Expected inflation and other determinants of Treasury yields

Gregory R. Duffee
Johns Hopkins University

First version April 2013
Current version September 2014

Abstract

A standard factor model is used to estimate the magnitude of inflation risk in nominal U.S. Treasury bonds. At a quarterly frequency, news about expected average inflation over a bond's life accounts for between 10 to 20 percent of shocks to nominal Treasury yields. This result is robust statistically, stable across time, and insensitive to the number of factors used in the model. Shocks to real rates and term premia account for the remainder of shocks to nominal yields.

Voice 410-516-8828, email duffee@jhu.edu. Address correspondence to 440 Mergenthaler Hall, 3400 N. Charles St., Baltimore, MD 21218. Thanks to seminar participants at Chicago's Booth School of Business, the Columbia Business School, the Shanghai Advanced Institute of Finance, the University of Lausanne, participants at the Bank of Canada "Advanced in Fixed Income Modeling" conference, Mike Chernov, Anna Cieslak, George Constantinides, Lars Lochstoer, and Jonathan Wright for helpful comments. This version is still preliminary. Comments are encouraged.

1 Introduction

A large and expanding literature explores the relation between nominal bond yields and inflation. Ang and Piazzesi (2003) make a particularly important contribution, introducing Gaussian macro-finance dynamic term structure models to determine the compensation investors require to face shocks to inflation and macroeconomic activity. The related literature has quickly branched out to include unspanned macro risks, non-Gaussian dynamics, and fundamental explanations for inflation risk premia that are grounded in investor preferences and New Keynesian macro models.

It is difficult to uncover from this literature any widely accepted conclusions about the joint dynamics of inflation and the nominal term structure. Ang, Bekaert, and Wei (2008) made the same point to motivate their own attempt to produce some basic facts. Research since Ang et al. has not exhibited any tendency to converge on their conclusions or any other set of core results. Thus it is not clear what branches of the macro-finance literature are likely to be fruitful and which should be abandoned.

This paper makes yet another effort to identify a robust empirical property that can be used to guide future research. How much inflation risk is embedded in nominal Treasury bonds? The term "inflation risk" is a little loose. The definition used here is derived from an accounting identity rather than from any assumptions about dynamics or risk premia. Inflation risk for a bond over a horizon such as a month or a quarter is the fraction of the variance of innovations in the bond's yield that is attributable to news about expected average inflation over the life of the bond. This fraction can be computed for all models that specify the joint dynamics of inflation and bond yields. Therefore if we can agree on a range of estimates of inflation risk as defined here, we can evaluate existing models and construct new ones based on their ability to generate plausible values.

Estimates of the quantity of inflation risk in nominal bonds vary wildly across prominent term structure models in the literature. Some estimated models imply that the variance of news about inflation exceeds the variance of yield shocks, so that inflation risk exceeds 100 percent. Others imply that almost none of the variance of yield shocks stems from news about inflation. Perhaps these differences are attributable to the variety of restrictions on risk premia dynamics that appear in the literature. These restrictions tie physical dynamics to cross-sectional covariances among yields. Not all of the restrictions can be correct, and false restrictions distort estimates of the physical dynamics.

To avoid imposing false restrictions, I estimate dynamic models of yields and expected inflation without imposing any restrictions on risk premia dynamics. In fact, I do not impose any no-arbitrage restrictions. The models explored here pin down physical dynamics but not equivalent-martingale dynamics. The focus is on the quantity of inflation risk, not the pricing of risk. Since yields depend on expected inflation rather than realized inflation, I follow the path pioneered by Pennacchi (1991) to link yields to survey forecasts of inflation.

Using 45 years of quarterly data, I conclude that empirically, shocks to expected inflation are a small part of shocks to the nominal bond yields. Roughly 10 to 20 percent of variances of quarterly shocks to long-term Treasury bond yields are attributable to news about expected inflation over the life of the bond. This conclusion is robust statistically and is satisfied for both parsimonious and highly flexible factor specifications. It also holds across subsamples. In particular, estimates of inflation risk are around 20 percent in a data sample that is dominated by the runup of inflation during the 1970s and the Federal Reserve's subsequent monetarist experiment. Estimates are around 10 percent for a sample that begins with Greenspan's arrival at the Fed.

The properties of the data that underlie this conclusion are easy to summarize. Although expectations of future inflation are highly persistent, they fluctuate little over time. For example, estimates of the standard deviation of quarterly shocks to average expected inflation over a ten-year horizon are in the neighborhood of 20 to 25 basis points. Estimates of the standard deviation of quartely shocks to the ten-year bond yield are around 55 to 60 basis points. Squaring and dividing produces variance ratio estimates between 10 and 20 percent.

Thus mechanically, innovations to expected short-term real rates and term premia are the primary drivers of yield shocks. There is insufficient information in the data to disentangle the relative contributions of these two components. Again, the relevant properties of the data are easy to summarize. Shocks to short-term real rates are large, and long-term nominal yields covary strongly with them. If short-term real rates are highly persistent, then the variation in long-term yields is explained by shocks to average expected future short-term real rates. If short-term real rates die out quickly, the variation is explained by term premia that positively covary with short-term real rates. Point estimates of the persistence are consistent with the latter version, but statistical uncertainty in these estimates cannot rule out the former version.

The next section describes how I measure the quantity of inflation risk and makes some preliminary calculations based on the existing literature. Section 3 describes the dynamic model I use to estimate inflation risk. The main results are in Section 4. Section 5 considers (and rejects) some possible objections to the modeling framework, such as informationally-sticky survey forecasts. Section 6 concludes.

2 The question and some preliminary evidence

Inflation risk is not a clearly defined concept, and there is no unique or best way to measure inflation risk in nominal bonds. The main problem is that shocks to yields cannot be divided cleanly between those that also alter the path of expected inflation and those that do not. This section advocates a measure that is intuitive, does not depend on modeling assumptions, and is straightforward to compute in many term structure models. The resulting estimates of inflation risk will help evaluate models and guide the development of theory.

The ambiguity of 'inflation risk' is inherent in both structural and reduced-form models. An example of the former is the New Keynesian model examined by Rudebusch and Swanson (2012). There are no exogeneous shocks to inflation. Thus in a limited sense, there is no inflation risk in the model. However, shocks to productivity, monetary policy, and government spending each affect the paths of expected inflation, real rates, and risk premia. All risk is inflation risk, in the sense that every shock contains news about expected future inflation. Models in which a monetary authority follows a Taylor rule typically have the same property. Outside of special cases, a shock to any variable that appears in the Taylor rule affects both yields and expected inflation.

Similarly, all shocks can be labeled as inflation shocks in standard reduced-form macro-finance models. To illustrate this property, consider a beta-type approach to measuring inflation risk, in which shocks to yields are regressed on shocks to expected inflation. (Ignore any difficulty in observing these shocks.) The fraction of yield shocks that is attributable to inflation risk could be measured by the R^2 of the regression.

In the canonical Gaussian macro-finance setting of Joslin, Le, and Singleton (2013) with inflation, straightforward manipulations allow a length-n latent state vector to be written entirely in terms of current and expected future inflation,

$$\text{state}_t = \begin{pmatrix} \pi_t & E_t\pi_{t+1} & \ldots & E_t\pi_{t+n-1} \end{pmatrix}'. \tag{1}$$

In these models, yield shocks are spanned by shocks to the state vector. Therefore if enough inflation expectations at different horizons are included in the regression, the R^2 will be one. The intuition is the same as it is with structural models. Unless restrictions are imposed on dynamics of the state vector, a shock to any element of the state vector affects the paths of all other state variables. There are no shocks to yields that do not affect inflation expectations at some horizon.

Distinguishing inflation shocks from other shocks is possible given sufficiently strong orthogonality assumptions. For example, the real business cycle term structure model of van Binsbergen, Fernández-Villaverde, Koijen, and Rubio-Ramírez (2012) exhibits money neutrality. Exogeneous shocks to inflation affect nominal yields but not real yields or risk premia. In a reduced-form setting, similar restrictions can be imposed on model dynamics. For example, Ang and Ulrich (2012) assume the existence of monetary policy shocks that affect nominal yields but are independent of inflation at all leads and lags.

Rather than adopt assumptions necessary for a model-based decomposition of shocks into inflation and non-inflation components, I measure inflation risk with an accounting approach that has its roots in the dividend/price decomposition of Campbell and Shiller (1988), as extended to returns by Campbell (1991).

2.1 Measuring inflation risk

I closely follow Campbell and Ammer (1993), who decompose unexpected bond returns into news about future real rates, news about future inflation, and news about future excess returns. The only mechanical difference is that I examine innovations in yields rather than innovations in returns. However, as Section 5.2.2 discusses, the conclusions I draw about the role of inflation contrast sharply with those of Campbell and Ammer.

Begin with some notation.

$y_t^{(m)}$: Continuously compounded yield, nominal zero-coupon bond maturing at $t+m$.

π_t : log change in the price level from $t-1$ to t.

π_t^e : Period-t expectation of next period's inflation, $\pi_t^e \equiv E_t(\pi_{t+1})$.

r_t : one-period real rate for nominal bonds, $r_t \equiv y_t^{(1)} - \pi_t^e$.

The one-period real rate r_t, also known as the ex-ante real rate, differs from the yield on a one-period real bond owing to both Jensen's inequality and the compensation investors require to face uncertainty in next period's price level. Inflation expectations should be thought of as investors' forecasts.

Yields on multiperiod bonds can be written as the sum of average expected short-term yields and term premia. Formally, the term premium on an m-period nominal bond, denoted $\tau_t^{(m)}$, satisfies

$$y_t^{(m)} \equiv \frac{1}{m}\sum_{j=0}^{m-1} E_t\left(y_{t+j}^{(1)}\right) + \tau_t^{(m)}. \tag{2}$$

Given the yield, the term premium depends on the expectations. We treat the expectation operator in (2) as investors' expectations and thus the term premium is investors' perception of the term premium.

Standard manipulations express this nominal yield in a variety of useful forms. One replaces the expected short-term nominal yield in (2) with its components:

$$y_t^{(m)} = \frac{1}{m}\sum_{j=0}^{m-1} E_t\left(r_{t+j}\right) + \frac{1}{m}\sum_{j=0}^{m-1} E_t\left(\pi_{t+j}^e\right) + \tau_t^{(m)}. \tag{3}$$

A nominal yield is the sum of expected average inflation and average real rates over the life of the bond, plus a term premium. I use some shorthand to refer to the first two components of this decomposition:

$$y_t^{(m)} = E_t(\overline{r}_{t,t+m-1}) + E_t(\overline{\pi}_{t+1,t+m}) + \tau_t^{(m)}. \tag{4}$$

Additional shorthand refers to the innovations in the elements of this decomposition:

$$
\begin{aligned}
\tilde{y}_t^{(m)} \equiv & y_t^{(m)} - E_{t-1} y_t^{(m)}, \\
\eta_{r,t}^{(m)} \equiv & E_t(\overline{r}_{t,t+m-1}) - E_{t-1}(\overline{r}_{t,t+m-1}, \\
\eta_{\pi,t}^{(m)} \equiv & E_t(\overline{\pi}_{t+1,t+m}) - E_{t-1}(\overline{\pi}_{t+1,t+m}), \\
\tilde{\tau}_t^{(m)} \equiv & \tau_t^{(m)} - E_{t-1}(\tau_t^{(m)}).
\end{aligned} \tag{5}
$$

This notation allows for a compact expression for shocks to yields. A yield shock is the sum of news about expected average real rates over the bond's life, news about expected average inflation over the bond's life, and a term premium shock:

$$
\tilde{y}_t^{(m)} = \eta_{r,t}^{(m)} + \eta_{\pi,t}^{(m)} + \tilde{\tau}_t^{(m)}. \tag{6}
$$

The variance of yield innovations is the sum of the variances of the individual components on the right side of (6) and twice their covariances:

$$
\begin{aligned}
\text{Var}\left(\tilde{y}_t^{(m)}\right) = & \text{Var}\left(\eta_{r,t}^{(m)}\right) + \text{Var}\left(\eta_{\pi,t}^{(m)}\right) + \text{Var}\left(\tilde{\tau}_t^{(m)}\right) \\
& + 2\text{Cov}\left(\eta_{r,t}^{(m)}, \eta_{\pi,t}^{(m)}\right) + 2\text{Cov}\left(\eta_{r,t}^{(m)}, \tilde{\tau}_t^{(m)}\right) + 2\text{Cov}\left(\eta_{\pi,t}^{(m)}, \tilde{\tau}_t^{(m)}\right).
\end{aligned} \tag{7}
$$

I define the direct measure of inflation risk as the ratio of the second term on the right to the total variance,

$$
\text{direct measure} \equiv \frac{\text{Var}\left(\eta_{\pi,t}^{(m)}\right)}{\text{Var}\left(\tilde{y}_t^{(m)}\right)}. \tag{8}
$$

The indirect measure of inflation risk is the sum of the first and third covariance terms divided by total variance,

$$
\text{indirect measure} \equiv \frac{2\left[\text{Cov}\left(\eta_{r,t}^{(m)}, \eta_{\pi,t}^{(m)}\right) + \text{Cov}\left(\eta_{\pi,t}^{(m)}, \tilde{\tau}_t^{(m)}\right)\right]}{\text{Var}\left(\tilde{y}_t^{(m)}\right)}. \tag{9}
$$

Survey data on yield and inflation expectations are not sufficiently comprehensive to construct model-free estimates of inflation risk. Thus (8) and (9) must be calculated from dynamic models. The models can be as simple as univariate autoregressive descriptions of inflation and an m-maturity yield. The variances of the residuals appear in the numerator and denominator of (8), and their covariance determines (9). The model presented in Section 3 is not much more complicated.

The focus of this paper is on the direct and indirect contributions of news about inflation to the variance decomposition (7). It is worth mentioning two related decompositions. The first is the decomposition of yields levels rather than yield shocks. Inflation and nominal yields roughly track each other over long horizons. This predictable variation shows up in the variance decomposition of yields rather than yield shocks. Versions of (8) and (9) for yields rather than yield shocks are reported at various points.

The second is the decomposition of conditional variances of yield shocks. Note that the variances in (7) are population variances of innovations. When shocks are heteroskedastic, conditional variances of a shock will generally differ from the population variance of the shock. Hence decompositions of conditional variances can differ from decompositions of population variances. For example, the amount of inflation risk may vary over time. However, the model I use for empirical work is completely Gaussian. Thus within the model there is no difference between population and conditional variances of innovations.

2.2 Some measurements from the literature

Dynamic term structure models that include inflation characterize fully the dynamics of nominal yields and inflation. Therefore given a parameterized model with a stationary distribution of shocks, the direct measure (8) can be calculated. When yields and inflation are both stationary, the corresponding variance ratio can also be calculated for levels of inflation and yields.

This subsection reports these measures for a few well-known contributions to the literature, beginning with Piazzesi and Schneider (2007). Piazzesi and Schneider construct a parsimonious benchmark equilibrium model of yields using a representative agent with recursive preferences.[1] Shocks are homoskedastic, thus term premia are constant over time. The point estimates of their model imply that long-run expected inflation unexpectedly jumps when expected long-run aggregate consumption growth unexpectedly falls. This pat-

[1]They also have a more sophisticated model with learning that is not examined here.

tern means that nominal yields are less volatile than is news about inflation; the latter is partially offset by opposing variation in expected short-term real rates.

The first row of Table 1 reports population properties of a five-year nominal bond yield.[2] The direct contribution ratio (8) at the quarterly horizon is reported in the column labeled "Innovations." The "Levels" column reports the the same ratio for unconditional variance of levels rather than innovations. Inflation risk—the variance ratio of news about inflation to yield shocks—is 1.5. The variance ratio for levels is almost identical.

The other models summarized in Table 1 are more flexible than that of Piazzesi and Schneider because they are not tied to the properties of aggregate consumption. Campbell and Viceira (2001) estimate two-factor Gaussian no-arbitrage models of nominal yields and inflation for two sample periods. Like Piazzesi and Schneider, their model rules out time-varying term premia. Ang, Bekaert, and Wei (2008) estimate a four-factor model with time-varying risk premia and an additional factor that captures changes in regimes. Chernov and Mueller (2012) estimate a variety of four-factor and five-factor Gaussian models, and Haubrich, Pennacchi, and Ritchken (2012) estimate a seven-factor model with stochastic volatility.[3]

None of these estimated models implies that inflation risk is as large as the value from Piazzesi and Schneider. Nonetheless, they disagree substantially about the relation between yields and expected inflation. Estimates of the direct contribution of news about inflation expectations range from almost none of the variance of yield innovations (0.03) to most of this variance (0.60). The range of estimates for the direct contribution to the level of yields is even greater.

One possible explanation for this wide range of estimates is that the data do not contain sufficient information to pin down the quantity of inflation risk. Another is that the differing restrictions embedded in these models override the information in the data. To distinguish between these possibilities we need a methodological approach that focuses on measuring inflation risk.

[2]The numbers are based on the "large information set" parameterization. Thanks to Monika and Martin for making their Matlab code available. The parameters used here are taken from their code, not from the published paper.

[3]For Ang et al., I use the parameters of their benchmark IV^C model. For Chernov and Mueller, I use the parameters of their $AO5$ model. Thanks to Mike and Philippe for helping me with their model.

3 A dynamic model of yields and expected inflation

This section describes how and why I model the joint dynamics of nominal yields and expected inflation. The first subsection presents the framework and the other subsections discuss the features of the model.

3.1 The framework

The dynamics of nominal yields and expected inflation are linked through their joint dependence on a state vector. Denote the length-n state vector by x_t. State-space models are standard in the dynamic term structure literature.[4] The state vector has Gaussian VAR(1) dynamics

$$x_{t+1} = \mu + Kx_t + \Sigma\epsilon_{t+1}, \qquad \epsilon_{t+1} \sim MVN(0, I). \tag{10}$$

Nominal yields are affine functions of the state vector. The notation for the yield on an m-maturity bond is

$$y_t^{(m)} = A_m + B_m' x_t + \eta_{m,t}, \tag{11}$$

where $\eta_{m,t}$ represents measurement error or some other deviation from an exact affine representation.

Like many other macro-finance researchers, I follow Pennacchi (1991) by drawing inferences about expected inflation from survey data. Substantial research supports the restriction that mean survey forecasts equal investor expectations. For example, survey forecasts are more accurate than model-based forecasts constructed using the history of inflation and other variables. Ang, Bekaert, and Wei (2007) conclude that mean inflation forecasts from surveys are more accurate than econometric forecasts. In addition, they find no evidence that using realized inflation in addition to survey forecasts helps reduce forecast errors. Faust and Wright (2009) and Croushore (2010) draw the same conclusion. Chernov and Mueller (2012) cannot reject the hypothesis that the subjective probability distribution of future inflation, as inferred from surveys, equals the true probability distribution. In a comprehensive handbook chapter, Faust and Wright (2012) concur: "...purely judgmental forecasts of inflation are right at the frontier of our forecasting ability."

[4]The first application of these models to interest rates is Hamilton (1985), although his motivation differs from that in the dynamic term structure literature.

Accordingly, the mean across survey respondents at t of predicted inflation at $t+j$ is assumed to equal the model-implied expectation of j-ahead inflation, possibly contaminated by some measurement error. Denoting mean survey forecasts with an "s" superscript, this restriction is

$$E_t^s(\pi_{t+j}) = E_t(\pi_{t+j}) + \eta_{\pi,j,t}, \tag{12}$$

Model-implied inflation expectations are affine in the state vector,

$$E_t(\pi_{t+j}) = A_{\pi,j} + B'_{\pi,j} x_t. \tag{13}$$

The absence of arbitrage is not imposed. Therefore the coefficients of (11) are unrestricted. By itself, the assumption of no-arbitrage is probably unimportant here. Joslin, Le, and Singleton (2013) show that when risk premia dynamics are not constrained, Gaussian no-arbitrage macro-finance models are close to factor-VAR models such as (10) and (11). No-arbitrage is typically imposed to allow researchers to impose economically-motivated restrictions on risk premia dynamics. All of the research discussed in the context of Table 1 imposes such additional restrictions.

Since these restrictions link equivalent-martingale dynamics to physical dynamics, they affect estimates of the magnitude of inflation risk in nominal yields. The evidence in Table 1 suggests that at least some of these restrictions produce estimates of inflation risk that are at odds with what we observe in the data. I choose to impose no pricing restrictions rather than attempt to determine which of these restrictions work better.

I break from the usual approach (including Pennacchi, Chernov and Mueller, and Haubrich et al.) by not including realized inflation among the observables. There are two reasons. First, including realized inflation is unlikely to improve the model's ability to describe expected inflation. Including realized inflation among the observables simply increases the number of free parameters and raises the likelihood of overfitting. Second, I am not interested in estimating the compensation investors require to face shocks to inflation. If investors are not risk-neutral with respect to the shock $\pi_t - E_{t-1}\pi_t$, then the one-period nominal rate will include a risk premium. This risk premium cannot be pinned down without observing realizations of inflation. However, risk premia are not the focus of this analysis.

Calculations of the direct and indirect measures of inflation risk use standard vector autoregression mathematics. Consider an m-period bond for which we have parameters of the affine mapping (11). 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

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

Write this 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

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

and the news at t about this average expected inflation is

$$\eta_{\pi,t} = B_{\pi,1}' W_{m,1} \Sigma \epsilon_t. \tag{14}$$

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

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

Term premia shocks are calculated by subtracting shocks to average expected inflation and real rates from the yield shock. Population variances and covariances among these shocks are computed using the population covariance matrix of state-vector shocks (the identity matrix).

3.2 A note on linear Gaussian dynamics

Overwhelming evidence documents stochastic volatility of both yields and inflation.[5] As mentioned at the end of Section 2.1, stochastic volatility drives a wedge between conditional and population variances of shocks. A completely Gaussian model is incapable of capturing time-varying inflation risk. There is also strong evidence of parameter instability for models of linear conditional means.[6] A linear model labels these nonlinearities as residuals and thus mismeasures innovations to yields and news about expected future real rates and inflation.

Rather than build a highly parameterized model capable of fitting these features, I explore their practical importance through subsample estimation. For example, the magnitude of shocks and their persistence were both different during the Great Moderation than they were during the Fed's monetarist experiment period in the late 1970s and early 1980s. How did these differences affect the relative importance of news about expected inflation? Similarly, does subsample estimation produce fitted time series of news about expected inflation and real rates that are similar to fitted time series produced by a single full-sample estimation? These informal tests help us judge the usefulness of the linear, completely Gaussian model.

4 Empirical evidence

This section presents results of estimating Section 3.1's model. Section 4.1 describes the data and Section 4.2 discusses the estimation procedure. Variance decompositions of shocks to bond yields are in Section 4.3 and evidence of their stability over time is in Section 4.4. These decompositions are closely linked to the persistence of shocks to expected inflation and real rates, a connection that Section 4.5 explores. Section 4.6 considers decomposing the term structure into level, slope, and curvature shocks. Finally, Section 4.7 compares variance decompositions of yields levels to those of yield shocks.

[5] The relevant literature is vast. A relatively early contribution is Schwert (1989). He estimates the variability and persistence of conditional volatility for a variety of important financial and macroeconomic time series.

[6] Again, the literature is vast. Early examples are Hamilton (1988), who examines regime shifts in the term structure, and Mishkin (1992), who examines instability in the relation between inflation and interest rates.

4.1 The data

Quarterly inflation forecasts are from the Survey of Professional Forecasters. During the first half of the second month in quarter t, respondents provide predictions of the GDP price level for quarters t, $t+1$, $t+2$, $t+3$, and $t+4$. These imply expected quarter-on-quarter inflation forecasts for quarters $t+1$ through $t+4$. I construct the mean cross-sectional prediction following the procedure of Bansal and Shaliastovich (2013), which discards outlier responses from individual forecasters. The first observation of the survey data is 1968Q4. The final observation used here is 2013Q4, for a total of 181 quarters. Intermittent observations of inflation forecasts for longer horizons are also available. I use these data for informal specification tests, comparing long-horizon forecasts from surveys to those implied by estimated models.

Treasury bond yields are observed in the middle of the second month of each quarter.[7] This choice approximately aligns the yield observations with dates that survey respondents make their predictions. Zero-coupon yields are for maturities of three months, one through five years, and ten years. The three-month yield is from the Federal Reserve Board's H15 release. The other yields are produced by Anh Le as described in Le and Singleton (2013).[8]

Figure 1 displays the term structures of forecasted inflation and yields. Over long horizons—say, decades—the two term structures move roughly together. At shorter horizons, they sometimes diverge substantially. For example, inflation forecasts jump in the mid 1970s without a corresponding increase in yields. Between 2000 and 2004, bond yields fall substantially, while inflation forecasts are stable. In the figure, cross-sectional differences in yields (i.e., measures of the slope of the term structure) are typically much larger than those for inflation forecasts, since the inflation forecasts do not look more than a year ahead.

Table 2 reports summary information about both levels and quarterly changes in inflation forecasts and yields. The latter are more relevant than the former when measuring inflation risk, since levels are highly persistent. The typical first-order serial correlation is about 0.98 at a quarterly frequency.

Two aspects of these statistics drive the main results in this paper. First, quarterly changes in inflation forecasts are substantially less volatile than quarterly changes in yields.

[7] Yields are for the 15th of the second month in the quarter. If the 15th is not a trading day, yields are observed on the last trading day prior to the 15th.

[8] Thanks very much to Anh for sharing the data.

Second, volatilities of changes in both inflation forecasts and yields decline with the horizon. The standard deviation of changes in inflation forecasts declines from 43 basis points at the one-quarter horizon to 31 basis points at the four-quarter horizon.

Given this pattern, it is plausible that standard deviations of changes in average expected inflation over a multi-year horizon, say five years, are well below 30 basis points. Yet standard deviations of changes in yields are around 60 to 70 basis points. If we treat changes as approximately equivalent to innovations, then these calculations imply that no more than one-fifth of the yield variance is explained by news about expected future inflation. Put differently, the direct measure of inflation risk (8) is no more than 0.2. Estimates of the formal model verify and refine this conclusion.

4.2 Estimation details

As Figure 1 shows, survey forecasts for different horizons track each other closely. I therefore estimate the formal model using just the one-quarter and three-quarter ahead forecasts. I use the three-quarter instead of the four-quarter forecast because the latter has five missing observations.

In the model description of Section 3.1, the state vector is latent and thus unidentified. The Appendix describes the normalizations imposed in estimation, and discusses at length a type of local underidentification that affects estimation and statistical inference. Certain parametric restrictions are adopted to address this local underidentification. The parameters are constrained to satisfy (13) for both the one-quarter-ahead and three-quarter-ahead inflation forecasts. Again, the Appendix describes the constraints.

The free parameters are estimated with the Kalman filter, which corresponds to maximum likelihood under the model's assumptions. Stationarity is imposed. Section 5.2 discusses this restriction in detail. The covariance matrix of parameter estimates is constructed with the outer product of first derivatives. Confidence bounds on nonlinear functions of the parameters are calculated using Monte Carlo simulations, randomly drawing parameter vectors from a multivariate Gaussian distribution with a mean equal to the parameter estimates. The Appendix contains some additional details.

I estimate the model using data for three periods. They are the full sample of 1968Q4 through 2013Q4, the subsample that ends with departure of Chairman Volcker (1968Q4 through 1987Q3), and the subsample that begins with Chairman Greenspan's appointment (1987Q4 through 2013Q4). The sample lengths are 181, 76, and 105 quarters respectively. Three-factor and four-factor models are estimated for all three periods. These models have 45 and 59 free parameters. A five-factor model, with 72 free parameters, is estimated for the full sample. The other samples are too short to estimate reliably the parameters of a five-factor model. For each sample, seven yields and two survey forecasts of inflation are observed.

The measures of inflation risk (8) and (9) can be calculated given a parameterized model. These measures—and more generally, variance decompositions of yield innovations—are discussed in detail in the next section. The model's parameters are not of direct interest, thus estimates are relegated to the Appendix.

Tables A3 through A5 report the full-sample estimated parameters, including standard deviations of cross-sectional deviations from an exact factor model. These standard deviations indicate that the three-month yield is not closely aligned with the other observables. The smallest standard deviation of its measurement error among the three sets of parameter estimates is 20 basis points. The largest is 47 basis points. The corresponding standard deviations for yields with maturities less than ten years are typically under 10 basis points, while estimates for the 10-year yield are around 10 to 30 basis points. Standard deviations of measurement error for survey forecasts of inflation are around 30 basis points for one-quarter-ahead inflation and less than 10 basis points for three-quarter-ahead inflation.

4.3 Variance decompositions

The most important message contained in these results is that inflation shocks account for a small fraction of the total variance of shocks to nominal yields. Table 3 presents detailed results for a four-factor model estimated over the entire sample 1968Q4 through 2013Q4. Variance decompositions are reported in Panel A for bonds with maturities of one, five, and ten years. Panel B sums the direct and indirect contributions of expected inflation.

There are four observations from the table worth emphasizing. First, for all the bonds, between 15 and 20 percent of the variance of nominal yield shocks is statistically explained by the direct contribution of news about inflation expectations. The two-sided 95 percent confidence bounds are tight. For each bond, we confidently conclude that quarterly inflation news accounts for between 5 and 35 percent of yield-shock variance.

Second, there is insufficient information in the data to decompose accurately the remaining variance of long-maturity yields into news about expected future real rates and news about term premia. The point estimates suggest that the former is more important at maturities less than five years and the latter is more important for longer maturities. However, the confidence bounds are huge, and do not rule out the possibility that either source dominates the other.

Third, point estimates imply a positive covariance between news about expected real rates and news about term premia. The estimates indicate that between 20 and 35 percent of the variance of yield shocks is attributable to this covariance. The confidence bounds are very large. Section 4.5 shows that this observation is closely related to the second observation above.

Fourth, sums of the direct and indirect measures of inflation risk (8) and (9) are also modest. The confidence bounds are wider than those for the direct contribution in Panel A because of the uncertainty in the covariances involving inflation expectations. Nonetheless, we can reject the hypothesis that the total contribution of inflation is outside of the range -0.5 to 0.5.

The result that inflation expectations contribute little to the variance of nominal yields is not dependent on the choice of a four-factor model. Table 4 reports estimates of the direct and total contributions of inflation expectations for three-factor and five-factor models estimated over the same sample. (To conserve space, other variance ratios are not reported in the table.) All of the reported point estimates of the direct contribution of inflation expectations are less than 0.20. Even the largest confidence bound in the table allows us to reject the hypothesis that more than 35 percent of the variance of a yield's shocks is attributable to news about inflation expectations. Simiarly, sums of the direct and indirect contributions of inflation expectations reported in Table 4 are very close to those reported in Table 3.

4.4 Stability over time

The sample period includes the monetarist experiment, the Great Moderation, and the financial crisis. It is well-known that inflation (and inflation expectations) were more volatile in the early part of the sample than in the latter. See, e.g., the discussion in Campbell, Shiller, and Viceira (2009). This pattern suggests that the variance ratios reported in Tables 3 and 4 might be unstable over time. I examine stability by contrasting results for the samples split at the transition from Chairman Volcker to Chairman Greenspan.

The point estimates for three-factor and four-factor models are displayed in Table 5. The role of inflation is indeed higher during the sample period that includes the monetarist experiment. However, the magnitude remains small. In the first subsample, inflation contributes directly about 20 percent of the variance of bond yields. In the second subsample, the direct contribution is about 10 percent. Confidence bounds, which are not reported in the table, are wide owing to the relatively short sample sizes.

How can the monetarist experiment period, characterized by volatile inflation expectations, have such a small fraction of yield innovations attributable to news about expected inflation? The reason is that all of the components of yields were more volatile during the monetarist experiment period. The evidence is in Figure 2.

Figure 2 displays filtered estimates of shocks to the three components of a five-year yield. The time series in the top row are based on a four-factor model estimated over the full sample. The time series in the bottom row are spliced together from separate estimation of three-factor models over the two subsamples. (Four-factor models over these short horizons produce highly unreliable estimates of both expected real rates and term premia.) Both rows provide visual evidence of the main empirical observation. Shocks to the five-year yield owing to news about expected inflation (the middle panels) are less volatile than either news about expected real rates or shocks to term premia.

Consistent with our intuition, the sample standard deviation of news about expected inflation in the early period is more than twice the corresponding standard deviation in the late period. According to the full-sample estimates, the volatilities of the other two components were similarly higher in early period. For the spliced sample estimates, the volatility of expected future real short rates did not vary much. The action is concentrated in changes in the volatility of term premia shocks, which is twice as large in the early period as it is in the late period. The differences between the top and bottom rows of this figure support an observation made in the previous subsection. The data do not allow us to distinguish clearly between the effects of average expected real rates and term premia.

4.5 Impulse responses

Impulse responses to shocks provide some additional information about the small contribution of inflation expectations to yield shocks. Figure 3 displays responses to a one standard deviation quarterly shock to the three-quarter-ahead inflation forecast, based on full-sample results for the four-factor model. Panel A reveals that the shock is small—about 30 basis points—and highly persistent. There is too much uncertainty in the parameter estimates to pin down the covariation between the shock to expected inflation and shocks to ex ante real rates, long-term yields, and term premia.

This figure contrasts sharply with Figure 4, which displays responses to a one standard deviation shock to the real short rate. (The shocks in Figures 3 and 4 are not orthogonalized.) The initial shock is large—almost a full percentage point—and the point estimates imply that it dies out quickly. The immediate response of the five-year yield is a little less than 50 basis points. Because the shock dies out so quickly, this response of the five-year yield substantially exceeds the shock to the average expected real rate over the next five years. Thus the term premium for the bond also immediately jumps by 20 basis points.

The pattern of these responses underlies an observation made in Section 4.3: the covariance between average expected real rates and term premia is large. The confidence bounds on these responses underlie another earlier observation. The data do not allow us to distinguish statistically between the roles played by average expected real rates and term premia. The point estimates imply that shocks to real rates die out quickly, but the confidence bounds in Panel A allow for the possibility that real rates are actually highly persistent. If real rates are highly persistent, then the immediate response of the five-year yield to the shock to real rates is in line with the shock to the average expected real rate over the next five years. Hence the confidence bounds on the response of the term premium includes the possibility that term premia do not react at all.

Another way to say this is that in the sample, shocks to real rates are volatile and not persistent. Long-term bond yields covary strongly with these shocks and these responses die out quickly as well. There are two ways to explain this pattern. One is that term premia are also volatile, covary strongly with real rates, and die out quickly. The other is that the sample pattern is at odds with the population properties of the data. If this explanation is correct, then shocks to real rates are truly highly persistent. Investors know this and price long-term bonds accordingly. This explanation implies that investors were surprised by the speed at which the shocks died out in the sample. There is not enough information in the sample to reject either explanation.

Neither theory nor the existing empirical literature offers much to help pin down the relative contributions of news about average expected real rates or shocks to term premia. If the shocks are primarily news about real rates, then shocks to real rates must be highly persistent. Such shocks arise naturally in settings where investors learn slowly about the dynamics of consumption, as in Johannes, Lochstoer, and Mou (2014). However, it is not clear that the amount of variation we see in real rates is consistent with learning. Other shocks may be more important. Hanson and Stein (2014) find that monetary policy shocks have substantial effects on long-term nominal and TIPS yields. They interpret these as term premia shocks rather than news about expected real rates, because standard theories of monetary policy do not allow policy shocks to have long-term effects on short-term real rates. Nakamura and Steinsson (2013) disagree about both the magnitude of shocks to long-maturity yields and their interpretation as primarily term premia.

A long-established empirical literature supports high real-rate persistence. Work that predates cointegration tools, most prominently Nelson and Schwert (1977), Garbade and Wachtel (1978), Mishkin (1981), and Fama and Gibbons (1982) finds that the real rate varies through time and is highly persistent. More recent research explicitly tests for unit roots and cointegration among inflation, nominal yields, and real rates. In their critical review of the literature, Neely and Rapach (2008) conclude that "...studies [of real rates] often report evidence of unit roots, or–at a minimum–substantial persistence." However, Neely and Rapach note that it is hard to tell whether *shocks* to real rates are persistent or whether the conditional mean of the real-rate process changes periodically.

4.6 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, thus it is not surprising that the same decomposition holds for shocks. This subsection asks whether this same decomposition holds for the components of yield shocks. 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 (14). The parameterized models imply a covariance matrix of inflation news for bonds with the seven maturities used in estimation. Figure 5 displays the first three principal components (PCs) of the covariance matrix constructed using the four-factor full-sample estimates. The first PC is the blue line in the 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 5. 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 percent 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. 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.

4.7 Unconditional properties of yields and expected inflation

Risk is created by shocks. Therefore the empirical analysis throughout this section emphasizes variance decompositions of yield shocks. At first glance, the small role of inflation in these decompositions is surprising. Our intuition is drawn largely from the way we interpret visual evidence such as Figure 1, which displays levels of yields and expected inflation. A glance at the figure shows that the series move roughly together. Correlations between yields and expected inflation over the period 1968Q4 through 2013Q4 range from about 0.7 to 0.8.

The first and most obvious mismatch between our intuition and the results of this section is the difference between levels and shocks. A second mismatch is that correlations ignore differences in scale. Note that the scales of the vertical axes in Figure 1 differ. Yields are more volatile than expected inflation, both in levels and in shocks. Thus variance decompositions of yield *levels* need not assign a large role to expected inflation, even though their correlations are high.

Table 6 reports a direct measure of inflation's contribution to yield levels. It is the level version of (8), the direct measure of inflation risk: the unconditional variance of average m-quarter inflation divided by the unconditional variance of the m-maturity bond yield. As with the variance decompositions of shocks, these variances are population values implied by model parameters.

There are two main conclusions to draw from the table. First, the estimates indicate that inflation accounts for less than half of the variation in yields. The variance ratios are a little less than 0.25. Second, the confidence bounds are much too large to draw strong conclusions. According to these bounds, the true contribution of inflation might be close to zero or greater than a half.

The large confidence intervals are another manifestation of the inability to pin down the persistence of shocks to real rates. Shocks to expected inflation are small and highly persistent. Shocks to real rates are large. If they are also highly persistent in population, then unconditional yield variability will be driven primarily by real rate variability. By contrast, if the population persistence of real rates matches their persistence in the sample, unconditional yield variability will be driven primarily by expected inflation variability.

5 Modeling alternatives

This section discusses in detail two alternatives to the modeling approach of Section 3.1. The first is that survey forecasts are sticky and the second is that inflation (and inflation expectations) are nonstationary.

5.1 Sticky information

The model treats mean survey forecasts as true expectations, albeit possibly contaminated with measurement error. Coibion and Gorodnichenko (2012a,b) 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 subsection reviews their model of forecasts. It concludes that a more plausible interpretation of the evidence is that forecasters update frequently, and any evidence of rigidities is an accident of the sample.

This discussion uses different notation than Section 3 to account for 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 the information is revealed and the second refers to the quarter over which inflation is measured. Respondents to the Survey of Professional Forecasters predict $\pi_{t|t-1}$ in the middle of calendar quarters $t-1, \ldots, t-5$.

Denote the full information rational expectations (FIRE) expectations with the usual expectations operator. A rational agent updates her expectation of $\pi_{t|t-1}$ every period. These updates are uncorrelated over time, 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 supercript 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

$$E_{t-j}\left(\pi_{t|t-1}\right) = \phi_{t-j}^{(t)} + \phi_{t-j-1}^{(t)} + \ldots \tag{15}$$

Coibion and Gorodnichenko (2012b) 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 $\phi_{t-j}^{(t)}$. 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,

$$\widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = (1-\rho)\phi^{(t)}_{t-j} + (1-\rho^2)\phi^{(t)}_{t-j-1} + (1-\rho^3)\phi^{(t)}_{t-j-2} + \ldots \tag{16}$$

Substituting (15) into (16) 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,

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \left\{\pi_{t|t-1} - E_{t-j}(\pi_{t|t-1})\right\} + \rho\sum_{i=0}^{\infty}\rho^i\phi^{(t)}_{t-j-i}. \tag{17}$$

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 (17) to produce the relation between forecast errors and forecast revisions,

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \frac{\rho}{1-\rho}\left(\widehat{E}_{t-j} - \widehat{E}_{t-j-1}\right)\pi_{t|t-1} + \left\{\pi_{t|t-1} - E_{t-j}\left(\pi_{t|t-1}\right)\right\}. \tag{18}$$

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 (18) can be interpreted as a regression equation.

Coibion and Gorodnichenko use a variant (18) to test for the presence of sticky information, focusing on annual forecasts of inflation. An implication of (18) 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 7 reports results of estimating the regression

$$\pi_{t|t-1} - \widehat{E}_{t-j}\left(\pi_{t|t-1}\right) = \beta_{0,j} + \beta_{1,i}\left(\widehat{E}_{t-j} - \widehat{E}_{t-j-1}\right)\pi_{t|t-1} + e_{t,j} \tag{19}$$

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.[9] For the full sample from 1969 through 2013, the point estimates are positive and significant, both economically and statistically. In this respect, the results support Coibion and Gorodnichenko's 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 sample 1985 through 2013, forecast revisions are only weakly associated with survey forecast errors, both economically and statistically. For example, none the regression R^2s exceed three percent. 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 during the sample period was concentrated in the subsample when inflation high and volatile. Casual intuition suggests that the opposite should be true.

An alternative and more perhaps 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 early 1980s surprised most forecasters, whether they were paying attention or not. For example, beginning with 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 (19)) were positive for all nine.

Just as unusual was the energy price shock in 1973 and 1974. Three-quarter-ahead

[9]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.

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.[10] The errors associated with predictions made during the first three quarters of 1973 cannot be attributed to forecaster inattention, since 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.

5.2 Nonstationary inflation

The dynamics (10) are consistent with unit roots in elements of the state vector, and thus consistent with unit roots in yields and expected inflation. In estimation I impose stationarity on (10). By contrast, the standard trend-cycle framework for inflation dynamics assumes a unit root. Prominent applications include Stock and Watson (2007) and Cogley, Primiceri and Sargent (2010). Dynamic term structure models typically assume stationarity. However, Fama (2006), Duffee (2011), and Christensen and Rudebusch (2012) all assume a unit root in yields.

This subsection examines issues related to stationarity, drawing two main conclusions. First, measures of inflation risk are largely unaffected by using a trend-cycle model of inflation. Second, the nonstationary inflation specification adopted by Campbell and Ammer (1993) is inappropriate because it produces forecasts of long-run inflation that are much more volatile than survey expectations of long-run inflation.

In-sample predictions of average ten-year inflation are displayed in Figure 6. The solid black line represents forecasts generated by Kalman filter estimation of the four-factor model over the full sample. The figure also displays mean survey forecasts of ten-year inflation. The forecasts are from Blue Chip Economic Indicators, the Livingston Survey, and the Survey of Professional Forecasters. Survey forecasts are available beginning in 1979Q4. Ignore the dashed-dotted line for the moment.

There are two notable differences between survey forecasts and model-based forecasts. First, survey forecasts are almost always higher. This is probably because the survey forecasts are primarily for CPI inflation while the model is parameterized using GDP inflation. Over

[10]Figure 1 displays the substantial swings in survey forecasts over this period. A couple of the residuals for (19) are more than four standard deviations above zero.

this sample, realized CPI inflation is about 50 annualized basis points higher than realized GDP inflation.

Second, survey forecasts rise and fall more sharply during the monetarist experiment. Survey forecasts reach a peak of 8.25 percent, while the model's forecasts do not reach 7.5 percent.[11] The model's long-run forecasts during this period are relatively low even though short-run survey inflation expectations are high. One-quarter-ahead and three-quarter-ahead forecasts reach 9.1 percent and 8.9 percent respectively.

This gap between survey and model long-run forecasts can be explained by different views of the persistence of inflation shocks. The model's estimates are inferred from the in-sample persistence of short-horizon inflation expectations under the maintained assumption of stationarity. When inflation expectations reach their peak in the sample, the model requires that expectations eventually return to their sample mean. Estimates of persistence are subject to the standard downward bias for a process that is close to a unit root.

If inflation persistence in the model is considerably lower than inflation persistence as perceived by investors, the model's estimates of inflation risk will be too small. All else equal, faster mean reversion of shocks to short-horizon inflation expectations implies less volatile news about long-run expected inflation.

I examine whether this effect is empirically important by comparing the model's estimates of inflation news to estimates from two methodologies that impose a unit root on inflation. The first is the trend-cycle model of inflation. The second is the approach of Campbell and Ammer (1993), which ignores any information in the level of inflation. I conclude that the trend-cycle model produces measures of inflation risk that are similar to those produced by the stationary model. The model of Campbell and Ammer produces substantially different measures. It does so via a highly volatile process for long-run inflation. The resulting estimates of long-run inflation are much too volatile relative to survey-based estimates.

5.2.1 Trend-cycle inflation

The trend-cycle model describes inflation as the sum of a random walk component and a transitory component,

$$\pi_t = \tau_t + \varphi_t, \tag{20}$$

[11] Blue Chip forecasts prior to 1983 are for GDP inflation. Therefore the gap is not explained by the difference between CPI and GDP inflation.

$$\tau_t = \tau_{t-1} + \xi_t, \tag{21}$$

$$\varphi_t = \theta\varphi_{t-1} + \upsilon_t. \tag{22}$$

I follow Nason and Smith (2014) by assuming shocks are homoskedastic and Gaussian,

$$\left(\begin{array}{cc} \xi_t & \upsilon_t \end{array} \right)' \sim MVN(0, \Omega\Omega'), \qquad \Omega = \left(\begin{array}{cc} \Omega_{11} & 0 \\ \Omega_{21} & \Omega_{22} \end{array} \right). \tag{23}$$

Other trend-cycle models appear in the literature. Stock and Watson (2007) use a variant of (20) through (23) in which the transitory shocks are not persistent and volatilities vary over time. Cogley, Primiceri and Sargent (2010) allow the persistence of transitory shocks to vary over time. I use (20) through (23) because the specification (other than the unit root) is similar to the model in Section 3.1.

I estimate the model using mean survey forecasts. Assuming the survey forecasts are formed in a manner consistent with the trend-cycle model, the link between the forecasts and the components of inflation is

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

Following the estimation procedure for Section 3.1's model, the two components of inflation are treated as unobserved, while the data are observed with measurement error. Parameter estimation uses the Kalman filter. Note that estimation does not involve any bond yields.

Panel A of Table 8 reports parameter estimates for the full sample, the subperiod that ends with the departure of Volcker, and the subperiod that begins with the arrival of Greenspan. In-sample predictions of average ten-year inflation are displayed with a dashed-dotted line in Figure 6. They are produced using the full-sample estimates. Panel B reports the model-implied standard deviation of shocks to average expected ten-year inflation. It reports the corresponding standard deviations for stationary models of yields and inflation that are discussed in Section 4.3.

The figure shows that during the monetarist experiment period, the long-horizon forecasts from trend-cycle model are very close to corresponding survey forecasts. Thus the fitted shocks to long-horizon inflation expectations over this period are larger for the trend-cycle model than the stationary models. Nonetheless, models estimated over a sample that includes this period largely agree on the standard deviation of quarterly shocks to ten-year inflation expectations. For the full sample, the trend-cycle estimate of this standard deviation, 22 basis points, equals the estimates of the three-factor and four-factor stationary models. The trend-cycle estimate for the pre-Greenspan sample is 28 basis points, between the estimates of the three-factor and four-factor stationary models.

Oddly, the only substantial difference appears for estimates based entirely on post-Volcker data. According to stationary model estimates, quarterly shocks to ten-year inflation expectations are small. Estimated of their standard deviation are less than 14 basis points. The estimated standard deviation for the trend-cycle model is noticeably larger, at 23 basis points. In results not detailed here, I find that first differences of the fitted trend produced by the Kalman filter are negatively serially correlated. This pattern is consistent with Nason and Smith (2014), who find that the trend-cycle model is insufficiently rich to capture the dynamics in survey forecasts of GDP inflation.

5.2.2 Campbell and Ammer (1993)

The measures of inflation risk described in Section 2.1 are borrowed from Campbell and Ammer's decomposition of excess bond returns. However, they assume a different process for inflation and conclude that shocks to average expected inflation account for the vast majority of shocks to nominal yields.

Their procedure assumes that expectations of future inflation and short-term nominal rates are determined by the stationary dynamics of a vector autoregression (VAR). The six variables included in the VAR are the ex-post real interest rate (short nominal rate at t less inflation at $t+1$), the change in the short nominal rate, the excess return to the aggregate stock market, the slope of the term structure, the dividend-price ratio, and the relative bill rate.[12] Inflation enters only in the form of a linear combination with the short nominal rate, and the level of yields is not included. Therefore yields and inflation are assumed to be nonstationary and cointegrated. Since neither the level of inflation nor the level of some yield appears by itself in the VAR, these variables are implicitly assumed to not have any stationary components. Real rates are assumed to be stationary.

Campbell and Ammer produce estimates of inflation risk that are substantially higher

[12] The slope is measured by the ten-year yield less the one-month bill rate. The relative bill rate is the one-month yield less the one-year backward moving average of the one-month yield.

than for either the stationary model of Section 3.1 or the trend-cycle model. They conclude that inflation risk accounts for effectively all of the variance of bond return shocks. To confirm and update their results, I apply their methodology to more recent data. I use their data definitions and frequency (monthly).

Table 9 reports the implied variance decompositions of shocks to a ten-year bond yield for the three sample periods studied in Section 4. The results are similar to those reported by Campbell and Ammer for different sample periods. The direct measure of inflation risk is roughly one for all three periods, including the post-Volcker era. In other words, the variance of news about expected average inflation over ten years is close to the variance of shocks to the ten-year yield. The confidence bounds all easily reject the null hypothesis that less than half of the yield variance is attributable to news about inflation.

When news about inflation is so large, expectations of long-run inflation must vary substantially over time. Figure 7 displays the model's monthly forecasts of average inflation over the next 120 months. The forecasts fluctuate substantially over time, from more than 12 percent in the early 1980s to less than -1 percent in 2011. Corresponding survey forecasts are also displayed; they are identical to those in Figure 5. It is clear from the figure that the model's estimates of long-run inflation are wildly at odds with both the survey forecasts and the forecasts produced by the models that allow some or all of the variability in inflation to be mean reverting.

6 Concluding comments

This paper studies the joint dynamics of nominal yields and inflation expectations from 1968 through 2013. For this sample as well as subsamples, quarterly shocks to nominal yields are primarily shocks to real rates and term premia. News about expected future inflation contributes relatively little to the variance of yield shocks.

This is a robust result that can be used to help evaluate dynamic term structure models. In particular, some estimated models in the literature imply that inflation shocks drive much of the variation in nominal yields. This feature of the models is inconsistent with the U.S. experience, even excluding the period of stable inflation during the Federal Reserve leadership of Greenspan and Bernanke.

References

Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do macro variables, asset markets or surveys forecast inflation better?, *Journal of Monetary Economics* 54, 1163-1212.

Ang, Andrew, Geert Bekaert, and Min Wei, 2008, The term structure of real rates and expected inflation, *Journal of Finance* 63, 797-849.

Ang, Andrew, and Monika Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745-787.

Ang, Andrew, and Maxim Ulrich, 2012, Nominal bonds, real bonds, and equity, Working paper, Columbia GSB.

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 24, 2738-2780.

Campbell, John Y., 1991, A variance decomposition for stock returns, *The Economic Journal* 101,157-179.

Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? A variance decomposition for long-term asset returns, *Journal of Finance* 48, 3-37.

Campbell, John Y., and Robert J. Shiller, 198, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195-228.

Campbell, John Y., Robert J. Shiller, and Luis M. Viceira, 2009, Understanding inflation-indexed bond markets, *Brookings Papers on Economic Activity*, Spring, 79-120.

Campbell, John Y., and Luis M. Viceira, 2001, Who should buy long-term bonds?, *American Economic Review* 91, 99-127.

Chernov, Mikhail, and Philippe Mueller, 2012, The term structure of inflation expectations, *Journal of Financial Economics* 106, 367-394.

Christensen, Jens H.E., and Glenn D. Rudebusch, 2012, The response of interest rates to U.S. and U.K. quantitative easing, *Economic Journal* 122, F385-F414.

Cogley, Timothy, Giorgio E. Primiceri and Thomas J. Sargent, 2010, Inflation-gap persistence in the U.S., *American Economic Journal: Macroeconomics* 2, 43-69.

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

Coibion, Olivier, and Yuriy Gorodnichenko, 2012b, Informational rigidity and the expectations formation process: A simple framework and new facts, Working paper, UT Austin and UC Berkeley.

Croushore, Dean, 2010, An evaluation of inflation forecasts from surveys using real-time data, *B.E. Journal of Macroeconomics* 10, 1-32.

Duffee, Gregory R., 2011, Forecasting with the term structure: The role of no-arbitrage restrictions, Working paper, Johns Hopkins University.

Fama, Eugene F., 2006, The behavor of interest rates, *Review of Financial Studies* 19, 359-379.

Fama, Eugene F., and Michael R. Gibbons, 1982, Inflation, real returns and capital investment, *Journal of Monetary Economics* 9, 297-323.

Faust, Jon, and Jonathan H. Wright, 2009, Comparing Greenbook and reduced form forecasts using a large real-time dataset, *Journal of Business and Economic Statistics* 27, 468-479.

Faust, Jon, and Jonathan H. Wright, 2012, Forecasting inflation, *Handbook of Forecasting*, forthcoming.

Garbade, Kenneth, and Paul Wachtel, 1978, Time variation in the relationship between inflation and interest rates, *Journal of Monetary Economics* 4, 755-765.

Hamilton, James D., 1985, Uncovering financial market expectations of inflation, *Journal of Political Economy* 93, 1224-1241.

Hamilton, James D., 1988, Rational-expectations econometric analysis of changes in regime: An investigation of the term structure of interest rates, *Journal of Economic Dynamics and Control* 12, 385-423.

Hanson, Samuel G., and Jeremy C. Stein, 2014, Monetary policy and long-term real rates, Working paper, Harvard University.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2012, Inflation expectations, real rates, and risk premia: evidence from inflation swaps, *Review of Financial Studies* 25, 1588-1629.

Johannes, Michael, Lars A. Lochstoer, and Yiqun Mou, 2014, Learning about consumption dynamics, *Journal of Finance*, forthcoming.

Joslin, Scott, Anh Le, and Kenneth J. Singleton, 2013, Why Gaussian macro-finance term structure models are (nearly) unconstrained factor-VARs, *Journal of Financial Economics*.

Le, Anh, and Kenneth J. Singleton, 2013, The structure of risks in equilibrium affine models of bond yields, Working paper, Stanford GSB.

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

Mishkin, Frederic S., 1981, The real rate of interest: an empirical investigation, *Carnegie-Rochester Conference Seies on Public Policy* 15, 151-200.

Mishkin, Frederic S., 1992, Is the Fisher effect for real? A reexamination of the relationship between inflation and interest rates, *Journal of Monetary Economics* 30, 195-215.

Nakamura, Emi, and Jón Steinsson, 2013, High frequency identification of monetary non-neutrality, Working paper, Columbia GSB.

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.

Neely, Christopher J., and David E. Rapach, 2008, Real rate persistence: evidence and implications, *Federal Reserve Bank of St. Louis Review* 90, 609-641.

Nelson, Charles R., and G. William Schwert, 1977, Short-term interest rates as predictors of inflation: On testing the hypothesis that the real rate of interest is constant, *American Economic Review* 67, 478-486.

Pennacchi, George G., 1991, Identifying the dynamics of real interest rates and inflation: evidence using survey data, *Review of Financial Studies* 4, 53-86.

Piazzesi, Monika, and Martin Schneider, 2007, "Equilibrium yield curves," NBER Macroeconomics Annual 2006. MIT Press, Cambridge, MA.

Rudebusch, Glenn D., and Eric Swanson, 2012, The bond premium in a DSGE model with long-run real and nominal risks, *American Economic Journal: Macroeconomics* 4, 105-143.

Schwert, G. William, 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115-1153.

Stock, James H., and Mark W. Watson, 2007, Why has U.S. inflation become harder to forecast? *Journal of Money, Credit, and Banking* 39, 3-33.

van Binsbergen, Jules H., Jesús Fernández-Villaverde, Ralph S.J. Koijen, and Juan F. Rubio-Ramírez, 2012, The term structure of interest rates in a DSGE model with recursive preferences, *Journal of Monetary Economics* 59, 634-648.

Table 1. Properties of nominal yields and expected inflation implied by term structure models

The table reports model-implied population properties of a five-year bond yield and inflation expectations. The source of the models and the sample period used to estimate the parameters are reported in the first two columns. The variance ratios are ratios of the variance of expected average inflation over five years divided by the variance of the five-year yield. The ratios are calculated for both quarterly innovations and levels.

	Sample	Variance ratios	
Source	Period	Innovations	Levels
Piazzesi and Schneider (2007)	1952–2005	1.50	1.51
Campbell and Viceira (2001)	1952–1996	0.63	0.99
Campbell and Viceira (2001)	1983–1996	0.13	0.01
Ang et al. (2008)	1952–2004	0.50	0.50
Chernov and Mueller (2012)	1971–2008	0.03	0.17
Haubrich et al. (2012)	1982–2010	0.17	0.15

Table 2. Summary statistics of inflation forecasts and nominal Treasury bond yields, 1968Q4 through 2013Q4

Inflation forecasts are means, across respondents, from the Survey of Professional Forecasters. Outlier responses are dropped, as discussed in the text. Yields are for zero-coupon nominal Treasury bonds. The column labeled AR(1) reports the first-order serial correlation. The units are annualized percentage points. There are 181 quarterly observations for all time series other than the four-quarter-ahead inflation forecast, for which there are 176 observations.

A. Inflation forecasts

Quarters Ahead	Mean	Std Dev	AR(1)	Std Dev, 1st Diff
1	3.57	1.98	0.98	0.43
2	3.57	1.89	0.98	0.36
3	3.55	1.80	0.99	0.33
4	3.55	1.77	0.98	0.31

B. Yields

Maturity (quarters)	Mean	Std Dev	AR(1)	Std Dev, 1st Diff
1	5.24	3.33	0.96	0.99
4	5.64	3.30	0.97	0.90
8	5.91	3.22	0.98	0.80
12	6.11	3.13	0.98	0.63
16	6.28	3.04	0.98	0.69
20	6.42	2.94	0.98	0.63
40	6.84	2.58	0.98	0.58

Table 3. Decompositions of population variances of yield innovations

The table reports model-implied standard deviations of quarterly shocks to nominal Treasury bond yields. Yields are expressed in percent per year. The table also reports decompositions of the corresponding variances. The model uses four factors to describe the joint dynamics of nominal yields and expected inflation. Yields are the sum of average expected inflation and short-term real rates over the life of the bond, plus a term premium. The contributions to total variance sum to one (aside from rounding). The sample period is 1968Q4 through 2013Q4. Brackets display [2.5% 97.5%] percentile bounds. Panel B reports the sum of the variance column and two covariance columns in Panel A that involve average expected inflation.

A. Full decomposition

Maturity (years)	Std dev	1. Average expected real rate	2. Average expected inflation	3. Term premium	2Cov([1], [2])	2Cov([1], [3])	2Cov([2], [3])
One	0.87	0.66	0.16	0.02	−0.06	0.25	−0.03
	[0.80 1.60]	[0.49 1.13]	[0.06 0.25]	[0.00 0.11]	[−0.54 0.09]	[−0.04 0.48]	[−0.15 0.03]
Five	0.61	0.29	0.18	0.22	0.12	0.35	−0.16
	[0.54 1.08]	[0.12 1.69]	[0.06 0.26]	[0.07 0.80]	[−0.44 0.29]	[−1.03 0.60]	[−0.48 0.10]
Ten	0.51	0.18	0.19	0.33	0.22	0.22	−0.15
	[0.44 0.81]	[0.05 2.26]	[0.05 0.35]	[0.15 1.60]	[−0.50 0.54]	[−2.43 0.52]	[−0.80 0.24]

B. Components related to inflation

	Point estimate	95th Percentile Bounds	
One-yr	−0.45	0.07	0.20
Five-yr	−0.41	0.14	0.22
Ten-yr	−0.45	0.26	0.36

Table 4. Decompositions of population variances of yield shocks, for different factor models

Two different factor models describe the joint dynamics of nominal yields and expected inflation. One uses three factors and the other uses five factors. The sample period is 1968Q4 through 2013Q4. The table reports a subset of the variance decomposition information reported in Table 3. It reports model-implied standard deviations of quarterly shocks to nominal Treasury bond yields, the fraction of variance attributable to expected inflation over the life of the bond, and the fraction of variance attributable to expected inflation and its covariance with expected future real short rates and term premia. Brackets display [2.5% 97.5%] percentile bounds.

A. Three factors

Maturity (years)	Std dev	Average expected inflation	All components related to expected inflation
One	0.86	0.16	0.08
	[0.80 1.46]	[0.07 0.21]	[−0.24 0.20]
Five	0.61	0.18	0.14
	[0.55 0.91]	[0.07 0.25]	[−0.29 0.25]
Ten	0.49	0.20	0.27
	[0.44 0.69]	[0.07 0.33]	[−0.28 0.40]

B. Five factors

Maturity (years)	Std dev	Average expected inflation	All components related to expected inflation
One	0.86	0.16	0.07
	[0.82 2.39]	[0.04 0.29]	[0.00 0.37]
Five	0.59	0.19	0.14
	[0.52 1.50]	[0.03 0.28]	[−0.04 0.31]
Ten	0.53	0.17	0.21
	[0.45 0.99]	[0.02 0.35]	[−0.01 0.30]

Table 5. Decompositions of population variances of yield shocks: subsamples

The table reports model-implied decompositions of sample variances of quarterly shocks to nominal Treasury bond yields. The model describes the joint dynamics of nominal yields and expected inflation. Yields are the sum of average expected inflation and short-term real rates over the life of the bond, plus a term premium. Rows sum to one (aside from rounding).

Sample Period	Number of Factors	Maturity (years)	1. Average expected real rate	2. Average expected inflation	3. Term premium	2Cov([1], [2])	2Cov([1], [3])	2Cov([2], [3])
1968Q4–1987Q3	3	One	0.46	0.18	0.08	−0.05	0.35	−0.02
		Five	0.15	0.25	0.42	0.13	0.34	−0.28
		Ten	0.09	0.23	0.55	0.20	0.19	−0.26
1968Q4–1987Q3	4	One	0.57	0.15	0.06	−0.06	0.35	−0.07
		Five	0.21	0.15	0.41	0.01	0.49	−0.27
		Ten	0.09	0.14	0.79	0.06	0.31	−0.40
1987Q4–2013Q4	3	One	0.66	0.10	0.01	0.14	0.10	−0.01
		Five	0.47	0.08	0.28	0.18	0.00	−0.01
		Ten	0.40	0.08	0.41	0.23	−0.13	0.00
1987Q4–2013Q4	4	One	0.81	0.14	0.02	0.14	−0.07	−0.04
		Five	0.92	0.09	0.47	0.27	−0.64	−0.11
		Ten	0.80	0.09	0.76	0.32	−0.87	−0.10

Table 6. The contribution of inflation to bond yield variance

The table reports the fraction of bond yield variance explained by the variance of expected inflation over the life of the bond. Variances are determined by the population properties of a dynamic multifactor model of yields and expected inflation. The models are estimated over the sample 1968Q4 through 2013Q4. They differ in the number of factors. Brackets display [2.5% 97.5%] percentile bounds.

	Maturity		
Factors	One year	Five years	Ten years
3	0.24	0.23	0.24
	[0.10 0.77]	[0.11 0.65]	[0.10 0.62]
4	0.25	0.24	0.24
	[0.04 1.55]	[0.07 1.14]	[0.08 1.22]
5	0.22	0.21	0.21
	[0.03 1.98]	[0.04 1.83]	[0.05 1.45]

Table 7. Explaining survey forecast errors with forecast revisions

The table reports estimates of (19) in the text, 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	Coef (Std err)	R^2
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)	
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 8. Estimated properties of a trend-cycle model of inflation

Panel A reports parameter estimates of a trend-cycle model of inflation. Inflation is the sum of two components. One follows a random walk and the other follows an AR(1). The shocks are drawn from a multivariate normal distribution. The Cholesky decomposition of the covariance matrix is denoted Ω. The model is estimated using quarterly mean survey forecasts of expected inflation one quarter ahead and three-quarters ahead. The two components are treated as unobserved and the observed survey forecasts are assumed to be contaminated by measurement error. The model is estimated with the Kalman filter over three sample periods. Panel B reports model-implied population standard deviations of shocks to the expected average ten-year inflation rate. For comparison, corresponding standard deviations are reported for four-factor dynamic models of yields and inflation that impose stationarity.

A. Parameter estimates

Parameter	1968Q4–2013Q4	1968Q4–1978Q3	1978Q4–2013Q4
AR(1) coef	0.658	0.668	0.453
Std dev, trend shock (b.p.)	21	27	23
Std dev, cycle shock (b.p.)	54	71	53
Shock correlation	0.31	0.50	−0.46
Std dev of measurement error,			
1-q-ahead forecast (b.p.)	0	0	0
Std dev of measurement error			
3-q-ahead forecast (b.p.)	16	20	3

B. Standard deviations of shocks to ten-year inflation (b.p.)

Model	1968Q4–2013Q4	1968Q4–1978Q3	1978Q4–2013Q4
Trend-cycle	22	28	23
Stationary, three factors	22	32	12
Stationary, four factors	22	26	13

Table 9. Variance decompositions using a model of nonstationary inflation and yields

The table reports model-implied variance decompositions of quarterly shocks to a ten-year nominal yield. The model is a first-order vector autoregression of monthly data. The observed data are described in Section 5.2.2. The model imposes cointegration on yields and inflation, while real rates are stationary. Yields are the sum of average expected inflation and short-term real rates over the life of the bond, plus a term premium. The contributions to total variance sum to one (aside from rounding). Brackets display [2.5% 97.5%] percentile bounds.

Sample period	1. Average expected real rate	2. Average expected inflation	3. Term premium	2Cov([1], [2])	2Cov([1], [3])	2Cov([2], [3])
10/1968–12/2013	0.02	1.27	0.60	−0.06	−0.02	−0.80
	[0.01 0.10]	[0.90 2.78]	[0.23 1.98]	[−0.40 0.14]	[−0.28 0.24]	[−3.45 0.23]
10/1968–9/1987	0.02	0.97	0.85	−0.15	0.21	−0.89
	[0.01 0.39]	[0.65 5.89]	[0.30 4.53]	[−2.10 0.05]	[−0.06 1.76]	[−9.210.21]
10/1987–12/2013	0.02	0.92	0.18	−0.07	−0.02	−0.02
	[0.01 0.17]	[0.60 2.65]	[0.07 1.09]	[−0.78 0.09]	[−0.23 0.37]	[−2.37 0.25]

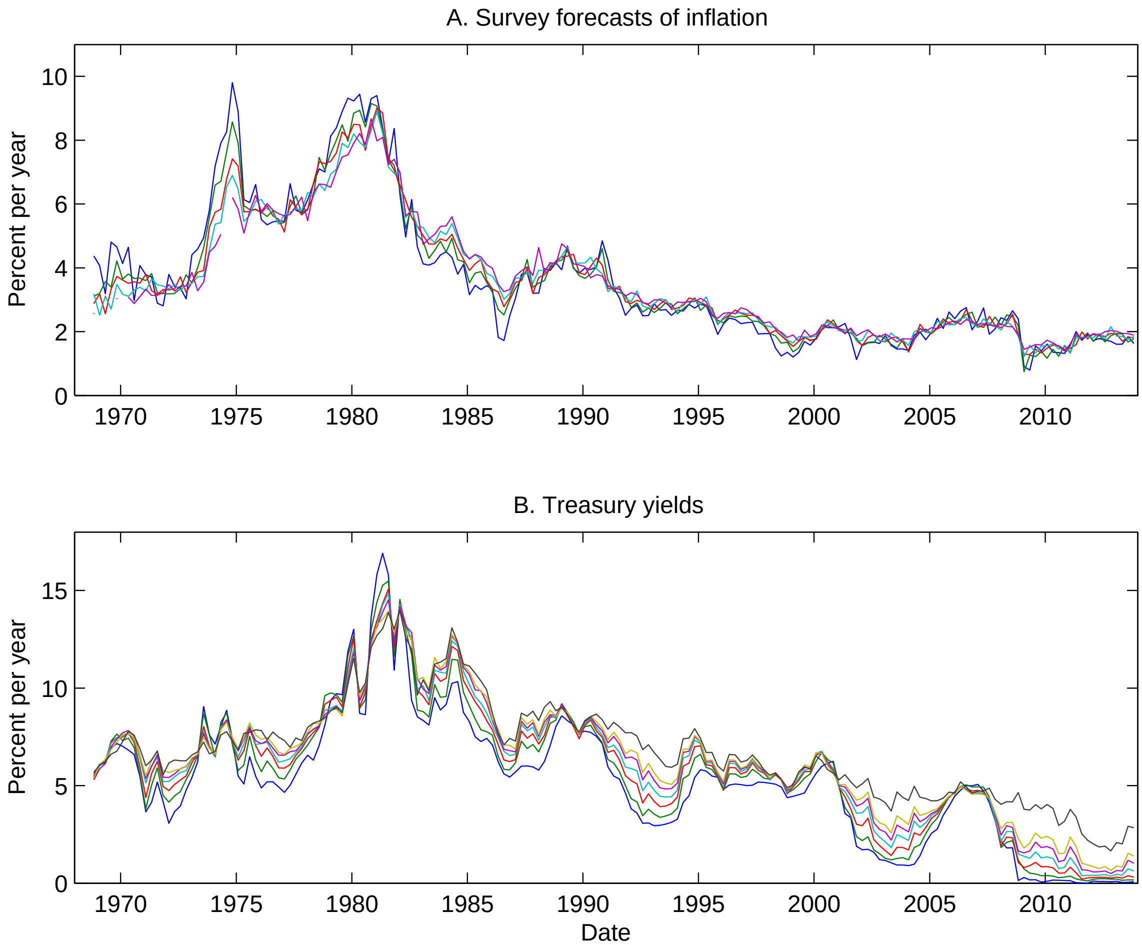

Figure 1. Expected inflation and Treasury bond yields

Panel A displays mean forecasts, from the Survey of Professional Forecasters, of the log change in the GDP price index for zero to four quarters ahead. Panel B displays yields on zero-coupon nominal Treasury bonds with maturities of three months, one through five years, and ten years.

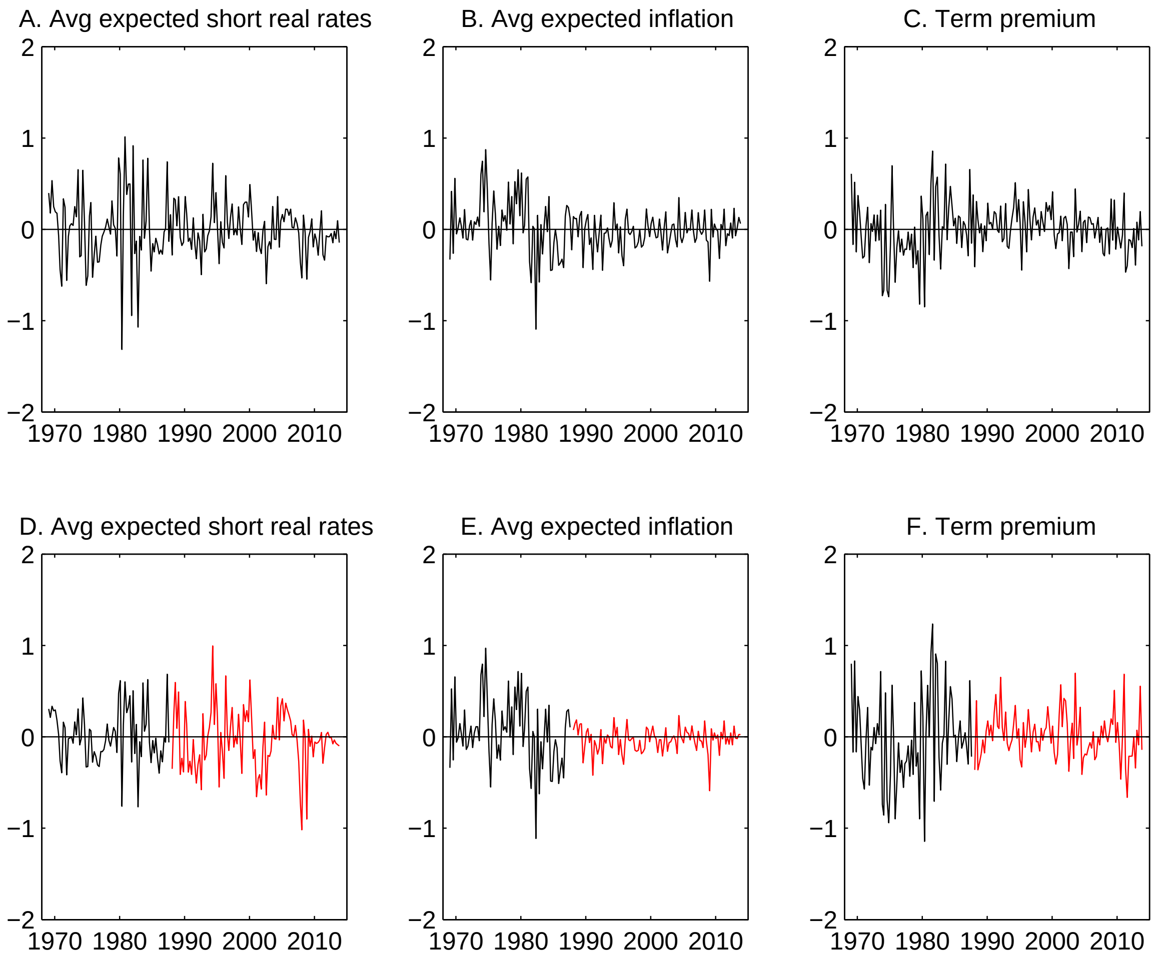

Figure 2. Shocks to components of the five-year nominal Treasury yield

Quarterly shocks to nominal yields are the sum of shocks to average expected real rates during the life of the bond, average expected inflation during the life of the bond, and term premia. The figure plots fitted shocks for a five-year zero-coupon Treasury bond, where the shocks are inferred from an estimated dynamic model of yields and expected inflation. Panels A through C use a four-factor model estimated over the sample 1968Q4 through 2013Q4. Panels D, E, and F splice two series from three-factor models. The earlier series uses data from 1968Q4 to 1987Q3 and the latter series use data from 1987Q4 through 2013Q4. Shocks are in percentage points of annualized yields.

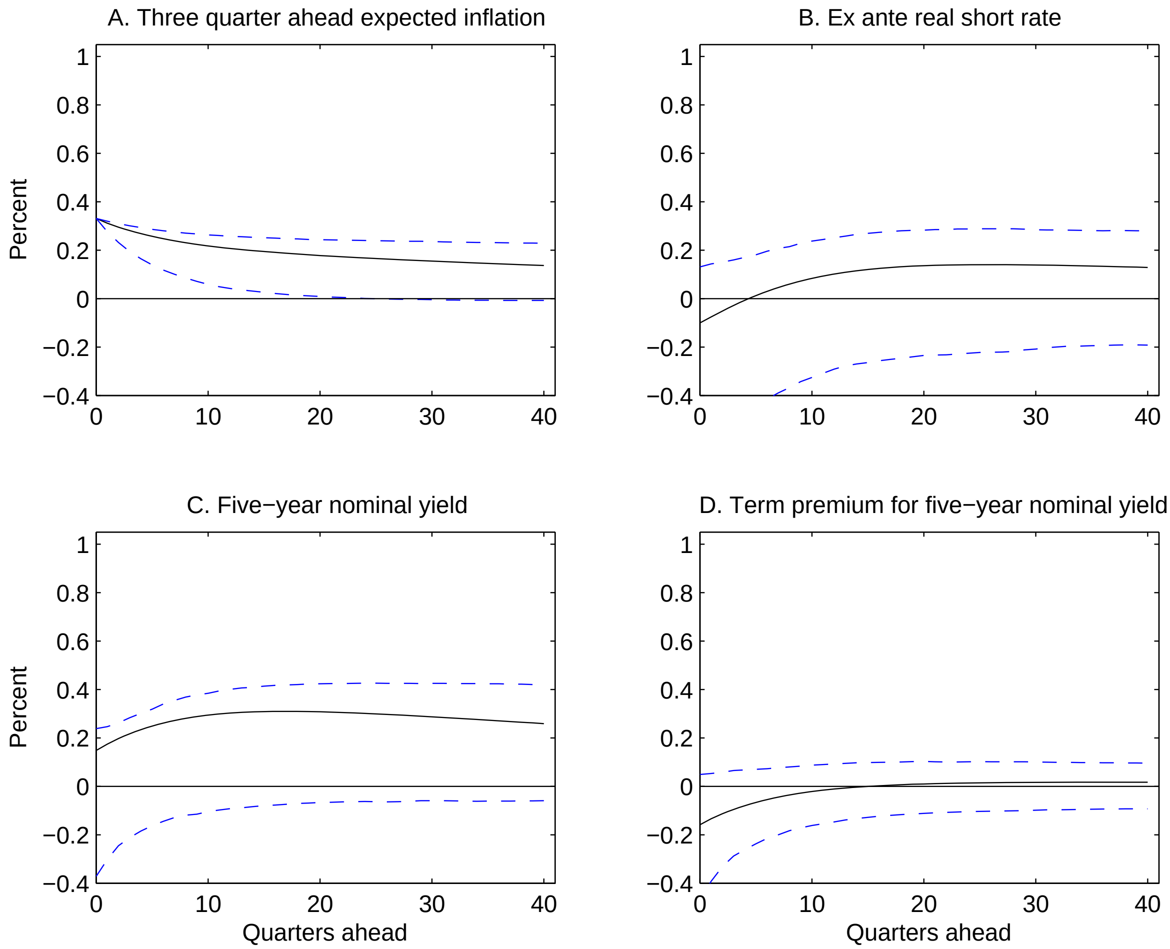

Figure 3. Impulse responses for a shock to expected inflation

A four-factor dynamic model of nominal yields and expected inflation is estimated over the sample 1968Q4 through 2013Q4. The figure displays model-implied impulse responses to a shock to three-quarter-ahead expected inflation. The initial shock is 33 basis points, which is the model-implied population standard deviation of the shock. The yield and term premium responses are for a five-year nominal bond. Also displayed are 95 percentile confidence bounds on the impulse responses.

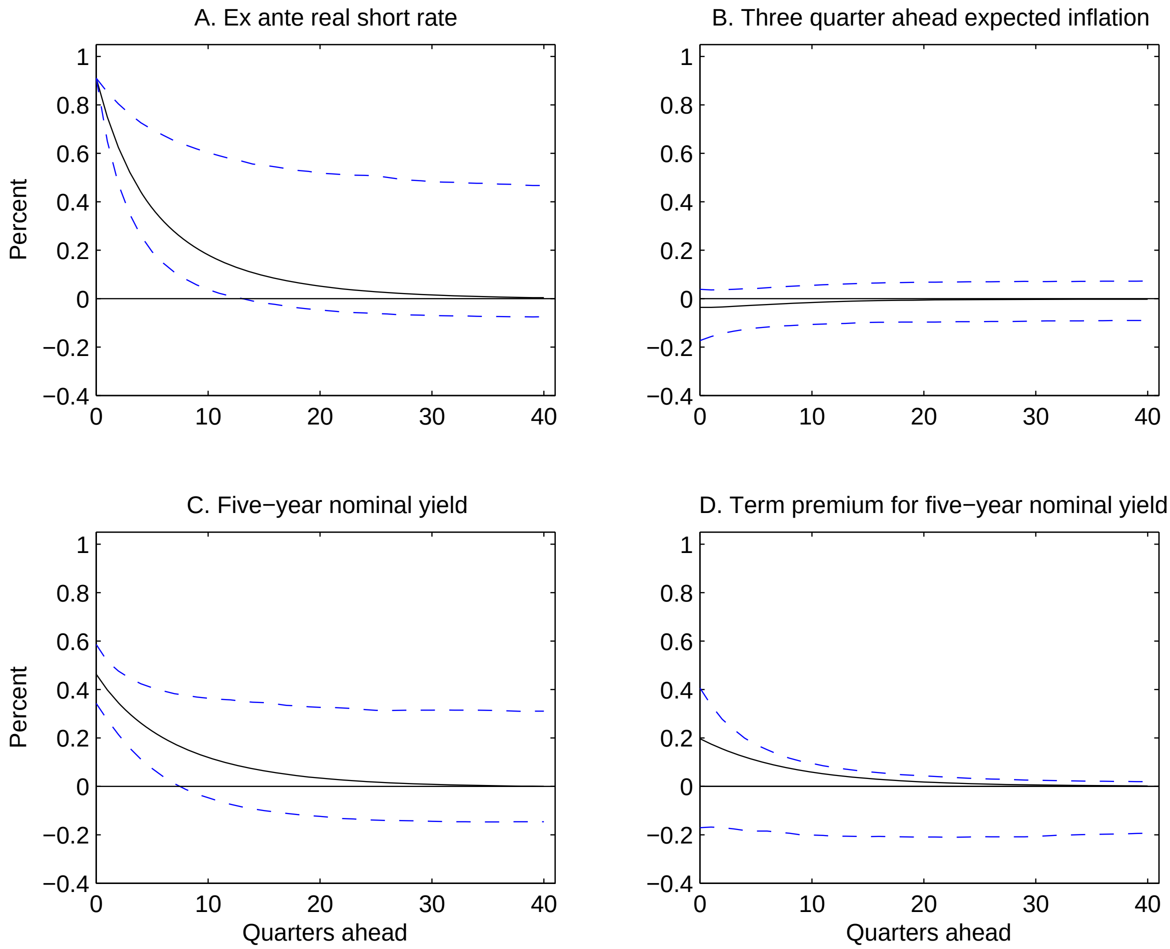

Figure 4. Impulse responses for a shock to the ex ante real rate

A four-factor dynamic model of nominal yields and expected inflation is estimated over the sample 1968Q4 through 2013Q4. The figure displays model-implied impulse responses to a shock to the ex ante real rate, defined as the three-month nominal yield less expected inflation during the next quarter. The initial shock is 91 basis points, which is the model-implied population standard deviation of the shock. The yield and term premium responses are for a five-year nominal bond. Also displayed are 95 percentile confidence bounds on the impulse responses.

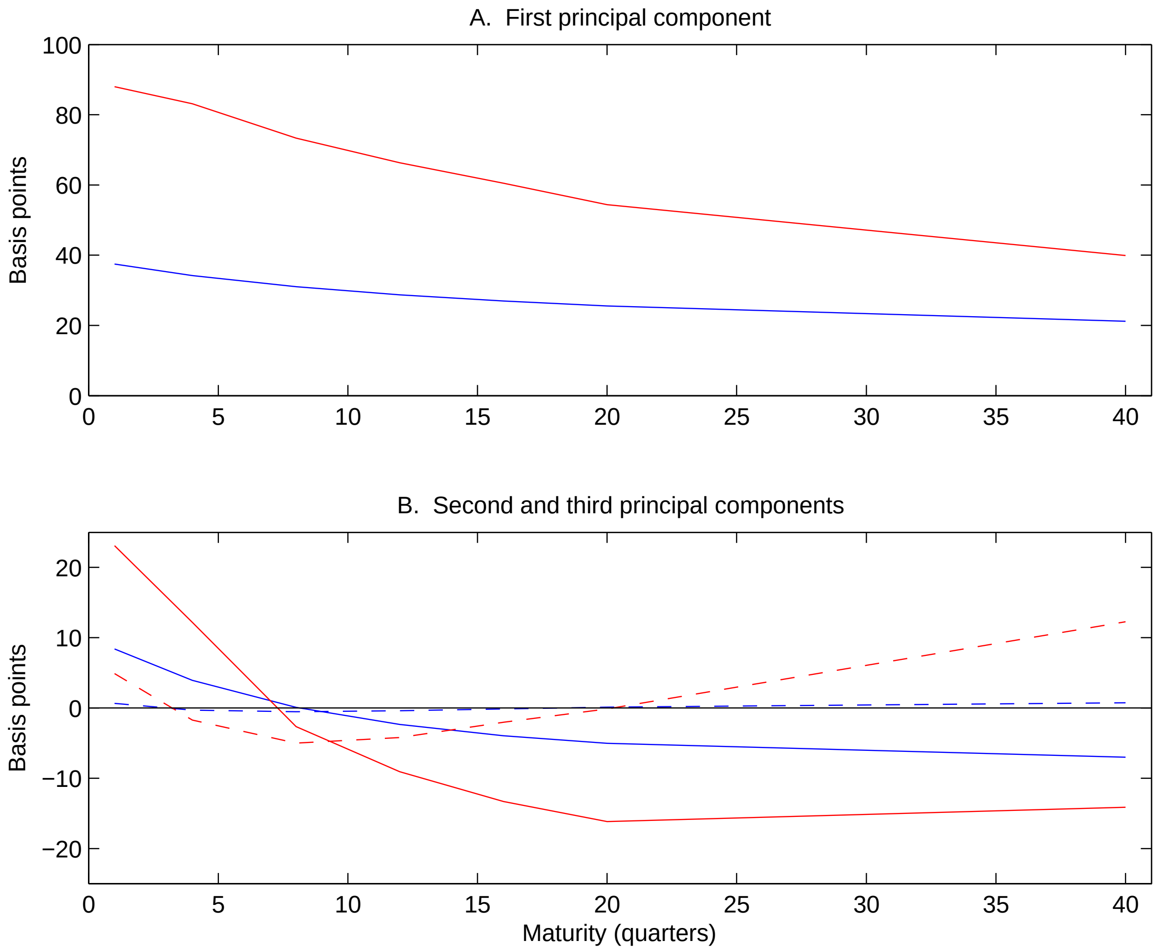

Figure 5. Principal component decompositions of shocks to bond yields

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 with 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.

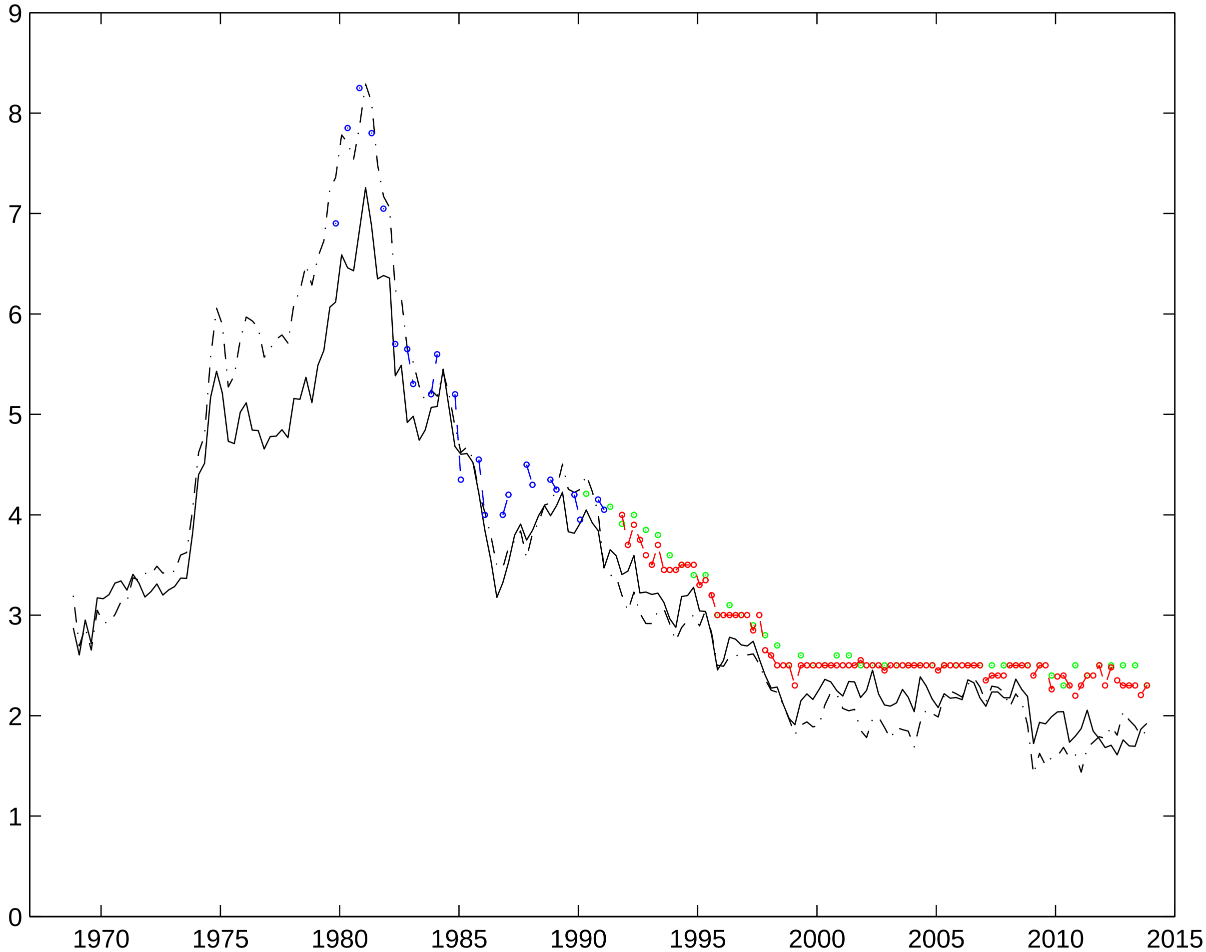

Figure 6. Ten-year inflation forecasts from surveys and models

A four-factor dynamic model of nominal yields and expected inflation is estimated over the sample 1968Q4 through 2013Q4. The solid black line represents the model's predictions of average GDP inflation during the next ten years. A trend-cycle model of expected inflation is estimated over the same period; its predictions are displayed with the dashed black line. The blue dots are mean ten-year inflation forecasts from Blue Chip Economic Indicators. Prior to 1983, Blue Chip forecasts are for GDP inflation. Subsequently they are for CPI inflation. The green dots are mean forecasts of CPI inflation from the Livingston Survey. The red dots are mean forecasts of CPI inflation from the Survey of Professional Forecasters.

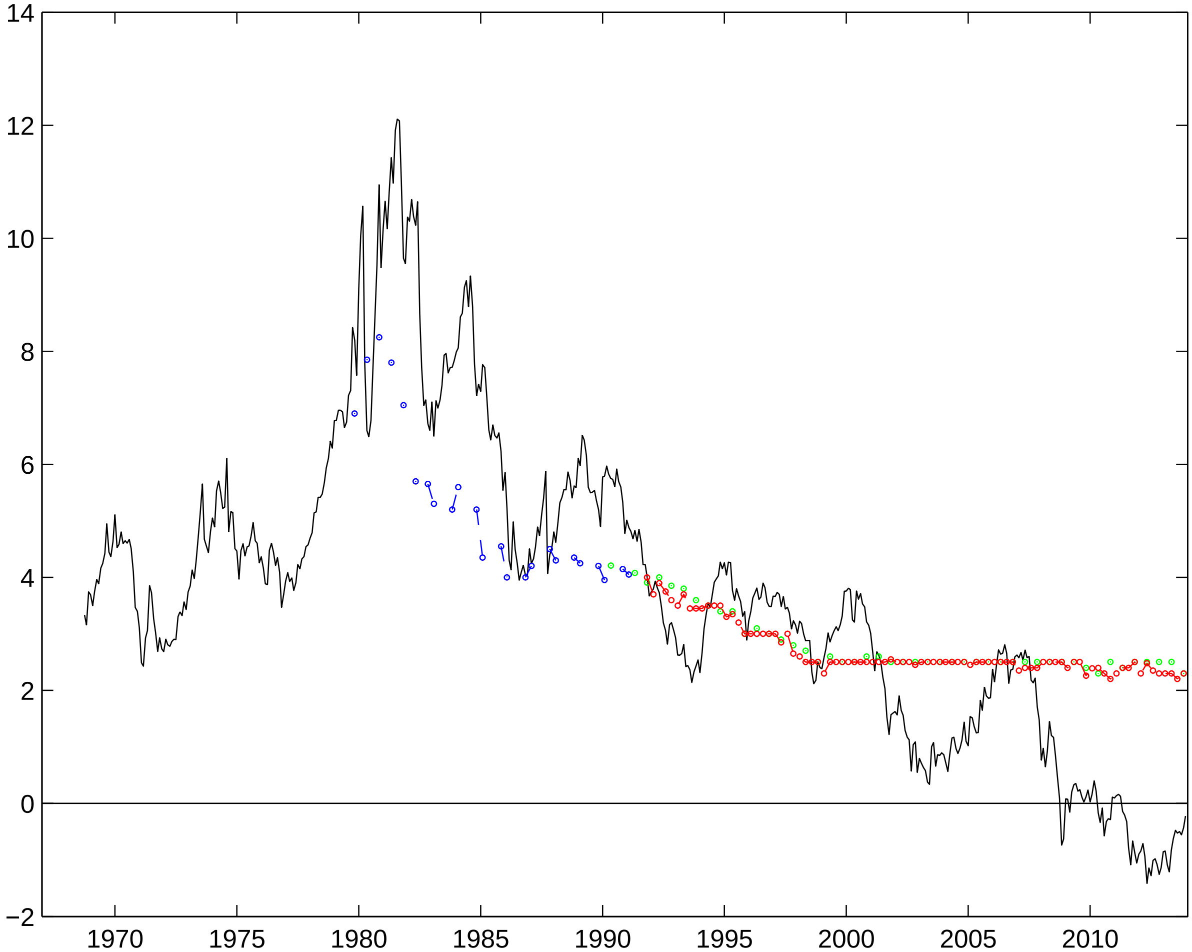

Figure 7. Ten-year inflation forecasts from surveys and a model of nonstationary inflation and yields

Section 5.2.2 describes a joint model of inflation and nominal yields that assumes no components of inflation are stationary. The solid black line represents the model's prediction of average CPI inflation during the next ten years. The model's parameters are estimated using the sample October 1968 through December 2013. The blue dots are mean ten-year inflation forecasts from Blue Chip Economic Indicators. Prior to 1983, Blue Chip forecasts are for GDP inflation. Subsequently they are for CPI inflation. The green dots are mean forecasts of CPI inflation from the Livingston Survey. The red dots are mean forecasts of CPI inflation from the Survey of Professional Forecasters.

Appendix

This appendix describes how the model is estimated and reports parameter estimates.

6.1 Normalizations

Model estimation uses a factor rotation of (10) 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

$$x_{t+1} = Dx_t + \Sigma\epsilon_{t+1}, \tag{24}$$

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 compex eigenvalues. Estimation of a version that allows for complex eigenvalues produced only real eigenvalues, thus (24) is not restrictive in practice.

Parameters of the mapping from factors to yields (11) are unrestricted. The corresponding parameters of (12) for the three-quarter-ahead inflation forecast are also unrestricted. For expectational consistency, the constant for the one-quarter-ahead forecast is set equal to the constant for the three-quarter-ahead forecast. Again for consistency, the factor loadings of the one-quarter-ahead forecast are set equal to

$$B_{\pi,1} = B_{\pi,3} \ ./ \ \mathrm{diag}(D),$$

where the operator is element-by-element division.

6.2 Local underidentification

This normalized model is identified econometrically with the time series properties of the observables as long as the eigenvalues of D are distinct. Factor i follows an autonomous process with AR(1) coefficient equal to D_{ii}. With distinct eigenvalues, linear combinations of these factors do not follow autonomous processes. Factor shocks are arbitrarily correlated across factors. Standard deviations of orthogonalized shocks are fixed.

Additional normalizations are required when the eigenvalues are not distinct. If, say, eigenvalues i and j are equal, then any linear combination of factors i and j also follows an autonomous AR(1) process. The model is locally unidentified (i.e., unidentified for these eigenvalues) because factor i can be replaced in (24) 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).

6.3 Local underidentification of unrestricted models

Table A1 reports maximum likelihood (ML) estimates of the eigenvalues and the Σ matrix for three-, four-, and five-factor models with dynamics (24). The sample period is 1968Q4 through 2013Q4. Likelihood function calculation and optimization uses Fortran with IMSL subroutines.[13]

The estimated models all have a single eigenvalue close to one—almost a unit root. The other estimated eigenvalues are sufficiently close to each other to produce local underidentification. Two of the estimated eigenvalues of the three-factor model differ by 4/1000, while two of the estimated eigenvalues of the five-factor model differ by 3/1000. The local underidentification of these two models is so severe that the covariance matrix of the parameter estimates is singular to machine precision.[14]

Underidentification is not quite as severe for the four-factor model. Two eigenvalues differ by about 2/100. For this model the covariance matrix of the parameter estimates is nonsingular. Nonetheless, there is not enough information in the sample to identify cleanly the factor rotation. As the Table documents, estimated standard errors on the elements of Σ range up to ten times the magnitude of any element of the matrix. More importantly, the distribution of parameter estimates implied by ML asymptotics is not sensible. The true distribution, whatever it is, is unquestionably nothing like the distribution implied by ML asymptotics. Monte Carlo simulations summarized in Table A2 make this clear.

Panel A's first row reports the estimated population standard deviation of quarterly shocks to the one-year yield. This standard deviation is is a function of only the model's parameter estimates. The estimate of 87 basis points is sensible, given the 90 basis point

[13]Previous versions of this paper used Matlab Simplex and derivative-based optimization routines. These routines often failed to locate the maximum likelihood, presumably because the local underidentification produces a nearly flat likelihood surface.

[14]The covariance matrix is estimated using the outer product of numerical first derivatives.

sample standard deviation of quarterly changes in the one-year yield reported in Table 2. The other elements of the row summarize Monte Carlo simulations of this population standard deviation. Each set of simulated parameters is drawn from a multivariate normal distribution with mean equal to the parameter estimates and a covariance matrix equal to the estimated covariance matrix of the parameter estimates.[15] The mean simulated standard deviation exceeds 1300 basis points. This is more than ten times both the point estimate and any plausible true value of the population standard deviation.

6.4 Choosing restricted models

Imposing restrictions on the Σ matrix reduces the problem of local underidentification. For the three-factor model, the natural restriction sets element (3,2) of the matrix to zero. This sets to zero the correlation between shocks to the second and third factors that are not shared with the first factor. Appropriate restrictions on the matrix for the four- and five-factor models are less clear. There is no clear line between restrictions imposed to reduce underidentification and restrictions that incorrectly imply extreme confidence in a point estimate. The former (latter) produce more (less) accurate covariance matrices of parameter estimates.

For the four-factor model, my method to restrict Σ begins with the observation that the single eigenvalue near one is not economically close to the other three eigenvalues. Therefore I treat it as well-identified and do not restrict the off-diagonal elements of the first column of Σ. Element (3,2) is set to zero because eigenvalue two is close to eigenvalue three. Estimation of this one-restriction model results in Monte Carlo distribution of parameter estimates that is still not well-behaved. The second row of Panel A in Table A2 reports the Monte Carlo mean standard deviation of quarterly shocks to the one-year yield exceeds 130 basis points and the upper 97.5 percentile upper bound is close to 340 basis points.

Therefore element (4,2) is also set to zero. The table reports Monte Carlo simulation results for this two-restriction model. The simulated properties of the standard deviation are better than those corresponding to the one-restriction model. Nonetheless, the mean standard deviation is 30 basis points greater than the point estimate, and the upper bound of nearly 300 b.p. is much too high. Finally, the remaining off-diagonal element (4,3) is set

[15] Simulations occasionally draw eigenvalues that exceed one. These are replaced with 0.999, effectively treating the draw as one that imposes a unit root.

to zero. This three-restriction model is used in the text. Both the mean standard deviation of 105 b.p. and the upper bound of 160 b.p. are much more realistic. Unfortunately, the lower bound of 80 basis points seems a little high.

A similar exercise for the five-factor model produces a similar model. The estimated model in the text has unrestricted off-diagonal elements for the first column of Σ. All other off-diagonal elements are set to zero. Monte Carlo simulations of the standard deviation of shocks to the one-year yield are summarized in Panel B of Table A2.

For the sake of consistency, models estimated over the two subsamples (pre-Greenspan and post-Volcker) use the same restrictions developed based on full-sample properties. Full-sample parameter estimates and standard errors for all three restricted models studied in the text are reported in Tables A3, A4, and A5.

It is important to recognize that restrictions on Σ have substantial effects on estimates of parameter uncertainty but almost no effect on point estimates of inflation risk. Panel C of Table A2 reports direct measures of inflation risk for the unrestricted models estimated in this Appendix. They are almost identical to those reported in Tables 3 and 4, which are calculated from the models with restricted versions of Σ.

Table A1. Partial parameter estimates for models estimated without restrictions on factor dynamics

The models described here factor dynamics described by (24) without any additional parameter restrictions. They are estimated over the entire sample 1968Q4 through 2013Q4. Standard errors are in parentheses for the four-factor model.

A. Three factors

	Column		
Parameter	1	2	3
Eigenvalues, D_{ii}	0.989	0.847	0.843
Σ	10		
	0.75	10	
	-12.64	-184.58	10

B. Four factors

	Column			
Parameter	1	2	3	4
Eigenvalues, D_{ii}	0.989	0.886	0.864	0.745
	(0.012)	(0.112)	(0.074)	(0.106)
Σ	10			
	(-)			
	1.44	10		
	(6.89)	(-)		
	0.00	-11.42	10	
	(8.83)	(114.01)	(-)	
	0.25	-1.20	-5.24	10
	(3.75)	(13.02)	(13.38)	(-)

C. Five factors

	Column				
Parameter	1	2	3	4	5
Eigenvalues, D_{ii}	0.989	0.890	0.887	0.729	0.708
Σ	10				
	2.25	10			
	-31.05	-141.66	10		
	-0.60	-6.72	-8.71	10	
	4.62	33.27	36.49	-51.63	10

Table A2. Features of estimated models

Unrestricted models have factor dynamics described by (24) without any parameter restrictions. Other models have restrictions on Σ in (24). All models are estimated over the entire sample 1968Q4 through 2013Q4. The appendix describes the standard deviation estimates in Panels A and B. Standard deviations of quarterly shocks to the one-year yield are measured in basis points of annualized yields In Panel C, inflation risk is the variance of quarterly news about average inflation over a bond's life as a fraction of the variance of quarterly yield shocks.

A. Four-factor estimates of the standard deviation of one-year yield shocks

Model	Point	Monte Carlo Distribution		
restrictions	estimate	2.5%	Mean	97.5%
None	87	73	1302	6535
$\Sigma_{32} = 0$	87	61	129	338
$\Sigma_{32} = \Sigma_{43} = 0$	88	58	118	289
$\Sigma_{32} = \Sigma_{43} = \Sigma_{42} = 0$	87	80	105	160

B. Five-factor estimates of the standard deviation of one-year yield shocks

Model	Point	Monte Carlo Distribution		
restrictions	estimate	2.5%	Mean	97.5%
None	87	-	-	-
$\Sigma_{32} = \Sigma_{52} = \Sigma_{53} = \Sigma_{54} = 0$	86	60	213	713
$\Sigma_{32} = \Sigma_{43} = \Sigma_{52} = \Sigma_{53} = \Sigma_{54} = 0$	86	57	131	423
$\Sigma_{32} = \Sigma_{42} = \Sigma_{43} = \Sigma_{52} = \Sigma_{53} = \Sigma_{54} = 0$	86	82	122	239

C. Estimates of inflation risk based on unrestricted models

Bond Maturity (yr)	3 factors	4 factors	5 factors
One	0.16	0.16	0.15
Five	0.18	0.18	0.19
Ten	0.20	0.19	0.17

Table A3. Full sample estimates of the restricted three-factor model

The sample period is 1968Q4 through 2013Q4. For estimation, all observables are expressed in annualized basis points. For example, 100 corresponds to one percent per year. In this table the estimated unconditional means of the observables are expressed in annualized percentage points.

A. Factor dynamics

	Column		
Parameter	1	2	3
Eigenvalues, D_{ii}	0.988	0.876	0.815
	(0.011)	(0.037)	(0.051)
Σ	10		
	(-)		
	2.80	10	
	(3.39)	(-)	
	-1.00	0.00	10
	(3.00)	(-)	(-)

B. Observables mapping

Observable	Mean (%)	Factor loadings			SD of meas error (b.p.)
1-q ahead forecast	2.71	2.58	1.68	1.88	32
	(-)	(-)	(-)	(-)	(2.5)
3-q ahead forecast	2.71	2.52	1.29	1.25	8
	(1.79)	(0.30)	(1.07)	(0.89)	(4.6)
3 mon yield	3.54	5.04	-5.48	6.72	47
	(3.86)	(1.44)	(4.16)	(3.07)	(3.5)
1 yr yield	3.90	5.15	-5.66	5.78	15
	(3.91)	(1.38)	(3.67)	(3.09)	(2.0)
2 yr yield	4.15	5.18	-5.57	4.21	7
	(3.94)	(1.27)	(2.83)	(3.12)	(0.8)
3 yr yield	4.38	5.10	-5.36	3.15	7
	(3.87)	(1.17)	(2.28)	(3.04)	(0.7)
4 yr yield	4.58	5.00	-5.14	2.28	6
	(3.80)	(1.09)	(1.83)	(2.94)	(0.5)
5 yr yield	4.76	4.86	-4.80	1.58	9
	(3.69)	(1.00)	(1.49)	(2.78)	(0.9)
10 yr yield	5.38	4.25	-3.74	0.11	32
	(3.22)	(0.75)	(0.87)	(2.28)	(2.3)

Table A4. Full sample estimates of the restricted four-factor model
The sample period is 1968Q4 through 2013Q4. For estimation, all observables are expressed in annualized basis points. For example, 100 corresponds to one percent per year. In this table the estimated unconditional means of the observables are expressed in annualized percentage points.

A. Factor dynamics

	Column			
Parameter	1	2	3	4
Eigenvalues, D_{ii}	0.988	0.900	0.860	0.735
	(0.012)	(0.070)	(0.032)	(0.083)
Σ	10			
	(-)			
	−3.70	10		
	(5.01)	(-)		
	0.95	0	10	
	(3.11)	(-)	(-)	
	−0.24	0	0	10
	(2.62)	(-)	(-)	(-)

B. Observables mapping

Observable	Mean (%)	Factor loadings				SD of meas error (b.p.)
1-q ahead forecast	2.76	2.65	0.10	2.38	0.96	32
	(-)	(-)	(-)	(-)	(-)	(2.7)
3-q ahead forecast	2.76	2.59	0.08	1.76	0.52	8
	(2.03)	(0.47)	(1.92)	(0.59)	(0.53)	(5.0)
3 mon yield	3.65	4.96	5.98	−1.32	6.85	35
	(4.45)	(2.05)	(2.86)	(4.92)	(4.92)	(2.2)
1 yr yield	4.01	5.08	6.20	−1.60	5.20	0
	(4.48)	(2.04)	(2.49)	(4.91)	(4.91)	(-)
2 yr yield	4.26	5.14	5.89	−2.03	3.21	7
	(4.48)	(1.92)	(2.18)	(4.53)	(4.53)	(0.6)
3 yr yield	4.49	5.07	5.40	−2.33	2.28	6
	(4.40)	(1.78)	(2.10)	(4.12)	(4.12)	(0.6)
4 yr yield	4.69	4.99	4.89	−2.58	1.67	7
	(4.31)	(1.65)	(2.11)	(3.69)	(3.69)	(0.6)
5 yr yield	4.88	4.86	4.35	−2.68	1.19	9
	(4.19)	(1.51)	(2.05)	(3.27)	(3.27)	(0.7)
10 yr yield	5.48	4.29	2.45	−2.97	1.07	16
	(3.72)	(1.08)	(2.16)	(1.89)	(1.89)	(2.5)

Table A5. Full sample estimates of the restricted five-factor model

The sample period is 1968Q4 through 2013Q4. For estimation, all observables are expressed in annualized basis points. For example, 100 corresponds to one percent per year. In this table the estimated unconditional means of the observables are expressed in annualized percentage points.

A. Factor dynamics

	Column				
Parameter	1	2	3	4	5
Eigenvalues, D_{ii}	0.988	0.921	0.867	0.751	0.623
	(0.015)	(0.044)	(0.062)	(0.095)	(0.096)
Σ	10				
	(-)				
	4.58	10			
	(9.07)	(-)			
	-1.76	0.00	10		
	(4.57)	(-)	(-)		
	-0.36	0.00	0.00	10	
	(2.41)	(-)	(-)	(-)	
	1.59	0.00	0.00	0.00	10
	(2.73)	(-)	(-)	(-)	(-)

B. Observables mapping

Observable	Mean (%)	Factor loadings					SD of meas error (b.p.)
1-q ahead forecast	2.89	2.60	0.93	1.70	0.70	1.55	30
	(-)	(-)	(-)	(-)	(-)	(-)	(2.8)
3-q ahead forecast	2.89	2.53	0.79	1.28	0.39	0.60	9
	(1.97)	(0.63)	(1.91)	(1.27)	(1.03)	(0.40)	(4.5)
3 mon yield	3.74	5.67	−4.47	3.44	6.89	−0.59	20
	(4.59)	(2.95)	(3.93)	(4.95)	(2.07)	(4.39)	(5.7)
1 yr yield	4.13	5.72	−4.41	3.58	5.27	−2.81	7
	(4.61)	(2.87)	(3.87)	(4.57)	(2.46)	(3.37)	(2.4)
2 yr yield	4.38	5.75	−4.37	2.96	3.45	−3.33	6
	(4.61)	(2.71)	(3.34)	(4.21)	(2.60)	(2.27)	(1.0)
3 yr yield	4.61	5.66	−4.22	2.30	2.58	−3.23	6
	(4.53)	(2.55)	(2.85)	(3.98)	(2.51)	(1.80)	(0.7)
4 yr yield	4.80	5.55	−4.05	1.69	2.00	−3.06	7
	(4.45)	(2.39)	(2.43)	(3.73)	(2.36)	(1.46)	(0.6)
5 yr yield	4.99	5.38	−3.76	1.18	1.54	−2.94	8
	(4.32)	(2.21)	(2.11)	(3.46)	(2.29)	(1.27)	(0.9)
10 yr yield	5.60	4.65	−2.57	−0.44	1.56	−3.09	10
	(3.81)	(1.58)	(1.36)	(2.51)	(2.26)	(1.39)	(3.8)