Internet Appendix for

"Expected Inflation and Other Determinants of Treasury Yields"

GREGORY R. DUFFEE*

* Citation format: Duffee, Gregory R., Internet Appendix for "Expected Inflation and Other Determinants of Treasury Yields," *Journal of Finance* [DOI STRING]. Please note: Wiley-Blackwell is not responsible for the content or functionality of any supporting information supplied by the author. Any queries (other than missing material) should be directed to the author of the article.

This appendix reports details, additional empirical analysis, and some discussion of the evidence in "Expected Inflation and Other Determinants of Treasury Yields." Section I contains details about measures of expected inflation. Section II discusses whether sticky information accounts for the observed low inflation variance ratios. Sections III and IV contain details about estimates of short-horizon and long-horizon inflation variance ratios, respectively. Section V presents additional results about inflation variance ratios in the model of Bansal and Shaliastovich (2013). Section VI presents additional results about inflation variance ratios in the models of Wachter (2006) and Ermolov (2015). Section VII describes the dynamic model in detail. Finally, Section VII discusses how news about expected inflation affects the term structure's level, slope, and curvature.

I. Details about Measuring Inflation Expectations

A. Point-in-Time and Average Price Levels

As described in the main text, this research uses inflation forecasts from two Blue Chip (BC) surveys and the Survey of Professional Forecasters (SPF). Respondents to the BC Economic Indicators Survey (EI) and the BC Financial Forecasts Survey (FF) forecast CPI inflation. Respondents to the SPF forecast the GDP price index.

Because of the difference between point-in-time price levels and average-over-time price levels, survey forecasts do not match up exactly with expectations in the accounting decomposition of the text's Section I.A. In the decomposition, expected inflation over the life of a bond is the expected log change in the price level from the date when a bond's yield is observed to the date when the bond matures. (This change is divided by the bond's maturity.) This is a log change in point-to-point prices. Forecasting in practice focuses on time-averaged inflation. For example, the SPF tracks forecasts of the GDP Price Index in future calendar quarters. This price index is an average of prices in the quarter. The BC survey has forecasts of the quarter-to-quarter percentage change in the CPI, where quarterly CPI is defined as the average of the three monthly CPI values in the quarter.

Figure IA1 contrasts the theory and practice. At date T_1 in 1997Q2, agents make predictions about the one-year bond yield at T_2, a date in 1997Q3 that is precisely three months after T_1. At T_1 agents also make predictions about average inflation during the life of the bond. The bond matures at T_3 in 1998Q3, precisely one year after T_2. Therefore, at T_1 agents forecast the log change in prices from T_2 to T_3. This time span is indicated in the figure with a line from T_2 to T_3. At T_2, the bond yield is realized and agents have updated predictions about inflation during the life of the bond. The difference between the inflation predictions at T_2 and T_1 is the news about expected inflation, as defined in the text's equation (4).

Researchers do not observe these T_2-to-T_3 inflation forecasts. In practice, surveys taken at dates T_1 and T_2 ask for predictions of quarter-to-quarter inflation for calendar quarters 1997Q4 through 1998Q3. Combining these calendar-quarter predictions creates forecasts of the log change from average prices during 1997Q3 to average prices during 1998Q3. The line in the figure labeled "Practice" connects the average of 1997Q3 to the average of 1998Q3.

Inflation news as defined by the theory will differ from inflation news as measured using surveys. The primary reason is that none of the point-to-point change in prices has been realized at T_2, but part of the average-to-average change *has* been realized at T_2. Therefore, investors at T_2 know more about the average-to-average measure of inflation than they do about the point-to-point measure.

A stark example is helpful. Imagine that expected inflation is always 5% per year. Variability in realized inflation is due entirely to permanent shocks to prices that are completely unforecastable. As a result, there is no news revealed between T_1 and T_2 about the expected log price change from T_2 to T_3. However, at date T_2 agents know the price-level shocks that have been realized since the beginning of 1997Q3. If these shocks are, on average, positive (negative), agents will predict an inflation rate from 1997Q3 to 1998Q3 that is greater (less) than 5%.

This effect of partially realized inflation raises the measured volatility of news about expected inflation relative to the true volatility. However, the effect is unlikely to be large even for short horizons. At long horizons it is even less important, since the distortion arises for only part of the first quarter of inflation news.

B. An Example of Constructing Short-Horizon Inflation Expectation Shocks

Panel A of Table IAI illustrates the construction of inflation news using BC-FF survey responses. The dates in this example correspond to those in Figure IA1. Consensus forecasts in the table are drawn from surveys published at the beginning of June and September 1997. The survey responses are gathered at the end of May and August, respectively, which are the dates reported in the table. The end-May survey has consensus forecasts of quarter-to-quarter inflation for five future quarters, from 1997Q3 through 1998Q3. The end-August survey has consensus forecasts for the five future quarters 1997Q4 through 1998Q4. The table reports consensus forecasts for the four future quarters that the surveys have in common.

These quarter-to-quarter forecasts are used to calculate, for each survey, the expected annualized log change in prices from 1997Q3 to each of the four quarters 1997Q4 through 1998Q3. The change from the end-May to end-August forecasts is defined as news about expected average inflation. In this example, forecasts of future average inflation drop by about 30 basis points at the one-quarter horizon. The magnitude of the decline shrinks to about 20 basis points at the four-quarter horizon.

C. Splicing Blue Chip Data Series

The text's Figure 1 displays monthly realizations of quarterly news about average expected inflation produced using the BC surveys. The data are from the BC-FF survey, augmented by the BC-EI survey for news realizations that cannot be produced with the BC-FF survey.

D. An Example of Constructing Short-Maturity Yield Shocks

The text compares two methods of forecasting short-maturity yields. One is a martingale forecast. The other is a forecasting regression, which is equation (9) in the text. Panel B of Table IAI illustrates the construction of short-maturity yield innovations. The first line reports end-May yields for bonds with maturities between one and four quarters. Equation (9) is estimated with OLS over the sample for which BC surveys are available. In-sample fitted one-quarter-ahead forecasts as of the end of May 1997 are reported in the second line of Panel B. Realized end-August yields are also reported, and the corresponding fitted innovations in yields.

Panel B shows that short-maturity yields were somewhat higher in August than anticipated in May, with yield innovations ranging from 3 to 30 basis points. Equation (5) in the text expresses these innovations as the sum of news about expected future inflation, expected future short-term real rates, and expected excess returns. Panel A shows that the news about expected future inflation is around negative 25 basis points at these maturities. Therefore, expected future real rates and/or expected excess returns must have increased unexpectedly.

II. Sticky Information?

The paper's model treats mean survey forecasts as true expectations, albeit possibly contaminated with i.i.d. measurement error. Coibion and Gorodnichenko (2012, 2015) argue that empirically, mean survey forecasts exhibit patterns consistent with informational rigidities. These rigidities imply that individual forecasters update their predictions infrequently, inducing sluggishness in mean forecasts. If true, expectations about inflation impounded in bond prices likely differ from those extracted from mean forecasts, since market prices are determined by active buyers and sellers. These agents are those most likely to have recently updated their information.

This section reviews and tests their model of rigidities using consensus forecasts from the SPF and real-time realizations of GDP inflation. I conclude that a more plausible interpretation of the evidence is that forecasters update frequently, and any evidence of rigidities is an accident of the sample.

Notation in this discussion differs from that used in the text, to accommodate the lag in reporting of quarterly inflation. Date t is the middle of calendar quarter t. At this time NIPA releases an estimate of inflation (GDP index) during quarter $t-1$. Denote this estimate by $\pi_{t|t-1}$, where the first part of the subscript refers to the quarter in which the information is revealed and the second part refers to the quarter over which inflation is measured. Respondents to the SPF predict $\pi_{t|t-1}$ in the middle of calendar quarters $t-1, \ldots, t-5$.

Denote the full-information rational expectations (FIRE) forecast with the usual expectations operator. A rational agent updates her expectation of $\pi_{t|t-1}$ every period. These updates are uncorrelated over time, and thus we can write the realization as the sum of news about $\pi_{t|t-1}$ at $t, t-1$, and so on:

$$\pi_{t|t-1} = \phi_t^{(t)} + \phi_{t-1}^{(t)} + \phi_{t-2}^{(t)} + \ldots, \qquad E_{t-j-1}\left(\phi_{t-j}^{(t)}\right) = 0.$$

(This equation ignores a constant term.) The superscript on the shock ϕ is the date of the inflation announcement and the superscript is the date the shock is revealed. The FIRE expectation of inflation of $t-j$ is

$$E_{t-j}\left(\pi_{t|t-1}\right) = \phi_{t-j}^{(t)} + \phi_{t-j-1}^{(t)} + \ldots. \tag{IA1}$$

Coibion and Gorodnichenko (2015) describe a simple model of sticky information. The discussion here adds some notation to their framework in order to consider forecast errors at different horizons. A fraction $1-\rho$ of respondents immediately update their prediction of $\pi_{t|t-1}$ in response to news at $t-j$. Therefore, the mean, across respondents, of the forecast of π_{t-1} adjusts by only $(1-\rho)$ of the true shock $\phi^{(t)}_{t-j}$. Of those who do not update immediately, $(1-\rho)$ update a period later, and so on. Hence, the mean survey forecast of $\pi_{t|t-1}$ made at $t-j$ is a sum of partial current and lagged shocks to the FIRE expectation. Denoting mean survey forecasts with hats,

$$\widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = (1-\rho)\phi^{(t)}_{t-j} + (1-\rho^2)\phi^{(t)}_{t-j-1} + (1-\rho^3)\phi^{(t)}_{t-j-2} + \dots. \tag{IA2}$$

Substituting (IA1) into (IA2) expresses mean survey forecasts as deviations from FIRE forecasts,

$$\widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = E_{t-j}\left(\pi_{t|t-1}\right) - \rho\sum_{i=0}^{\infty}\rho^i\phi^{(t)}_{t-j-i}.$$

Therefore, the error in the mean survey forecast made at $t-j$ is the sum of the FIRE forecast error and the expectational error,

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \left\{\pi_{t|t-1} - E_{t-j}(\pi_{t|t-1})\right\} + \rho\sum_{i=0}^{\infty}\rho^i\phi^{(t)}_{t-j-i}. \tag{IA3}$$

These forecast errors are closely related to lagged revisions in mean survey expectations. The revision in the mean survey expectation from $t-j-1$ to $t-j$ is

$$\left(\widehat{E}_{t-j} - \widehat{E}_{t-j-1}\right)\pi_{t|t-1} = (1-\rho)\sum_{i=0}^{\infty}\rho^i\phi^{(t)}_{t-j-i}.$$

Plug this expression for forecast revisions into the survey forecast error (IA3) to produce the relation between forecast errors and forecast revisions,

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \frac{\rho}{1-\rho}\left(\widehat{E}_{t-j} - \widehat{E}_{t-j-1}\right)\pi_{t|t-1} + \left\{\pi_{t|t-1} - E_{t-j}\left(\pi_{t|t-1}\right)\right\}. \tag{IA4}$$

The term in curly brackets is unforecastable (by both FIRE and survey respondents) as of $t-j$, and therefore is orthogonal to the first term on the right. Hence, (IA4) can be interpreted as a regression equation.

Coibion and Gorodnichenko (2012) use a variant of (IA4) to test for the presence of sticky information, focusing on annual forecasts of inflation. An implication of (IA4) explored here is that the coefficient on the forecast revision is independent of j. According to the model, the length of time between the forecast $t-j$ and the realization t does not affect the coefficient on the forecast revision.

Table IAII reports results of estimating the regression

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \beta_{0,j} + \beta_{1,i}\left(\widehat{E}_{t-j} - \widehat{E}_{t-j-1}\right)\pi_{t|t-1} + e_{t,j} \tag{IA5}$$

for various choices of quarterly lags j and two sample periods. The inflation measure is real-time GDP inflation, available from the Federal Reserve Bank of Philadelphia.[1] For the full sample (1969 through 2013), the point estimates are positive and significant, both economically and statistically. In this respect, the results support Coibion and Gorodnichenko's (2012) interpretation.

However, two aspects of these results cast considerable doubt on this story. First, the regression coefficients rise substantially as the forecast horizon increases. The estimate for $j = 1$ corresponds to a value of ρ less than 0.3, while the estimate for $j = 4$ corresponds to $\rho = 0.65$. Yet the theory does not accommodate shorter periods of inattention for near-term inflation.

Second, the statistical significance disappears after 1984. For the 1985 through 2013 subsample, forecast revisions are only weakly associated with survey forecast errors, both economically and statistically. For example, none of the regression R^2s exceed 3%. Nason and Smith (2014) test the sticky information hypothesis for inflation expectations and arrive at a similar conclusion about the role of the sample period. If inattention to news about inflation accounts for these results, inattention concentrates in the subsample when inflation is high and volatile but casual intuition suggests that the opposite should be true.

An alternative and perhaps more plausible story is that the observed relation between survey forecasts of inflation and realized inflation is not representative of the population relation. It is well known that the steady increase in inflation during the late 1970s and

[1] Because of a Federal government shutdown, the estimate of GDP inflation for 1995Q4 was not published by the Bureau of Economic Analysis during 1996Q1. For these regressions, the published value in 1996Q2 is treated as known by participants in 1996Q1.

early 1980s surprised most forecasters, whether they were paying attention or not. For example, beginning in 1978Q2, mean survey forecasts of three-quarter-ahead inflation were revised upwards for nine straight quarters. The corresponding realized forecast errors (the left side of (IA5)) were positive for all nine quarters.

Just as unusual was the energy price shock in 1973 and 1974. Three-quarter-ahead inflation forecasts were revised upwards for nine straight quarters beginning in 1973Q1. The corresponding three-quarter-ahead and four-quarter-ahead forecast errors were almost entirely positive and occasionally extremely large.[2] The errors associated with predictions made during the first three quarters of 1973 cannot be attributed to forecaster inattention, as the Yom Kippur War did not begin until October 1973. Subsequent forecast errors are more reasonably associated with difficulties in predicting the peak of energy prices rather than inattention to the OPEC oil embargo.

III. Details for Short-Horizon Estimates of Variance Ratios

Standard errors reported in the text's Table I are estimated using Generalized Method of Moments (GMM). The two moments for the martingale yield forecasts are the unconditional variance of yield shocks and the unconditional inflation variance ratio. For yield forecasts produced with the text's equation (9), four additional moments are the OLS orthogonality conditions. Persistent fluctuations in conditional variances are accommodated through the use of 12 quarters of Newey-West lags for the variance moments.

The text briefly mentions measurement error and claims that the role of such error is negligible. A simple model supports this claim. Consider a straightforward model with i.i.d. measurement error. Observed variance ratios satisfy

$$\text{observed } VR = \frac{(\text{true } \pi \text{ news variance}) + (\pi \text{ measurement error variance})}{(\text{true yield innov variance}) + (\text{yield measurement error variance})}.$$

Since we observe both the numerator and the denominator, we can make assumptions about the variances of measurement error to back out inflation variance ratios uncontaminated by measurement error.

Assume, for example, that both are contaminated by i.i.d. measurement error with a standard deviation of 10 basis points. Then the true variance ratio for the one-year

[2] A couple of the residuals for (IA5) are more than four standard deviations above zero.

yield using BC forecasts is 0.064 rather than 0.074 as reported in Table I. If yields are contaminated but inflation news is observed without error, the true variance ratio is 0.075. In sum, measurement error is irrelevant to these results.

As noted in the text, variance ratios based on the SPF differ somewhat from those based on BC surveys. One reason for the differences is that forecasts of CPI inflation are more volatile than forecasts of GDP inflation. A clear example is the post-Lehman revision in expected inflation. Figure 1 in the text shows that news about one-quarter-ahead expected CPI inflation is negative 2%. In the text's Figure 2, the corresponding news for expected GDP inflation is approximately negative 1%. News about expected CPI inflation is more volatile than news about expected GDP inflation for all maturities during the common sample periods in the text's Table II (1983Q1 to 2008Q2 and 2008Q3 to 2013Q4).

A more surprising reason for the differences is that yield innovations are somewhat more volatile when measured at the end of the month (aligning them with BC surveys) than when measured in the middle of the month (aligning them with SPF surveys). Part of this is undoubtedly just sampling variation, and there may be month-end effects as well.

IV. Details for Long-Horizon Estimates

A. Parameterization and Cointegration

Repeating the model's structure from the text,

$$\pi_t = \tau_t + \varphi_t + \phi_t, \tag{IA6}$$

$$\tau_t = \tau_{t-1} + \xi_t, \qquad E_{t-1}(\xi_t) = 0, \tag{IA7}$$

$$\varphi_t = \theta\varphi_{t-1} + \upsilon_t, \qquad E_{t-1}(\upsilon_t) = 0, \tag{IA8}$$

$$E_{t-1}(\phi_t) = 0. \tag{IA9}$$

Yield innovations are given by

$$y_t^{(m)} = y_{t-1}^{(m)} + \tilde{y}_t^{(m)}, \quad E_{t-1}\left(y_t^{(m)}\right) = 0, \tag{IA10}$$

for some long maturity m. The unconditional covariance matrix of shocks is

$$\mathrm{Var}\left(\left(\begin{array}{cccc} \phi_t & \xi_t & \upsilon_t & \tilde{y}_t^{(m)} \end{array}\right)'\right)^{1/2} = \left(\begin{array}{cccc} \sigma_\phi^2 & 0 & 0 & 0 \\ 0 & \Omega_{11} & 0 & 0 \\ 0 & \Omega_{21} & \Omega_{22} & 0 \\ 0 & \Omega_{31} & \Omega_{32} & \Omega_{33} \end{array}\right). \tag{IA11}$$

The square root in (IA11) refers to a Cholesky decomposition. The completely transitory shock to inflation is orthogonal to all other shocks. Its realization does not affect bond prices since it does not affect expectations of future inflation. Bond yields comove with the permanent and transitory shocks to inflation. Yield innovations also have a component orthogonal to inflation.

Forecasts of future inflation are formed with

$$E_t(\pi_{t+j}) = \tau_t + \theta^j \varphi_t.$$

The innovation from t to $t+1$ in the expectation of inflation in period $t+j$ is

$$\left(E_{t+1} - E_t\right) \pi_{t+j} = \xi_{t+1} + \theta^{j-1} \upsilon_{t+1}.$$

Similarly, the innovation from t to $t+1$ in the expectation of average inflation from $t+2$ to $t+m+1$ is

$$\left(E_{t+1} - E_t\right) \frac{1}{m} \sum_{i=1}^{m} \pi_{t+i+1} = \xi_{t+1} + \frac{1}{m} \theta \left(\frac{1-\theta^m}{1-\theta}\right) \upsilon_{t+1}. \tag{IA12}$$

The inflation variance ratio is measured by the variance of (IA12) divided by the variance of the yield innovation in (IA11).

Unit roots raise the issue of cointegration. Recall that both the trend τ_t and the bond yield $y_t^{(m)}$ are modeled as martingales. Cointegration requires that they have identical shocks. If they did not (say, the shock to the yield at time t was the sum of the trend-inflation shock and some other shock), then over time yields would diverge from trend inflation, as the other shocks accumulated without dying out. Therefore, cointegration imposes

$$\text{cointegration:} \qquad \Omega_{11} = \Omega_{31}, \quad \Omega_{32} = \Omega_{33} = 0.$$

This assumption implies an inflation variance ratio of one for long maturities. Inflation news for large m reduces to the martingale shock (take the limit of (IA12)), which is the same as the yield shock.

B. Measurement Error and Estimation

Recall that all observables other than the nowcast are assumed measured with error. The measurement error for a particular observable is assumed to be i.i.d. For example, denoting survey consensus forecasts with "s" superscripts, we have

$$E_t^s\left(\pi_{t+k}\right) = E_t\left(\pi_{t+k}\right) + \chi_{t,k}^s, \quad E\left(\chi_{t,k}^s\right) = 0, \ \mathrm{Cov}(E_t(\pi_{t+k}), \chi_{t,k}^s) = 0. \tag{IA13}$$

Thus, survey forecasts are noisy observations of true expectations. In principle, an econometrician could use the structure of (IA13) to produce more accurate forecasts than surveys. But in practice, estimation of (IA13) introduces sampling error that degrades forecast performance.

An alternative modeling assumption reverses the location of noise. Rather than assuming that surveys contain noise, we could assume that true expectations contain noise. Formally,

$$E_t\left(\pi_{t+k}\right) = E_t^s\left(\pi_{t+k}\right) + \chi_{t,k}^s, \quad E\left(\chi_{t,k}^s\right) = 0, \ \mathrm{Cov}(E_t^s(\pi_{t+k}), \chi_{t,k}^s) = 0. \tag{IA14}$$

In (IA13) the measurement error is orthogonal to the true expectation, while in (IA14) it is orthogonal to the survey forecast.

Dynamics of survey forecasts based on (IA13) differ from those based on (IA14). The latter equation implies that survey expectations react slowly to information about future inflation. Equation (IA14) says that the survey forecast at t of inflation at $t+k$ does not incorporate the information in $\chi_{t,k}^s$. Yet we can combine (IA14) with the law of iterated expectations to produce

$$E_t\left(E_{t+1}^s\left(\pi_{t+k}\right)\right) = E_t^s\left(\pi_{t+k}\right) + \chi_{t,k}^s. \tag{IA15}$$

Equation (IA15) says that the survey forecast at $t+1$ of inflation at $t+k$ incorporates the information in $\chi_{t,k}^s$. In this sense survey expectations respond to information with a lag. Evidence in this appendix supports the conclusion that survey forecasts do not react slowly.

I use assumption (IA13) in empirical estimation both because of this evidence and because the covariance assumption of (IA13) works when filtering the state vector from observables.

Exact filtering requires information about the conditional joint distribution of observables and factors. The model described by (IA6) through (IA11) does not specify conditional distributions—it only specifies conditional first moments and unconditional second moments. Since conditional second moments are neither known nor of direct interest here, I use a homoskedastic Gaussian model as a misspecified model in GMM estimation. For this misspecified model, the unconditional covariance matrix of shocks (IA11) is also the conditional covariance matrix of shocks.

Model parameters are estimated with exactly identified GMM. The moment vector is the score vector from maximum likelihood (ML) estimation of the homoskedastic Gaussian model. The assumption underlying this choice is that as the number of observations goes to infinity, the ML estimate of the covariance matrix (IA11) converges to the true unconditional covariance matrix of shocks. It is worth emphasizing that this is an assumption, not a conclusion validated by either Monte Carlo simulation or asymptotic arguments. The Kalman filter recursion allows for computation of the homoskedastic Gaussian likelihood function. The filter is also a straightforward method to handle the many missing observations of consensus survey forecasts.

The score vector is identically zero at the ML estimates, and thus the GMM parameter estimates are the ML parameter estimates. However, the covariance matrix of the estimates is not the ML covariance matrix. I use the GMM covariance matrix with a Newey-West adjustment for 12 quarterly lags of serially correlated moments.

The standard deviation of the yield's measurement error is fixed at 10 basis points, a value consistent with the analysis of Bekaert, Hodrick, and Marshall (1997). The standard deviations of the inflation expectations' measurement errors are set to a common value that is a free parameter.

Table IAIII reports parameter estimates and standard errors for the trend-cycle model of inflation and bond yields in the text's Section II.B. I only report results for the 10-year bond yield.

V. Time-Varying Risk Premia in the Bansal-Shaliastovich (2013) model

The text has an example that links news about expected excess returns to the mean slope of the term structure. Some formal derivations are below. Assume, as in Bansal and Shaliastovich (2013), that the joint conditional distribution of the nominal stochastic discount factor (SDF) and nominal bond prices is log-normal. Standard no-arbitrage arguments imply that the expected excess log return to an m-maturity bond is

$$E_{t-1}\left(ex_t^{(m)}\right) = -\frac{1}{2}\text{Var}_{t-1}\left(ex_t^{(m)}\right) - \text{Cov}_{t-1}\left(\log SDF_t, ex_t^{(m)}\right). \qquad \text{(IA16)}$$

Consider the case in which both the conditional variance and covariance in (IA16) are proportional to a single state variable V. The main example of this case is when the conditional standard deviations of both log returns and the log SDF are linear in the square root of the state variable, while the conditional correlation between log returns and the log SDF is constant. This is a simpler setting than that studied in Bansal and Shaliastovich (2013).

In this case we can rewrite (IA16) as

$$E_{t-1}\left(ex_t^{(m)}\right) = k_{1,m}V_{t-1} + k_{2,m}V_{t-1} = k_m V_{t-1}. \qquad \text{(IA17)}$$

Now assume that the state variable that drives variances is stationary, so unconditional expectations exist. Multiply and divide (IA17) by its unconditional expectation to produce

$$E_{t-1}\left(ex_t^{(m)}\right) = E\left(ex_t^{(m)}\right)\frac{V_{t-1}}{E(V)}. \qquad \text{(IA18)}$$

We can iterate this equation forward to produce multi-step forecasts,

$$E_{t-1}\left(ex_{t+i}^{(m)}\right) = E\left(ex_t^{(m)}\right)\frac{E_{t-1}\left(V_{t+i-1}\right)}{E(V)}, \qquad i \geq 0. \qquad \text{(IA19)}$$

Plug (IA19) into the equation for news about expected excess returns, the text's equation (4). The result is

$$\eta_{ex,t}^{(m)} = \frac{1}{m}\sum_{i=1}^{m} E\left(ex_{t+i}^{(m-i+1)}\right)\frac{(E_t - E_{t-1})V_{t+i-1}}{E(V)}. \qquad \text{(IA20)}$$

Equation (IA20) tells us that the news about expected average excess returns depends on the size of the variance shock relative to the average level of variance, the persistence of the variance shock, and the average compensation to holding nominal bonds. To get a sense of an upper bound on the magnitudes of news, consider a variance shock at t that does not begin to die out until after the bond has matured. Denoting this shock with a tilde, the news about expected average excess returns is

$$\eta_{ex,t}^{(m)} = \frac{\tilde{V}}{E(V)} \frac{1}{m} \sum_{i=1}^{m} E\left(ex_{t+i}^{(m-i+1)}\right). \tag{IA21}$$

Finally, plug the text's equation (1) into (IA21), replacing the sum of expected excess log returns with the yield spread:

$$\eta_{ex,t}^{(m)} = \left(\frac{\tilde{V}}{E(V)}\right) E\left(y_t^{(m)} - y_t^{(1)}\right). \tag{IA22}$$

The key observation from (IA22) is that the mean yield spread determines the magnitude of news about expected excess returns. This is important because mean yield spreads are small. Over the 1969 to 2013 sample, the mean yield spread between the five-year (10-year) yield and the three-month yield is 117 (159) basis points.

VI. Inflation Variance Ratios in Habit Formation Models

Table IAIV reports, for Wachter's (2006) model and Ermolov's (2015) model, standard deviations of yield shocks and news about expected inflation. Two sets of numbers are reported for Wachter's (2006) model. The second sets the correlation between consumption and inflation shocks to zero. In the original set this correlation is negative. The increase in variance ratios associated with the change in this coefficient indicates the sensitivity of the variance ratio to this parameter.

In Ermolov's (2015) model, the correlation is positive. The model-implied variance ratios are about the same as those for Wachter's (2006) model at short maturities, but are much larger at long maturities.

VII. The Dynamic Model

This section contains details about the model in Section IV of the main text.

A. Matrix Calculations of News

Consider an m-period bond for which we have parameters of the affine mapping, equation

(16) in the text. Shocks to the yield (excluding measurement error) are then

$$\tilde{y}_t^{(m)} = B_m' \Sigma \epsilon_t.$$

As of time t, the average expected value of the state vector from t to $t+m-1$ is

$$\frac{1}{m}\sum_{j=0}^{m-1} E_t(x_{t+j}) = \frac{1}{m}\left(\sum_{j=1}^{m-1}\sum_{i=1}^{j} K^{i-1}\right)\mu + \frac{1}{m}\left(\sum_{j=0}^{m-1} K^j\right) x_t.$$

In the case of stationary dynamics, this can be written as

$$\frac{1}{m}\sum_{j=0}^{m-1} E_t(x_{t+j}) = \left(I - \frac{1}{m}(I-K^m)(I-K)^{-1}\right)(I-K)^{-1}\mu + \frac{1}{m}\left(I-K^m\right)\left(I-K\right)^{-1} x_t.$$

Write this (either the general or the stationary version) as

$$\frac{1}{m}\sum_{j=0}^{m-1} E_t(x_{t+j}) = W_{m,0} + W_{m,1} x_t.$$

Then average expected inflation from $t+1$ to $t+m$ is

$$\frac{1}{m}\sum_{j=0}^{m-1} E_t\left(\pi_{t+1+j}\right) = A_\pi + B_\pi' W_{m,0} + B_\pi' W_{m,1} x_t,$$

and the news at t about this average expected inflation is

$$\eta_{\pi,t} = B_\pi' W_{m,1} \Sigma \epsilon_t. \tag{IA23}$$

Similarly, news at t about the average expected real short rate over the life of the bond is

$$\eta_{r,t} = \left(B_1 - B_\pi\right)' W_{m,1} \Sigma \epsilon_t.$$

Shocks to expected excess returns are calculated by subtracting shocks to average expected inflation and real rates from the yield shock. Variances and covariances among these shocks are computed using the covariance matrix of state-vector shocks.

B. Measurement Error

All four consensus forecasts and seven bond yields are assumed to be observed with i.i.d. measurement error. As with the trend-cycle model, (IA13) describes the errors in survey forecasts. The four standard deviations of measurement errors of the consensus forecasts are free parameters. Standard deviations of measurement errors of yields is fixed to 10 basis points for all maturities.

C. Normalizations

Model estimation uses a factor rotation of the text's equation (15) that sets the unconditional means of the factors to zero, diagonalizes the matrix K, and restricts the diagonal of Σ. The rotated factors satisfy the dynamics

$$x_{t+1} = Dx_t + \Sigma\epsilon_{t+1}, \tag{IA24}$$

where D is diagonal. The matrix Σ is lower triangular with fixed diagonal elements. (In implementation, they are fixed at 10.) The diagonal elements of D are its eigenvalues, which are ordered from largest to smallest. These dynamics rule out complex eigenvalues. Estimation of a version that allows for complex eigenvalues produces only real eigenvalues, and thus (IA24) is not restrictive in practice.

The A_m, A_π, B_m, and B_π parameters of the text's equations (28) and (29) are unrestricted.

D. Local Underidentification

When the diagonal elements of D are distinct, this normalized model is identified econometrically. In particular, the model is invariant to factor rotations. This invariance follows from the property that factor i follows an autonomous process with the AR(1) coefficient equal to D_{ii}. When the eigenvalues of D (the diagonal elements) are distinct, other linear combinations of these factors do not follow autonomous processes, and hence other linear combinations are ruled out.

Additional normalizations are required when the diagonal elements are not distinct.[3] If, say, the diagonal elements D_{ii} and D_{jj} are equal, then any linear combination of factors i

[3] Recall that the model rules out complex eigenvalues. Thus, this discussion is unrelated to the algebraic multiplicity of eigenvalues discussed in the appendix to Joslin, Singleton, and Zhu (2011).

and j also follows an autonomous AR(1) process. The model is locally unidentified (i.e., unidentified for $D_{ii} = D_{jj}$) because factor i can be replaced in (IA24) with such a linear combination. This is equivalent to a factor rotation that changes the covariances among the factors (elements of Σ) and the loadings of the observables on the factors (B).

Local underidentification occurs when estimated diagonal elements of D are sufficiently close for the likelihood function to be close to flat along the dimension of such a factor rotation. Even when the diagonal matrix is estimated with high precision, estimates of Σ and the factor loadings of yields and inflation are estimated with extremely low precision. Effectively, the likelihood of one factor rotation cannot be distinguished from another.

Unrestricted estimates of the normalized model reveal substantial local underidentification. This underidentification has no effect on point estimates of the inflation variance ratio, but makes it impossible to produce reliable confidence bounds. I therefore impose restrictions on the Σ matrix to pin down factor rotations in the presence of nearly identical (statistically) eigenvalues of D. Recall that the matrix is lower triangular and has diagonal elements fixed at 10. Additional underidentification restrictions result in the form

$$\Sigma = \begin{pmatrix} 10 & 0 & 0 & 0 \\ \Sigma_{21} & 10 & 0 & 0 \\ \Sigma_{31} & 0 & 10 & 0 \\ \Sigma_{41} & 0 & 0 & 10 \end{pmatrix}. \tag{IA25}$$

The eigenvalues of the factors are ordered in D from largest to smallest. This form therefore limits the covariances of shocks among the less-persistent factors, forcing a particular factor rotation even when they share the same eigenvalues.

E. Parameter Estimates and Related Information

Parameter estimates are reported in Table IAV. Asymptotic standard errors are in parentheses. Table IAVI gives a little more information about the local underidentification problem. It reports information for both the estimated model that uses (IA25) and an unrestricted model. For this model, the lower triangular off-diagonal elements of Σ are all free parameters. This adds three free parameters relative to the restricted model. The unrestricted model is estimated with ML.

Table IAVI shows that the two sets of point estimates produce almost identical values of the population standard deviation of quarterly shocks to yields. Differences are less than a basis point of annualized yields. Uncertainty in the estimate of Σ is much larger with the unrestricted model. The greater uncertainty makes it difficult to draw statistical inferences about population variance ratios.

A typical way to draw such inferences is to randomly draw parameters from the distribution implied by the point estimates and the covariance matrix of the estimates. We then calculate the desired population properties, such as variance ratios, for this draw. We repeat this procedure thousands of times, then construct confidence bounds. Because the covariance matrix of shocks is the square of Σ, greater uncertainty in the parameters of Σ leads to higher mean simulated variances. When uncertainty is small, this effect is also small.

However, Table IAVI shows that this effect is extremely large for the unrestricted model. The mean, across simulations, of the population standard deviation of quarterly shocks to the three-month yield exceeds 1800%. Even for the 10-year yield, the mean simulated standard deviation of quarterly shocks is almost 600%. Therefore, simulating parameters from the unrestricted model is effectively useless—at least for drawing inferences about variances and variance ratios. By contrast, Table IAVI also shows that mean simulated standard deviations for the restricted model are much better behaved.

VIII. A Reinterpretation of Level, Slope, and Curvature

Litterman and Scheinkman (1991) show that almost all of the cross-sectional variation in bond returns can be characterized by level, slope, and curvature factors. Returns are closely related to yield shocks, and thus it is not surprising that the same decomposition holds for shocks. In this section I ask whether this same decomposition holds for the components of yield shocks. To do so, I use the dynamic model of yields and expected inflation presented in Section IV of the main text. Can news about average expected inflation also be summarized by level, slope, and curvature? What about the innovations in yields not attributable to news about inflation?

News about expected inflation for an m-maturity bond is defined by the text's equation (31). The model of Section IV implies a covariance matrix of inflation news for bonds with the seven maturities used in estimation. Figure IA2 displays the first three principal components (PCs) of the covariance matrix constructed using the estimated model. The first PC is the blue line in Panel A. The second and third PCs are in Panel B, illustrated with a blue solid line and a blue dashed line, respectively.

The same decomposition could be produced separately for news about average expected real rates and term premia shocks. However, since these shocks are difficult to distinguish statistically, I sum these two shocks and produce a single PC decomposition. The first three PCs are displayed as red lines in Figure IA2. All of the PCs are scaled to represent the effect of a unit standard deviation shock.

The figure shows that both sets of PCs can be described as level, slope, and curvature. For both sets, the first PC accounts for about 96% of the total variance. The most obvious difference between the two sets is that shocks are smaller for news about expected inflation than for combined shocks to expected future real rates and term premia.

The similarities motivate a more complicated description of yield shocks than we typically infer from Litterman and Scheinkman (1991). There are two types of level shocks to yields. The smaller type is a shock to average expected inflation. The larger is a shock to the combination of expected real rates and term premia. Similarly, there are two types of slope and curvature shocks.

REFERENCES

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1-33.

Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On biases in tests of the expectations hypothesis of the term structure of interest rates, *Journal of Financial Economics* 44, 309-348.

Coibion, Olivier, and Yuriy Gorodnichenko, 2012, What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120, 116-159.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Informational rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644-2678.

Ermolov, Andrey, 2015, Time-varying risk of nominal bonds: How important are macroeconomic shocks? Working paper, Columbia GSB.

Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu, 2011, A new perspective on Gaussian dynamic term structure models, *Review of Financial Studies* 24, 926-970.

Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 54-61.

Nason, James M., and Gregor W. Smith, 2014, Measuring the slowly evolving trend in U.S. inflation with professional forecasts, CAMA–Australian National University Working Paper 7/2014.

Wachter, Jessica A., 2006, A consumption-based model of the term structure of interest rates, *Journal of Financial Economics* 79, 365-399.

Table IAI

Predictions and Realizations of Inflation and Bond Yields, mid-1997

Panel A reports, for May and August 1997, consensus forecasts of four future quarter-to-quarter CPI inflation rates from the Blue Chip Financial Forecast. It also reports implied forecasts of average inflation over horizons of one, two, three, and four quarters. Differences between the August and May forecasts are the news about expected future inflation.

Panel B reports yields on short-term zero-coupon Treasury bonds for the same two months. The row labeled "Forecasted change" contains predicted changes in yields from May to August. Yield forecasts are from OLS regressions using yields on three-month, 12-month, and 18-month bonds as predictive variables. Yield innovations are realized changes from May to August less predicted changes.

Quarter-to-quarter forecasts are in annualized percent without compounding. Average inflation and bond yields are in annualized percent using continuous compounding.

Panel A. Inflation expectations

Month-end, 1997	Value	Quarter 1997Q4	1998Q1	1998Q2	1998Q3
May	q-to-q inflation forecast	2.99	2.98	2.91	2.86
May	average inflation forecast	2.94	2.94	2.92	2.89
August	q-to-q inflation forecast	2.70	2.69	2.71	2.76
August	average inflation forecast	2.65	2.66	2.68	2.69
August	average inflation news	−0.28	−0.28	−0.25	−0.21

Panel B. Bond yields

Month-end, 1997	Value	Maturity 3 mon	6 mon	9 mon	12 mon
May	level	4.85	5.42	5.53	5.72
May	forecasted change	−0.02	−0.14	−0.20	−0.25
August	level	5.13	5.30	5.40	5.66
August	yield innovations	0.30	0.03	0.07	0.19

Table IAII

Explaining Survey Forecast Errors with Forecast Revisions

The table reports estimates of (IA5), which regresses inflation forecast errors on lagged revisions in mean survey forecasts. Mean survey forecasts of GDP inflation are from the Survey of Professional Forecasters. Forecast errors are the difference between the quarter-to-quarter GDP inflation rate observed at t and the mean survey forecast as of $t-j$. Revisions in forecasts equal the mean survey forecast at $t-j$ less the forecast at $t-j-1$. Asymptotic standard errors are adjusted for $j-1$ lags of moving average residuals. A few observations of survey forecasts are missing.

Forecast lag j (quarters)	Number of Obs	Coeff (Std Err)	R^2
Panel A. 1969Q1 to 2013Q4			
1	179	0.41	0.05
		(0.17)	
2	178	0.84	0.09
		(0.40)	
3	177	1.22	0.09
		(0.48)	
4	171	1.83	0.13
		(0.47)	
Panel B. 1985Q1 to 2013Q4			
1	115	−0.29	0.02
		(0.28)	
2	114	0.10	0.00
		(0.31)	
3	113	0.07	0.00
		(0.48)	
4	112	0.75	0.03
		(0.45)	

Table IAIII

Detail of Trend-Cycle Estimation of 10-year Inflation Uncertainty

The model is described by the main text's equations (10), (11), (12), (13), and (14). The bond's maturity is 10 years. The bond yield and consensus forecasts of future inflation are assumed to be contaminated by i.i.d. measurement error. The standard deviation of the yield's measurement error is fixed at 10 basis points. The common standard deviation of consensus forecasts is a free parameter, denoted SD-M in the table.

This table reports the parameter estimates from GMM estimation. Asymptotic standard errors are in parentheses. All parameters and standard errors are expressed in annualized basis points except θ, which is persistence at a quarterly frequency.

Survey	Sample	θ	Ω_{11}	Ω_{21}	Ω_{22}	Ω_{31}	Ω_{32}	Ω_{33}	$\text{Var}(\phi_t)^{1/2}$	SD-M
BC	1980Q1–2013Q4	0.60	21.3	−24.2	80.2	13.6	24.2	57.3	56.3	10.1
		(0.064)	(3.1)	(16.4)	(51.8)	(51.3)	(102.1)	(11.4)	(47.7)	(1.2)
BC	1980Q1–1982Q4	0.49	47.4	−13.3	230.7	39.1	0.30	124.8	159.5	15.6
BC	1983Q1–2008Q2	0.70	20.2	−28.6	45.6	14.1	20.4	48.6	32.4	8.9
BC	2008Q3–2013Q4	0.53	10.0	−0.8	79.6	6.2	9.0	36.8	48.1	10.9
SPF	1968Q4–2013Q4	0.74	20.5	5.5	48.2	16.4	4.3	53.7	28.4	14.9
		(0.017)	(3.4)	(7.1)	(8.8)	(7.7)	(4.5)	(6.5)	(3.1)	(1.5)
SPF	1968Q4–1979Q2	0.72	23.1	20.1	73.4	−1.4	2.1	42.2	39.5	19.7
SPF	1979Q3–1982Q4	0.84	16.8	63.0	8.3	61.2	−3.0	108.2	46.1	23.7
SPF	1983Q1–2008Q2	0.75	16.6	−5.7	31.9	22.6	10.5	41.9	21.9	11.3
SPF	2008Q3–2013Q4	0.78	7.7	−12.6	33.5	−39.9	10.3	15.3	16.8	10.9

Table IAIV
Volatilities of Shocks to Bond Yields and Expected Inflation Implied by Models with Habit Formation

The table reports population standard deviations of quarterly shocks to yields and expected inflation over the same horizon as the bond maturities. The population values are implied by parameterized models described in the indicated source. The inflation variance ratio is measured by the ratio of the latter variance to the former variance. The revised version of Wachter's model sets the correlation between shocks to consumption growth and inflation to zero. Using Wachter's notation, ρ is changed from -0.205 to 0. No other parameters are altered. Thanks very much to Andrey for calculating the numbers for his paper.

Source	Maturity (quarters)	Std Dev of Inflation News	Std Dev of Yield Innovations	Variance Ratio
Wachter (2006)	1	70	93	0.57
	4	64	89	0.52
	20	42	78	0.28
	40	27	77	0.12
Wachter (revised)	1	70	85	0.68
	4	64	81	0.63
	20	42	70	0.35
	40	27	69	0.15
Ermolov (2015)	1	85	106	0.65
	4	77	94	0.68
	20	47	58	0.66
	40	29	41	0.50

Table IAV
Parameter Estimates for the Dynamic Model of Yields and Inflation Expectations

The model is described by equations (IA24), (IA25), and the main text's equations (16) and (17). The observables are four consensus survey forecasts of future GDP inflation and yields of Treasury bonds with maturities of three months, one through five years, and 10 years. The units of all observables are in basis points at annualized rates. For example, a yield of 400 corresponds to 5% per year and a survey forecast of 300 corresponds to an inflation forecast of 3% per year. Each observable is assumed to be contaminated with i.i.d., normally distributed measurement error. This table reports the parameter estimates from maximum likelihood estimation. Asymptotic standard errors are in parentheses. Note that the factor dynamics are stationary, with the largest eigenvalue set to 0.999. Therefore, unconditional means of the observables exist. However, since the dynamics are almost a unit root, the standard errors on these means are extremely large.

Panel A. Factor dynamics

	Index			
Parameter	1	2	3	4
D_i (eigenvalues)	0.999	0.885	0.854	0.704
	–	(0.055)	(0.036)	(0.082)
$\Sigma_{1,i}$	10			
	(–)			
$\Sigma_{2,i}$	−2.801	10		
	(4.495)	(–)		
$\Sigma_{3,i}$	−1.342	0	10	
	(3.605)	(–)	(–)	
$\Sigma_{4,i}$	−0.061	0	0	10
	(2.371)	(–)	(–)	(–)

Table IAV, continued

Panel B. Observables						
		Factor Loading				SD of
Observable	Constant	1	2	3	4	Meas Error
1-q-ahead expected π	285.9	-2.428	0.670	-2.437	1.252	26.84
	(1858.4)	(0.298)	(3.321)	(1.087)	(0.641)	(2.24)
2-q-ahead expected π	Implied by 1-q-ahead, factor dynamics					15.88
						(1.57)
3-q-ahead expected π	Implied by 1-q-ahead, factor dynamics					14.04
						(1.13)
4-q-ahead expected π	Implied by 1-q-ahead, factor dynamics					21.33
						(1.49)
1q yield	384.7	-4.804	-7.204	-2.225	5.345	10
	(3712.1)	(1.432)	(3.406)	(9.090)	(1.697)	(–)
1y yield	422.3	-4.914	-7.254	-2.202	2.639	10
	(3795.8)	(1.432)	(3.406)	(9.090)	(1.697)	(–)
2y yield	448.1	-4.992	-6.813	-1.545	0.835	10
	(3857.6)	(1.442)	(3.306)	(8.972)	(1.726)	(–)
3y yield	470.9	-4.954	-6.326	-0.969	0.150	10
	(3829.0)	(1.358)	(2.591)	(8.336)	(1.577)	(–)
4y yield	490.6	-4.900	-5.855	-0.451	-0.225	10
	(3788.5)	(1.263)	(1.958)	(7.726)	(1.480)	(–)
5y yield	508.1	-4.795	-5.320	-0.064	-0.518	10
	(3708.2)	(1.164)	(1.425)	(7.109)	(1.369)	(–)
10y yield	565.7	-4.332	-3.622	1.182	-0.145	10
	(3353.4)	(1.067)	(1.049)	(6.444)	(1.273)	(–)

Table IAVI
Model-Implied Population Standard Deviations
of Quarterly Shocks to Bond Yields

This table reports properties of two parameterized versions of the model in the main text's Section IV. The restricted version is described by equations (IA24), (IA25), and the text's equations (16) and (17). The unrestricted version replaces (IA25) with a more general lower-triangular matrix, in which all of the off-diagonal elements are free parameters. Both models are estimated with maximum likelihood (ML), as described in the notes to Table IAV.

The table presents the population standard deviation of quarterly shocks to bond yields implied by the ML point estimates. It also reports mean population standard deviations from Monte Carlo draws of the model parameters. For a given model, the distribution of the parameters is centered at the point estimates with a covariance matrix equal to the outer-product estimate of the covariance matrix of the estimates. The units are basis points of annualized yields.

	Restricted		Unrestricted	
Maturity	Point Estimate	Mean from Simulations	Point Estimate	Mean from Simulations
Three months	95.8	137.9	95.2	1842.2
One year	84.4	129.9	84.0	1793.3
Five years	62.8	90.6	62.5	1056.7
Ten years	51.6	66.4	51.2	560.9

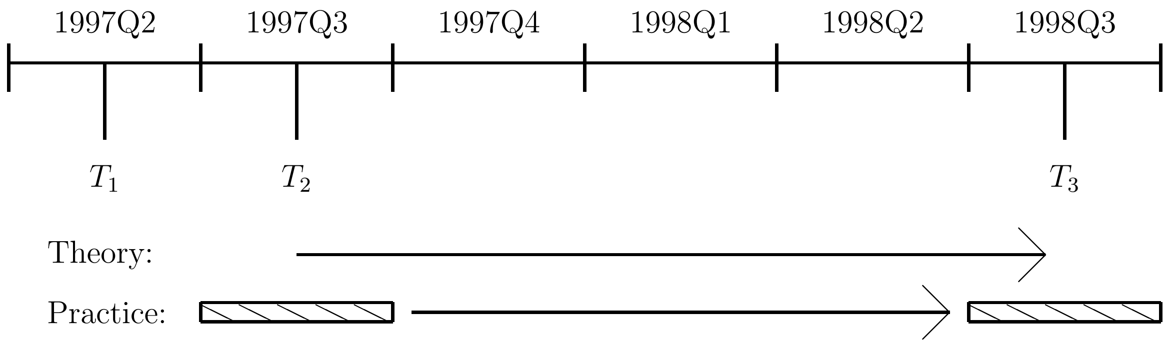

Figure IA1. Time line for news about expected inflation over the life of a one-year bond. The theory of the news decomposition assumes that we observe agents' predictions at dates T_1 and T_2 about the log change in the price level from date T_2 to date T_3. In practice, we observe survey forecasts as of T_1 and T_2 of the log change from the average price level in the quarter containing T_2, indicated by the first shaded box, to the average price of the quarter containing T_3, indicated by the second shaded box.

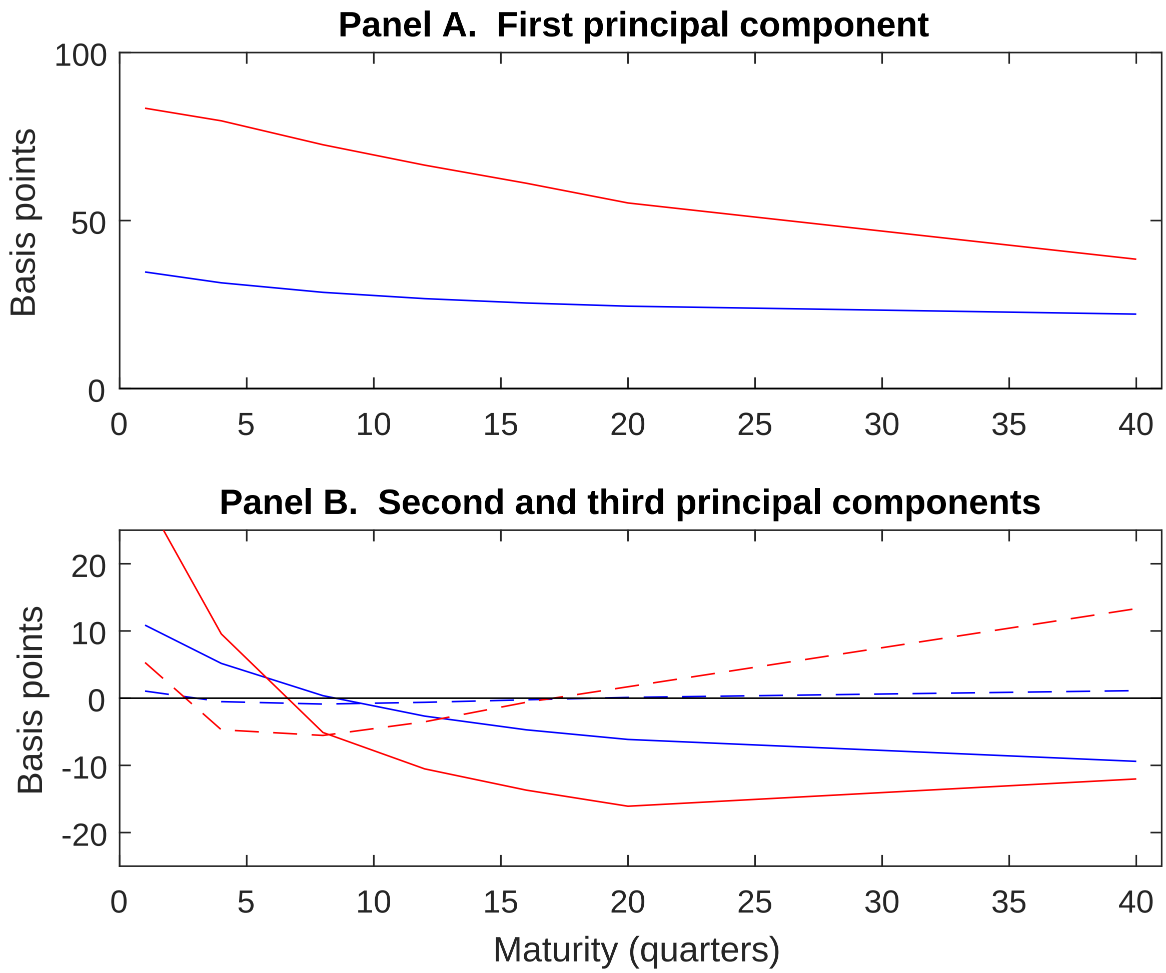

Figure IA2. Principal component decompositions of shocks to bond yields. Shocks to nominal yields are decomposed into shocks to average expected inflation over the life of the bond and a remainder, which is the sum of a shock to average expected real rates and a term premium shock. The decomposition and the resulting principal components of the shocks are calculated using a four-factor model of yields and inflation expectations estimated over the period 1968Q4 through 2013Q4. Principal components for inflation shocks are in blue and principal components for the remainder are in red. The third principal component is displayed with a dashed line. The principal components all correspond to a one-standard-deviation shock to the respective component.