between smoothing the variation in the real rate and obtaining a timely proxy for the persistent inflation component. In unreported results, we establish that for plausible parameters of the real rate dynamics, a moving average of past

¹¹The SPF survey is conducted quarterly, Livingston and BCEI surveys are semi-annual. To obtain monthly data, we linearly interpolate between quarterly/semi-annual observations (e.g. Kiley, 2008).

inflation provides a better gauge of the persistent component than the moving average of past yields.

IV. Link to other bond return predictors

IV.A. The Cochrane-Piazzesi factor

Cochrane and Piazzesi (2005) show that a single factor, which they make observable through a linear combination of forward rates, captures almost complete variation in expected excess returns on bonds with different maturities. Given its prominence in the empirical term structure research, it is useful to connect our findings with the CP factor. To this end, let us run the predictive regression of an average holding period excess return $\overline{r\bar{x}}_{t+1}$ on a set of m forward rates with maturities 1 to m years at time t :

$$\overline{r\bar{x}}_{t+1} = \gamma_0 + \sum_{i=1}^m \gamma_i f_t^{(i)} + \bar{\varepsilon}_{t+1} \quad (18)$$

$$= \gamma_0 + \gamma' \mathbf{f}_t + \bar{\varepsilon}_{t+1}. \quad (19)$$

Cochrane and Piazzesi (2005) construct the return forecasting factor as $\gamma' \mathbf{f}_t$. From decomposition (9) and the definition of the forward rate $f_t^{(i)}$, we have:

$$\overline{r\bar{x}}_{t+1} = \tilde{\gamma}_0 + \tau_t^{CPI} \left(\sum_{i=1}^m \tilde{\gamma}_i \right) + \sum_{i=1}^m \tilde{\gamma}_i c_t^{(i)} + \bar{\varepsilon}_{t+1}, \quad (20)$$

$$= \tilde{\gamma}_0 + \tilde{\gamma}' \mathbf{1} \tau_t^{CPI} + \tilde{\gamma}' \mathbf{c}_t + \bar{\varepsilon}_{t+1}, \quad (21)$$

where $\mathbf{1}$ is an m -dimensional vector of ones, $\tilde{\gamma}, \tilde{\gamma}$ are respective vectors of loadings that are explicit functions of $b_\tau^{(i)}$ and γ_i coefficients, and $\mathbf{c}_t = \left(c_t^{(1)}, \dots, c_t^{(m)} \right)'$.

A common empirical property of regression (19) is that the γ_i coefficients have a neutralizing effect on each other: Independent of the data set used or the particular shape of the loadings, γ_i 's (and so $\tilde{\gamma}_i$'s) sum to a number close to zero. If the long-run Fisher effect holds, this result is intuitive since only the cyclical part of yield variation matters for risk premia and thus for forecasting $\overline{r\bar{x}}$. Equation (20) tells us that the OLS tries to remove the common

τ_t^{CPI} from forward rates, while preserving a linear combination of the cycles. For this reason, forecasting returns with forward rates embeds an implicit restriction on the slope coefficients: γ_i 's are constrained by the dual role of removing the persistent component and minimizing the prediction error of excess returns using the cycles.

This interpretation can be tested by allowing the excess returns in (20) to load with separate coefficients on $\bar{\gamma}'\mathbf{1}\tau_t$ and $\tilde{\gamma}'\mathbf{c}_t$. Effectively, we can split the forward factor into two components, and estimate: $\bar{r}\bar{x}_{t+1} = a_0 + a_1(\bar{\gamma}'\mathbf{1}\tau_t) + a_2(\tilde{\gamma}'\mathbf{c}_t) + \varepsilon_{t+1}$. This exercise gives an \bar{R}^2 of 30%, similar to 26% obtained with $\gamma'\mathbf{f}_t$. As expected, the predictability comes from the strongly significant $\tilde{\gamma}'\mathbf{c}_t$ term (Newey-West t-statistic of 5.9). The persistent component $\gamma'\mathbf{1}\tau_t$ is not significantly different from zero, and thus has a minimal contribution to the overall dynamics of the CP factor.

These results suggest an interpretation of the CP factor as a constrained linear combination of the cycles. By the presence of the persistent component in forward rates, the factor is restricted in its ability to extract information about premia. Using just forward rates, and with no information about long-horizon inflation expectations, τ_t^{CPI} , this is the best predictability one can achieve.

IV.B. Level and slope

It is well-documented that the predictability of bond returns by the level factor (PC1) is close to zero. Instead, the predictability by the slope (PC2) or the term-spread gives R^2 's of about 15% for long maturities. In this section, we interpret these results in the context of our decomposition of the yield curve.

Given the representation of yields in equation (1), the level factor should reflect premia unless they are precisely offset by the short rate expectations. To see the connection between the level factor lvl_t and the return forecasting factor \widehat{cf}_t , we note that the level, lvl_t , is proportional to the average yield across maturities.¹² Therefore, by projecting lvl_t onto

¹²The level factor is obtained as $lvl_t = q\mathbf{1}'\mathbf{y}_t$ where $\mathbf{y}_t = (y_t^{(1)}, \dots, y_t^{(m)})'$, $\mathbf{1}$ is a $m \times 1$ vector of ones, q is a constant. $q\mathbf{1}$ is the eigenvector corresponding with the largest eigenvalue in the singular-value decomposition of the yield covariance matrix.

τ_t^{CPI} , we obtain the average cycle as the cointegrating residual. We denote this residual by c_t^{lv} , noting that c_t^{lv} captures essentially the same source of variation in the yield curve as the average cycle \bar{c}_t , and their correlation exceeds 99%. Therefore, we can form the forecasting factor in equation (16) using c_t^{lv} instead of \bar{c}_t . The correlation between the two, $\text{corr}(\widehat{cf}_t^{lv}, \widehat{cf}_t)$, is 99.9% so the return predictability remains unaffected. This is consistent with the notion that term premia should contribute to the level of yields.

Principal components rotate short rate expectations and risk premia in yields into factors. Figure 3 illustrates this point by plotting the contribution of risk premia \widehat{cf}_t to the variation in the level factor (average yield curve) and the slope ($y_t^{(10)} - y_t^{(1)}$).

[Figure 3 here.]

The plot follows from performing a variance decomposition of the level and the slope in terms of $\tau_t^{CPI}, c_t^{(1)}, \widehat{cf}_t$.¹³ We note that the level factor is predominantly related to short rate expectations τ_t^{CPI} and $c_t^{(1)}$, that explain 85% and 9% of its variation, respectively, while the premium component \widehat{cf}_t accounts for a small portion of less than 5% of overall variance of the level. This low contribution of \widehat{cf}_t explains why the level does not predict bond returns. The slope of the yield curve, instead, contains information about (transitory) short rate expectations and risk premia, with $c_t^{(1)}$ and \widehat{cf}_t explaining 75% and 23% of its variance, respectively. This relatively larger contribution of \widehat{cf}_t is the reason why the slope carries some degree of predictability for future returns. However, because of the dominant role of the transitory short rate expectations $c_t^{(1)}$, predicting bond returns with the slope suffers from an error-in-variables problem.

V. Macroeconomic fundamentals and \widehat{cf}

This section studies the link between the return forecasting factor and macroeconomic fundamentals. We find that \widehat{cf}_t comprises the predictability of many macro-finance variables.

¹³The regression of the level on those three factors gives an $R^2 = 1$ by construction, the regression of the slope on the three factors gives an $R^2 = 0.995$. The contributions of the factors to the explained variance are obtained as the partial R^2 's. We note that these conclusions are unchanged if we use the first two principal components rather than the average yield and the term spread.

Conditional on that factor, the additional predictive power of macroeconomic risk pertains to bonds with short maturities, which we associate with the influence of monetary policy on this segment of the curve.

V.A. Do macro variables predict returns beyond \widehat{cf} ?

Including macro-finance variables in predictive regressions together with the CP factor or with yield principal components usually leads to an increase in R^2 . Ludvigson and Ng (2009) summarize information in 132 macro-finance series and find that real activity and inflation factors remain highly significant and increase the forecasting power relative to the CP factor. Cooper and Priestley (2009) reach a similar conclusion considering the output gap.

It is natural to ask whether and how these conclusions may change when we take \widehat{cf}_t as our benchmark for predictability. Specifically, we estimate the regression:

$$rx_{t+1}^{(n)} = b_0 + b_1\widehat{cf}_t + b_2'\text{Macro}_t + \varepsilon_{t+1}^{(n)}, \quad (22)$$

where Macro_t represents the additional macro-finance information. This regression allows us to assess which macroeconomic variables are reflected in the movements of bond risk premia.

Table VII displays estimates of (22) with eight macro-finance factors, \widehat{F}_t , constructed according to Ludvigson and Ng (2009), and indicates the domains that these factors capture. We use data from 1971:11 through 2007:12. The end of the sample is dictated by the availability of the macro series. Alone, \widehat{F}_t explain more than 20% of variation in bond excess returns. Although we do not report the details of the separate regression of rx on \widehat{F}_t only, in Table VII we indicate significant factors at the 1%, 5% and 10% level with superscripts H, M, L, respectively. If we exclude \widehat{cf} from regression (22), the significant factors involve financial spreads, stock market returns, inflation, and monetary conditions.

[Table VII here.]

In the presence of \widehat{cf}_t , however, most macro variables lose predictive power. Their contribution to R^2 , denoted as “ ΔR^2 ” in the table, does not exceed 2%. The only exception is the two-year bond for which inflation and, to a lesser degree, the real activity factor remain significant yielding ΔR^2 of 5%.¹⁴

We also use output gap to represent macro information in equation (22). Following Cooper and Priestley (2009), we obtain gap_t from the unrevised data on industrial production by applying a quadratic time trend.¹⁵ Also here, the estimates (not reported) suggest that gap_t does not provide additional information beyond that conveyed by \widehat{cf}_t .

Out of eight factors considered in Table VII, only \widehat{F}_{2t} is statistically significant for intermediate and long maturities. To the extent that \widehat{F}_{2t} is related to different financial spreads, as shown by Ludvigson and Ng (2009), it seems to reflect the variation in funding liquidity. To explore this predictability channel, we construct several liquidity proxies such as spreads on commercial papers, swap rates, Baa corporate bonds, three-month T-bill over Fed’s target, and the TED. We also consider the on-the-run liquidity factor recently proposed by Fontaine and Garcia (2011) (henceforth, FG factor).¹⁶ Exact descriptions of the variables are in Appendix B. We evaluate the joint significance of \widehat{cf}_t and each of those variables over the sample period 1987:04 through 2007:12 when all series are available. The FG factor and the Moodys Baa spread turn out to be the only variables that, albeit weakly and only at short maturities, continue to contribute to the predictability achieved with \widehat{cf}_t .¹⁷

At this point, it is worth summarizing two properties of \widehat{cf} revealed by our analysis up to now: \widehat{cf} has a non-trivial effect on the cross-section of yields, and its predictive power increases with the maturity of the bond. The current section suggests that while macroeconomic risks have different effect on term premia across maturities, their contribution is strongest at the short maturity range. Specifically, using macroeconomic information

¹⁴We do not report analogous estimates with the CP factor for our sample, and just note that they conform with the conclusions of Ludvigson and Ng (2009). Using the CP factor as a benchmark, increases the contribution of macroeconomic information in (22) in that most \widehat{F}_t variables preserve their significance.

¹⁵We construct gap_t using the industrial production going back to 1948:01 as in Cooper and Priestley (2009).

¹⁶We thank Jean-Sébastien Fontaine for providing the data on their liquidity factor.

¹⁷The details of these results are available upon request.

beside \widehat{cf} could improve investors' forecast of the return on the two-year bond, but not on bonds with longer maturities.

V.B. Predicting economic activity

The ability of the slope of the yield curve to predict future economic activity is one of the well-documented facts in macroeconomics (e.g. Estrella and Hardouvelis, 1991; Harvey, 1989). We have shown that the slope reflects both short rate expectations and the risk premium. In this section, we explore how each of these components is related to economic growth.

We forecast future economic activity, measured by the Chicago Fed National Activity Index (CFNAI),¹⁸ by estimating the following regression:

$$\text{CFNAI}_{t+k} = \beta' X_t + u_{t+k}, \quad k = 1, 2, \dots, 48 \text{ months}, \quad (23)$$

where as regressors X_t we consider the following variables: (i) slope $y_t^{(10)} - y_t^{(1)}$, (ii) the negative of the transitory short rate expectations $-c_t^{(1)}$,¹⁹ (iii) the return forecasting factor \widehat{cf}_t , and (iv) jointly $-c_t^{(1)}$ and \widehat{cf}_t . To study the impact on growth of the different components we run each of these regressions separately.

Figure 4 summarizes the R^2 's, β coefficients, and t-statistics of these regressions. The key observation is that the term premium and the expected short rate differ in their respective abilities to predict economic activity in terms of the magnitude and forecast horizon. Moreover, separating these two components leads to a stronger growth predictability relative to the slope, because it alleviates the errors-in-variables problem. Lower expected short rate (higher $-c_t^{(1)}$) forecasts higher growth up to 24 months ahead, with a peak around 12 months

¹⁸The CFNAI is published on the Chicago Fed web page as a weighted average of 85 existing monthly indicators of national economic activity including (i) production and income, (ii) employment, unemployment, and hours, (iii) personal consumption and housing, and (iv) sales, orders, and inventories. The index is constructed to have an average value of zero and a standard deviation of one. Since economic activity tends toward trend growth rate over time, a positive index reading corresponds to growth above trend and a negative index reading corresponds to growth below trend. Thus, we can interpret CFNAI as a growth indicator. We use the monthly unsmoothed reading of the index.

¹⁹The minus sign in front of $c_t^{(1)}$ makes it positively correlated with the slope, and is for the ease of exposition.

ahead. This effect is analogous to the one captured by the slope. A new effect comes with the term premium: Initially, a higher term premium (higher \widehat{cf}_t) predicts lower growth up to one year ahead, but this effect is marginally significant and tapers off relatively fast. Over longer time horizon, higher bond risk premium predicts higher economic activity—this effect becomes strongest about two and a half years ahead, precisely when the impact of short rate expectations wears down to approximately zero.²⁰ This last results provides an analogy to the consumption-based CAPM literature (Parker and Julliard, 2005; Malloy, Moskowitz, and Vissing-Jorgensen, 2009). Similar to our findings, this literature finds that a positive relationship between future consumption growth and stock excess returns becomes visible at long horizons of about three years.

[Figure 4 here.]

VI. Additional results

VI.A. Link to unspanned factors

The return predictability that we document crucially depends on the new information we provide about long-horizon inflation expectations. That is, relative to the forward rate regressions, the extra predictive power comes from extending the information set of an econometrician. It is therefore natural to relate our findings to the literature on unspanned (or hidden) factors in term premia (Joslin, Pribsch, and Singleton, 2010; Duffee, 2011). One way to assess the additional information that τ_t^{CPI} provides, and that is not in the cross section of yields, is to estimate the residual of the regression:

$$\tau_t^{CPI} = b_0 + b_1' \mathbf{y}_t + u_t. \tag{24}$$

²⁰The relation between the yield curve and economic activity is often discussed by policy makers. “(...) the decline in term premiums in the Treasury market of late may have contributed to keeping long-term interest rates relatively low and, consequently, may have supported the housing sector and consumer spending more generally” (D. Kohn, speech on Jul 25, 2005). “Easier financial conditions will promote economic growth.” (B. Bernanke, Nov 4, 2010, The Washington Post). These quotes appear to refer to the part of the relationship between risk premia and growth that we find at shorter horizons.

We estimate (24) using six yields with maturities one, two, five, seven, ten and 20 years as regressors.²¹ The adjusted R^2 from this regression is 89.5%. While high, the R^2 less than one suggest a possibility of important information being left in \hat{u}_t .

By construction the fitted \hat{u}_t is orthogonal to the vector of yields \mathbf{y}_t , i.e. it is a hidden factor. Two additional features of the residual are worth emphasizing: First, \hat{u}_t exhibits a non-zero correlation with factors in our decomposition $(\tau_t^{CPI}, c_t^{(1)}, \widehat{cf}_t)$. The correlation between \widehat{cf}_t and \hat{u}_t is the strongest at -0.71 . Other correlations are also non-negligible reaching 0.32 for τ_t^{CPI} and 0.47 for $c_t^{(1)}$. Second, while \hat{u}_t predicts bond excess returns, its predictive power is on average 45% lower than that of \widehat{cf}_t . In a predictive regression of bond excess returns on \hat{u}_t , the R^2 values range between 19% and 31% across maturities. This evidence indicates that \hat{u}_t does not capture the same term premium effect that \widehat{cf}_t does.

The regression (24) assumes that yields are measured perfectly. We now analyze the extent to which the empirical properties of \hat{u} are sensitive to this assumption.²² To this end, we set up a simulation experiment from an affine three-factor model under the null hypothesis that three factors are perfectly spanned by the cross section of yields.²³ In such a model, in absence of the measurement error, the R^2 in regression (24) equals one, and the theoretical correlation between the residual \hat{u} and underlying factors equals zero. We find that already a small measurement error introduces large biases in the regression coefficients. Those biases tend to decrease the R^2 away from the theoretical value of one and to increase the correlation between the residual \hat{u} and factors away from the theoretical value of zero.

The experiment suggests that the empirical values of R^2 in regression (24) and the correlation between \widehat{cf} and \hat{u} found in the data could arise under the measurement error of about

²¹Including more yields or different maturities does not change the results. Including more yields raises the concern of multicollinearity.

²²The perfect observability of yields can be violated for several reasons, e.g. the splining of the zero curve. For instance the measurement error implied by different constructions of the zero curve (GSW, FB, CMT) ranges between 6 and 25 basis points, with an average of 14 basis points. Similarly, using the CRSP data set, we estimate the average bid-ask spread in yields across maturities is to equal 8 bps (1971–1980), 5 bps (1981–1990), 2 bps (1991–2000), and 1 bps (2001–2010).

²³Details of the simulation design and the results are available upon request.

five basis points. These results grant some caution in interpreting \widehat{cf} as the hidden factor. We leave a further investigation of these issues for future research.

VI.B. Out-of-sample predictability

Out-of-sample predictability presents an important check for the stability of our findings. Suppose an investor perceives the process generating the slow variation in yields as being driven by the long-run inflation expectations, and estimates the persistent factor τ_t^{CPI} using core CPI. In doing so, they exploit inflation information that is available only up to time t , and update the estimates of τ_t^{CPI} with the gain parameter of 0.9868.

We consider three out of sample periods starting in 1978:01, 1985:01 and 1995:01, ending in 2009:12. For each of the samples, we obtain the initial estimates based on the period from 1971:11 until 1977:01, until 1984:01 and until 1994:01, respectively. With information up to this point, say t_0 , we obtain cycles as in equation (12), and run regression (13) predicting excess returns realized up to t_0 using cycles up to t_0 less 12 months. At the estimated parameters, we then predict excess returns 12 months ahead, i.e. realized at t_0 plus 12 months. We extend the sample month-by-month, and repeat these steps until we reach the maximum sample length. The performance of cycles is compared to that of forward rates and the slope.

Our out-of-sample tests involve three measures (see online appendix for details). We start with the encompassing test (ENC-NEW) proposed by Clark and McCracken (2001). Following discussion in Section IV.A, we treat cycles as an unrestricted model and forward rates as a restricted one. The null hypothesis of the ENC-NEW test is that the restricted model (forwards) encompasses all the predictability in bond excess returns, and it cannot be further improved by the unrestricted model (cycles).²⁴

The second measure is the ratio of mean squared errors implied by the unrestricted versus restricted model, MSE_{cyc}/MSE_{fwd} . A number less than one indicates that the unrestricted model is able to generate lower prediction errors.

²⁴Clark and McCracken (2005) show that the ENC-NEW test statistic has a non-standard distribution under the null, therefore we obtain the critical values by bootstrapping.

The third measure is the out-of-sample R^2 of Campbell and Thompson (2008), R_{OOS}^2 . R_{OOS}^2 compares the forecasting performance of a given predictor toward a “naive” forecast obtained with the historical average return. The statistic is analogous to the in-sample R^2 : Its positive value indicates that the predictive model has a lower mean-squared prediction error than the “naive” forecast.

Throughout, for forward rate regressions, we use forward rates with maturities of one, two, five, seven, ten and 20 years as predictors. For cycle regressions (except the ENC test), we use the short maturity cycle $c_t^{(1)}$ and the average cycle \bar{c}_t as predictors in bivariate regressions.²⁵ For the sake of comparability with the forward rate regressions, in the ENC test we simply employ six cycles with the same maturities as the forward rates. For slope regressions, we forecast excess return on the n -year bond with the corresponding spot forward spread, $f_t^{(n)} - y_t^{(1)}$. The slope regressions provide a useful benchmark out of sample because, other than forward rate regressions, they do not require an estimation of a large number of coefficients.

The panels of Table IX report the results for 1971–2009, 1985–2009, and 1995–2009, respectively. The ENC-NEW test rejects the null hypothesis for all maturities at the 95% confidence level: The cycles’ model significantly improves the predictive performance over forwards. The MSE ratio, $\text{MSE}_{\text{cyc}}/\text{MSE}_{\text{fwd}}$, is reliably below one for all maturities. In the recent sample, the $\text{MSE}_{\text{cyc}}/\text{MSE}_{\text{fwd}}$ ratio is substantially lower than in the period 1971–2009. Indeed, while in the last two decades the performance of forward rates deteriorates compared to the full sample, the performance of the cycles remains relatively stable. With one exception, R_{OOS}^2 values obtained with cycles are large and positive for all maturities across all sample periods. In summary, the out-of-sample statistics support the previous in-sample evidence, indicating the relevance of the economic mechanism that the cycles capture.

[Table IX here.]

²⁵The results remain almost identical if we first construct \widehat{cf} as in Section III.B, and then use it for predicting returns out of sample.

Additionally, in the online appendix we show that the out-of-sample results are only weakly influenced when varying the learning parameter between 0.975 (fast updating) and 0.995 (very slow updating). Based on the literature which we have summarized above, this range for the gain parameter can be viewed as covering the extremes.

VII. Robustness

VII.A. Actual non-splined bond portfolio returns across horizons

One may be worried that the predictive results depend on the use of splined (artificial) bond returns. To address this concern, we use the actual bond returns for different maturity segments available from the CRSP Fama files. The data have a monthly frequency and are an equal weighted average of unadjusted holding period returns for each bond in the portfolio. From the monthly files we compute returns in excess of Tbill rate over investment horizons of one, three, six, nine and 12 months.

Table VIII summarizes the predictive regressions showing that \widehat{cf}_t is highly significant across all horizons, and the R^2 increases with the horizon. For 12-month holding period the results broadly agree with the ones we have reported in Section III, Table V. Therefore, our conclusions are robust to the way we measure bond returns.

[Table VIII here.]

VII.B. Other econometric concerns

Predictive regressions like those we estimate pose econometric problems, especially at longer forecast horizons and when working with the available sample sizes. Therefore, we provide additional evidence based on reverse regression and non-overlapping annual returns.

Let us go back to regression (13), whose results we report in Table IV. The Wald test strongly rejects that all coefficient on c_t 's are zero, using both the Hansen-Hodrick and the Newey-West method. However, since both tests are known to overreject the null hypothesis in small samples (e.g., Ang and Bekaert, 2007), in Table IV we additionally

provide a conservative test based on the reverse regression proposed by Hodrick (1992) and recently extended by Wei and Wright (2010). This approach amounts to regressing short-horizon (one-month) returns on the long-run (twelve-month) mean of the cycles, and is less prone to overrejection.²⁶ While the reverse regression test statistics are by design more moderate, we consistently reject the null of no predictability by the cycles at the conventional significance levels. We compare the standard errors obtained with the cycles to those of the forward rate regressions. In both samples and across all maturities, cycles give much stronger evidence of predictability than do forward rates.

We also study the predictability of nonoverlapping annual bond returns by the single factor \widehat{cf} . Our conclusions do not change, both in terms of statistical and economic significance. Based on a non-overlapping sample, the predictive R^2 range from 40% for the two-year, through 56% for the ten-year and 52% for the twenty-year bond.

VII.C. Alternative ways of constructing \widehat{cf}

There are alternative ways of constructing the single forecasting factor. We could proceed in analogy to Cochrane and Piazzesi (2005) and obtain the forecasting factor as a fitted value of the linear combination of several cycles. Also, we could estimate the factor in one step via non-linear least squares, or by means of an eigenvalue decomposition of the covariance matrix of expected returns. As an equally valid approach, in the second step regression (16), we could only use contemporaneous time- t variables on the left- and right-hand sides of the equation. As it turns out, the single factor amounts to orthogonalizing the average longer-maturity cycle \bar{c}_t with respect to $c_t^{(1)}$.²⁷ Those different approaches to constructing \widehat{cf}_t produce essentially an identical return forecasting factor (correlation exceeds 99% among the alternatives), and lead to very similar conclusions for return predictability. These alternatives are discussed in the online appendix.

²⁶Wei and Wright (2010) extend the reverse regressions proposed by Hodrick (1992) beyond testing the null hypothesis of no predictability, and allowing for highly persistent regressors. In constructing one-month excess returns on bonds, we follow Campbell and Shiller (1991), approximating the log price of a $(n - 1/12)$ -maturity bond as $-(n - 1/12)y_t^{(n)}$.

²⁷As the online appendix shows, we could estimate the return forecasting factor as the residual from the regression: $\bar{c}_t = \gamma_1 c_t^{(1)} + \bar{\varepsilon}_{t+1}$, $\widehat{cf}_t = \bar{c}_t - \hat{\gamma}_1 c_t^{(1)}$.

VII.D. Other data sets

One may be concerned that the return predictability we document depends to the CMT yields, and the way we construct the zero curve. To show that our results are robust to these choices, we perform the predictive exercise on other two commonly used data sets constructed by Fama and Bliss (FB) and Gürkaynak, Sack, and Wright (2006, GSW). We remain conservative on two fronts. First, we focus on the range of maturities from one to five years, as dictated by the FB data. Second, to assess the sensitivity of our results to the recent crisis, we consider two samples: (i) excluding the crisis (1971–2006), and (ii) including the crisis (1971–2009). The data sets we consider differ not only in the way of constructing the zero curve, but also in the choice of the underlying yields. For instance, CMT yields are based on the on-the-run securities while GSW yields are off-the-run. We are in this way able to assess if our conclusions are driven by the liquidity premium pertaining to the on-the-run curve.

The forecasting ability of the cycles is confirmed across all data sets. Despite using a restricted number of maturities, the R^2 's obtained with the cycles are in the 50% range.²⁸ Using yields and forward rates, or spreads leads to much lower predictability, diminishing the R^2 's by half. The gap between cycles and other predictors becomes even more apparent when we include the crisis years. While the recent turmoil leads to a weakened performance across all regressors, with forward rates explaining just about 17% of variation in \overline{rx}_{t+1} , the predictive power of the cycles still remains confidently above 45%.

VIII. Conclusions

This paper is concerned with the role of frequencies in the yield curve and how they encode different economic forces at work. In a first step, we separate two effects in yields: a slow adjustment related to the changing long-run mean of inflation, and transitory fluctuations (cycles) around that smooth component reflecting current macro-finance conditions. Consistent with the long-run Fisher hypothesis, cycles present stationary deviations from the

²⁸We report these results in the online appendix. Unreported results show that the one to five year yield maturity range that we use here can be restrictive. For example, in the Greenspan's subperiod we find that the inclusion of cycles with longer maturities improves the forecasting performance.

long-term relationship between yields and long-horizon inflation expectations. The cycles across different maturities combine the term structure of transitory short rate expectations with the term structure of risk premia. Using their cross-sectional composition, in the second step, we distill those two elements into distinct factors. This leaves us with three observable variables: persistent short rate expectations, transitory short rate expectations, and the term premium factor, \widehat{cf} . We confirm our findings by various robustness checks including different measures of long-horizon inflation expectations.

References

- ANG, A., AND G. BEKAERT (2007): “Stock Return Predictability: Is It There?,” *Review of Financial Studies*, 20, 651–707.
- BARSKY, R. B., AND J. B. DELONG (1991): “Forecasting Pre-World War I Inflation: The Fisher Effect and the Gold Standard,” *Quarterly Journal of Economics*, 106, 815–836.
- CAMPBELL, J. Y., AND P. PERRON (1991): “Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots,” *NBER Macroeconomics Annual*, 6, 141–201.
- CAMPBELL, J. Y., AND R. J. SHILLER (1991): “Yield Spreads and Interest Rate Movements: A Bird’s Eye View,” *Review of Economic Studies*, 58, 495–514.
- CAMPBELL, J. Y., AND S. THOMPSON (2008): “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?,” *Review of Financial Studies*, 21, 1509–1531.
- CARCELES-POVEDA, E., AND C. GIANNITSAROU (2007): “Adaptive Learning in Practice,” *Journal of Economic Dynamics and Control*, 31, 2659–2697.
- CARLSON, J. A. (1977): “A Study of Price Forecasts,” *Annals of Economic and Social Measurement, NBER*, 6, 33–63.
- CLARK, T., AND M. MCCrackEN (2001): “Tests of Equal Forecast Accuracy and Encompassing for Nested Models,” *Journal of Econometrics*, 105, 85–110.
- (2005): “Evaluating Direct Multi-Step Forecasts,” *Econometric Reviews*, 24, 369–404.
- COCHRANE, J. H., AND M. PIAZZESI (2005): “Bond Risk Premia,” *American Economic Review*, 95, 138–160.
- COOPER, I., AND R. PRIESTLEY (2009): “Time-Varying Risk Premiums and the Output Gap,” *Review of Financial Studies*, 22, 2801–2833.
- CROUSHORE, D. (2008): “Revisions to PCE Inflation Measures: Implications for Monetary Policy,” Working paper FRB Philadelphia.
- CROUSHORE, D., AND T. STARK (1999): “A Real-Time Data Set for Macroeconomists,” Working paper, FRB Philadelphia.
- CROWDER, W. J., AND M. E. WOHar (1999): “Are Tax Effects Important in the Long-Run Fisher Relationship? Evidence from the Municipal Bond Market,” *Journal of Finance*, 54, 307–317.
- DELONG, J. (1997): *America’s Only Peacetime Inflation: The 1970s* Chicago: University of Chicago Press.
- DEWACHTER, H., AND L. IANIA (2010): “An Extended Macro-Finance Model with Financial Factors,” Working paper, Katholieke Universiteit Leuven.
- DEWACHTER, H., AND M. LYRIO (2006): “Learning, Macroeconomic Dynamics and the Term Structure of Interest Rates,” Working paper, Katholieke Universiteit Leuven, Erasmus University of Rotterdam.

- DUFFEE, G. R. (2011): “Information in (and Not in) the Term Structure,” *Review of Financial Studies*, forthcoming.
- ENGLE, R., AND C. W. GRANGER (1987): “Co-integration and Error Correction: Representation, Estimation, and Testing,” *Econometrica*, 55, 251–276.
- ESTRELLA, A., AND G. A. HARDOUVELIS (1991): “The Term Structure as a Predictor of Real Economic Activity,” *Journal of Finance*, 46, 555–576.
- EVANS, G. W., AND S. HONKAPOHJA (2009): “Learning and Macroeconomics,” *Annual Review of Economics*, 1, 421–449.
- FAMA, E. (2006): “The Behavior of Interest Rates,” *Review of Financial Studies*, 19, 359–379.
- FONTAINE, J.-S., AND R. GARCIA (2011): “Bond Liquidity Premia,” *Review of Financial Studies*, forthcoming.
- GOYAL, A., AND I. WELCH (2008): “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 21, 1455–1508.
- GÜRKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2006): “The U.S. Treasury Yield Curve: 1961 to the Present,” Working paper, Federal Reserve Board.
- HARVEY, C. R. (1989): “Forecasts of Economic Growth from the Bond and Stock Markets,” *Financial Analysts Journal*, September–October, 38–45.
- HODRICK, R. J. (1992): “Dividend Yields and Expected Stock Returns: Alternative Procedures for Inference and Measurement,” *Review of Financial Studies*, 5, 357–386.
- HUANG, J.-Y., AND Z. SHI (2010): “Determinants of Bond Risk Premia,” Working paper, Penn State University.
- JARDET, C., A. MONFORT, AND F. PEGORARO (2010): “No-Arbitrage Near-Cointegrated VAR(p) Term Structure Models, Term Premia and GDP Growth,” Working paper, Banque de France, CNAM, CREST.
- JOSLIN, S., M. PRIEBSCHE, AND K. SINGLETON (2010): “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” Working paper, MIT Sloan School of Management and Stanford University.
- KILEY, M. (2008): “Monetary Policy Actions and Long-Run Inflation Expectations,” Working paper, Federal Reserve Board.
- KOIJEN, R. S., O. VAN HEMERT, AND S. VAN NIEUWERBURGH (2009): “Mortgage Timing,” *Journal of Financial Economics*, 93, 292–324.
- KOZICKI, S., AND P. TINSLEY (1998): “Moving Endpoints and the Internal Consistency of Agents’ Ex Ante Forecasts,” *Computational Economics*, 11, 21–40.
- (2001a): “Shifting Endpoints in Term Structure of Interest Rates,” *Journal of Monetary Economics*, 47, 613–652.

- (2001b): “Term Structure Views of Monetary Policy under Alternative Models of Agent Expectations,” *Journal of Economic Dynamics & Control*, 25, 149–184.
- (2005): “Permanent and Transitory Policy Shocks in an Empirical Macro Model with Asymmetric Information,” *Journal of Economic Dynamics and Control*, 29, 1985–2015.
- (2006): “Survey-Based Estimates of the Term Structure of Expected U.S. Inflation,” Working paper, Bank of Canada.
- KÜNSCH, H. R. (1989): “The Jackknife and the Bootstrap for General Stationary Observations,” *Annals of Statistics*, 17, 1217–1241.
- LUDVIGSON, S. C., AND S. NG (2009): “Macro Factors in Bond Risk Premia,” *Review of Financial Studies*, 22, 5027–5067.
- MALLOY, C. J., T. J. MOSKOWITZ, AND VISSING-JORGENSEN (2009): “Long-Run Stockholder Consumption Risk and Asset Returns,” *Journal of Finance*, 64, 2427–2479.
- MALMENDIER, U., AND S. NAGEL (2009): “Learning from Inflation Experiences,” Working paper, UC Berkeley and Stanford University.
- MISHKIN, F. (1992): “Is the Fisher effect for real?,” *Journal of Monetary Economics*, 30, 195–215.
- MUNDELL, R. (1963): “Inflation and Real Interest,” *Journal of Political Economy*, 71, 280–283.
- ORPHANIDES, A., AND M. WEI (2010): “Evolving Macroeconomic Perceptions and the Term Structure of Interest Rates,” Working paper, Board of Governors of the Federal Reserve System.
- PARKER, J. A., AND C. JULLIARD (2005): “Consumption Risk and the Cross Section of Expected Returns,” *Journal of Political Economy*, 113, 185–222.
- PIAZZESI, M., AND M. SCHNEIDER (2011): “Trend and Cycle in Bond Premia,” Working paper, Stanford University and NBER.
- RUDEBUSCH, G. D., AND T. WU (2008): “A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy,” *The Economic Journal*, 118, 906–926.
- STOCK, J. H., AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit, and Banking*, 39, 3–33.
- SUMMERS, L. (1982): “The Non-Adjustment of Nominal Interest Rates: A Study of the Fisher Effect,” NBER working paper.
- WEI, M., AND J. H. WRIGHT (2010): “Reverse Regressions and Long-Horizon Forecasting,” Working paper, Federal Reserve Board and Johns Hopkins University.

IX. Figures

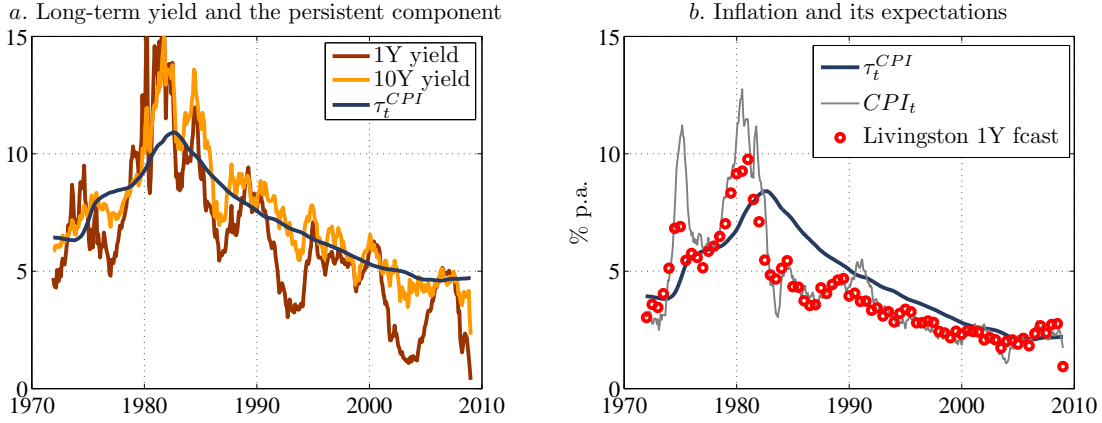


Figure 1: The persistent inflation expectations factor, τ_t^{CPI}

Panel *a* superimposes the one- and ten-year yield with τ_t^{CPI} . τ_t^{CPI} is constructed as the discounted moving average of the core CPI in equation (8), with sums truncated at $N = 120$ months and the discount factor $v = 0.9868$. τ_t^{CPI} is fitted to yields so that all variables match in terms of magnitudes. Panel *b* plots the one-year ahead median inflation forecasts from the Livingston survey and realized core CPI inflation.

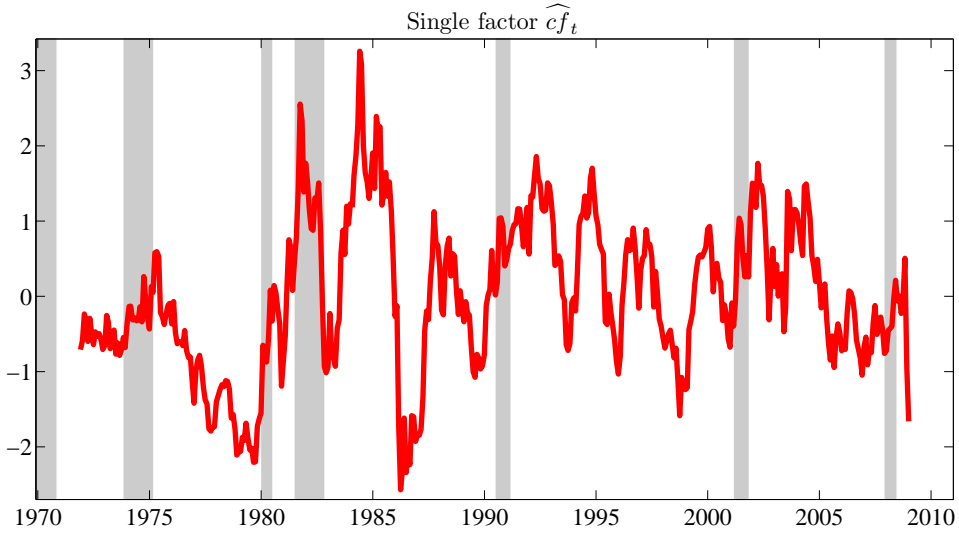


Figure 2: Single factor

The figure displays the return forecasting factor \widehat{cf}_t formed with equation (16). Shaded areas mark the NBER recessions. The series has been standardized.

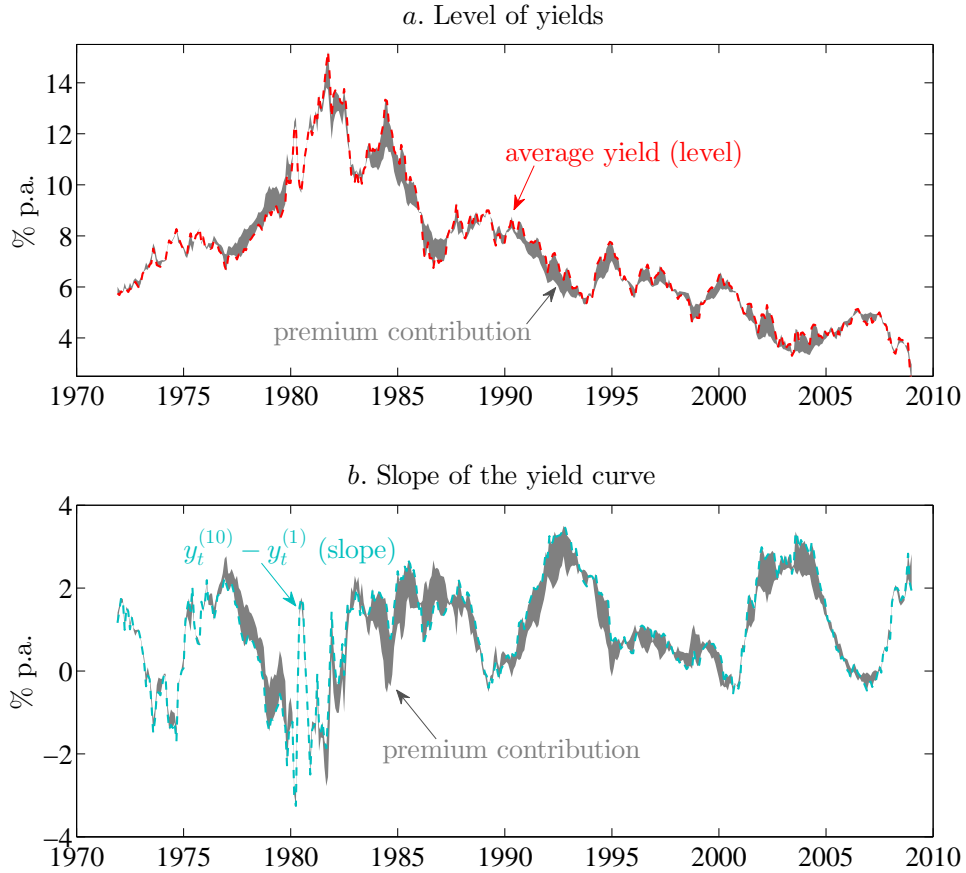


Figure 3: Contributions of \widehat{cf}_t to level and slope

The figure plots the contributions of \widehat{cf}_t to the variation in the level and the slope (term spread) of the yield curve. We run the regression of the level and slope, respectively, on three factors $\tau_t^{CPI}, c_t^{(1)}, \widehat{cf}_t$. The three factors explain 100% of variation in the level and 99.5% of variation in the term spread. Then, the shadowed area shows the difference between the level and the slope, respectively, and their value fitted by two factors τ_t^{CPI} and $c_t^{(1)}$, i.e. excluding \widehat{cf}_t .

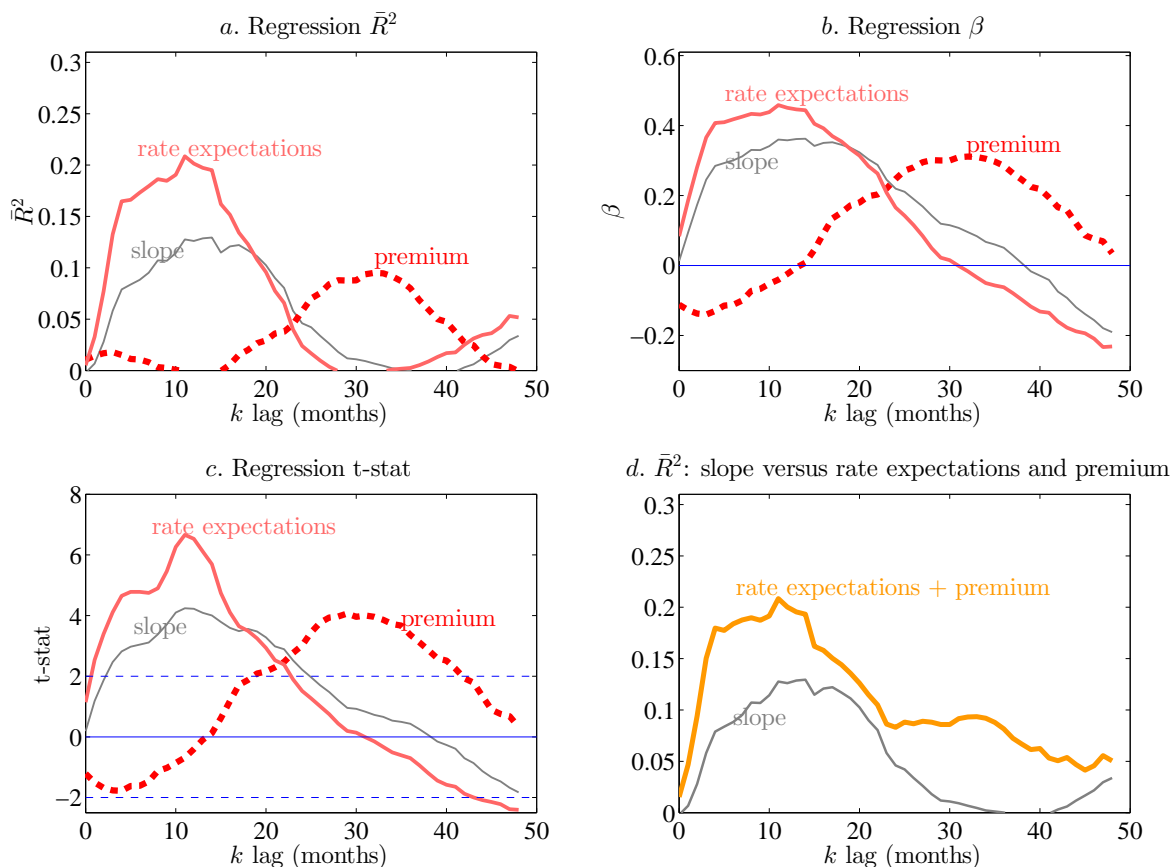


Figure 4: Predicting economic activity

The figure shows R^2 , β coefficients, and t-statistics from regressions of CFNAI (a real activity measure) on yield curve variables. The regressions are of the type:

$$\text{CFNAI}_{t+k} = \beta' X_t + u_{t+k},$$

where we use four different sets of regressors in X_t : (i) risk premium \widehat{cf}_t (labeled “premium”), (ii) transitory short rate expectations $-c_t^{(1)}$ (labeled “rate expectations”), (iii) both rate expectations and premium $(-c_t^{(1)}, \widehat{cf}_t)$ (labeled “rate expectations + premium”), and (iv) slope of the yield curve $y_t^{(10)} - y_t^{(1)}$ (labeled “slope”). Transitory short rate expectations $c_t^{(1)}$ enter the regressions with a minus so that they are positively correlated with the slope of the yield curve.

X. Tables

Table II: Estimates of the vector error correction model

The table reports the estimated coefficients from the error correction model on monthly frequency:

$$\Delta y_t^{(n)} = a_c c_{t-\Delta t}^{(n)} + a_y \Delta y_{t-\Delta t}^{(n)} + a_\tau \Delta \tau_{t-\Delta t} + a_0 + \varepsilon_t, \quad \Delta t = 1 \text{ month}$$

Reported t-statistics use Newey-West adjustment with 12 lags. For ease of comparison, all variables are standardized. $\Delta \bar{y}_t$ in the last column denotes the average yield change across maturities.

Regressor	Dependent variable						$\Delta \bar{y}_t$
	$\Delta y_t^{(1)}$	$\Delta y_t^{(2)}$	$\Delta y_t^{(5)}$	$\Delta y_t^{(7)}$	$\Delta y_t^{(10)}$	$\Delta y_t^{(20)}$	
$c_{t-\Delta t}^{(n)}$	-0.19 (-2.37)	-0.20 (-2.84)	-0.21 (-3.54)	-0.20 (-3.78)	-0.20 (-3.94)	-0.20 (-4.06)	-0.20 (-3.64)
$\Delta y_{t-\Delta t}^{(n)}$	0.22 (3.02)	0.22 (4.41)	0.18 (4.22)	0.14 (3.15)	0.14 (3.28)	0.12 (2.71)	0.19 (3.99)
$\Delta \tau_{t-\Delta t}$	0.07 (1.16)	0.07 (1.36)	0.09 (1.66)	0.08 (1.62)	0.09 (1.75)	0.06 (1.22)	0.09 (1.65)
\bar{R}^2	0.06	0.06	0.06	0.05	0.05	0.04	0.06

Table III: Bond excess returns: summary statistics

Panel A reports summary statistics for bond excess returns. **Panel B** reports the mean and standard deviation of duration-standardized excess returns to remove the effect of duration. AR(1) denotes the first order autocorrelation coefficient. **Panel C** reports the correlation between excess returns of different maturities. Bond excess returns are computed at an annual horizon as $rx_{t+1}^{(n)} = p_{t+1}^{(n-1)} - p_t^{(n)} - y_t^{(1)}$ and multiplied by 100.

Panel A. Bond excess returns						
	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
Mean	0.69	1.61	1.97	1.99	2.57	3.05
Stdev	1.99	6.29	8.60	11.69	16.99	22.59
AR(1)	0.94	0.93	0.93	0.93	0.93	0.92

Panel B. Duration standardized excess returns						
	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
Mean	0.34	0.32	0.28	0.20	0.17	0.15
Stdev	1.00	1.26	1.23	1.17	1.13	1.13

Panel C. Correlation of excess returns						
	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
$rx^{(2)}$	1.00	–	–	–	–	–
$rx^{(5)}$	0.95	1.00	–	–	–	–
$rx^{(7)}$	0.91	0.99	1.00	–	–	–
$rx^{(10)}$	0.86	0.96	0.99	1.00	–	–
$rx^{(15)}$	0.81	0.93	0.96	0.99	1.00	–
$rx^{(20)}$	0.76	0.87	0.91	0.94	0.97	1.00

Table IV: First look at predictive regressions of bond returns

The table reports the results of predictive regressions in equation (13). In the first row, we provide adjusted R^2 values. To assess the small sample (SS) properties of \bar{R}^2 , the next three rows give its 5%, 50% and 95% percentile values obtained with the block bootstrap (see Appendix E). The $\chi^2(6)$ tests if the coefficients (excluding the constant) are jointly equal to zero. We report the Hansen-Hodrick (HH) and the Newey-West (NW) correction, using 12 and 15 lags, respectively. “LS” means that the statistics were estimated using the full sample. The row “ $\chi^2(6)$ (SS 5%)” states the lower 5% bound on the values of the χ^2 -test (using NW adjustment) obtained with the bootstrap. We also provide conservative standard errors obtained using the reverse regression delta method (rev.reg.) of Wei and Wright (2010) and the corresponding p-values. The last five rows summarize the corresponding results for the forward rate regressions. Cycles \mathbf{c}_t and forward rates \mathbf{f}_t are of maturities one, two, five, seven, ten, and 20 years. Sample is 1971:11–2009:12. The asymptotic 1%, 5%, and 10% critical values for $\chi^2(6)$ are 16.81, 12.59, and 10.64, respectively.

Statistic	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
Cycle regressions: $rx_{t+1}^{(n)} = \delta_0 + \delta' \mathbf{c}_t + \varepsilon_{t+1}^{(n)}$						
\bar{R}^2	0.42	0.49	0.52	0.55	0.56	0.57
\bar{R}^2 (SS,5%)	0.31	0.37	0.40	0.44	0.45	0.44
\bar{R}^2 (SS,50%)	0.47	0.53	0.56	0.59	0.59	0.59
\bar{R}^2 (SS,95%)	0.61	0.65	0.67	0.68	0.68	0.69
$\chi^2(6)$ (LS, HH)	46.98	117.38	149.18	182.40	181.22	149.69
$\chi^2(6)$ (LS, NW)	61.59	131.12	150.20	172.52	166.63	125.39
$\chi^2(6)$ (SS,5%)	48.01	86.10	95.26	116.47	116.08	87.13
$\chi^2(6)$ (rev. reg.)	13.19	24.21	28.56	35.80	36.66	30.02
pval	0.04	0.00	0.00	0.00	0.00	0.00
Forward-rate regressions: $rx_{t+1}^{(n)} = d_0 + d' \mathbf{f}_t + \varepsilon_{t+1}^{(n)}$						
\bar{R}^2	0.21	0.22	0.24	0.24	0.26	0.31
$\chi^2(6)$ (LS, HH)	23.82	26.07	23.13	22.77	23.20	20.77
$\chi^2(6)$ (LS, NW)	25.38	28.76	28.13	28.59	28.84	26.95
$\chi^2(6)$ (rev. reg.)	9.14	12.73	13.15	13.48	13.63	14.22
pval	0.17	0.05	0.04	0.04	0.03	0.03

Table V: Predicting returns with the single forecasting factor

Panel A reports the estimates of equation (15). Rows denoted as “LS” give the full sample t-statistics and adjusted R^2 's. Rows denoted as “SS” summarize the small sample distributions of the statistics obtained with the block bootstrap. **Panel B** shows the predictability of individual bond returns with the single factor. Again, the full sample (LS) and small sample (SS) distributions are provided. Row “ $\Delta \bar{R}^2$ ” gives the difference in \bar{R}^2 values between the corresponding unconstrained predictive regressions using six cycles in Table IV and the regressions using the single factor, \widehat{cf}_t . HH denotes Hansen-Hodrick adjustment in standard errors, NW denotes the Newey-West adjustment. We use 12 and 15 lags, respectively. Bootstrapped t-statistics use the NW adjustment with 15 lags to ensure a positive definite covariance matrix in all bootstrap samples. To facilitate comparisons, in panel B all left- and right-hand variables have been standardized. Sample is 1971:11–2009:12.

Panel A. $\overline{rx}_{t+1}^{(n)} = \gamma_0 + \gamma_1 c_t^{(1)} + \gamma_2 \bar{c}_t + \bar{\varepsilon}_{t+1}$					
	$\hat{\gamma}_1$	t-stat (HH,NW)	$\hat{\gamma}_2$	t-stat (HH,NW)	\bar{R}^2
LS	-0.85	(-5.69, -6.52)	1.18	(9.99, 10.24)	0.54
SS (5%, 50%, 95%)		[-11.05, -7.52, -5.17]		[7.36, 10.70, 14.52]	[0.40, 0.56, 0.66]

Panel B. $rx_{t+1}^{(n)} = \beta_0 + \beta_1 \widehat{cf}_t + \varepsilon_{t+1}^{(n)}$, where $\widehat{cf}_t = \hat{\gamma}_0 + \hat{\gamma}_1 c_t^{(1)} + \hat{\gamma}_2 \bar{c}_t^{(n)}$						
Statistic	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
β	0.62	0.68	0.71	0.73	0.74	0.72
tstat (LS, HH)	5.77	8.66	9.37	10.21	10.07	9.12
tstat (LS, NW)	6.60	9.45	9.92	10.55	10.26	9.06
tstat (SS,5%)	4.58	6.37	6.95	7.35	7.50	6.89
tstat (SS,50%)	7.54	10.28	10.69	11.09	10.77	9.56
tstat (SS,95%)	12.80	15.40	14.95	15.17	14.24	12.70
R^2 (LS)	0.38	0.46	0.50	0.53	0.55	0.52
$\Delta \bar{R}^2$ (LS)	0.03	0.02	0.03	0.02	0.01	0.05
R^2 (SS, 5%)	0.22	0.30	0.35	0.39	0.41	0.39
\bar{R}^2 (SS, 50%)	0.40	0.48	0.52	0.55	0.56	0.53
\bar{R}^2 (SS, 95%)	0.55	0.61	0.63	0.65	0.66	0.63

Table VI: Return predictability with different measures of long-horizon inflation expectations

This table reports the R^2 's from predictive regression using different measures of the long-horizon inflation expectations. Because the survey samples are of different lengths, in panels A and B, the first row reports R^2 's when \widehat{cf}_t is constructed using τ^{CPI} , as in the previous sections. The second row reports R^2 's obtained when the alternative proxy for the persistent component from surveys or from realized inflation. The third row reports R^2 's for the CP factor. Panel C contains only results obtained using alternative realized CPI measures, as the sample length for \widehat{cf} and CP factors overlaps with the one reported in previous tables.

	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
Panel A: Ten-year ahead inflation expectations						
Livingston and BCEI, 1979–2009						
$\bar{R}^2 \widehat{cf}, \tau^{CPI}$	0.28	0.41	0.45	0.50	0.51	0.49
$\bar{R}^2 \widehat{cf}, \tau^{surv}$	0.43	0.43	0.47	0.47	0.48	0.48
$\bar{R}^2 CP$	0.12	0.13	0.15	0.16	0.18	0.19
Panel B: One-year ahead inflation expectations						
B1. GDP deflator, SPF, 1971–2009						
$\bar{R}^2 \widehat{cf}, \tau^{CPI}$	0.38	0.46	0.50	0.53	0.55	0.52
$\bar{R}^2 \widehat{cf}, \tau^{surv}$	0.29	0.26	0.30	0.31	0.32	0.30
$\bar{R}^2 CP$	0.19	0.20	0.23	0.23	0.26	0.30
B2. Median CPI, BCFF, 1984–2009						
$\bar{R}^2 \widehat{cf}, \tau^{CPI}$	0.26	0.43	0.50	0.54	0.55	0.51
$\bar{R}^2 \widehat{cf}, \tau^{surv}$	0.22	0.32	0.38	0.40	0.44	0.46
$\bar{R}^2 CP$	0.24	0.28	0.32	0.34	0.36	0.35
Panel C: Realized inflation						
C1. Current realized core CPI, 1971–2009						
$\bar{R}^2 \widehat{cf}, CPI$	0.19	0.18	0.21	0.23	0.25	0.23
C2. 12-month MA of realized core CPI, 1971–2009						
$\bar{R}^2 \widehat{cf}, CPI(MA12)$	0.19	0.20	0.23	0.26	0.28	0.27

Table VII: Marginal predictability of bonds excess returns by macro factors

The table reports predictive regressions of bond excess returns on the single return forecasting factor \widehat{cf}_t and eight macro factors proposed by Ludvigson and Ng (2009), $\widehat{F}_{1t}, \dots, \widehat{F}_{8t}$. “ $\Delta \bar{R}^2$ ” denotes the **gain** in adjusted R^2 from adding all eight macro factors to the predictive regression with \widehat{cf}_t . “ \bar{R}^2 (\widehat{F}_t only)” reports the adjusted R^2 values from regressing the excess returns on $\widehat{F}_t, \dots, \widehat{F}_{8t}$. Macro factors are constructed from 132 macroeconomic and financial series. The sample period is 1971:11–2007:12. Superscripts H, M, L at t-statistics indicate variables that are significant in the macro-only regression of rx on \widehat{F}_t at 1%, 5% and 10%, respectively. In parentheses, t-statistics use the Newey-West adjustment with 15 lags. All variables are standardized.

Macro factors: $rx_{t+1}^{(n)} = b_0 + b_1 \widehat{cf}_t + \mathbf{b}'_2 \widehat{\mathbf{F}}_t + \varepsilon_{t+1}^{(n)}$, sample 1971-2007						
Regressor	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
\widehat{cf}_t	0.52	0.60	0.63	0.67	0.69	0.68
	4.89	6.53	6.72	7.18	7.16	6.37
\widehat{F}_{1t} (real)	0.17	0.08	0.04	0.01	-0.01	0.00
	(1.66) ^M	(0.86)	(0.45)	(0.15)	(-0.15)	(-0.03)
\widehat{F}_{2t} (financial spreads)	0.05	0.07	0.09	0.09	0.09	0.06
	(1.03) ^M	(1.53) ^M	(2.00) ^H	(2.02) ^H	(1.91) ^H	(1.14) ^M
\widehat{F}_{3t} (inflation)	-0.03	-0.01	0.00	0.01	0.01	0.01
	(-1.58)	(-0.44)	(0.17)	(0.31)	(0.51)	(0.47)
\widehat{F}_{4t} (inflation)	-0.15	-0.06	-0.02	0.01	0.05	0.06
	(-2.12) ^H	(-0.84) ^M	(-0.23) ^M	(0.17) ^L	(0.64)	(0.77)
\widehat{F}_{5t}	0.07	0.02	0.01	0.00	-0.01	0.00
	(1.17)	(0.35)	(0.19)	(-0.03)	(-0.18)	(-0.01)
\widehat{F}_{6t} (monetary)	-0.09	-0.10	-0.10	-0.10	-0.10	-0.09
	(-1.00) ^H	(-1.06) ^H	(-1.05) ^H	(-1.17) ^H	(-1.14) ^H	(-0.92) ^H
\widehat{F}_{7t} (bank reserves)	-0.04	-0.08	-0.08	-0.08	-0.09	-0.09
	(-0.74) ^M	(-1.24) ^H	(-1.26) ^H	(-1.45) ^H	(-1.56) ^H	(-1.47) ^H
\widehat{F}_{8t} (stock market)	0.01	0.02	0.02	0.01	0.01	0.03
	(0.23)	(0.72) ^M	(0.66) ^M	(0.43) ^M	(0.52) ^H	(1.18) ^H
\bar{R}^2	0.45	0.49	0.52	0.55	0.57	0.54
R^2 (\widehat{F}_t only)	0.25	0.22	0.22	0.22	0.22	0.20
$\Delta \bar{R}^2 = \bar{R}^2 - \bar{R}^2(\widehat{cf}_t)$	0.05	0.02	0.01	0.02	0.02	0.02

Table VIII: Predicting actual bond portfolio returns

The table reports the predictive regressions of actual bond portfolio excess returns from CRSP on \widehat{cf}_t at different investment horizons, $rx_{t+h/12}$, $h = 1, 3, 6, 9, 12$ months. The column labels reflect the maturity bin for each bond portfolio: from below two years to above ten years. The sample is from 1972:01 (where return for the $> 120m$ portfolio become available) through 2010:12. We use log returns; returns are in excess of Tbill rates of corresponding maturity obtained from H.15 Fed tables. Tbill rates are converted from discount to continuous basis. T-statistics use Newey-West adjustment with 12 lags. All variables are standardized.

	12m < rx < 24m	rx < 36m	rx < 48m	rx < 60m	rx < 120m	rx > 120m
Panel A. $h = 1$ month						
\widehat{cf}_t	0.17 (3.60)	0.17 (3.85)	0.19 (4.29)	0.19 (4.60)	0.20 (4.86)	0.22 (5.62)
\bar{R}^2	0.03	0.03	0.03	0.03	0.04	0.05
Panel B. $h = 3$ months						
\widehat{cf}_t	0.30 (4.71)	0.31 (4.98)	0.33 (5.45)	0.34 (5.82)	0.36 (6.07)	0.38 (6.83)
\bar{R}^2	0.09	0.09	0.10	0.11	0.12	0.14
Panel C. $h = 6$ months						
\widehat{cf}_t	0.44 (5.70)	0.45 (6.05)	0.47 (6.56)	0.48 (7.03)	0.51 (7.43)	0.54 (8.36)
\bar{R}^2	0.19	0.20	0.22	0.23	0.25	0.29
Panel D. $h = 9$ months						
\widehat{cf}_t	0.54 (6.92)	0.56 (7.41)	0.58 (7.95)	0.59 (8.49)	0.62 (9.00)	0.67 (9.93)
\bar{R}^2	0.29	0.31	0.33	0.35	0.39	0.44
Panel E. $h = 12$ months						
\widehat{cf}_t	0.62 (7.47)	0.64 (8.38)	0.67 (9.09)	0.68 (9.92)	0.72 (10.31)	0.76 (11.60)
\bar{R}^2	0.39	0.41	0.44	0.46	0.51	0.58

Table IX: Out-of-sample tests

The table reports the results of out-of-sample tests for the period 1978–2009 (**panel A**), 1985–2009 (**panel B**), and 1995–2009 (**panel C**). Row (1) in each panel contains the ENC-NEW test. The null hypothesis is that the predictive regression with forward rates (restricted model) encompasses all predictability in bond excess returns. The null is tested against the alternative that cycles (unrestricted model) improve the predictability achieved by the forward rates. For forwards and cycles we use maturities of one, two, five, seven, ten and 20 years. Row (2) reports bootstrapped critical values (CV) for the ENC-NEW statistic at the 95% confidence level. Row (3) shows the ratio of mean squared errors for the unrestricted and restricted models, MSE_{cyc}/MSE_{fwd} . Rows (4), (5) and (6) report the out-of-sample R^2 , R_{OOS}^2 , defined in equation (41), for cycles, forwards and the yield curve slope, respectively. For forwards we use six maturities as above, for cycles we use $c_t^{(1)}$ and \bar{c}_t . The slope for predicting the bond return with maturity n is constructed as $f_t^{(n)} - y_t^{(1)}$. Implementation details for the out-of-sample tests are collected in Appendix F.

Test	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
Panel A. Out-of-sample period: 1978–2009						
(1) ENC-NEW	131.41	140.12	150.39	168.35	167.68	149.99
(2) Bootstrap 95% CV	74.77	65.22	64.04	61.39	61.72	63.27
(3) MSE_{cyc}/MSE_{fwd}	0.65	0.56	0.54	0.49	0.50	0.63
(4) R_{OOS}^2 cyc	0.18	0.29	0.34	0.40	0.39	0.31
(5) R_{OOS}^2 fwd	-0.25	-0.26	-0.22	-0.24	-0.23	-0.10
(6) R_{OOS}^2 slope	0.10	0.09	0.08	0.08	0.10	0.10
Panel B. Out-of-sample period: 1985–2009						
(1) ENC-NEW	108.16	111.04	121.87	137.93	138.71	127.53
(2) Bootstrap 95% CV	42.36	41.10	44.50	43.01	45.24	46.74
(3) MSE_{cyc}/MSE_{fwd}	0.59	0.52	0.51	0.49	0.51	0.59
(4) R_{OOS}^2 cyc	0.12	0.34	0.40	0.44	0.44	0.41
(5) R_{OOS}^2 fwd	-0.48	-0.26	-0.18	-0.14	-0.09	-0.00
(6) R_{OOS}^2 slope	0.07	0.07	0.06	0.09	0.10	0.11
Panel C. Out-of-sample period: 1995–2009						
(1) ENC-NEW	53.74	53.64	61.99	74.83	80.02	75.91
(2) Bootstrap 95% CV	20.35	22.01	24.12	24.57	29.71	39.26
(3) MSE_{cyc}/MSE_{fwd}	0.48	0.47	0.41	0.40	0.37	0.31
(4) R_{OOS}^2 cyc	-0.06	0.10	0.18	0.20	0.27	0.37
(5) R_{OOS}^2 fwd	-1.20	-0.92	-1.00	-1.00	-0.99	-1.06
(6) R_{OOS}^2 slope	-0.20	-0.06	-0.08	-0.04	-0.04	-0.05

A. Cointegration

In Section II.C, we invoke cointegration to argue that cycles should predict bond returns. This Appendix provides unit root tests for yields, τ_t^{CPI} and residuals from the cointegrating regression (9).

Table A-I reports values of the augmented Dickey-Fuller (ADF) test. We consider changes in respective variables up to lag 12 as indicated in the first column. Tests in panel A are specified with a constant since all series have nonzero mean. Tests in panel B are specified without a constant since the cointegration residuals are zero mean by construction. Each panel provides the corresponding critical values. Additionally, we also apply the Phillips-Perron test and find that it conforms very closely with the ADF test. Therefore, we omit these results for brevity. The tests indicate that: (i) we cannot reject the hypothesis that both yields and τ_t have a unit root, (ii) that cointegration residuals (cycles) are stationary.

Table A-I: Unit root test

Panel A reports values of the ADF test for τ_t^{CPI} and yields with different maturities. τ_t^{CPI} is specified in equation (8). In the last column, \bar{y}_t is the average of yields across maturities: $\bar{y}_t = \frac{1}{20} \sum_{i=1}^{20} y_t^{(i)}$. For all variables the test contains a constant since yields and τ_t are both nonzero mean. **Panel B** reports the values of the ADF test for the cointegrating residuals $c_t^{(n)}$ from the regression of $y_t^{(n)}$ on τ_t^{CPI} (the regression includes a constant). We specify the test without a constant since $c_t^{(n)}$ is zero mean by construction. \bar{c}_t in the last column is obtained as the residual from a regression of \bar{y}_t on τ_t^{CPI} . The null hypothesis states that a variable has a unit root. Corresponding critical values are reported separately in each panel.

Panel A. ADF test for τ_t^{CPI} and yields								
# lags	τ_t^{CPI}	$y_t^{(1)}$	$y_t^{(2)}$	$y_t^{(5)}$	$y_t^{(7)}$	$y_t^{(10)}$	$y_t^{(20)}$	\bar{y}_t
1	-2.75	-1.90	-1.68	-1.34	-1.09	-0.95	-1.14	-1.09
3	-1.15	-1.42	-1.25	-1.02	-0.90	-0.80	-0.90	-0.84
6	-1.07	-1.21	-1.11	-1.03	-0.95	-0.90	-1.19	-0.94
12	-0.87	-1.63	-1.54	-1.38	-1.27	-1.15	-1.31	-1.22
Critical values: -3.46 (1%), -2.87 (5%), -2.59 (10%)								
Panel B. ADF test for cointegrating residual								
# lags	$c_t^{(1)}$	$c_t^{(2)}$	$c_t^{(5)}$	$c_t^{(7)}$	$c_t^{(10)}$	$c_t^{(20)}$	\bar{c}_t	
1	-4.05	-4.22	-4.44	-4.28	-4.29	-4.30	-4.35	
3	-3.38	-3.58	-3.88	-3.93	-3.97	-3.81	-3.84	
6	-3.12	-3.39	-3.97	-4.10	-4.27	-4.58	-4.14	
12	-4.05	-4.41	-4.97	-5.09	-5.21	-5.10	-5.08	
Critical values: -2.58 (1%), -1.96 (5%), -1.63 (10%)								

B. Data

This section describes the construction of data series and compares bond excess returns obtained from different data sets: Gürkaynak, Sack, and Wright (2006, GSW), Fama-Bliss (FB) and constant maturity Treasury rates (CMT).

Interest rate data:

- *CMT rates.* We use constant maturity Treasury rates (CMT) compiled by the US Treasury, and available from the H.15 Fed’s statistical release. The maturities comprise one, two, three, five, seven, ten and 20 years. Our sample period is November 1971 through December 2009. The beginning of our sample coincides with the end of the Bretton Woods system in August 1971. This is also when the GSW data for long-term yields become available. Data on 20-year CMT yield are not available for the period from January 1987 through September 1993. We fill this gap by computing the monthly yield returns of the 30-year CMT yield and using them to write the 20-year CMT yield forward. To compute the zero curve, we treat CMT rates as par yields and apply the piecewise cubic Hermite polynomial.
- *Short maturity rate.* The three and six month T-bill rates are from the H.15 tables. We use secondary market quotes, and convert them from the discount to the continuously compounded basis.
- *Zero curve.* For comparison, we also use the GSW and Fama-Bliss zero yields. GSW data set is compiled by the Fed. The GSW data are available at Fed website.²⁹ Fama-Bliss data are obtained from the CRSP database.

Macroeconomic variables:

- *Inflation.* CPI for all urban consumers less food and energy (core CPI) is from Bureau of Labor Statistics, downloaded from the FRED database. We define core CPI inflation as the year-on-year simple growth rate in the core CPI index. We construct the cyclical component of inflation CPI_t^c as the difference between the core CPI inflation and permanent component τ_t^{CPI} computed according to equation (8).
- *Unemployment.* UNEMPL is the year-on-year log growth in the unemployment rate provided by the Bureau of Labor Statistics. The series is downloaded from the FRED database.

Financial variables:

- *Commercial paper spread.* Commercial paper spread is defined as the difference between the yield on a three-month commercial paper and the yield on a three-month T-bill.
- *Swap spread.* Swap spread is the difference between ten-year swap rate and the corresponding CMT yield.
- *Moody’s Baa spread.* Moody’s Baa spread is the difference between the Moody’s Baa corporate bond yield and the 30-year CMT yield. To compute the yield, Moody’s includes bonds with remaining maturities as close as possible to 30 years.
- *TED spread.* The TED spread is the difference between the three-month LIBOR and the yield on three-month Treasury bill.
- *T-bill3M spread.* T-bill3M spread is the difference between the three-month T-bill and the Fed funds target rate.
- *Fed funds rate.* The Federal funds denotes the monthly effective Fed funds rate. Monthly Fed funds rates are obtained as the average of daily values.

All financial data series are obtained from the FRED database, the only exception are the swap and LIBOR rates which are downloaded from Datastream.

²⁹<http://www.federalreserve.gov/econresdata/researchdata.htm>.

Survey data:

- *Blue Chip Financial Forecasts.* Blue Chip Financial Forecasts (BCFF) survey contains monthly forecasts of yields, inflation and GDP growth given by approximately 45 leading financial institutions. The BCFF is published on the first day of each month, but the survey itself is conducted over a two-day period, usually between the 23rd and 27th of each month. The exception is the survey for the January issue which generally takes place between the 17th and 20th of December. The forecasts are for the current quarter, the next quarter out to five quarters ahead. Inflation survey data four quarters ahead and are available for the period 1984:07–2009:12.
- *Livingston Survey.* Livingston survey was started in 1946, it covers the forecasts of economists from banks, government and academia. The survey contains semi-annual forecasts of key macro and financial variables such as inflation, industrial production, GDP, unemployment, housing starts, corporate profits and T-bills. It is conducted in June and December each year. The survey contains forecast out to ten years ahead for some variables. However, the inflation forecasts for ten years ahead start only in 1990.
- *Survey of Professional Forecasters (SPF).* Conducted quarterly; respondents provide estimates of the one- and ten-year inflation, among other variables. One-year inflation forecasts start in 1981:Q3, and the ten-year forecasts begin in 1991:Q4.
- *Long-horizon inflation expectations surveys.* Philadelphia Fed provides a combined file of ten-year inflation forecasts³⁰ compiled from Blue Chip Economic Indicators (BCEI) and from the Livingston survey, both of which are conducted semi-annually. The sample covers the period 1979:12–2009:12.

C. Basic expression for the long-term yield

It is straightforward to express an n -period yield as the expected sum of future short rates plus the term premium. For completeness, we briefly provide the argument. The price of an n -period nominal bond P_t^n satisfies:

$$P_t^{(n)} = E_t \left(M_{t+1} P_{t+1}^{(n-1)} \right), \quad (25)$$

where M_{t+1} is the nominal stochastic discount factor. Let lowercase letters $(m_t, p_t^{(n)})$ denote natural logarithms of the corresponding variables. Under conditional joint lognormality of M_{t+1} and the bond price, from (25) we obtain the recursion:

$$p_t^{(n)} = E_t \left(p_{t+1}^{(n-1)} + m_{t+1} \right) + \frac{1}{2} \text{Var}_t \left(p_{t+1}^{(n-1)} + m_{t+1} \right),$$

where r_t is the short rate: $r_t = y_t^{(1)}$. By recursive substitution, we can express $p_t^{(n)}$ as:

³⁰<http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/historical-data/inflation-forecasts.cfm>

$$\begin{aligned}
p_t^{(n)} &= -E_t(r_t + r_{t+1} + \dots + r_{t+n-1}) + E_t \left[\frac{1}{2} \text{Var}_t(p_{t+1}^{(n-1)}) + \text{Cov}_t(p_{t+1}^{(n-1)}, m_{t+1}) \right. \\
&\quad + \frac{1}{2} \text{Var}_{t+1}(p_{t+2}^{(n-2)}) + \text{Cov}_{t+1}(p_{t+2}^{(n-2)}, m_{t+2}) + \\
&\quad \left. \dots + \frac{1}{2} \text{Var}_{t+n-2}(p_{t+n-1}^{(1)}) + \text{Cov}_{t+n-2}(p_{t+n-1}^{(1)}, m_{t+n-1}) \right].
\end{aligned}$$

Let $rx_{t+1}^{(n)} = \ln \frac{P_{t+1}^{(n-1)}}{P_t^{(n)}} - r_t$ and $y_t^{(n)} = -\frac{1}{n} p_t^{(n)}$. For an n -maturity yield, since $E_t(rx_{t+1}^{(n)}) = -\text{Cov}_t(m_{t+1}, p_{t+1}^{(n-1)}) - \frac{1}{2} \text{Var}_t(p_{t+1}^{(n-1)})$, we obtain:

$$y_t^{(n)} = \frac{1}{n} E_t \left(\sum_{i=0}^{n-1} r_{t+i} \right) + \underbrace{\frac{1}{n} E_t \left(\sum_{i=0}^{n-2} rx_{t+i+1}^{(n-i)} \right)}_{:= rpy_t^{(n)}}. \quad (26)$$

D. Long-run inflation expectations: τ_t^{CPI} parameters

Our τ_t^{CPI} variable can be interpreted as an endpoint of inflation expectations, i.e. the local long-run mean to which current inflation expectations converge. In this appendix, we show how τ_t^{CPI} can be embedded within a simple model of the term structure of inflation expectations. To obtain the gain parameter v that is consistent with inflation forecasts we estimate the model using survey data for CPI.

D.1. Model of the term structure of inflation expectations

The model of the term structure of inflation expectations follows Kozicki and Tinsley (2006). Let the realized inflation CPI_t follow an AR(p) process, which we can write in a companion form as:

$$CPI_{t+1} = e_1' z_{t+1} \quad (27)$$

$$z_{t+1} = Cz_t + (I - C)\mathbf{1}\mu_\infty^{(t)} + e_1\varepsilon_{t+1}, \quad (28)$$

where $z_t = (CPI_t, CPI_{t-1}, \dots, CPI_{t-p+1})'$, $e_1 = (1, 0, \dots, 0)'$ with dimension $(p \times 1)$, $\mathbf{1}$ is a $(p \times 1)$ vector of ones, and companion matrix C is of the form

$$C = \begin{pmatrix} c_1 & c_2 & \dots & c_{p-1} & c_p \\ 1 & 0 & & 0 & 0 \\ 0 & 1 & & 0 & 0 \\ & & \ddots & & \vdots \\ 0 & \dots & 0 & 1 & 0 \end{pmatrix}. \quad (29)$$

Then,

$$CPI_{t+1} = e_1' Cz_t + e_1'(I - C)\mathbf{1}\mu_\infty^{(t)} + e_1'\varepsilon_{t+1}. \quad (30)$$

In the above specification, inflation converges to a time varying, rather than constant, long-run mean:

$$\mu_\infty^{(t)} = \lim_{k \rightarrow \infty} E_t(CPI_{t+k}), \quad (31)$$

which itself follows a random walk:

$$\mu_\infty^{(t+1)} = \mu_\infty^{(t)} + v_{t+1}. \quad (32)$$

The expected inflation j -months ahead is given as:

$$E_t(CPI_{t+j}) = e'_1 C^j z_t + e'_1 (I - C^j) \mathbf{1} \mu_\infty^{(t)}. \quad (33)$$

Thus, survey expectations can be expressed as:

$$s_{t,k} = \frac{1}{k} \sum_{j=1}^k E_t^s(CPI_{t+j}), \quad (34)$$

where the survey is specified as the average inflation over k periods, and E^s denotes the survey expectations. We treat the survey data as expected inflation plus a normally distributed measurement noise:

$$s_{t,k} = \frac{1}{k} \sum_{j=1}^k E_t(CPI_{t+j}) + \eta_{t,k}. \quad (35)$$

It is convenient to cast the model in a filtering framework with the state equation given as:

$$\mu_\infty^{(t+1)} = \mu_\infty^{(t)} + v_{t+1}, \quad v_{t+1} \sim N(0, Q), \quad (36)$$

and the measurement equation:

$$m_t = Az_t + H\mu_\infty^{(t)} + w_t, \quad w_t \sim N(0, R) \quad (37)$$

where $m_t = (CPI_{t+1}, s_{t,k_1}, s_{t,k_2}, \dots, s_{t,k_n})$ and $w_t = (\varepsilon_{t+1}, \eta_{t,k_1}, \eta_{t,k_2}, \dots, \eta_{t,k_n})'$, where k_i is the forecast horizon of a given survey i . We assume that the covariance matrix R is diagonal, and involves only two distinct parameters: (i) the variance of the realized inflation shock, and (ii) the variance of the measurement error for $s_{k,t}$, which is assumed identical across different surveys. From equations (33) and (34), A and H matrices in (37) have the form:

$$A = \begin{pmatrix} e'_1 C \\ e'_1 \frac{1}{k_1} \sum_{j=1}^{k_1} C^j \\ e'_1 \frac{1}{k_2} \sum_{j=1}^{k_2} C^j \\ \dots \\ e'_1 \frac{1}{k_n} \sum_{j=1}^{k_n} C^j \end{pmatrix}, \quad H = \begin{pmatrix} e'_1 (I - C) \mathbf{1} \\ e'_1 \left(I - \frac{1}{k_1} \sum_{j=1}^{k_1} C^j \right) \mathbf{1} \\ e'_1 \left(I - \frac{1}{k_2} \sum_{j=1}^{k_2} C^j \right) \mathbf{1} \\ \dots \\ e'_1 \left(I - \frac{1}{k_n} \sum_{j=1}^{k_n} C^j \right) \mathbf{1} \end{pmatrix}. \quad (38)$$

We consider two versions of the endpoint process, $\mu_\infty^{(t)}$. In the first version, we treat $\mu_\infty^{(t)}$ as a random walk as in equation (32). We estimate the model by maximum likelihood combined with the standard Kalman filtering of the latent state, $\mu_\infty^{(t)}$.

In the second version, we obtain the endpoint as the discounted moving average of past inflation, as we do in the body of the paper, i.e.

$$\mu_\infty^{(t)} := \tau_t^{CPI}(v, N). \quad (39)$$

In the expression above, we explicitly stress the dependence of τ_t^{CPI} on the parameters. This case allows us to infer the gain parameter v that is consistent with the available survey data. We estimate the model with maximum likelihood. Since v and N are not separately identified, we fix the window size at $N = 120$ months, and estimate the v parameter for this window size. In the subsequent section, we provide an extensive sensitivity analysis of the predictive results for bond returns to both v and N parameters.

D.2. Data

We combine two inflation surveys compiled by the Philadelphia Fed that provide a long history of data and cover different forecast horizons:

- Livingston survey: Conducted bi-annually in June and December; respondents provide forecasts of the CPI level six and twelve months ahead. Following Kozicki and Tinsley (2006) and Carlson (1977) we convert the surveys into eight- and 14-month forecasts to account for the real time information set of investors. We use data starting from 1955:06.
- Survey of Professional Forecasters: Conducted quarterly; respondents provide estimates of the one- and ten-year inflation. One-year forecasts start in 1981:Q3, and the ten-year forecasts begin in 1991:Q4.

We use the median survey response. We match the data with the realized CPI (all items) because it underlies the surveys. The estimation covers the period from 1957:12–2010:12. In the adaptive learning version of the model, we use data from 1948:01 to obtain the first estimate of τ_t^{CPI} .

D.3. Estimation results

We estimate the model assuming an AR(12) structure for inflation. In the adaptive learning version of the model, we estimate the gain parameter at $v = 0.9868$. The BHHH standard error of 0.0025 suggest that v is highly significant. Its value implies that when forming their long-run inflation expectations, each month agents attach the weight of about 1.3% to the current inflation. Other parameters are not reported for brevity. Panel *a* of Figure D-1 displays the realized inflation and the estimates of its long-run expectations considering the two specifications. The series labeled as “random walk” is the filtered $\mu_\infty^{(t)}$ state. The series marked as “adaptive” shows discounted moving average of inflation τ_t^{CPI} constructed at the estimated parameter $v = 0.9686$, and assuming $N = 120$ months. Panel *b* of Figure D-1 plots the survey data used in estimation. Comparing the filtered series in panel *a*, we note that while both estimates trace each other closely, the random walk specification points to a faster downward adjustment in inflation expectation during the disinflation period compared to the adaptive learning proxy. This finding is intuitive in that in the first part of the sample until early 1990s, the long-horizon survey information is unavailable. Thus, the filtered inflation endpoint $\mu_\infty^{(t)}$ is tilted towards the realized inflation, and short horizon survey forecasts.

E. Small sample standard errors

We use the block bootstrap (e.g., Künsch, 1989) to assess the small sample properties of the test statistics and to account for the uncertainty about $c_t^{(n)}$. This appendix provides the

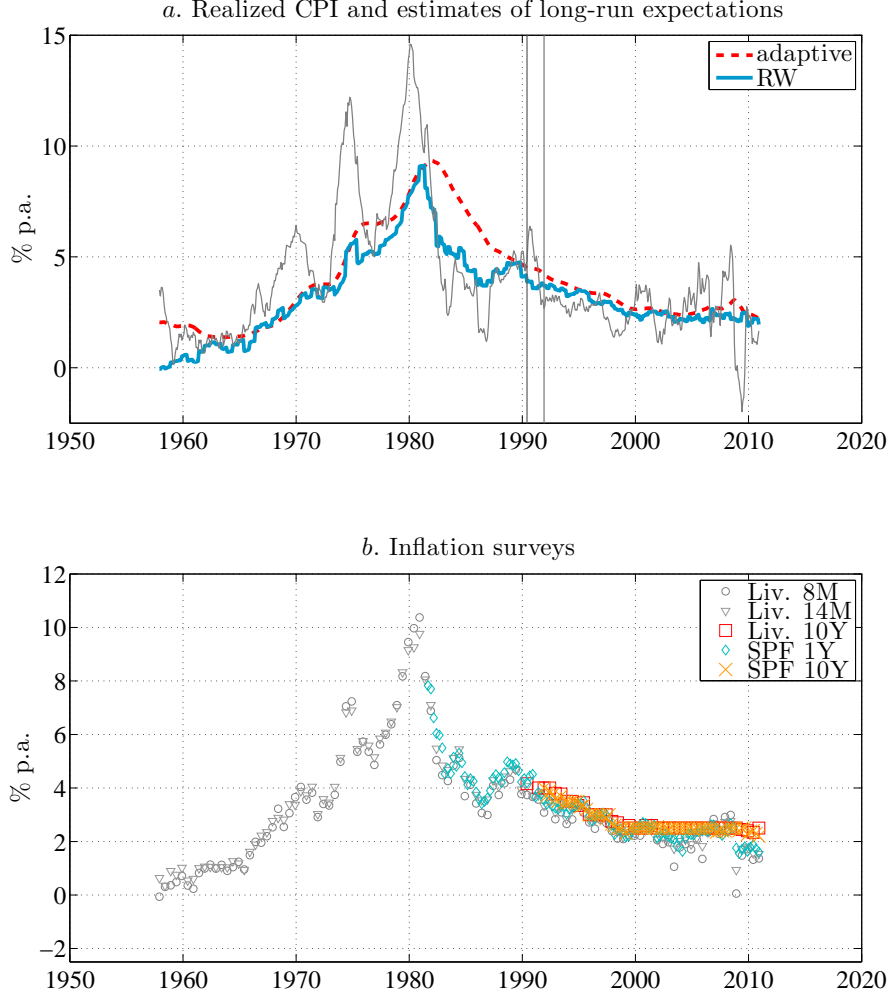


Figure D-1: Long-horizon inflation expectations

Panel *a* shows the realized CPI and the estimated long-run inflation expectations specified as a random walk and discounted moving average of past CPI data (adaptive). Vertical lines mark the dates on which 10-year inflation forecasts from Livingston and SPF surveys become available, respectively. Panel *b* shows the CPI surveys used in estimation.

details of the bootstrap procedure for regressions reported in Table V, which use the single factor to forecast individual bond returns. Small sample inference in other regressions is analogous.

The estimation consists of the steps as described in Section III.B, i.e:

Step 1. Project yields on the persistent component τ_t^{CPI} to obtain the cycles, $c_t^{(n)}$.

Step 2. Construct the single forecasting factor, \widehat{cf}_t from the cycles as in equation (16).

Step 3. Forecast individual excess bond returns with \widehat{cf}_t .

Let Z be a $T \times p$ data matrix with the t -th row: $Z_t = (\mathbf{y}_t', \tau_t^{CPI}, rx_{t+1}^{(i)})'$, and $\mathbf{y}_t = (y_t^{(1)}, y_t^{(2)}, \dots, y_t^{(m)})'$.

We split Z into blocks of size $bs \times p$, where $bs = \sqrt{T}$ ($bs = 21$ for the 1971–2009 sample).

Specifically, we create $(T - bs + 1)$ overlapping blocks consisting of observations: $(1, \dots, bs)$, $(2, \dots, bs + 1)$, \dots , $(T - bs + 1, \dots, T)$. In each bootstrap iteration, we select T/bs blocks with replacement, out of which we reconstruct the sample in the order the blocks were chosen. We perform steps 1 through 3 on the newly created sample, store the coefficients, t-statistics and adjusted R^2 values. For the statistics of interest, we approximate the empirical distribution using 1000 bootstrap repetitions, and obtain its 5% and 95% percentile values.

F. Out-of-sample tests

Below we describe the implementation of the bootstrap procedure to obtain the critical values for the ENC-NEW test. The test statistic for maturity n is given by:

$$\text{ENC-NEW}^{(n)} = (T - h + 1) \frac{\sum_{t=1}^T \left(u_{t+12}^{2,(n)} - u_{t+12}^{(n)} \varepsilon_{t+12}^{(n)} \right)}{\sum_{t=1}^T \varepsilon_{t+12}^{2,(n)}}, \quad (40)$$

where T is the number of observations in the sample, $\varepsilon_t^{(n)}$ and $u_t^{(n)}$ denote the prediction error from the unrestricted and restricted model, respectively, and h measures the forecast horizon, in our case $h = 12$ months. Note that the time step in (40) is expressed in months.

Our implementation of bootstrap follows Clark and McCracken (2005) and Goyal and Welch (2008). To describe the dynamics of yields and to obtain shocks to the state variables generating them, we assume that the yield curve is described by four principal components following a VAR(1). Persistent component τ_t is assumed to follow an AR(12) process. We account for the overlap in bond excess returns by implementing an MA(12) structure of errors in the predictive regression. Imposing the null of predictability by the linear combination of forward rates, we estimate the predictive regression, the VAR(1) for yield factors and VAR(12) for τ_t by OLS on the full sample. We store the estimated parameters and use the residuals as shocks to state variables for the resampling. We sample with replacement from residuals and apply the estimated model parameters to construct the bootstrapped yield curve, the persistent component and bond excess returns. To start each series, we pick a random date and take the corresponding number of previous observations to obtain the initial bootstrap observation. In our case, the maximum lag equals 12, hence we effectively sample from $T - 12$ observations. We construct 1000 bootstrapped series, run the out-of-sample prediction exercise and compute the ENC-NEW statistic for each of the constructed series. We repeat this scheme for different maturities. The critical value is the 95-th percentile of the bootstrapped ENC-NEW statistics.

The out-of-sample R^2 proposed by Campbell and Thompson (2008) is defined as:

$$R_{\text{OOS}}^{2,(n)} = 1 - \frac{\sum_{t=1}^{T-12} \left(rx_{t+12}^{(n)} - \widehat{rx}_{cyc,t+12}^{(n)} \right)^2}{\sum_{t=1}^{T-12} \left(rx_{t+12}^{(n)} - \overline{rx}_{t+12}^{(n)} \right)^2}, \quad (41)$$

where the time step t and sample size T is expressed in months. $\widehat{rx}_{cyc,t+12}^{(n)}$ is the forecast of annual excess return based on time t cycles, where the parameters of the predictive model are estimated

using cycles up to time $t - 12$ and returns realized up to time t . $\overline{r\bar{x}}_{t+12}^{(n)}$ is the return forecast using historical average excess return estimated through time t .

Online appendix for Understanding bond risk premia

— Not intended for publication —

ANNA CIESLAK and PAVOL POVALA

December 24, 2011

A. Comparison of excess returns from different data sets

Realized bond excess returns are commonly defined on zero coupon bonds. Since the computation of returns can be sensitive to the interpolation method, we compare returns obtained from CMTs to those from the GSW and FB data. Table I presents the regressions of one-year holding period CMT excess returns on their GSW and FB counterparts with matching maturities. Figure 1 additionally graphs selected maturities. Excess returns line up very closely across alternative data sets. The R^2 's from regressions of CMT excess returns on GSW and FB consistently exceed 99%, except for the ten-year bond for which the R^2 drops to 98% due to one data point in the early part of the sample (1975). Beta coefficients are not economically different from one. We conclude that any factor that aims to explain important features of excess bond returns shall perform similarly well irrespective of the data set used. Therefore, our key results are not driven by the choice of the CMT data.

Table I: Comparison of one-year holding period excess returns: CMT, GSW and FB data

The table reports β 's and R^2 's from regressions of excess returns constructed from CMT data on GSW (**panel A**) and FB (**panel B**) counterparts. We consider a monthly sample 1971:11–2009:12 with maturities from two to ten (five) years for GSW (FB) data. Excess returns are defined over a one-year holding period.

	$rx^{(2)}$	$rx^{(3)}$	$rx^{(4)}$	$rx^{(5)}$	$rx^{(6)}$	$rx^{(7)}$	$rx^{(8)}$	$rx^{(9)}$	$rx^{(10)}$
Panel A. Regressions of rx from CMT on GSW									
β	1.04	1.03	1.04	1.05	1.06	1.05	1.04	1.04	1.04
R^2	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.98
Panel B. Regressions of rx from CMT on FB									
β	1.04	1.01	1.02	1.05	—	—	—	—	—
R^2	0.99	0.99	0.99	0.99	—	—	—	—	—

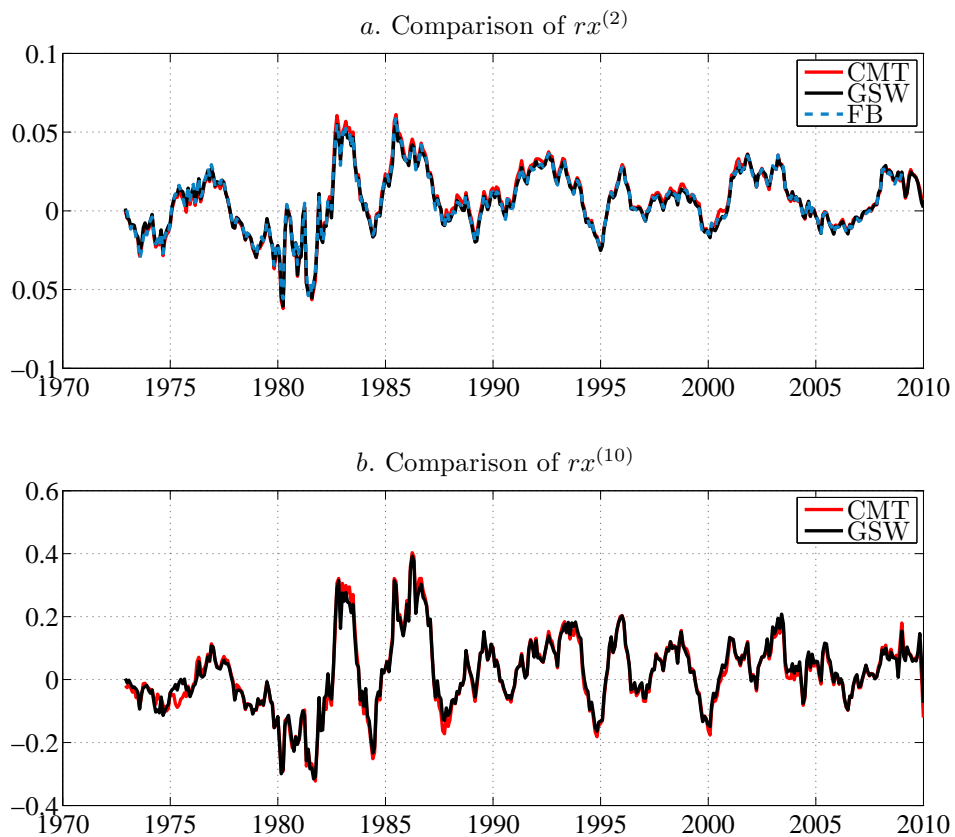


Figure 1: Comparison of realized excess returns across data sets

The figure plots one-year holding period returns on zero bonds constructed from three data sets: CMT, GSW and FB over the period 1971:11–2009:12. Upper panel provides a comparison for the excess returns on a two-year bond, the bottom panel compares the excess returns on the ten-year bond.

B. Constructing the single factor

This appendix introduces alternative approaches to constructing the single factor discussed in Section III.B.

B.1. Exploiting no information about future returns

The baseline construction of the return forecasting factor in equation (15) uses $\overline{r\bar{x}}$ on the left-hand side. We now use only time- t variables to obtain \widehat{cf}_t , i.e. unlike the CP factor, this way of obtaining the return forecasting factor does not involve future returns. We can estimate the return forecasting factor as the residual from the regression of \bar{c}_t onto $c_t^{(1)}$:

$$\bar{c}_t = \gamma_1 c_t^{(1)} + \bar{\varepsilon}_{t+1}, \quad (42)$$

$$\widehat{cf}_t^{CS} = \bar{c}_t - \hat{\gamma}_1 c_t^{(1)}, \quad (43)$$

and $\bar{c}_t = \frac{1}{m-1} \sum_{i=2}^m c_t^{(i)}$. The superscript ‘‘CS’’ indicates that we provide only cross sectional information to construct the forecasting factor.

B.2. One-step NLS estimation

We form a single factor as a linear combination of c_t 's:

$$\widehat{cf}_t^{NLS} = \lambda' \mathbf{c}_t, \quad (44)$$

and estimate the restricted system:

$$\mathbf{r}\mathbf{x}_{t+1} = A \begin{pmatrix} 1 \\ \lambda' \mathbf{c}_t \end{pmatrix} + \varepsilon_{t+1}, \quad (45)$$

where $\mathbf{r}\mathbf{x}_{t+1}$ is a $(m-1) \times 1$ vector of individual returns with maturities from two to m years, $\mathbf{r}\mathbf{x}_{t+1} = (rx_{t+1}^{(2)}, rx_{t+1}^{(3)}, \dots, rx_{t+1}^{(m)})'$, \mathbf{c}_t is a vector of cycles, and A is a matrix parameters:

$$A = \begin{pmatrix} \alpha_0^{(2)} & \alpha_1^{(2)} \\ \alpha_0^{(3)} & \alpha_1^{(3)} \\ \vdots & \vdots \\ \alpha_0^{(m)} & \alpha_1^{(m)} \end{pmatrix}. \quad (46)$$

We perform non-linear least squares (NLS) estimation, by minimizing the sum of squared errors:

$$(\hat{A}, \hat{\lambda}) = \min_{A, \lambda} \sum_{t=1}^T \left(\mathbf{r}\mathbf{x}_{t+1} - A \begin{pmatrix} 1 \\ \lambda' \mathbf{c}_t \end{pmatrix} \right)' \left(\mathbf{r}\mathbf{x}_{t+1} - A \begin{pmatrix} 1 \\ \lambda' \mathbf{c}_t \end{pmatrix} \right). \quad (47)$$

For identification, we set $\alpha_1^{(7)} = 1$. This choice is without loss of generality. The loss function (47) is minimized iteratively until its values are not changing between subsequent iterations. In application, being interested in the dynamics of the single factor \widehat{cf}_t^{NLS} , we additionally standardize excess returns cycles prior to estimation.

B.3. Common factor by eigenvalue decomposition

Alternatively, in constructing the single factor we can exploit the regression of an individual excess return on $c_t^{(1)}$ and the cycle of the corresponding maturity:

$$rx_{t+1}^{(n)} = \alpha_0^{(n)} + \alpha_1^{(n)} c_t^{(1)} + \alpha_2^{(n)} c_t^{(n)} + \varepsilon_{t+1}^{(n)}. \quad (48)$$

We form a vector \mathbf{erx}_t of expected excess returns obtained from this model:

$$\mathbf{erx}_t = E_t \left(rx_{t+1}^{(2)}, rx_{t+1}^{(3)}, \dots, rx_{t+1}^{(m)} \right)'. \quad (49)$$

The single factor is obtained as the first principal component of the covariance matrix of \mathbf{erx}_t :

$$\widehat{cf}_t^{PC} = U'_{(:,1)} \mathbf{erx}_t, \quad (50)$$

where $Cov(\mathbf{erx}_t) = ULU'$, and $U_{(:,1)}$ denotes the eigenvector associated with the largest eigenvalue in L . Using returns from two to 20 years, the first principal component explains 94% of common variation in \mathbf{erx}_t .

B.4. Comparing the results

We compare the single factor obtained with the different procedures. To distinguish between approaches, we use the notation: \widehat{cf}_t^{TS} for the construction involving future returns, \widehat{cf}_t^{NLS} for the one-step NLS estimation, \widehat{cf}_t^{PC} for the factor obtained with the eigenvalue decomposition of expected returns, and \widehat{cf}_t for the simple approach introduced in the body of the paper in Section III.B.

First, panel A of Table II presents the correlations among the four measures. Clearly, while the methods differ, they all identify virtually the same dynamics of the single factor. The correlation between the constructed factors reaches 98% or more.

Second, the way we obtain the single factor is inconsequential for the predictability we report. As a summary, panel B of Table II displays the adjusted R^2 values obtained by regressing individual excess returns on \widehat{cf}_t^{NLS} , \widehat{cf}_t^{PC} and \widehat{cf}_t , respectively. The difference between the three measures is negligible.

Table II: Comparing alternative constructions of the single factor

The table reports correlations between alternative approaches to constructing the single forecasting factor (**panel A**), as well as \bar{R}^2 values for predictability of individual excess returns (**panel B**). \widehat{cf}_t is used in the body of the paper, and defined in equation (16); \widehat{cf}_t^{CS} differs from the baseline specification in that it involves no information about future returns as discussed in Section B.1; \widehat{cf}_t^{PC} is obtained from the eigenvalue decomposition of expected excess returns in Section B.3 of this Appendix; \widehat{cf}_t^{NLS} is obtained in a one-step estimation in Section B.2.

Panel A. Correlations				
	\widehat{cf}_t	\widehat{cf}_t^{TS}	\widehat{cf}_t^{PC}	\widehat{cf}_t^{NLS}
\widehat{cf}_t	1	0.999	0.998	0.979
\widehat{cf}_t^{CS}	·	1	0.999	0.979
\widehat{cf}_t^{PC}	·	·	1	0.986
\widehat{cf}_t^{NLS}	·	·	·	1

Panel B. \bar{R}^2 from predictive regressions						
	$rx^{(2)}$	$rx^{(5)}$	$rx^{(7)}$	$rx^{(10)}$	$rx^{(15)}$	$rx^{(20)}$
\widehat{cf}_t	0.38	0.46	0.50	0.53	0.55	0.52
\widehat{cf}_t^{CS}	0.38	0.46	0.50	0.53	0.55	0.52
\widehat{cf}_t^{PC}	0.39	0.47	0.51	0.54	0.55	0.52
\widehat{cf}_t^{NLS}	0.40	0.48	0.52	0.55	0.56	0.53

C. Sensitivity of predictive results to τ_t^{CPI}

We analyze the sensitivity of our predictive results towards the specification of the persistent component τ_t^{CPI} . One concern is that these results could be highly dependent on the weighting scheme (v parameter) or the length of the moving average window (N) used to construct τ_t . For this reason, in Figure 2 we plot in-sample and out-of-sample R^2 's varying v between 0.975 and 0.995 and N between 100 and 150 months. While the predictability is stable across a wide range of parameter combinations, it weakens for values of v approaching one and for long window sizes. The deterioration in this region of the parameter space is intuitive: When combined with a long moving window, v close to one oversmooths the CPI data and leads to a less local estimate of the persistent component.³¹

In Figure 3, we use a simple moving average of past core CPI to isolate how the in- and out-of-sample predictive results depend on the window size, N . We consider N between 10 and 150 months. The predictive results are relatively stable for windows between 40 and 100 months, and taper off at the extremes. A very short moving window tilts τ_t^{CPI} to current realized inflation, a very long window, in turn, oversmooths the data. Both cases provide a poor measurement of the current long-run inflation mean, thus the predictability of bond returns weakens.

Beside the core CPI, Figure 3 considers two alternative variables used in the literature to construct proxies of the local mean reversion in interest rates: (i) the effective fed funds rate, and (ii) the one-year yield. For usual and economically plausible window sizes, neither alternative delivers predictability of bond returns at the level documented with the CPI. The question that underlies the difference in predictability is how each variable captures the persistent movement in interest rates. Using the moving average of the short rate one is faced with a tradeoff between smoothing the business cycle frequency in the short rate and contemporaneously measuring the generational inflation factor. Apart from the statistical fit, the advantage of τ_t^{CPI} lies in its direct link to an economic quantity, rather than to bond prices themselves. The benefit of an economic interpretation is also revealed in the estimates of a simple Taylor rule that we entertain in the Introduction to this paper. Indeed, considering $r_t = \gamma_0 + \gamma_c CPI_t^c + \gamma_y UNEMPL_t + \gamma_\tau \tau_t^i + \varepsilon_t$ with different proxies for $\tau_t, i = \{CPI, FFR\}$ shows that it is the CPI that provides highly stable coefficients across different subsamples (not reported).

³¹As $v \rightarrow 1$, the discounted moving average converges to a simple moving average with the corresponding window.

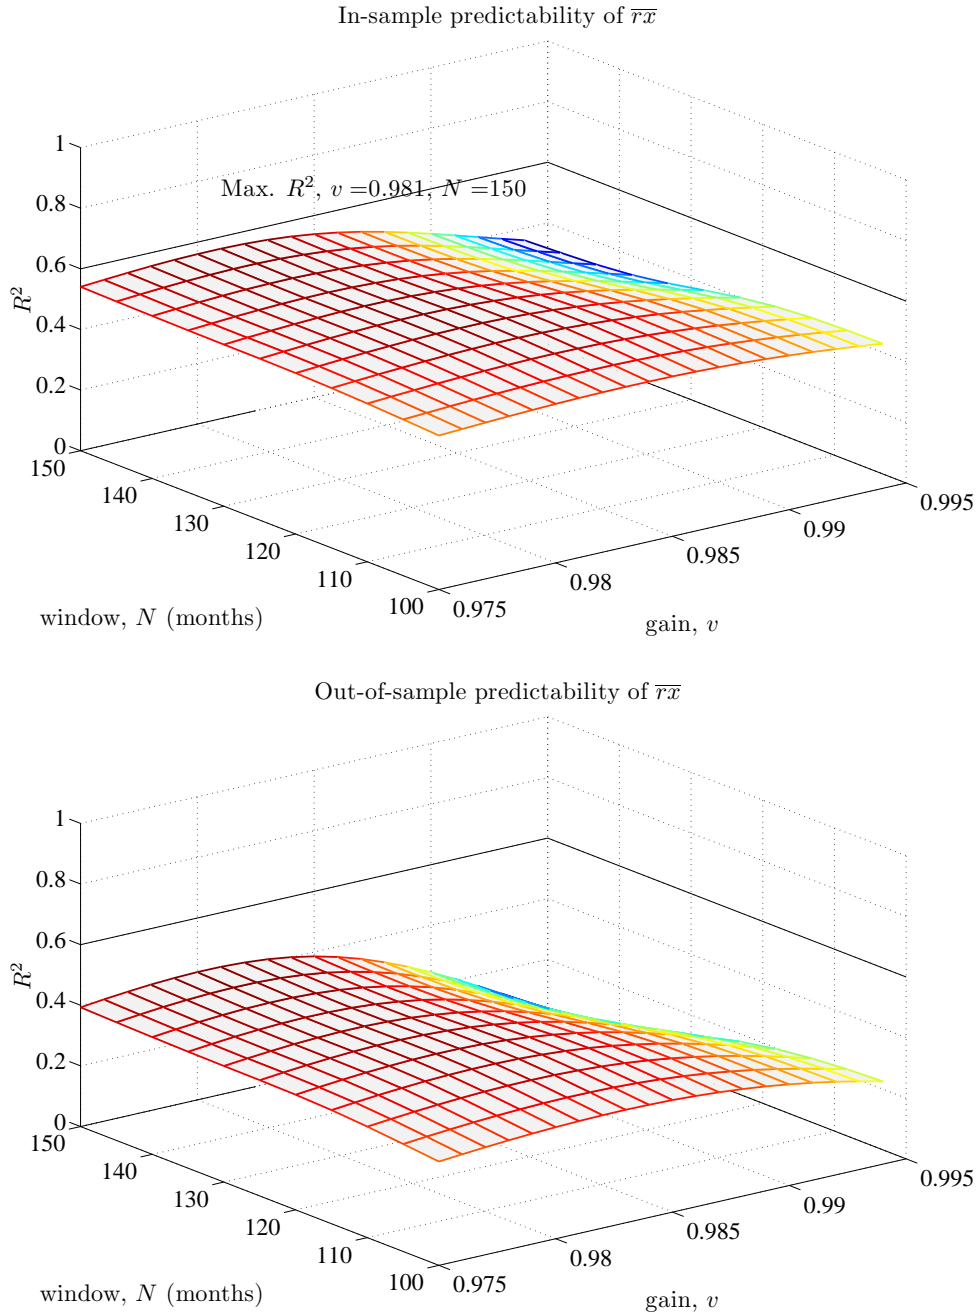


Figure 2: Sensitivity of the predictability evidence to v and N parameters

The figure studies the sensitivities of predictive \bar{R}^2 's to the values of N and v parameters used to construct τ_t^{CPI} . We predict the average bond return across maturities by regressing it on $c_t^{(1)}$ and \bar{c}_t as in the body of the paper. We consider the gain parameter v between 0.975 and 0.995, and the window size between 100 and 150 months. Panel *a* and *b* correspond to in-sample and out-of-sample results, respectively.

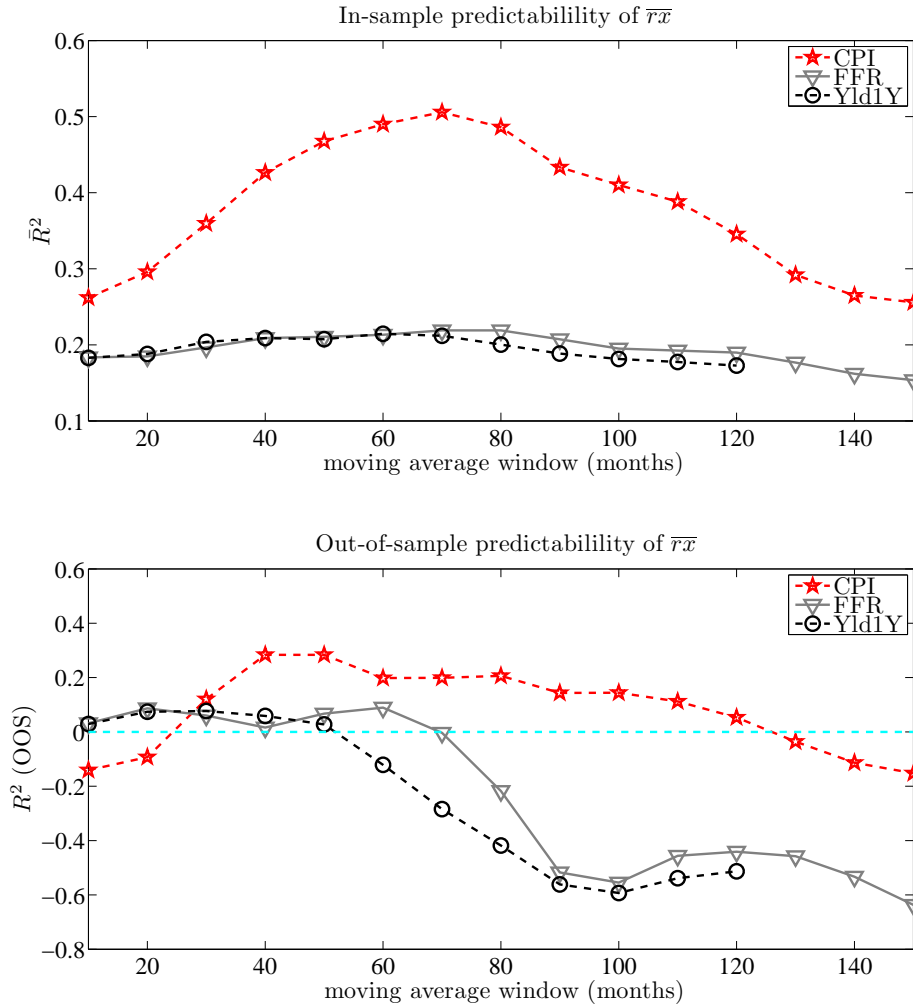


Figure 3: Sensitivity of predictability evidence to the window size

The figure depicts the sensitivity of the predictive results, in- and out-of-sample, to the length of the moving average window. All results are based on a simple (i.e. undiscounted) moving average. The window varies from ten to 150 months. We consider predictability of \bar{r}_x , and use three different proxies for the persistent component using moving average of: (i) past core CPI, (ii) past fed funds rate and (iii) past one-year rate.

D. Different datasets

This appendix, Table III, reports details of the predictive regressions reported in Section VII.D.

Table III: Comparing predictive R^2 in different data sets

The table compares the predictive adjusted R^2 's for three different zero curves obtained from: Fama-Bliss (FB), Gürkaynak, Sack, and Wright (2006, GSW), and Treasury constant maturity (CMT) rates. The dependent variable is:

$$\overline{rx}_{t+1} = \frac{1}{4} \sum_{i=2}^5 rx_{t+1}^{(i)}, \quad (51)$$

and regressors are indicated in the first column. In both panels, row (1) uses two cycles with maturity one and five years, row (2): five cycles with maturities from one through five years, row (3): two yields with maturity one and five years, row (4): five forward rates with maturity one through five years, row (5): spread between five- and one-year cycle, row (6): spread between five- and one-year yield. The column “sample” provides adjusted R^2 's for each regression; the column “bootstrap” gives the 5%, 50% and 95% percentile values for the adjusted R^2 's obtained with the block bootstrap (Appendix E).

Regressor	CMT		GSW		FB	
	Sample	Bootstrap	Sample	Bootstrap	Sample	Bootstrap
Panel A. Pre-crisis: 1971–2006						
(1) $c^{(1)}, c^{(5)}$	0.53	[0.39, 0.53, 0.66]	0.53	[0.40, 0.54, 0.66]	0.51	[0.38, 0.51, 0.64]
(2) $c^{(1)}, \dots, c^{(5)}$	0.54	[0.42, 0.55, 0.68]	0.53	[0.42, 0.55, 0.67]	0.56	[0.44, 0.57, 0.69]
(3) $y^{(1)}, y^{(5)}$	0.22	[0.10, 0.25, 0.44]	0.19	[0.08, 0.23, 0.42]	0.19	[0.08, 0.22, 0.42]
(4) $f^{(1)}, \dots, f^{(5)}$	0.27	[0.18, 0.32, 0.49]	0.21	[0.12, 0.27, 0.46]	0.30	[0.21, 0.34, 0.49]
(5) $c^{(5)} - c^{(1)}$	0.13	[0.02, 0.13, 0.29]	0.11	[0.01, 0.12, 0.27]	0.11	[0.01, 0.12, 0.27]
(6) $y^{(5)} - y^{(1)}$	0.13	[0.02, 0.14, 0.30]	0.11	[0.01, 0.12, 0.29]	0.11	[0.01, 0.12, 0.28]
Panel B. Post crisis: 1971–2009						
(1) $c^{(1)}, c^{(5)}$	0.47	[0.32, 0.49, 0.62]	0.46	[0.31, 0.49, 0.61]	0.44	[0.30, 0.47, 0.60]
(2) $c^{(1)}, \dots, c^{(5)}$	0.47	[0.34, 0.51, 0.65]	0.47	[0.34, 0.50, 0.63]	0.48	[0.35, 0.53, 0.65]
(3) $y^{(1)}, y^{(5)}$	0.15	[0.05, 0.19, 0.37]	0.13	[0.04, 0.17, 0.35]	0.13	[0.04, 0.16, 0.35]
(4) $f^{(1)}, \dots, f^{(5)}$	0.17	[0.10, 0.25, 0.42]	0.14	[0.08, 0.21, 0.38]	0.21	[0.13, 0.27, 0.43]
(5) $c^{(5)} - c^{(1)}$	0.11	[0.00, 0.10, 0.25]	0.09	[0.00, 0.09, 0.23]	0.09	[0.00, 0.09, 0.23]
(6) $y^{(5)} - y^{(1)}$	0.11	[0.00, 0.11, 0.26]	0.09	[0.00, 0.10, 0.25]	0.09	[0.00, 0.09, 0.24]