

Expected Inflation and Other Determinants of Treasury Yields

GREGORY R. DUFFEE*

ABSTRACT

Shocks to nominal bond yields consist of news about expected future inflation, expected future real short rates, and expected excess returns—all over the bond's life. I estimate the magnitude of the first component for short- and long-maturity Treasury bonds. At a quarterly frequency, variances of news about expected inflation account for between 10% to 20% of variances of yield shocks. Standard dynamic models with long-run risk imply variance ratios close to 1. Habit formation models fare somewhat better. The magnitudes of shocks to real rates and expected excess returns cannot be determined reliably.

A LARGE AND EXPANDING LITERATURE explores the relation between nominal bond yields and inflation. In a particularly important contribution, Ang and Piazzesi (2003) introduce Gaussian macro-finance dynamic term structure models to determine the compensation investors require to face shocks to inflation and macroeconomic activity. Subsequent studies have 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. Yet it is difficult to uncover from this literature any widely accepted conclusions about the joint dynamics of inflation and the nominal term structure. Motivated by this point, Ang, Bekaert, and Wei (2008) attempt to produce some basic facts. More recent research does not converge on their conclusions or any other set of core results. Thus, it remains unclear which branches of the macro-finance literature are likely to be fruitful and which should be abandoned.

In this paper, I make an additional attempt to identify a robust empirical property that can be used to guide future research. I focus on the question “How large are shocks to expected inflation relative to shocks to nominal bond yields?” Embedded in this question is an accounting identity. Campbell and Ammer (1993) show that the shock to a nominal bond's yield equals the sum of news about expected inflation, expected short-term real rates, and expected excess returns, all over the life of the bond. The “inflation variance ratio,” as

*Gregory R. Duffee is with Johns Hopkins University. I thank seminar participants at many schools and conferences, Ravi Bansal, Mike Chernov, Anna Cieslak, George Constantinides, Lars Lochstoer, Kenneth Singleton (Editor), Jonathan Wright, and three anonymous referees for helpful comments. I thank especially the discussants Anh Le and Scott Joslin. I have read the *Journal of Finance's* disclosure policy and have no conflicts of interest to disclose.

defined here, is the variance of expected inflation news relative to the variance of the yield shock.

Three observations motivate a focus on this ratio. First, it can be estimated without much structure, using survey forecasts of inflation to identify revisions in inflation expectations. Second, ratios based on quarterly U.S. data are reliable, in the sense that standard errors are tight and estimates are reasonably stable over time. Third, ratios inferred from these data are strongly at odds with corresponding values from both endowment economy long-run risk models and standard New Keynesian dynamic models.

Results for the 10-year horizon are representative. During the past 35 years, the standard deviation of quarterly shocks to 10-year average expected inflation is in the neighborhood of 20 basis points. The standard deviation of quarterly shocks to the 10-year bond yield is much larger—around 60 basis points. Squaring and dividing produces a variance ratio estimate close to 10%.

The challenge for economists is easier to see when viewing the result from a different angle. In the data, shocks to nominal yields are large and driven primarily by a combination of news about expected short-term real rates and expected excess returns. In our benchmark macro-finance models, channels for both types of news are small.

Bansal and Yaron (2004) develop one of these benchmark models, combining a representative agent, recursive utility preferences, and persistent fluctuations in the endowment growth rate. Short-term real rates are driven largely by fluctuations in expected consumption growth. The long-run risk literature follows Bansal and Yaron (2004) by relying on a high elasticity of intertemporal substitution and fairly small shocks to expected growth. The combination of these properties results in low volatility of news about expected short-term real rates.

These conditionally log-normal models generate news about expected excess returns only through shocks to conditional volatilities of macroeconomic shocks. The amount of news this mechanism produces depends on the average level of bond risk premia, which are much too small to allow for sizable volatilities of news about expected excess returns.

Dynamic New Keynesian models also produce low volatilities of real rate news because the dynamics are not sufficiently persistent. Since these models do not have conditional log-normal dynamics, nonlinearities can potentially create more news about expected excess returns than can models in the tradition of Bansal and Yaron (2004). But for parameterized models in the literature, the nonlinearities are not sufficient to generate realistic volatilities of shocks to yields.

Habit formation preferences in the spirit of Campbell and Cochrane (1999) break the link between expected consumption growth and short-term real rates, creating another mechanism for generating real rate news. In addition, the models' nonlinearities can create substantial news about expected excess returns. We can therefore choose parameterizations of these models that are consistent with observed inflation variance ratios. However, the evidence is much too tentative to warrant the conclusion that nominal bond dynamics are

best understood through the lens of habit formation. In particular, parameterizations that are successful in reproducing inflation variance ratios exhibit other properties that appear implausible.

Empirically, innovations to expected short-term real rates and expected excess returns are the primary drivers of yield shocks. Unfortunately, there is insufficient information in the data to disentangle the relative contributions of these two components, at least without imposing restrictive assumptions. 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 shocks to expected excess returns—also called shocks to 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 paper is organized as follows. Section I describes how I measure inflation variance ratios and discusses the data used to construct shocks to inflation expectations. Section II documents the low level of the ratio in the data. Section III discusses volatilities of components of yield shocks in various macro-finance equilibrium models. Section IV attempts to determine the relative roles of news about expected future short rates and expected future returns. An Internet Appendix, available in the online version of the article on the *Journal of Finance* website, contains detailed discussions of various issues.

I. Inflation Variance Ratios

“Inflation news” is not a clear-cut concept—there is no unique or best way to measure how shocks associated with an inflation process affect nominal bond yields. The New Keynesian model examined by Rudebusch and Swanson (2012; hereafter R/S) helps illustrate the ambiguity. In the model, there are no exogenous shocks to inflation. Thus, in a narrow sense, there is no inflation news. However, shocks to productivity, monetary policy, and government spending each affect the paths of expected inflation, real rates, and nominal bond risk premia. Thus, in a broad sense, all news is inflation news, as every shock conveys information about expected future inflation. Models in which a monetary authority follows a Taylor rule typically have the same property. Outside the special cases, a shock to any variable that appears in the Taylor rule affects both yields and expected future inflation.

Rather than adopt a specific model’s interpretation of inflation shocks, in this paper I measure the magnitude of inflation news using an accounting approach that has its roots in the dividend/price decomposition of Campbell and Shiller (1988), as extended to returns by Campbell (1991). The measure is straightforward to estimate with available data. Any dynamic model of both inflation and bond yields—a class of which includes a wide variety of dynamic macro models—implies a value of a bond’s inflation variance ratio.

A. An Accounting Identity

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

Begin with notation. All yields are continuously compounded and expressed per period. For example, using quarterly periods, a yield of 0.02 corresponds to 8% per year.

- $y_t^{(m)}$: yield on a nominal zero-coupon bond maturing at $t + m$.
- π_t : log change in the price level from $t - 1$ to t .
- r_t : ex ante real rate, the yield on a one-period nominal bond less expected inflation, $r_t \equiv y_t^{(1)} - E_t(\pi_{t+1})$.

Note that the ex ante real rate is not the rate on a one-period real bond. In no-arbitrage complete-market models, the ex ante real rate differs from the yield on a one-period real bond owing to both a Jensen's inequality term associated with price shocks and the compensation investors require to face uncertainty in next period's price level. Investors who disagree about next period's expected inflation will also disagree about the level of the ex ante real rate. I discuss inflation expectations in more detail in Section I.C.

The log return to holding an m -period nominal bond from t to $t + 1$ in excess of the log return to a one-period nominal bond is

$$ex_{t+1}^{(m)} = \left(m y_t^{(m)} - (m - 1) y_{t+1}^{(m-1)} \right) - y_t^{(1)}. \quad (1)$$

An accounting identity decomposes the m -maturity yield into future average inflation, ex ante real rates, and excess log returns:

$$y_t^{(m)} = \frac{1}{m} \sum_{i=1}^m E_{t+i-1}(\pi_{t+i}) + \frac{1}{m} \sum_{i=1}^m r_{t+i-1} + \frac{1}{m} \sum_{i=1}^m ex_{t+i}^{(m-i+1)}. \quad (2)$$

The accounting identity formalizes observations such as, holding constant a bond's yield, higher expectations of inflation over the life of the bond must correspond to either lower ex ante real rates or lower excess returns. Future expectations of inflation appear in (2) rather than realized inflation because the short rate in the excess return definition (1) is nominal rather than real. This identity holds regardless of how inflation expectations are calculated.

The time- t expectation of (2) decomposes the m -period yield into expectations of average inflation, average ex ante real rates, and average excess returns over the life of the bond. Using iterated expectations, the bond yield is

$$y_t^{(m)} = \frac{1}{m} \sum_{i=1}^m E_t(\pi_{t+i}) + \frac{1}{m} \sum_{i=1}^m E_t(r_{t+i-1}) + \frac{1}{m} \sum_{i=1}^m E_t(ex_{t+i}^{(m-i+1)}). \quad (3)$$

Again, this equation is an identity regardless of the process for calculating expectations. The third sum on the right of (3) is often described as the bond's term premium. The accounting identity puts no structure on the term premium. In a frictionless no-arbitrage setting, the term premium is determined by the risk premium investors require to hold the bond and a Jensen's inequality component associated with the log transformation. In models with frictions, the term premium may also include a safety or convenience component.

Using this accounting framework, express the innovation in the m -maturity yield from $t-1$ to t as the sum of news about expected average inflation, ex ante real rates, and excess returns. Denote the news by

$$\begin{aligned} \tilde{y}_t^{(m)} &\equiv y_t^{(m)} - E_{t-1}y_t^{(m)}, \\ \eta_{\pi,t}^{(m)} &\equiv E_t\left(\frac{1}{m} \sum_{i=1}^m \pi_{t+i}\right) - E_{t-1}\left(\frac{1}{m} \sum_{i=1}^m \pi_{t+i}\right), \\ \eta_{r,t}^{(m)} &\equiv E_t\left(\frac{1}{m} \sum_{i=1}^m r_{t+i-1}\right) - E_{t-1}\left(\frac{1}{m} \sum_{i=1}^m r_{t+i-1}\right), \\ \eta_{ex,t}^{(m)} &\equiv E_t\left(\frac{1}{m} \sum_{i=1}^m ex_{t+i}^{(m-i+1)}\right) - E_{t-1}\left(\frac{1}{m} \sum_{i=1}^m ex_{t+i}^{(m-i+1)}\right). \end{aligned} \quad (4)$$

A yield shock is then the sum of news, or

$$\tilde{y}_t^{(m)} = \eta_{\pi,t}^{(m)} + \eta_{r,t}^{(m)} + \eta_{ex,t}^{(m)}. \quad (5)$$

This paper uses (5) to study the relative contributions of different types of news to yield innovations. The unconditional variance of yield innovations is the sum of the unconditional variances of the individual components on the right side of (5) and twice their unconditional covariances:

$$\begin{aligned} \text{Var}(\tilde{y}_t^{(m)}) &= \text{Var}(\eta_{\pi,t}^{(m)}) + \text{Var}(\eta_{r,t}^{(m)}) + \text{Var}(\eta_{ex,t}^{(m)}) \\ &\quad + 2\text{Cov}(\eta_{\pi,t}^{(m)}, \eta_{r,t}^{(m)}) + 2\text{Cov}(\eta_{\pi,t}^{(m)}, \eta_{ex,t}^{(m)}) + 2\text{Cov}(\eta_{r,t}^{(m)}, \eta_{ex,t}^{(m)}). \end{aligned} \quad (6)$$

Divide (6) by the variance on the left to express the fraction of the variance explained, in an accounting sense, by news about expected inflation, expected real rates, and expected excess returns. I use the term inflation variance ratio

to refer to one of these ratios, namely, the variance of inflation news to the variance of yield shocks. The unconditional inflation variance ratio is

$$\text{inflation variance ratio} \equiv VR_{\pi}^{(m)} = \frac{\text{Var}\left(\eta_{\pi,t}^{(m)}\right)}{\text{Var}\left(\tilde{y}_t^{(m)}\right)}. \quad (7)$$

It is important to understand what is, and what is not, measured by the inflation variance ratio. The magnitude of expected inflation news is not a summary measure of the difference in risk between nominal and real bonds because it does not capture all of the uncertainty that investors bear when they bet on future inflation. For example, a shock to the conditional variance of inflation will alter the risk of nominal bonds and thereby change bond prices. In this accounting framework, such a shock appears in news about expected excess returns, not news about expected inflation. Also note that inflation news is not necessarily orthogonal to other news. For example, in the New Keynesian model discussed above, inflation news and real rate news are correlated. More broadly, the inflation variance ratio is not a fundamental measure of inflation risk derived from a macro-finance model.

It is better to view the inflation variance ratio as an informative moment of the data rather than a fundamental measure of inflation risk. Campbell (1991) is the obvious analogy, both formally and intuitively. Campbell's conclusion that news about future cash flows accounts for less than half of the variation in aggregate stock returns strongly challenges macroeconomic models of equity prices. Similarly, the conclusions here are a strong empirical challenge to macroeconomic models of nominal yields.

B. Conditional and Unconditional Ratios

The terms “conditional” and “unconditional” can create some confusion because the shocks defined in (4) use conditioning information, while the variance ratio (7) does not. The focus here is on unconditional second moments of one-step-ahead shocks. Unconditional variance ratios are ratios of average conditional variances,

$$VR_{\pi}^{(m)} = \frac{E\left(\text{Var}_{t-1}\left(\eta_{\pi,t}^{(m)}\right)\right)}{E\left(\text{Var}_{t-1}\left(\tilde{y}_t^{(m)}\right)\right)}. \quad (8)$$

(This equation uses the fact that conditional means of shocks are identically zero.) Sample inflation variance ratios are calculated using sample variances of one-step-ahead shocks. The sample variance ratios are then compared with corresponding ratios of unconditional variances implied by workhorse macro-finance models.

Conditional second moments can be used in (6) instead of unconditional second moments. For example, we could calculate inflation variance ratios conditioned on time- t information. I do not focus on conditional variance ratios,

although they are worth a detailed study. Rigorous analysis of conditional moments requires an explicit model of conditioning information. Balduzzi and Lan (2014) take a conditional approach to interpreting the news content of shocks to the 10-year bond yield. Cram (2016), building on an earlier version of this research, models the dynamics of conditional inflation variance ratios using specific conditioning information. Here I get substantial mileage out of unconditional ratios without attempting to characterize conditional variances.

I do, however, estimate sample inflation variance ratios for interesting sub-periods. Subperiod results shed (model-free) light on the time-variation in conditional variance ratios, which in turn helps us evaluate the economic significance of the wedge between full-sample variance ratios and model-implied unconditional variance ratios. For example, if a model is incapable of matching full-sample results but is better able to match results from the 1970s and 1980s, we might conclude that the model helps us understand a high inflation regime in the United States.

Bauer and Rudebusch (2017) also extend the results of this research by modifying conditioning information. They generalize the shocks defined in (4) to h -period-ahead shocks, while retaining the focus on unconditional second moments. As h gets large, the numerator of the inflation variance ratio converges to the unconditional variance of average expected inflation, while the denominator converges to the unconditional variance of the bond's yield. Cieslak and Povala (2015) discuss the long-run relation between inflation expectations and bond yields.

C. Measuring Innovations in Inflation Expectations

Like many other researchers beginning with Pennacchi (1991), I infer inflation expectations from surveys of market practitioners. Consensus forecasts—in other words, cross-sectional means—from these surveys are close in spirit to the subjective expectations of a sophisticated investor, although no agent's beliefs may correspond exactly to consensus forecasts.

Substantial research concludes that forecasts from econometric models of inflation dynamics are not more accurate than consensus survey forecasts. Ang, Bekaert, and Wei (2007) document that survey forecasts are more accurate than model-based forecasts constructed using the history of inflation and other nonsurvey information. In addition, they find no evidence that using realized inflation in addition to survey forecasts helps reduce survey-based 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 (2013, p. 21) concur: "... purely judgmental forecasts of inflation are right at the frontier of our forecasting ability."¹

¹ An alternative view, advocated by Coibion and Gorodnichenko (2012, 2015), is that a variety of consensus forecasts are sticky owing to inattentive respondents. The Internet Appendix evaluates and rejects their argument, at least for inflation expectations.

This earlier work supports the interpretation of consensus forecasts as expectations of both market participants and researchers. (We may want to allow for measurement error, an issue that is discussed in the next section.) This research uses inflation forecasts from two types of Blue Chip (BC) surveys and the Survey of Professional Forecasters (SPF). The BC data are monthly beginning with March 1980. The SPF data are quarterly beginning in 1968Q4. The data samples used here run through 2013. The BC consensus forecasts are means across respondents. The SPF consensus forecast is the mean across respondents, dropping outliers.²

Survey forecasts are concentrated at relatively short horizons. The length of the cross-section from BC surveys varies across observations, up to a maximum of seven quarters ahead. The SPF has forecasts for only four future quarters. Section II.B uses an econometric model to extend the information in survey forecasts to longer horizon forecasts.

D. Estimating Yield Innovations

Shocks to bond yields as defined by (4) are realizations less the previous period's forecast. Survey forecasts of Treasury yields are available for a variety of maturities. Unlike inflation forecasts, survey forecasts of yields are not superior to—or even as accurate as—less subjective forecasts. Cieslak (2017) and Giacomelli, Laursen, and Singleton (2015) show that the martingale assumption produces forecasts that have lower root mean squared errors than consensus survey forecasts. Therefore, the denominator of the inflation variance ratio (7) will be larger when evaluated using survey forecasts than when using martingale forecasts.

Since an important message of this paper is that the ratio (7) is quite small, I make the conservative choice to not use consensus survey forecasts of yields. Instead, I use methods advocated in the empirical term structure literature. Research beginning with Duffee (2002) documents that martingale forecasts of Treasury bond yields typically have lower root mean squared errors in pseudo out-of-sample forecasting than do forecasts produced by parameterized models. Thus, the benchmark forecasts in this paper are martingale forecasts.

I also explore using the shape of the short end of the yield curve to predict future changes in yields. The evidence of Campbell and Shiller (1991) supports this approach. In practice, as the results in the next section document, this choice does not have much of an effect on measures of inflation variance ratios.

Yields are taken from two sources. The one-quarter yield is from the Federal Reserve Board's H15 release. Yields on zero-coupon bonds with maturities from two to six quarters, as well as 5 and 10 years, are produced by Anh Le as described in Le and Singleton (2013).³ I use both month-end yields and mid-month yields, depending on whether the yields are to be matched with

²I follow the procedure of Bansal and Shaliastovich (2013) to discard outliers from the SPF. Data limitations prevent me from dropping BC outliers.

³Thanks very much to Anh Le for sharing the data.

BC forecasts or SPF forecasts.⁴ All yields are continuously compounded and expressed at an annual rate.

II. Measuring Inflation Variance Ratios

In this section I estimate the inflation variance ratio measure (7) at both short and long horizons. Survey data allow the model-free construction of news about average expected inflation over short horizons. Longer horizon forecasts require a dynamic model of inflation expectations. I use a simple model drawn from the literature on inflation expectations to estimate inflation variance ratios at multiyear horizons.

A. Short-Horizon Forecasts

A survey at quarter t reports k -quarter-ahead consensus predictions of inflation for $k = 1, \dots, K_{max}$. Discard the one-quarter-ahead prediction, and convert the others to predictions of average log inflation from quarter $t + 1$ to quarter $t + k$, $k = 2, \dots, K_{max}$. Use the survey at quarter $t + 1$ to calculate predictions of average log inflation over the same $K_{max} - 1$ horizons. The latter predictions minus the former predictions are the consensus innovations at quarter $t + 1$ of average expected inflation from $t + 1$ to $t + k$, $k = 2, \dots, K_{max}$. No model is necessary to construct these measures of news about expected future inflation.⁵

Figure 1 displays time series of quarterly news about average expected inflation produced with consensus forecasts from BC surveys. The horizons range from one to six quarters.⁶ A glance at the figure reveals that news is highly correlated across horizons. Part of the correlation is mechanical, since expected inflation over the next k quarters is embedded in average expected inflation over the next $k + 1$ quarters. Also note that the magnitude of news declines with the forecast horizon. For example, news about expected inflation over the next year is less volatile than news about expected inflation over the next quarter.

Figure 2 displays similar news produced with consensus forecasts from SPF. The time series is longer and the cross-section is shorter. The horizons range from one to three quarters. The long time series reveals substantial heteroskedasticity of news, with volatility peaking during the late 1970s.

I use two methods to construct corresponding yield innovations. One imposes the martingale assumption, so that current yields equal expected future yields. The other uses in-sample forecasts produced by a regression that

⁴ For comparison with the quarterly data of the SPF, yields are observed on 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.

⁵ Additional details are in the Internet Appendix. This appendix also discusses the imperfect mapping from survey forecasts of average quarterly inflation to the accounting identity's forecasts of point-to-point inflation.

⁶ The displayed realizations contain overlapping information because of the monthly survey frequency and quarterly news interval.

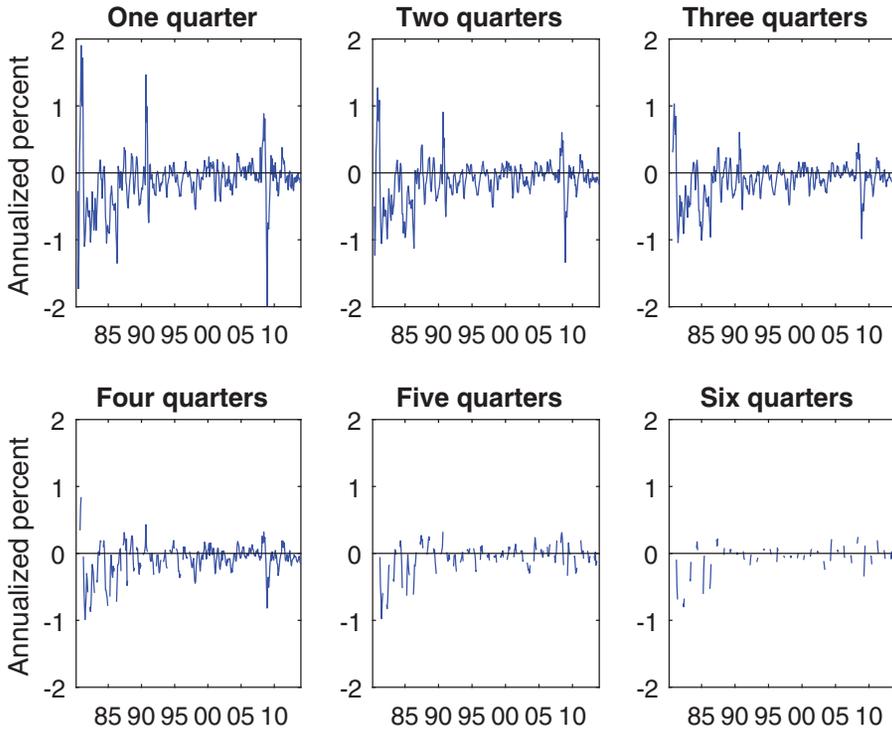


Figure 1. Realized quarterly news about average expected inflation from Blue Chip surveys. Quarterly changes in consensus forecasts from Blue Chip surveys are used to construct innovations in forecasts of expected average inflation over horizons ranging from one to six quarters. The sample range is May 1980 through December 2013. (Color figure can be viewed at wileyonlinelibrary.com)

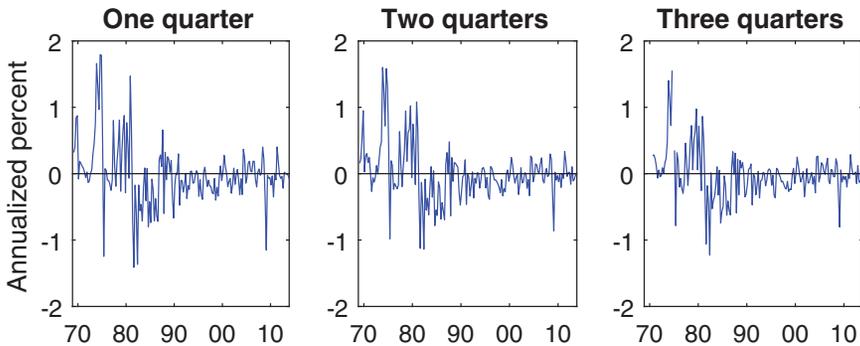


Figure 2. Realized quarterly news about average expected inflation from the Survey of Professional Forecasters. Quarterly changes in consensus forecasts from SPF surveys are used to construct innovations in forecasts of expected average inflation over horizons ranging from one to three quarters. The sample range is 1968Q4 through 2013Q4. (Color figure can be viewed at wileyonlinelibrary.com)

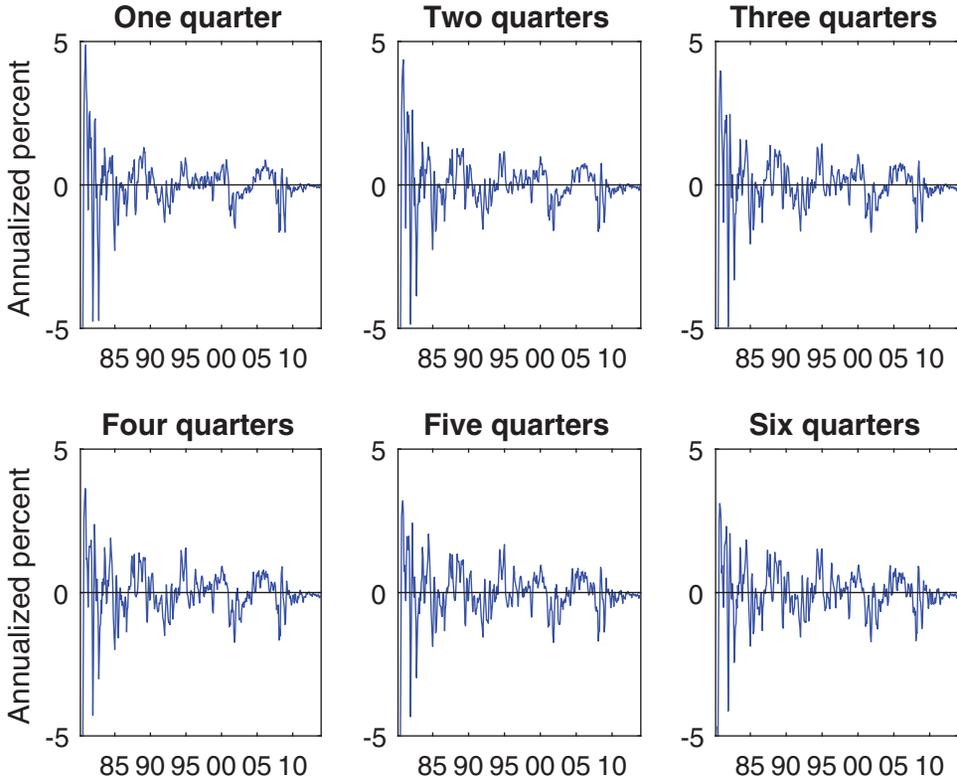


Figure 3. Realized quarterly innovations of Treasury yields, 1980 to 2013. Ordinary least squares (OLS) forecasting regressions are used to construct quarterly innovations of yields. The forecasting variables are yields for maturities of one, four, and six quarters. The sample range is May 1980 through December 2013. (Color figure can be viewed at wileyonlinelibrary.com)

predicts future changes in short-maturity yields using the shape of the short end of the yield curve. The regression equation is

$$y_{t+1}^{(m)} - y_t^{(m)} = b_{0,m} + b'_{1,m} \begin{pmatrix} y_t^{(1)} & y_t^{(4)} & y_t^{(6)} \end{pmatrix}' + \hat{y}_{t+1}^{(m)}. \tag{9}$$

In (9), both time and maturities are measured in quarters. The parameters $b_{0,m}$ and $b_{1,m}$ are a scalar and a length-three vector, respectively. In words, quarterly changes in bond yields are predicted using yields on one-quarter, four-quarter, and six-quarter bonds. The residuals are the innovations.

Figure 3 displays quarterly innovations in yields from the regressions. The sample period and the maturity range match those for inflation expectation news displayed in Figure 1. Figure 4 also displays quarterly innovations, with a sample period and maturity range matching those in Figure 2. (The regressions are estimated separately for the two sample periods.) As with news about average inflation, yield innovations are highly correlated across horizons. One

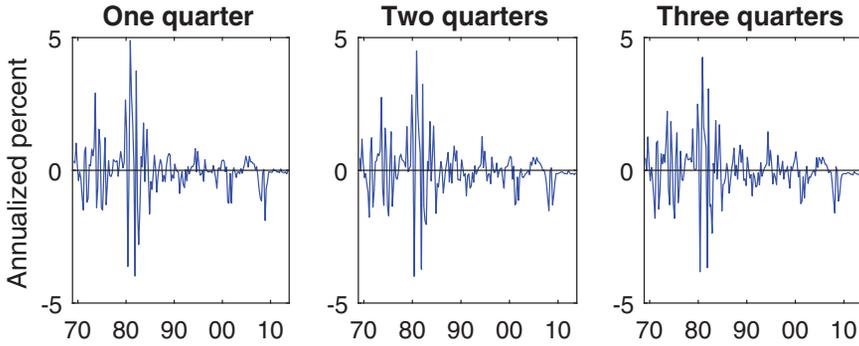


Figure 4. Realized quarterly innovations of Treasury yields, 1968 to 2013. Ordinary least squares (OLS) forecasting regressions are used to construct quarterly innovations of yields. The forecasting variables are yields for maturities of one, four, and six quarters. The sample range is 1968Q4 through 2013Q4. (Color figure can be viewed at wileyonlinelibrary.com)

important difference is that volatilities of yield innovations decline only slightly with maturity. This difference implies that inflation variance ratios decline with maturity.

A glance at the scales of the vertical axes of these four figures reveals the main result of this paper. The typical magnitude of yield innovations is more than twice as large as the typical magnitude of corresponding news about average expected inflation. An immediate implication is that the ratio of the two variances is less than a quarter. We need to make some numerical calculations, extend the analysis to longer maturities, and evaluate statistical significance. But nothing that follows is surprising, given the evidence in Figures 1 through 4.

A.1. Unconditional Inflation Variance Ratios

Point estimates of unconditional inflation variance ratios are displayed in Table I. Robust standard errors are estimated using generalized method of moments (GMM). I describe the GMM estimation in the Internet Appendix.

The table's main result is that unconditional inflation variance ratios are small, in the neighborhood of 0.1 to 0.2. Estimates in Panel A are for the 1980Q1 through 2013Q4 period, with inflation news from BC surveys.⁷ Because of the sparseness of data for five- and six-quarter horizons (see the gaps in Figure 1), results are displayed only for horizons up to four quarters. In Panel A, the largest variance ratio—for the one-quarter yield—is only 0.2. Estimates of variance ratios for three- and four-quarter yields are less than 0.1. Estimates in Panel B are for the 1968Q4 through 2013Q4 period, with inflation news from SPF forecasts of the GDP deflator. The estimates are all less than a quarter and decline with maturity.

⁷ Although the BC data are monthly, estimation uses only surveys from the end of the second month of each quarter (the surveys published at the beginning of March, June, September, and December) to avoid overlapping observations.

Table I
Short-Horizon Estimates of Inflation Variance Ratios

Quarterly shocks to average expected inflation over various maturities are from surveys of forecasters. Corresponding shocks to bond yields are constructed using forecasts from either in-sample estimates of a forecasting model or the assumption that yields follow martingales. Standard deviations are in annualized percent per quarter. Inflation variance ratios are the ratio of the variance of shocks to average expected inflation to the variance of yield shocks. GMM standard errors are in parentheses. A robust correction to the covariance matrix of moments is used with 12 lags.

Statistic	Maturity (quarters)			
	1	2	3	4
Panel A. 1980Q1 through 2013Q4, Blue Chip Forecasts of CPI				
<i>SD</i> of news about expected inflation	0.45	0.35	0.31	0.26
<i>SD</i> of yield innovations				
Regression forecasts	1.01	1.01	1.01	0.97
Martingale forecasts	1.05	1.04	1.02	0.98
Variance ratio				
Regression forecasts	0.202 (0.079)	0.121 (0.044)	0.096 (0.033)	0.074 (0.022)
Martingale forecasts	0.184 (0.071)	0.116 (0.042)	0.096 (0.031)	0.073 (0.021)
Panel B. 1968Q4 through 2013Q4, SPF Forecasts of GDP Deflator				
<i>SD</i> of news about expected inflation	0.46	0.40		0.36
<i>SD</i> of yield innovations				
Regression forecasts	0.95	0.94		0.92
Martingale forecasts	0.99	0.95		0.93
Variance ratio				
Regression forecasts	0.234 (0.090)	0.185 (0.064)		0.150 (0.045)
Martingale forecasts	0.218 (0.086)	0.180 (0.064)		0.148 (0.046)

Standard errors for these variance ratios are sufficiently small that we can reliably conclude that population variance ratios are less than a half for the one-quarter yield and less than a quarter for the longer maturities. Standard errors for the SPF estimates are slightly larger, but we can reliably conclude that population variance ratios are less than 0.5 for the one-quarter yield and less than 0.25 for the one-year yield.

The estimates in Table I assume that both inflation news and yield innovations are measured without error. This assumption is too strong. Consensus forecasts are cross-sectional sample means, and hence the period- t forecast depends on the particular makeup of the panel at t . Bond yields are also measured with error. All but the one-quarter yield are interpolated from yields on coupon bonds. Bekaert, Hodrick, and Marshall (1997) estimate that, for maturities around one year, standard deviations of interpolation measurement error are in the range of eight to nine basis points.

Table II
Short-Horizon Estimates of Inflation Variance Ratios: Subsamples

Quarterly shocks to average expected inflation over various maturities are from surveys of forecasters. Corresponding shocks to bond yields are constructed using the assumption that yields follow martingales. Standard deviations are in annualized percent per quarter. Inflation variance ratios are the ratio of the variance of shocks to average expected inflation to the variance of yield shocks.

Period	Survey	Statistic	Maturity (Quarters)			
			1	2	3	4
1968Q4 to 1979Q2	SPF	<i>SD</i> , inflation news	0.67	0.59	0.51	
		<i>SD</i> , yield innovations	0.94	0.96	0.95	
		Variance ratio	0.50	0.38	0.28	
1979Q3 to 1982Q4	SPF	<i>SD</i> , inflation news	0.81	0.66	0.61	
		<i>SD</i> , yield innovations	2.88	2.66	2.51	
		Variance ratio	0.08	0.06	0.06	
1980Q1 to 1982Q4	Blue Chip	<i>SD</i> , inflation news	0.98	0.77	0.73	0.70
		<i>SD</i> , yield innovations	3.21	3.11	3.03	2.75
		Variance ratio	0.09	0.06	0.06	0.06
1983Q1 to 2008Q2	Blue Chip	<i>SD</i> , inflation news	0.33	0.27	0.24	0.19
		<i>SD</i> , yield innovations	0.58	0.60	0.62	0.65
		Variance ratio	0.33	0.20	0.15	0.09
	SPF	<i>SD</i> , inflation news	0.29	0.26	0.24	
		<i>SD</i> , yield innovations	0.53	0.56	0.57	
		Variance ratio	0.29	0.22	0.17	
2008Q3 to 2013Q4	Blue Chip	<i>SD</i> , inflation news	0.50	0.35	0.27	0.14
		<i>SD</i> , yield innovations	0.37	0.31	0.30	0.30
		Variance ratio	1.77	1.23	0.85	0.63
	SPF	<i>SD</i> , inflation news	0.30	0.23	0.22	
		<i>SD</i> , yield innovations	0.37	0.27	0.25	
		Variance ratio	0.65	0.74	0.73	

I discuss measurement error in detail in the Internet Appendix. Here it is sufficient to note that modest amounts of measurement error in both inflation news and yields will likely artificially increase observed inflation variance ratios as measured in Table I, but the effect will be negligible.

A.2. Inflation Variance Ratios across Monetary Regimes

To reiterate, the evidence in Table I refers to unconditional inflation variance ratios. A glance at Figures 1 through 4 reveals considerable heteroskedasticity of both inflation news and yield shocks. Table II documents that conditional inflation variance ratios also vary over time.

Results are presented for four subperiods. A common view is that monetary policy was accommodative during the pre-Volcker years, ending in 1979Q2. Accommodation increases the responsiveness of expected inflation to macroeconomic shocks. In addition, as argued by Clarida, Galí, and Gertler (2000), it can produce sunspots in inflation expectations, creating another channel for

news about expected inflation. The second period is disinflation, beginning in 1979Q3. Following Lubik and Schorfheide (2004) and others, I treat 1982Q2 as the end of the disinflation period. The aggressive monetary policy period begins with 1983Q1. The financial crisis/zero lower bound (ZLB) subperiod begins with the Lehman failure in 2008Q3 and continues through the end of the sample, 2013Q4.

Before turning to detailed results, note that in Table I there is little difference between variance ratios calculated with regression-based yield forecasts and martingale yield forecasts. In other words, quarterly changes in yields are largely unpredictable with the shape of the term structure. With subsamples, there is greater danger of overfitting. Therefore, in Table II all the results use martingale yield forecasts.

Three results in Table II are worth highlighting. First, outside of the crisis/ZLB period, none of the reported inflation variance ratios exceed a half. Second, they are higher during the period of accommodative monetary policy than during the period of aggressive monetary policy. Variance ratios during the former period are roughly 1.6 times corresponding variance ratios during the latter period. Third, extreme realizations of inflation variance ratios are observed during the extreme periods of the sample. Variance ratios are less than 0.1 during the disinflation period and well above a half during the crisis/ZLB period. I discuss these results in Section II.B, after presenting estimation of inflation variance ratios for long-maturity bonds.

Table II also documents that holding the sample period fixed, variance ratios based on the SPF differ from those based on the BC surveys. The latter are typically larger than the former. I briefly discuss the differences, which are not economically substantial, in the Internet Appendix.

B. Longer Horizon Estimates

Survey data on long-horizon expectations of inflation are sparse. BC and SPF participants are occasionally asked to predict inflation at the 5- to 10-year horizon, but neither the frequency of the responses nor the precision of the inflation horizon allows for model-free calculation of innovations to inflation expectations. Empirical implementation of (3) for bond maturities at horizons greater than a year requires a model.

B.1. A Trend-Cycle Model

A conclusion of Faust and Wright (2013) motivates the modeling approach. The authors describe a method that produces intermediate-range inflation forecasts “close to the frontier of predictive performance.” Simply use a glide path to connect survey forecasts of current inflation to survey forecasts of distant inflation. A corollary to their conclusion, verified below, is that longer horizon forecasts can be extrapolated from the glide path on which short-horizon forecasts lie.

A trend-cycle model captures the intuition of the glide path approach. The model assumes a unit root in inflation, consistent with models of inflation such as Stock and Watson (2007) and Cogley, Primiceri, and Sargent (2010). Period- t inflation is the sum of three components. One follows a martingale, another follows a persistent stationary process, and the third is a serially uncorrelated shock. The equations are

$$\pi_t = \tau_t + \varphi_t + \phi_t, \quad (10)$$

$$\tau_t = \tau_{t-1} + \xi_t, \quad E_{t-1}(\xi_t) = 0, \quad (11)$$

$$\varphi_t = \theta\varphi_{t-1} + \nu_t, \quad E_{t-1}(\nu_t) = 0, \quad (12)$$

$$E_{t-1}(\phi_t) = 0. \quad (13)$$

Following the discussion of Section I.D, long-term bond yields are assumed to follow martingales, as in

$$y_t^{(m)} = y_{t-1}^{(m)} + \tilde{y}_t^{(m)}, \quad E_{t-1}(y_t^{(m)}) = 0, \quad (14)$$

for some long-maturity m . Similar models appear in Nason and Smith (2014), Stock and Watson (2007), and Cogley, Primiceri, and Sargent (2010). More information about the model, including formulas for calculating inflation variance ratios, is provided in the Internet Appendix.

In this model, both inflation expectations and long-term yields have unit roots. This raises the question of cointegration. The earliest comprehensive empirical analysis of cointegration among nominal yields is in Campbell and Shiller (1987). If both yields and inflation have unit roots but they are not cointegrated, then either real rates or term premia must also have a unit root. Campbell and Ammer (1993) assume that inflation and yields are cointegrated with a unit cointegrating vector. This assumption, imposed on the trend-cycle model, requires that the martingale component of inflation must move one-for-one with the long-maturity yield (which is also a martingale). Hence, inflation variance ratios must converge to 1 for long maturities. Further details can be found in the Internet Appendix.

I do not impose cointegration, but rather allow the data to determine the magnitude of inflation variance ratios. This choice is consistent with results in the applied cointegration literature, which typically produces a double negative: we cannot reject the hypothesis that inflation and nominal yields are both nonstationary and not cointegrated. Examples include Lardic and Mignon (2004) and Hjalmarsson and Österholm (2010). In their critical review of the literature, Neely and Rapach (2008, p. 609–610) conclude that “. . . studies [of real rates] often report evidence of unit roots, or—at a minimum—substantial persistence.”⁸

⁸ Neely and Rapach (2008) 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.

Some readers may be uncomfortable with the assumption of unit roots in inflation and nominal yields. In Section IV, I present and estimate a stationary model of inflation and nominal yields. This choice has a minimal effect on measures of inflation variance ratios.

B.2. Estimation

Estimation uses current and expected future inflation from surveys, as well as observations of a five-year bond yield. I then repeat the exercise using a 10-year bond yield.⁹ Since realized quarter- t inflation is announced well after the end of quarter t (and after the bond yield is determined), I use survey consensus forecasts of current-quarter inflation (nowcasts) rather than actual inflation. Survey consensus forecasts of future inflation are used up to the maximum available horizon. This maximum is seven quarters for the BC survey and four quarters for the SPF.

With more observables than state variables, a stochastic singularity problem arises if variables are assumed to be measured without error. I therefore assume that all observables other than the inflation nowcast are contaminated by measurement error. There is nothing special about the nowcast's accuracy, but it is impossible to untangle measurement error from the purely transitory shock to inflation. The measurement error for a particular observable is assumed to be i.i.d.

I estimate model parameters using exactly identified GMM. Robust standard errors adjust for conditional heteroskedasticity of shocks to yields and expected inflation. The Internet Appendix contains more details about measurement error, estimation, and standard errors.

B.3. Results

Table III reports estimated inflation variance ratios. The Internet Appendix contains all other estimates associated with the model.

These long-horizon results are consistent with the short-horizon results in Tables I and II. For both 5- and 10-year bonds, full-sample point estimates of unconditional inflation variance ratios are all less than 0.15. The standard errors allow us to reject at standard confidence levels the hypothesis that a variance ratio exceeds 0.25. This conclusion is reinforced graphically in Figure 5, which displays, for the BC data, time series of the 10-year bond yield and filtered estimates of expected inflation over the next 10 years. The former series is much more volatile than the latter. Only a small part of the variation in bond yields is attributable to variation in expected future inflation over the life of the bond.

It is important to verify that the long-run inflation forecasts from this dynamic model are accurate, in the sense that they capture investor expectations

⁹This is inefficient. I do not impose the restriction that inflation dynamics should be the same for both bond yields.

Table III
Estimates of Inflation Variance Ratios for Long-Term Bonds

Quarterly shocks to average expected inflation over 5- and 10-year horizons are estimated from a model that assumes inflation is the sum of a martingale and an AR(1) component. The model is fit to survey consensus forecasts of inflation over horizons from zero to up to seven quarters ahead. Shocks to the yields are calculated assuming they are martingales. All consensus forecasts and yields are assumed to be contaminated with i.i.d. measurement error. The inflation variance ratio is the ratio of the variance of shocks to average expected inflation over the life of the bond to the variance of yield shocks. Standard deviations are in annualized percent per quarter. Standard errors in parentheses use 12 lags to adjust for serial correlation.

Period	Survey	Bond Maturity	<i>SD</i> Inflation News	<i>SD</i> Yield Shocks	Variance Ratio
1980Q1 to 2013Q4	Blue Chip	5 years	0.20	0.73	0.079 (0.033)
		10 years	0.21	0.64	0.105 (0.039)
1968Q4 to 2013Q4	SPF	5 years	0.23	0.61	0.136 (0.041)
		10 years	0.21	0.56	0.141 (0.042)
1968Q4 to 1979Q2	SPF	5 years	0.27	0.50	0.290
		10 years	0.25	0.42	0.345
1979Q3 to 1982Q4	SPF	5 years	0.33	1.28	0.065
		10 years	0.25	1.24	0.040
1980Q1 to 1982Q4	Blue Chip	5 years	0.46	1.46	0.101
		10 years	0.47	1.31	0.131
1983Q1 to 2008Q2	Blue Chip	5 years	0.18	0.62	0.080
		10 years	0.19	0.55	0.118
	SPF	5 years	0.16	0.55	0.089
		10 years	0.16	0.49	0.113
2008Q3 to 2013Q4	Blue Chip	5 years	0.11	0.43	0.061
		10 years	0.10	0.38	0.070
	SPF	5 years	0.08	0.42	0.034
		10 years	0.07	0.44	0.027

of long-run inflation. Visual evidence is in Figure 6. The circles are semiannual BC survey consensus forecasts of CPI inflation over the period beginning 5 years and ending 10 years from the survey date. These data are unavailable prior to 1984. The x's are BC survey forecasts of GNP inflation over the same horizon. The solid line in Figure 6 displays filtered estimates from the trend-cycle model of expected inflation over the same future horizon.¹⁰ The dashed line is explained in Section II.C.

A glance at the figure reveals that the BC and model-implied forecasts closely correspond. In the early part of the sample, the GNP inflation survey forecasts are about 50 basis points lower than the model-implied CPI forecasts. This is consistent with the mean difference between CPI inflation and GNP (and GDP)

¹⁰ The model is expressed in terms of continuously compounded inflation. To simplify comparison with survey forecasts, expected continuously compounded inflation is converted to simple (annually compounded) terms. This transformation ignores Jensen's inequality term.

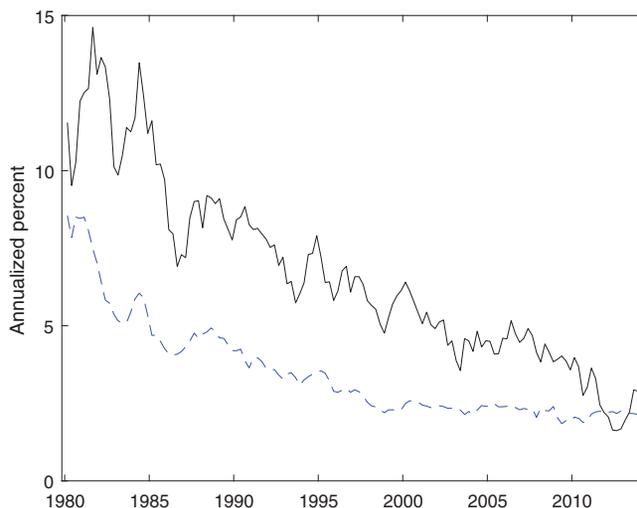


Figure 5. Ten-year bond yields and model-implied inflation forecasts. The solid line is the yield on a 10-year Treasury zero-coupon bond. The dashed line is an estimate of expected inflation over the next 10 years. The model of expected inflation assumes that inflation is the sum of a random walk, an AR(1) process, and white noise. The parameters are estimated using Blue Chip survey forecasts of CPI inflation up to seven quarters ahead. (Color figure can be viewed at wileyonlinelibrary.com)

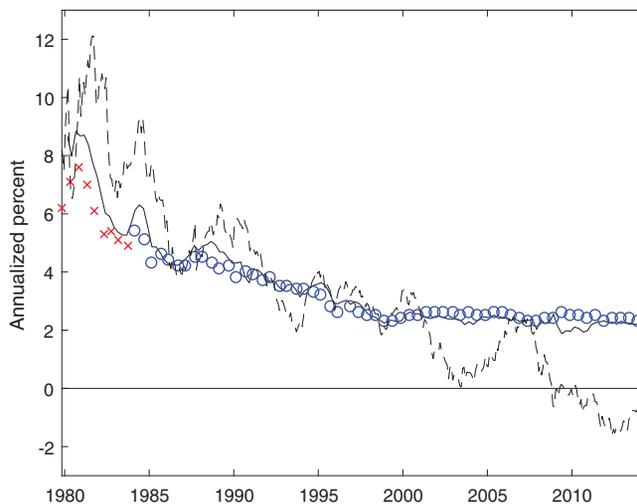


Figure 6. Long-horizon inflation forecasts from surveys and models. The solid and dashed lines are fitted forecasts, as of date t on the horizontal axis, of average CPI inflation from $t + 5$ years to $t + 10$ years. The model underlying the solid line assumes that inflation is the sum of a random walk, an AR(1) process, and white noise. The model underlying the dashed line assumes that inflation and yields are cointegrated with a cointegrating vector of 1's. The circles are Blue Chip survey forecasts of CPI inflation over the same horizon. The x's are Blue Chip survey forecasts of GDP inflation over the same horizon. Neither the circles nor the x's are used in model estimation. (Color figure can be viewed at wileyonlinelibrary.com)

inflation. After 1985, the largest differences between the forecasts occur during the financial crisis. During the crisis the model's forecasts are modestly more volatile than are the survey forecasts.

Table III also displays subsample results. The samples are the same as those for which short-horizon inflation variance ratios are estimated in Table II. The joint results can be summarized as follows. Inflation variance ratios are highest during the passive monetary policy period of the 1970s. Yet even in this period, news about expected inflation accounts for less than half of the variation in yields. Inflation variance ratios are around one-third for maturities greater than three months. During both the Volcker disinflation and the Great Moderation, inflation variance ratios are in the range of 0.05 to 0.15. (Again, inflation variance ratios are somewhat higher at the three-month horizon.) During the financial crisis/ZLB period, inflation variance ratios are very high—well above a half—at short maturities. But at long maturities, they are similar to those observed during both disinflation and the Great Moderation.

C. Revisiting Campbell and Ammer (1993)

The variance measure in Section I is borrowed from Campbell and Ammer's (1993) decomposition of excess bond returns. However, they conclude that shocks to average expected inflation account for the vast majority of shocks to nominal yields. Here I show that their conclusion relies on the assumption of cointegration and produces strongly counterfactual estimates of expected long-run inflation.

The authors assume 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.¹¹ Inflation enters only in the form of a linear combination with the short nominal rate, and the level of yields is not included.

This setup implies that yields and inflation are nonstationary and cointegrated with a cointegrating vector of ones. In addition, 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. By contrast, real rates are assumed to be stationary. Thus, their assumptions differ sharply with the trend-cycle model of inflation in Section II.B.

Campbell and Ammer (1993) produce estimates of inflation variance ratios for long-term bonds that are substantially higher than those reported in Section II.B. They conclude that news about expected inflation 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

¹¹ The slope is measured by the 10-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.

Table IV
Variance Decompositions Using a Model of Cointegrated
Inflation and Yields

The table reports model-implied variance decompositions of quarterly shocks to a 10-year nominal yield. The model is the first-order vector autoregression of monthly data. The observed data are described in Section II.C. The model imposes cointegration with a unit vector on yields and inflation, while real rates are stationary. The three components of yield shocks are news about expected future real rates, expected future inflation, and expected future excess returns. The contributions to total variance sum to 1, aside from rounding.

Statistic	10/1968 to 12/2013	10/1968 to 2/1987	3/1987 to 12/2013
Var(real rate news)	0.02	0.02	0.02
Var(inflation news)	1.27	1.02	0.89
Var(return news)	0.60	0.90	0.20
2Cov(real rate news, inflation news)	-0.06	-0.15	-0.06
2Cov(real rate news, return news)	-0.02	-0.21	-0.02
2Cov(inflation news, return news)	-0.80	-0.98	-0.03

their monthly frequency. The sample is split after February 1987, which is the last month in the sample examined by Campbell and Ammer (1993). Their methodology makes it easy to infer all of the components of (6), and thus each is reported in Table IV. I do not bother calculating standard errors.

Table IV confirms the results of Campbell and Ammer (1993), in the sense that estimated inflation variance ratios are about one for each sample. The full-sample point estimate is about 1.3. The point estimates for the two subsamples are slightly less than 1. Put differently, their VAR implies that the variance of news about expected average inflation over 10 years is close to the variance of shocks to the 10-year yield.

When news about inflation is so large, expectations of long-run inflation must vary substantially over time. Figure 6 displays with a dashed line the full-sample monthly forecasts of average inflation from year 5 to year 10. The forecasts fluctuate substantially over time, from more than 12% in the early 1980s to less than -1% in 2011. It is clear from the figure that the VAR estimates of long-run inflation are wildly at odds with both survey forecasts and forecasts produced by the models that allow some or all of the variability in inflation to be mean-reverting.

A straightforward implication of Figure 6 is that we should not impose the assumption of cointegration on the joint dynamics of inflation and yields. When we allow for more flexibility in their joint dynamics, estimated unconditional inflation variance ratios at long horizons are similar to model-free estimates for short horizons. At both short and long horizons, the estimates are no more than 0.2.

The results of this section naturally lead to two follow-up questions that I address in the paper's next sections. First, how tightly do these results bind on our macroeconomic models? Second, since inflation news does not drive bond yields, what does?

III. Volatilities in Standard Dynamic Models

This section explains why standard long-run risk models have substantial difficulty matching the volatilities of yield shocks and the volatilities of news about expected inflation documented in Section II. It also explains why habit formation models are less restrictive along the dimensions important for matching the documented behavior.

A. A Recursive Utility Starting Point

Bansal and Yaron (2004) combine recursive utility with long-run consumption growth risk in an endowment economy. At first glance, it might seem that the evidence of Section II has no implications for models in the style of Bansal and Yaron (2004). Their approach specifies the dynamics of the real economy without reference to an inflation process. Thus, in principle, we can overlay any inflation process on a standard long-run risk model—just pick one that matches the inflation variance ratios of Section II.

However, the important modeling challenge posed by Section II is matching both the numerators and denominators of inflation variance ratios. Put differently, can a dynamic model produce both inflation variance ratios well below one *and* volatilities of yield shocks similar to those documented in Section II? Such a model must generate substantial news about either expected future short-term real rates or expected future excess bond returns. The claim here is that long-run risk models have difficulty generating much news of either type.

The first extensive exploration of term structure dynamics in a long-run risk framework is Piazzesi and Schneider (2007; hereafter P/S). They add an exogenous inflation process to a model with recursive utility and long-run endowment risk. The elasticity of intertemporal substitution (EIS) is fixed at 1. Their benchmark model is set in a log-normal framework with homoskedastic shocks. Therefore, there is no news about expected excess returns. Nominal yields react only to news about expected future short-term real rates and expected future inflation. Standard deviations of these news components are readily calculated analytically.¹²

The top panel of Figure 7 displays these model-implied standard deviations. At all horizons, the magnitude of news about expected average inflation exceeds the magnitude of news about expected average real rates. Standard deviations of quarterly inflation news range from more than 60 basis points at the one-quarter horizon to 30 basis points at the 10-year horizon. Corresponding standard deviations for real rate news range from about 50 basis points at the short end to less than 7 basis points at the long end.

¹² P/S also study a model with learning, in which shocks (according to the investors' filtration) are heteroskedastic. In the learning model, the economy converges to full information in the long run. Since I study population properties of models here, I skip examination of their learning model. The calculations were made easier because the authors graciously provide their Matlab code on Monika Piazzesi's website.

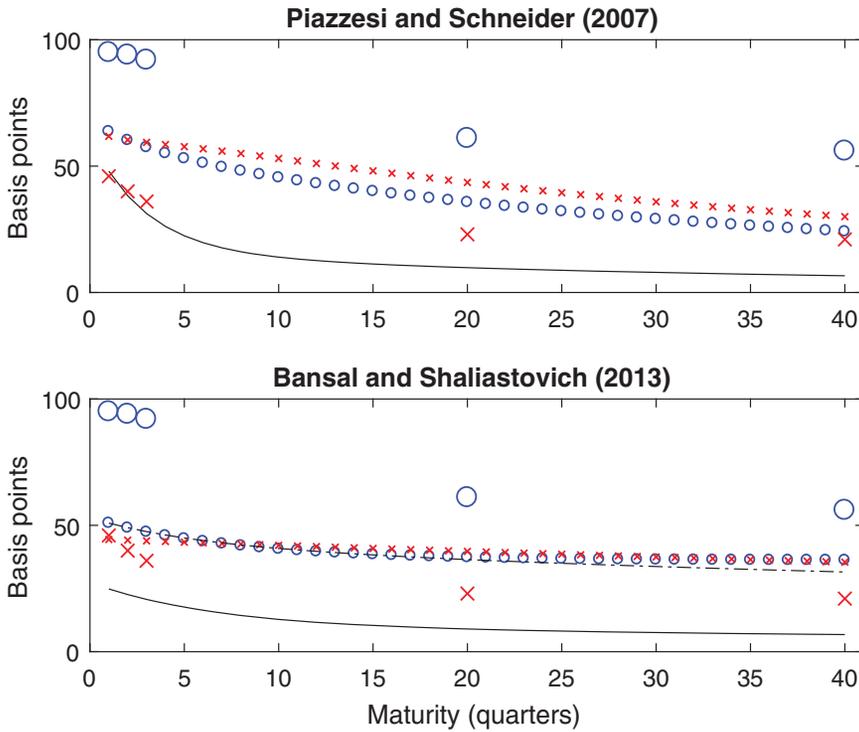


Figure 7. Standard deviations of the components of shocks to nominal yields for two long-run risk models. The figure plots sample and model-implied unconditional standard deviations of quarterly shocks. Sample standard deviations of yields (large O's) and news about expected inflation (large X's) are for the period 1968Q4 through 2013Q4 and come from Tables I and III. Unconditional model-implied standard deviations of yield shocks (small o's), news about expected inflation (small x's), and news about expected real rates (solid line) are determined using the “Benchmark” estimated model in Piazzesi and Schneider (2007) and the estimated model in Bansal and Shaliastovich (2013). The lower panel also displays model-implied standard deviations of the sum of news about expected inflation and expected real rates (dotted-dashed). (Color figure can be viewed at wileyonlinelibrary.com)

The parameterized model exhibits a standard property of long-run risk models: low volatility of real rate shocks. Two key features of the long-run risk approach are (a) shocks to expected consumption growth are small, and (b) the EIS is high. In combination, these features imply that short-term real rates do not vary much over time. From quarter to quarter, news about expected future real rates is small. Thus, shocks to nominal bonds must be driven primarily by shocks to expected inflation.

In fact, Figure 7 shows that model-implied standard deviations of yield shocks—given by the sum of the two types of news—are almost all *less* than the corresponding standard deviations of news about expected inflation. Real rate news and inflation news are sufficiently negatively correlated that model-implied inflation variance ratios are almost all greater than 1.

Negative news correlations are a consequence of macroeconomic dynamics that exhibit stagflation: negatively correlated shocks to expected consumption growth and expected inflation. Stagflation dynamics are consistent with the positive average risk premia earned by nominal bonds because nominal bonds unexpectedly decline in value when the future looks gloomy. Since real rates move in lockstep with expected consumption growth, news about real rates and expected inflation are also negatively correlated. Model-implied news correlations range from about -0.3 at the one-quarter horizon to -0.9 at the 10-year horizon.

The same panel of Figure 7 also displays sample standard deviations of quarterly shocks to yields and expected inflation for the 1968Q4 through 2013Q4 period, which come from Tables I and III. Model-implied volatilities of inflation news are much larger than the sample values, while model-implied volatilities of yield shocks are much smaller. A successful model requires smaller shocks to inflation expectations and much larger shocks to either real rates or expected excess returns.

B. Adding Stochastic Volatility

Bansal and Shaliastovich (2013; hereafter B/S) extend the approach of P/S by including time-varying conditional variances. This generalization opens the channel for shocks to expected excess returns, since shocks to conditional variances produce shocks to current and expected future risk premia.

The top line of Table V summarizes a few important properties of the model estimated by B/S. The unconditional standard deviation of log-differenced quarterly consumption is a little more than 1% (annualized). The correlation between consumption growth at quarters t and $t + 4$ is 0.16. The standard deviation matches, by construction, the sample standard deviation in postwar U.S. data. The fourth-order serial correlation is a little high relative to the postwar U.S. value of 0.1. The model-implied mean real yield curve slopes down because real bonds are a hedge. The nominal yield curve slopes up on average because stagflation risk outweighs the hedging properties of real bonds.¹³

The bottom panel of Figure 7 displays model-implied unconditional standard deviations of quarterly shocks, calculated using the point estimates of B/S. This panel plots one more function than does the P/S panel. In P/S, yield shocks are identical to the sum of news about expected real rates and expected inflation. In B/S, these differ because yield shocks include shocks to risk premia. The panel plots both standard deviation of yield shocks and standard deviation of the sum of news about expected real rates and expected inflation. The magnitude of news about expected excess returns determines the wedge between these standard deviations.

The figure shows that the B/S model, like the P/S model, produces inflation variance ratios close to 1. The B/S model shares with the P/S model the

¹³ Numbers in the table differ slightly from those reported in B/S, presumably because of rounding error in reported parameters.

Table V

Properties of Versions of the Bansal and Shaliastovich (2013) Model

The table reports properties of the dynamic representative-agent endowment economy modeled in Bansal and Shaliastovich (2013). They are calculated for the parameters reported in B/S and for alternative parameterizations. The standard deviation of log-differenced quarterly consumption is multiplied by 2 to put it in annual terms. The fourth autocorrelation of the series is denoted ρ_4 . Mean bond yields are expressed in annualized percent. Using the notation of B/S, Version 2 changes σ_{xc} from 1.09×10^{-3} to 4.0×10^{-3} . Version 3 uses the high value of σ_{xc} and changes ρ_c from 0.81 to 0.96. Version 4 uses the high value of σ_{xc} and changes $\sigma_{\omega c}$ from 1.85×10^{-7} to 3.7×10^{-6} . Version 5 changes the elasticity of intertemporal substitution from 1.81 to 0.55 and the coefficient of relative risk aversion from 20.9 to 170. All the alternatives also reduce the persistence of inflation, ρ_π , from 0.988 to 0.95.

Version	Log-Differenced Quarterly Consumption		Mean Yields (%)			
	SD (%)	ρ_4	Real		Nominal	
			1 year	10 years	1 year	10 years
Original	1.05	0.16	2.5	1.9	6.3	8.5
Version 2: high volatility of shocks to expected growth	1.65	0.30	2.0	0.7	5.7	4.8
Version 3: high volatility of expected growth shocks, high persistence	3.05	0.78	-6.3	-20.4	-2.5	-15.5
Version 4: high volatility of expected growth shocks, high stochastic volatility	1.65	0.30	-1.6	-10.6	2.1	-6.6
Version 5: high preference for smoothing	1.0	0.08	-26,487	-26,565	-26,483	-26,534

problematic features of news about expected real rates: there is not much news, and the news that exists is negatively correlated with news about expected future inflation. Perhaps surprisingly, there is also very little news about expected future excess returns. Standard deviations of yield shocks are almost identical to standard deviations of sums of news about expected real rates and expected inflation. At the 10-year maturity, these standard deviations differ by less than five basis points.

Why is the time-varying risk premium channel so small? The short answer is that quarterly shocks to conditional variances are small, and the effects of these shocks on bond risk premia are proportional to average bond risk premia—which are also small. To take an extreme example, imagine that all conditional variances are at their means, and they suddenly double. Investors believe this doubling of conditional variances is permanent. Conditional covariances therefore double. Conditional risk premia are proportional to conditional covariances, and therefore they also double. This logic holds for any asset, including bonds.

Asset prices must fall to provide investors these higher conditional risk premia. The magnitude of the decrease depends on the amount by which risk premia increase. At one extreme is an asset with a zero average risk premia. The asset's price will not budge. Equities are at the other end of the spectrum. For nominal bonds, the permanent doubling of risk premia requires a doubling of mean yield spreads between long- and short-maturity bonds. (The math can be found in the Internet Appendix.) For the 1969 through 2013 sample, the mean yield spread between the five-year yield and the three-month yield is 117 basis points. Therefore, a permanent unexpected doubling in conditional variances raises the five-year yield by 117 basis points.

This hypothetical five-year yield shock of 117 basis points is large relative to the 61 basis point sample standard deviation of quarterly five-year yield shocks plotted in Figure 7. But in the parameterized B/S model, a doubling of conditional variances over a quarter is close to a seven-standard-deviation event.¹⁴ In addition, the conditional variance shocks are not permanent. Moreover, shocks to risk premia are negatively correlated with shocks to real rates because higher conditional volatilities drive a larger precautionary demand for saving. The bottom line is that time-varying risk premia contribute little to quarterly shocks to bond yields.

C. Alternative Parameterizations

Parameter estimates in B/S maximize the joint likelihood of observed nominal yields, survey forecasts of inflation, and survey forecasts of economic growth. The flexibility of the B/S model makes it easy to choose parameters that are instead consistent with observed volatilities of shocks to yields and expected inflation. This section shows that such parameterizations require extremely unrealistic properties of consumption growth and/or bond yields.

At short maturities, some combination of a higher volatility to shocks to expected consumption growth and a lower EIS is necessary to generate high volatilities of yield shocks and low volatilities of news about expected inflation. Matching volatilities at longer maturities requires either a volatile and persistent short-term real rate process or volatile risk premia. The latter can be generated with either more stochastic volatility or higher average risk premia (e.g., a larger coefficient of relative risk aversion).

The panel "Version 2" in Figure 8 uses a higher volatility of shocks to expected consumption growth. It is generated by the B/S model using an unconditional standard deviation of shocks to expected consumption growth that is almost four times the B/S value. Version 2 also uses a lower persistence of expected inflation to better match the empirical pattern of declining volatilities of expected inflation shocks. (The parameter values are in Table V.) These choices produce highly volatile short-term real rates, combined with low volatility inflation shocks.

¹⁴ For both expected consumption growth and expected inflation, a 1 *SD* shock to conditional variances is approximately 0.15 of the mean conditional variance.

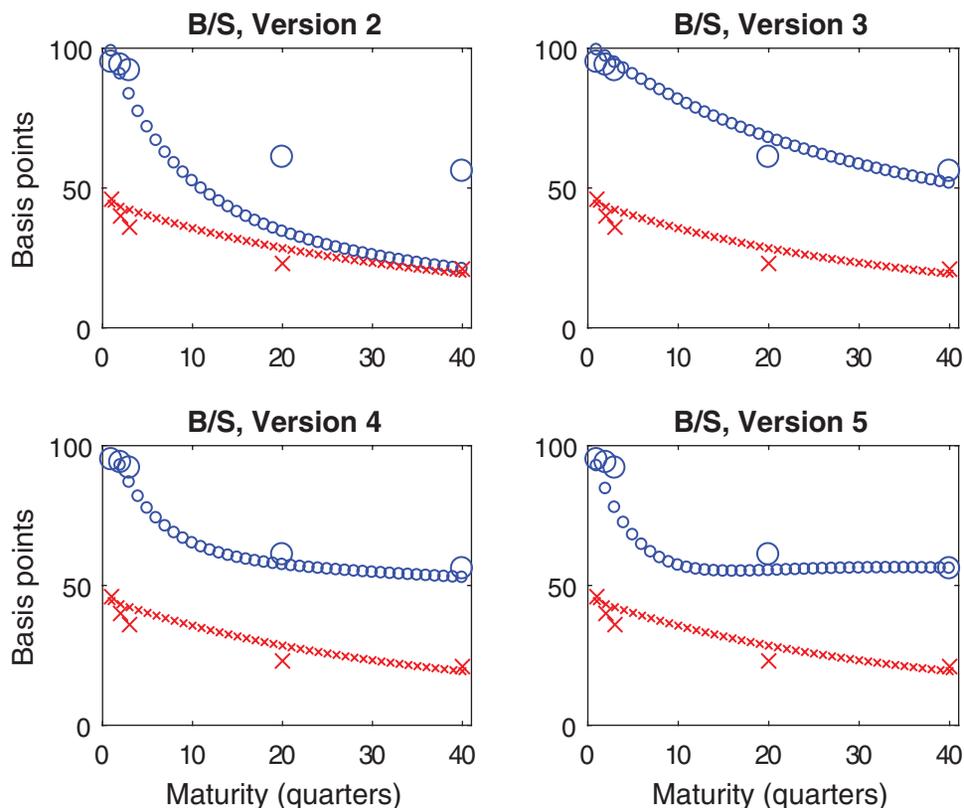


Figure 8. Standard deviations of the components of shocks to nominal yields for variants of the Bansal and Shaliastovich (2013) long-run risk model. See the notes for Figure 7 for a description of the points and lines. Unconditional model-implied standard deviations of yield shocks (solid) and news about expected inflation (dashed) are determined using various parameterizations of the estimated model in Bansal and Shaliastovich (2013). Version 2 uses a high volatility of shocks to expected consumption growth. Version 3 combines Version 2 with a highly persistent process for expected consumption growth. Version 4 combines Version 2 with high volatility of conditional volatility. Version 5 uses both a low elasticity of intertemporal substitution and a high coefficient of relative risk aversion. See the text for details. (Color figure can be viewed at wileyonlinelibrary.com)

The most obvious problem with the parameterization of Version 2 is excessive volatility of consumption. Table V reports that the standard deviation of quarterly consumption growth is 1.5 times its B/S baseline. The fourth-order serial correlation doubles. Less obvious is the model’s inability to match the positive mean slope of the nominal yield curve. Higher volatility of expected consumption growth raises the hedging demand for long-term real bonds. Therefore, the slope of the real yield curve is more negative with Version 2 than with the original parameters. In addition, Version 2’s reduction in expected inflation uncertainty makes the uncertainty in nominal bonds closer to the uncertainty

in real bonds. In combination, these effects produce a negatively sloped mean nominal yield curve.

The panel “Version 3” adds to Version 2 a higher persistence of the process for expected consumption growth. This change raises the persistence of short-term real rates, thus producing more news about expected future real rates at long horizons. The chosen parameters roughly match the volatilities of both yield shocks and expected inflation shocks across the range of maturities from three months to 10 years.

However, the fitting problems noted with Version 2 are greatly magnified with Version 3. Table V shows that the standard deviation of quarterly consumption growth exceeds 3% and the fourth-order serial correlation is nearly 0.8. Marginal utility is more volatile, which both lowers short-term rates through the precautionary savings motive and reduces the slope of the term structure through greater hedging demand for real bonds. Mean nominal yields are negative (there is no money in the economy to create a ZLB) and are much lower for long maturities than short maturities.

“Version 4” matches volatilities at long maturities with news about expected excess returns rather than news about expected real rates. This requires an extremely large amount of stochastic volatility relative to the B/S baseline. The standard deviation of shocks to the conditional variance of expected consumption growth is 20 times the B/S value. This choice drives larger swings in risk premia, and thus larger volatilities of news about expected excess returns. Figure 8 shows that this parameterization fits observed volatilities of shocks to yields and expected inflation. But the problems noted with Versions 2 and 3 carry over here. As with Version 2, consumption growth is still too volatile and persistent. Large variation in marginal utility associated with stochastic volatility creates substantial hedging demand for bonds, and thus, as with Version 3, the mean slope of the nominal yield curve is negative and large.

Versions 2 through 4 alter the consumption process parameterized in B/S. Version 5 instead adjusts the preference parameters. Therefore, the consumption properties of Version 5 are plausible, as reported in Table V.¹⁵ Matching volatilities at the short end requires an EIS close to one-half. Matching volatilities at long maturities requires substantial news about expected excess returns, which is achieved by having agents who are extremely sensitive to shocks. The coefficient of relative risk aversion must be 170 to match these volatilities. Model-implied volatilities are again shown in Figure 8.

Bansal and Yaron (2004) emphasize the importance of an EIS greater than 1 in matching asset price behavior using long-run risk models. Moreover, the large coefficient of relative risk aversion creates an astronomically large precautionary savings effect. The model-implied mean yields in Table V are below $-25,000\%$, illustrating the difficulty in matching the joint behavior of bond yields and expected inflation in a long-run risk framework.

¹⁵ Version 5, like the other alternatives, uses a relatively low persistence of inflation expectations. This indirectly affects the consumption process, resulting in slightly lower consumption volatility and serial correlation than in the B/S version.

D. Nonlinear New Keynesian Models

The endowment economies of P/S and B/S use conditionally log-normal joint dynamics for the stochastic discount factor (SDF) and asset prices. Backus, Boyarchenko, and Chernov (2016) explore ways in which departures from log-normality, such as jumps, lead to more flexible risk premia dynamics in a recursive utility setting. Aside from dynamic New Keynesian asset pricing models, this is largely unexplored territory.

Swanson (2016) emphasizes that New Keynesian models exhibit empirically relevant nonlinearities when preferences are sufficiently curved. Conditional heteroskedasticity in both the SDF and asset returns are created by these nonlinearities, even when shocks to fundamental variables such as productivity and government spending are homoskedastic. Thus, risk premia, and conditional expectations of excess returns to assets, will vary over time.

A leading example of New Keynesian asset pricing models is Kung (2015). Households have recursive utility preferences. Transitory productivity shocks have long-run effects that are endogenously generated by investment in both physical capital and R&D. Stagflation also arises endogenously. Positive shocks to productivity raise investment, raise expected productivity growth, and lower marginal costs, so that monopolistically competitive firms cut prices in order to capture business. Relative to an endowment economy, the endogenous capital economy has much higher short-run volatility of short-term real rates, as illustrated in the top panel of Figure 9.¹⁶ The positive shock to productivity initially lowers consumption because investment demand is high. Expected consumption growth is steep for a few quarters (producing high real rates), then tapers off and turns negative.

Details of Kung's (2015) model differ from those of P/S and B/S, but Figure 9 shows that basic properties of yield shocks are unchanged. Inflation variance ratios in Kung's (2015) model all exceed one. Model-implied volatilities of inflation news are too large relative to sample volatilities, while model-implied volatilities of yield shocks are too small. (The figure displays standard deviations of yield shocks for only a few maturities because there is no analytic expression for these yields.) The main mechanisms are the same as in the endowment economy models. News about expected real rates is small and negatively correlated with news about expected inflation.

Kung's (2015) model also shares with P/S and B/S the negligible contribution of news about expected excess returns. With an EIS of 2 and a coefficient of relative risk aversion of 10, apparently the model is not sufficiently nonlinear to drive a wedge between its risk premium dynamics and those of conditionally log-normal models.

The much simpler New Keynesian model in R/S helps illustrate both potential role of nonlinearities and boundaries on their contributions. The baseline recursive utility model of R/S does not have long-run risk dynamics. Since real shocks are less persistent in R/S than in Kung's (2015) long-run risk model,

¹⁶ Howard Kung graciously provided his Dynare++ code.

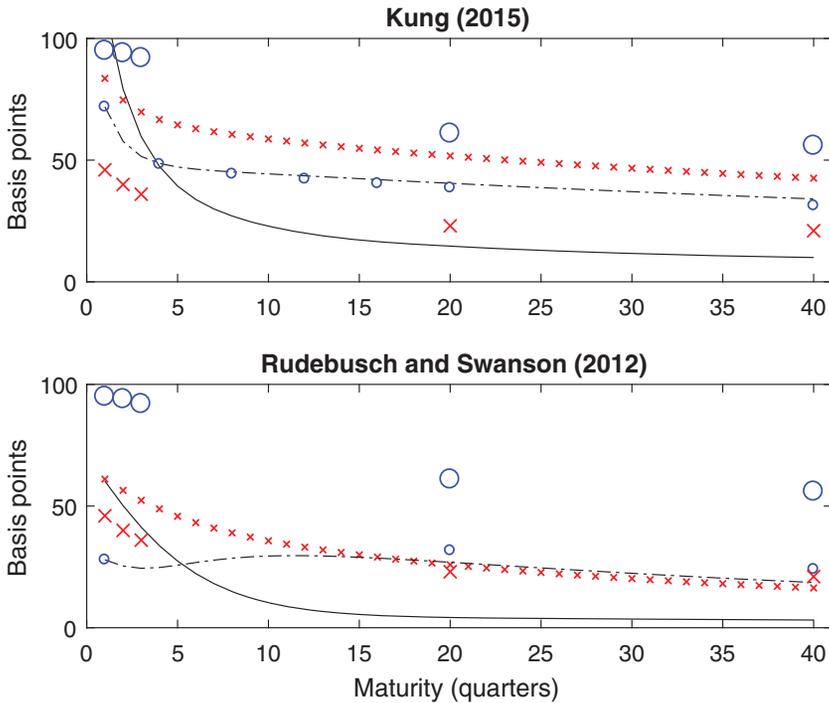


Figure 9. Standard deviations of the components of shocks to nominal yields for two New Keynesian models. See the notes for Figure 7 for a description of the points and the lines. The estimated models are those of Kung (2015) and the baseline model of Rudebusch and Swanson (2012). (Color figure can be viewed at wileyonlinelibrary.com)

Figure 9 shows that volatilities of both news about expected inflation and expected real rates are much smaller than in Kung's (2015) model.¹⁷

R/S's representative agent, with a coefficient of relative risk aversion of 75, has a much greater desire to smooth consumption than does the agent in Kung's (2015) model.¹⁸ This high desire to smooth creates strong nonlinearities that in turn create noticeable conditional heteroskedasticity of the SDF. For example, the model-implied unconditional standard deviation of quarterly shocks to the 10-year yield is 24 basis points. (This is displayed in Figure 9, although it is a little difficult to distinguish this standard deviation from three others that are plotted in the same neighborhood.) The model-implied unconditional standard deviation of the sum of news about expected inflation and expected real rates is 19 basis points, about two basis points greater than the standard deviation

¹⁷Eric Swanson graciously produced for me a long time series of data from the model's data generating process.

¹⁸This coefficient is not the usual coefficient γ in the exponent $(1 - \gamma)$ on consumption used to convert consumption into utility. Swanson (2018) explains how to calculate relative risk aversion in a recursive utility framework with utility derived from both consumption and labor. The coefficient γ in R/S is approximately 150.

of news about expected inflation. Therefore, the inflation variance ratio of 0.46 is driven largely by the volatility of news about excess returns. Although this value is not in the 0.1 to 0.2 range documented in Section II, it is well below the variance ratios implied by the other models studied here.

At this point it is worth restating the challenge posed by the evidence in Section II. Can a model match *both* sample inflation variance ratios and the sample volatilities of yield shocks? In the R/S model, news about expected inflation raises the standard deviation of 10-year yield shocks to 24 basis points from 19 basis points, but the sample standard deviation of yield shocks is near 60 basis points. A successful model requires much more news about either expected real rates or expected excess returns.

Perhaps the most obvious way to increase the volatilities of yield shocks in the R/S model is to introduce long-run productivity risk.¹⁹ As in Kung's (2015) model, this will magnify news about expected real rates and expected inflation. But it may well shrink, rather than magnify, news about expected excess returns. As noted by Swanson (2016), models with a low quantity of risk, such as the benchmark model of R/S, require high curvature to generate plausible mean risk premia and thus a plausible average slope of the nominal term structure. Models with higher quantities of risk, such as in Kung (2015), require smaller curvature to fit the average slope. A corollary to Swanson's (2016) point is that, since smaller curvature reduces the conditional heteroskedasticity of the SDF, a model with more plausible levels of yield shock volatilities will also have less time-variation in conditional risk premia.

E. A Preference Shock Interpretation

This research adds to a voluminous literature documenting that much of the variation in asset prices is unexplained by macroeconomic dynamics. Responding to this general empirical failure, Albuquerque et al. (2016) argue that shocks to preferences rather than shocks to the macroeconomy can help explain a variety of asset pricing puzzles. They model agents' time rate of preference as an exogenous persistent stochastic process. Shocks to the process create shocks to current short-term real rates. The persistence of the process implies that these shocks also affect expected future short-term real rates. Therefore, an "animal spirits" state variable generates news about expected future short-term real rates.

Cram (2016) builds on an earlier version of the current research by asking whether the approach of Albuquerque et al. (2016) can explain the evidence of Section II and simultaneously match other features of the term structure, such as mean yields and yield forecastability. Given a state variable expressly designed to generate news about real rates, it is easy to match observed inflation variance ratios. Taking the inflation process as exogenous, inflation variance ratios at various maturities pin down the volatility and persistence of the

¹⁹ An extension in R/S considers and rejects such a model on grounds not directly relevant to this discussion.

process for the time rate of preference. Testable restrictions come from matching other aspects of bond yields. Cram (2016) argues that the parameterized model can fit properties such as conditional expected excess bond returns.

Of course, attributing the central puzzle of these results to animal spirits is not particularly satisfying. It is worth exploring how well another standard consumption-based framework can explain the decomposition of shocks to nominal yields.

F. Habit Formation

Habit formation preferences provide a consumption-based setting that breaks the link between real rates and expected consumption growth. Asset pricing investigations using habit formation usually follow the path of Campbell and Cochrane (1999), who emphasize the role of surplus consumption. Their specification of preferences includes surplus consumption in a linear specification of the log SDF. Because of the formal similarity of this approach to a preference shock specification, preference shock models such as Albuquerque et al. (2016) have sometimes been labeled habit formation models. The empirical content of the Campbell-Cochrane (1999) view of habit formation relative to a general preference shock view comes from three restrictions. Shocks to surplus are perfectly correlated with shocks to consumption, surplus is mean-reverting because agents grow accustomed to their consumption level, and surplus is more volatile when it is low.

Wachter (2006) studies bond pricing in an extension of Campbell and Cochrane's (1999) model. She chooses parameter estimates to fit specific properties of the joint dynamics of nominal yields, inflation, and aggregate consumption growth. Although the estimates are not expressly picked to match observed inflation variance ratios, the estimated model's properties are much more encouraging than those of long-run risk models.

Figure 10 displays the same information for Wachter's estimated model that is displayed in Figure 7 for long-run risk models.²⁰ Comparing the figures reveals two important differences between the long-run risk and habit formation models. First, the habit formation model generates large standard deviations of real rate news at all horizons. Second, the model generates large standard deviations of news about expected excess returns to long-maturity nominal bonds. In combination, these properties of habit formation produce high volatilities of shocks to yields accompanied by low inflation variance ratios. The ratios range from about 0.6 at the one-quarter maturity to only 0.12 at the 10-year maturity.

Taken at face value, these results suggest that nominal term structure dynamics are more consistent with habit formation models than with long-run risk models. The first of two large caveats is that the results are qualitatively sensitive to plausible variations in the parameters. In particular, Wachter's parameterized model implies that news about expected real rates and expected

²⁰ Jessica Wachter graciously provided her Matlab code and helped me understand some of its features.

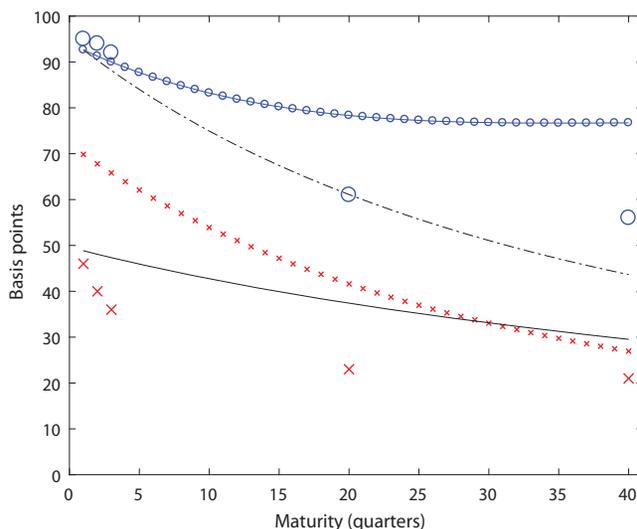


Figure 10. Standard deviations of the components of shocks to nominal yields for a habit formation model. The figure plots sample and model-implied unconditional standard deviations of quarterly shocks. Sample standard deviations of yields (large O's) and news about expected inflation (large X's) are for the period 1968Q4 through 2013Q4 and come from Tables I and III. Unconditional model-implied standard deviations of yield shocks (small o's), news about expected inflation (small x's), and news about expected real rates (solid line) are determined using the estimated model in Wachter (2006). The figure also displays model-implied standard deviations of the sum of news about expected inflation and expected real rates (dotted-dashed). (Color figure can be viewed at wileyonlinelibrary.com)

inflation are positively correlated. This positive correlation magnifies shocks to nominal yields. It also increases the magnitude of shocks to expected excess returns because assets with high volatility return shocks are more exposed to risk preferences than are assets with low volatility return shocks.

However, negative correlations between these two types of news are easy to defend. Le, Singleton, and Dai (2010) note that Wachter's parameters imply that the average term structure of yield volatilities is upward-sloped, a property that is inconsistent with the data. In the Internet Appendix, I show that positive correlations contribute to the upward-sloped term structure of volatilities. The results of Ermolov (2015) illustrate how these correlations affect inflation variance ratios. He studies term structure behavior in the habit formation model of Bekaert and Engstrom (2017). The parameter estimates imply a negative correlation between news about expected real rates and expected inflation. Results reported in the Internet Appendix document that the resulting inflation variance ratios are much larger than those of Wachter's model.

The second caveat is that evaluating the conditional properties of habit formation models is beyond the scope of this paper. Yet such models' unconditional properties critically rely on their nonlinear conditional properties. For example, Wachter's model implies that conditional inflation variance ratios are much

lower in recessions than in booms. In recessions, surplus is low and both real rates and expected excess returns are relatively volatile. If these conditional properties are not supported empirically, the ability of the model to fit unconditional properties is irrelevant.

IV. A Dynamic Model of Yields and Expected Inflation

What drives shocks to bond yields? Section II confirms that it is not news about inflation but does not help disentangle other possibilities. In an accounting sense, yields must be driven primarily by news about expected future short-term real rates, expected future excess returns, or both. This section attempts to answer the question. To preview the results, there is not enough information in the data to tell.

A. The Framework

Yield shocks are the sums of news about expected future inflation, expected future short-term real rates, and expected future excess returns. Section II.B uses a dynamic model of short-horizon inflation expectations to infer inflation news. Here we expand the dynamic model to include short- and long-term nominal rates. News about expected future short-term real rates is inferred from the model. News about expected future excess returns is produced as a residual. The model necessarily requires much more structure than the trend-cycle model in Section II.B.

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. The first application of these models to interest rates is Hamilton (1985), although his motivation differs from that in the dynamic term structure literature. The state vector has homoskedastic Gaussian VAR(1) dynamics

$$x_{t+1} = \mu + Kx_t + \Sigma\epsilon_{t+1}, \quad \epsilon_{t+1} \sim MVN(0, I). \quad (15)$$

We know from Figures 1 through 4 that shocks to neither yields nor inflation expectations are homoskedastic. I adopt the assumption both because it simplifies the model considerably and because it is a conservative assumption here. The main conclusion drawn from the results that follow is that there is not enough information in the data to explain why yield shocks are so large. This conclusion holds even when we impose a counterfactually strong assumption on conditional second moments.

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

$$y_t^{(m)} = A_m + B'_m x_t + \chi_{m,t}, \quad (16)$$

where $\chi_{m,t}$ represents measurement error or some other deviation from an exact affine representation. Similarly, the one-period-ahead expectation of future inflation is an affine function of the state,

$$E_t(\pi_{t+1}) = A_\pi + B'_\pi x_t. \quad (17)$$

There is no measurement error in (17). The measurement error is included in the consensus forecast. Expectations of j -period-ahead inflation combine (17) with the state dynamics (15),

$$E_t(\pi_{t+j}) = A_\pi + B'_\pi E_t(x_{t+j-1}). \quad (18)$$

Many researchers use similar frameworks to study the joint dynamics of inflation and bond yields. Notable examples include Campbell and Viceira (2001), who estimate two-factor Gaussian no-arbitrage models of nominal yields and inflation; Ang, Bekaert, and Wei (2008), who estimate a four-factor model with time-varying risk premia and an additional factor that captures changes in regimes; Chernov and Mueller (2012), who estimate a variety of four-factor and five-factor Gaussian models; and Haubrich, Pennacchi, and Ritchken (2012), who estimate a seven-factor model with stochastic volatility.

This paper diverges from the earlier literature in both its objective and its estimation procedure. The earlier work investigates the risk compensation investors require to face shocks, and thus they can put a price on the nominal risk embedded in nominal Treasury bonds. Here, the absence of arbitrage is not imposed. No arbitrage places restrictions on the coefficients of (16). By itself, the assumption of no-arbitrage is unimportant in this VAR setting. 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 (15) and (16). No-arbitrage restrictions matter only when they are coupled with restrictions on risk premia dynamics. Such restrictions can certainly affect inference about the variance decomposition (6). But I do not want to inadvertently impose restrictions on (6) that force a conclusion onto the data.

Calculations of news about expected future inflation and expected future short-term real rates use standard VAR mathematics. More details can be found in the Internet Appendix.

B. The Data and Estimation Details

Statistical inference with this highly parameterized model is improved with a longer time series. I therefore use the SPF consensus forecasts. The sample is 1968Q4 through 2013Q4, a total of 181 quarters. The observables are the consensus forecasts of GDP inflation from one to four quarters ahead, as well as Treasury bond yields. Treasury bond yields are observed in the middle of the second month of each quarter, roughly aligned with the SPF forecast dates. Zero-coupon yields are for maturities of three months, 1 through 5 years,

and 10 years. Sources of yield data are described in Section I.D. The Internet Appendix contains information about measurement error assumptions.

The length of the state vector is not specified in the model of Section IV.A. Results for a four-factor model are discussed below. I explore versions with three and five factors, but the results were not sufficiently novel to present. The state vector is latent and thus unidentified. Normalizations described in the Internet Appendix are imposed in estimation to eliminate global and local underidentification.

The likelihood function is given by the Kalman filter and the parameters are estimated by maximizing the likelihood. The largest eigenvalue of K is set to 0.999. Therefore, the model is stationary but indistinguishable from one with a unit root. In practice there is no way to distinguish statistically a unit root in the K matrix from an extremely persistent process. Allowing all eigenvalues to be free parameters produces less persistent shocks to yields and expected inflation, but does not affect appreciably estimates of inflation variance ratios.

After imposing the eigenvalue restriction and other restrictions described in the Internet Appendix, the four-factor model has 50 free parameters. 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 individual parameter estimates are not of direct interest and thus are reported in the Internet Appendix.

C. Variance Decompositions

The clearest message contained in these results is unsurprising, given the evidence of Section III. Inflation shocks account for a small fraction of the total variance of shocks to nominal yields. Table VI presents detailed results. Variance decompositions are reported for bonds with maturities of one quarter and 1, 5, and 10 years. For all of the bonds, the inflation variance ratio ranges from roughly 0.15 to 0.2. The two-sided 95th percentile confidence bounds are tight. For each bond, we confidently conclude that the inflation variance ratio is less than 0.3.

Table VI documents two new results. First, there is insufficient information in the data to accurately decompose the remaining variance of long-maturity yields into news about expected future real rates and expected excess returns. The point estimates suggest that at the 5-year maturity, their relative contributions are equal, while expected excess return news is more important at the 10-year maturity. However, the confidence bounds are huge, nesting point estimates that allow either source to dominate the other.

Second, the point estimates in the table imply a positive covariance between news about expected real rates and expected excess returns. The estimates indicate that between 15% and 30% of the variance of yield shocks is attributable to this covariance. The confidence bounds are again very large. We next consider what features of the data produce these two results.

Table VI
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. The three components of yield shocks are news about expected future real rates, expected future inflation, and expected future excess returns. The contributions to total variance sum to 1, aside from rounding. The sample period is 1968Q4 through 2013Q4. Brackets display [2.5%, 97.5%] confidence bounds.

Statistic	Maturity (Years)			
	1/4	1	5	10
<i>SD</i> (yield shocks)	0.96 [0.93 2.94]	0.84 [0.81 2.86]	0.63 [0.59 1.78]	0.52 [0.47 1.13]
Var(real rate news)	0.86 [0.75 1.19]	0.72 [0.53 0.97]	0.32 [0.13 1.06]	0.24 [0.07 1.28]
Var(inflation news)	0.14 [0.06 0.20]	0.14 [0.05 0.18]	0.16 [0.03 0.20]	0.22 [0.03 0.29]
Var(return news)	0.00 [0.00 0.00]	0.03 [0.01 0.10]	0.22 [0.09 0.66]	0.29 [0.13 1.28]
2Cov(real rate news, inflation news)	0.00 [-0.34 0.12]	-0.02 [-0.27 0.10]	0.14 [-0.22 0.28]	0.31 [-0.18 0.56]
2Cov(real rate news, return news)	0.00 [0.00 0.00]	0.16 [-0.02 0.38]	0.30 [-0.50 0.55]	0.15 [-1.39 0.53]
2Cov(inflation news, return news)	0.00 [0.00 0.00]	-0.03 [-0.09 0.03]	-0.14 [-0.38 0.01]	-0.21 [-0.73 0.08]

D. Some Impulse Responses

Figure 11 displays responses to a one-standard-deviation shock to the real short rate. The shock affects current and expected future real short rates, inflation, the 10-year bond yield, and the term premium on the 10-year bond. The initial real rate shock is about 90 basis points. The point estimates imply that the shock dies out quickly, as illustrated in Panel A. Panel B shows that expected inflation is unaffected by shocks to the real short rate. Panel C shows that the point estimate of the immediate response of the 10-year yield is a little more than 30 basis points and statistically distinguishable from zero. Because the effects of the shock on expected future real rates and inflation are so small, this response of the 10-year yield cannot be explained by either real rate news or inflation news. It must therefore be explained by an increase in the term premium, as shown in Panel D. The term premium jumps by more than 15 basis points.

The responses to real rate shocks account for the positive estimated covariance between average expected real rates and term premia. In addition, the confidence bounds on these responses account for the inability 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

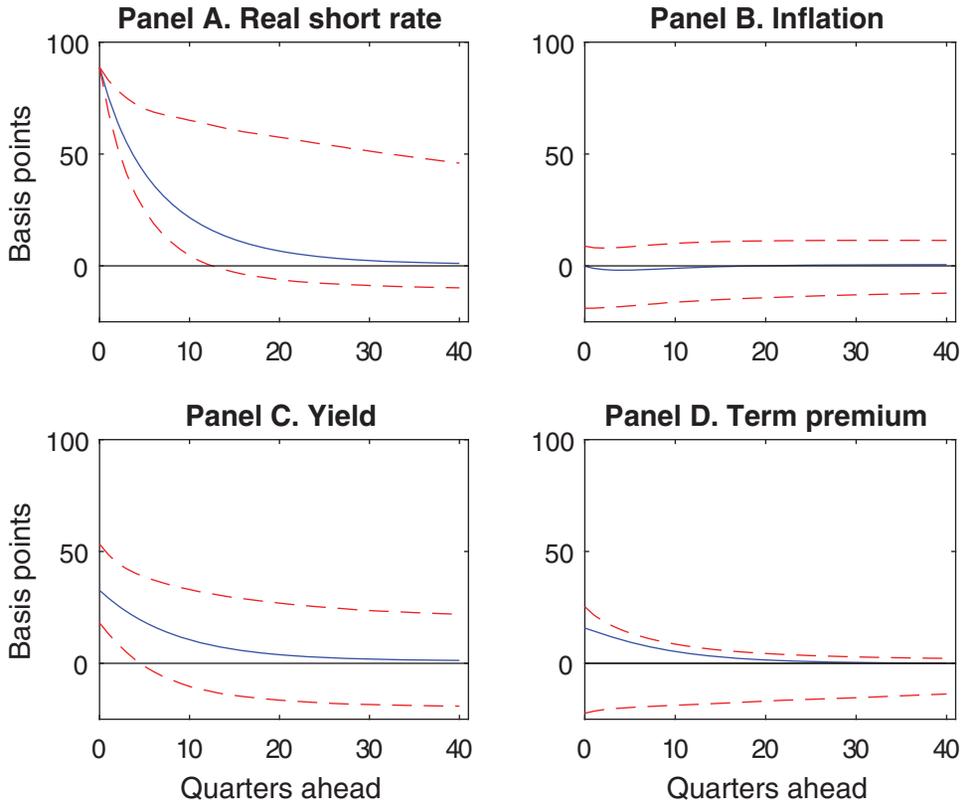


Figure 11. Impulse responses for a shock to the ex ante real rate. A four-factor dynamic model of nominal yields and expected inflation is estimated over the sample 1968Q4 through 2013Q4. The figure displays model-implied impulse responses to a shock to the ex ante real rate, defined as the three-month nominal yield less expected inflation during the next quarter. The initial shock is 89 basis points, which is the model-implied population standard deviation of the shock. The yield and term premium responses are for a 10-year nominal bond. Also displayed are 95 percentile confidence bounds on the impulse responses. (Color figure can be viewed at wileyonlinelibrary.com)

actually highly persistent. If real rates are highly persistent, then the immediate response of the 10-year yield to the shock to real rates is in line with the shock to the average expected real rate over the next 10 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 with effects that die out fairly quickly. 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 have highly persistent effects.

Investors know this and price long-term bonds accordingly. This explanation implies that investors were surprised by the speed with which the shocks died out in the sample. There is not enough information in the sample to reject either explanation.

Neither theory nor existing empirical literature offers much to help pin down the relative contributions of news about average expected real rates or shocks to expected excess returns. 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 (2016). 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 (2015) find that monetary policy shocks have substantial effects on long-term nominal and inflation indexed 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 (2018) disagree about both the magnitude of shocks to long-maturity yields and their interpretation as primarily term premia.

The evidence in Table VI indirectly helps motivate this paper's focus on the inflation variance ratio. Yield shocks consist of news about expected future inflation, expected future real rates, and expected excess returns. Therefore, as in equation (6), three variances and three covariances contribute to the variance of yield shocks. Each of these six elements is economically interesting, but only one—the variance of news about expected inflation—can be pinned down in the data.

V. Conclusion

This paper studies the joint dynamics of nominal yields and inflation expectations from 1968 through 2013. Unconditionally, news about expected future inflation contributes relatively little to the variance of yield shocks. Quarterly shocks to nominal yields are primarily shocks to real rates and term premia. This result holds for long-maturity bonds during a variety of subperiods: the passive monetary policy regime of the 1970s, the Volcker disinflation, the Great Moderation, and the financial crisis/ZLB subperiod.

This robust result can help evaluate dynamic equilibrium macroeconomic models. For example, long-run risk models in the literature imply that inflation shocks drive almost all of the variation in nominal yields. Plausible parameterizations of the models generate little news about either expected future short-term real rates or expected future excess bond returns.

Initial submission: October 4, 2014; Accepted: September 1, 2017
Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

REFERENCES

Albuquerque, Rui, Martin Eichenbaum, Victor Xi Luo, and Sergio Rebelo, 2016, Valuation risk and asset pricing, *Journal of Finance* 71, 2861–2903.

- 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.
- Backus, David, Nina Boyarchenko, and Mikhail Chernov, 2016, Term structures of asset prices and returns, Working paper, UCLA.
- Balduzzi, Pierluigi, and Chunhua Lan, 2014, Survey forecasts and time-varying second moments of stock and bond returns, Working paper, Boston College.
- Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *Review of Financial Studies* 26, 1–33.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Bauer, Michael D., and Glenn D. Rudebusch, 2017, Interest rates under falling stars, Working paper, Federal Reserve Bank of San Francisco.
- Bekaert, Geert, and Eric Engstrom, 2017, Asset return dynamics under habits and bad environment–good environment fundamentals, *Journal of Political Economy* 125, 713–760.
- Bekaert, Geert, Robert J. Hodrick, and David A. Marshall, 1997, On biases in tests of the expectations hypothesis of the term structure of interest rates, *Journal of Financial Economics* 44, 309–348.
- 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 John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Robert J. Shiller, 1987, Cointegration and tests of present value models, *Journal of Political Economy* 95, 1062–1088.
- Campbell, John Y., and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- Campbell, John Y., and Robert J. Shiller, 1991, Yield spreads and interest rate movements: A bird's eye view, *Review of Economic Studies* 58, 495–514.
- 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.
- Cieslak, Anna, 2017, Short-rate expectations and unexpected returns in Treasury bonds, *Review of Financial Studies* (forthcoming).
- Cieslak, Anna, and Pavol Povala, 2015, Expected returns in Treasury bonds, *Review of Financial Studies* 28, 2859–2901.
- Clarida, Richard, Jordi Galí, and Mark Gertler, 2000, Monetary policy rules and macroeconomic stability: Evidence and some theory, *Quarterly Journal of Economics* 115, 147–180.
- 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, 2012, What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120, 116–159.
- Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Informational rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.
- Cram, Roberto Gómez, 2016, How important are inflation expectations for the nominal yield curve? Working paper, Wharton.

- 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., 2002, Term premia and interest rate forecasts in affine models, *Journal of Finance* 57, 405–443.
- Ermolov, Andrey, 2015, Time-varying risk of nominal bonds: How important are macroeconomic shocks? Working paper, Columbia GSB.
- 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, 2013, Forecasting inflation, in Graham Elliot, Clive Granger, and Allan Timmerman, eds: *Handbook of Economic Forecasting*, Vol. 2 (North Holland, Amsterdam).
- Giacoletti, Marco, Kristoffer Laursen, and Kenneth Singleton, 2015, Belief heterogeneity, learning, and risk premiums in an arbitrage-free term structure model, Working paper, Stanford University.
- Hamilton, James D., 1985, Uncovering financial market expectations of inflation, *Journal of Political Economy* 93, 1224–1241.
- Hanson, Samuel G., and Jeremy C. Stein, 2015, Monetary policy and long-term real rates, *Journal of Financial Economics* 115, 429–448.
- 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.
- Hjalmarsson, Erik, and Pär Österholm, 2010, Testing for cointegration using the Johansen methodology when variables are near integrated: Size distortions and partial remedies, *Empirical Economics* 39, 51–76.
- Johannes, Michael, Lars A. Lochstoer, and Yiqun Mou, 2016, Learning about consumption dynamics, *Journal of Finance* 71, 551–600.
- 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* 109, 604–622.
- Kung, Howard, 2015, Macroeconomic linkages between monetary policy and the term structure of interest rates, *Journal of Financial Economics* 115, 42–57.
- Lardic, Sandrine, and Valérie Mignon, 2004, Fractional cointegration and the term structure, *Empirical Economics* 29, 723–736.
- Le, Anh, and Kenneth J. Singleton, 2013, The structure of risks in equilibrium affine models of bond yields, Working paper, Stanford GSB.
- Le, Anh, Kenneth J. Singleton, and Qiang Dai, 2010, Discrete-time affine^Q term structure models with generalized market prices of risk, *Review of Financial Studies* 23, 2184–2227.
- Lubik, Thomas A., and Frank Schorfheide, 2004, Testing for indeterminacy: An application to U.S. monetary policy, *American Economic Review* 94, 190–217.
- Nakamura, Emi, and Jón Steinsson, 2018, High frequency identification of monetary non-neutrality: The information effect, *Quarterly Journal of Economics* (forthcoming).
- Nason, James M., and Gregor W. Smith, 2014, Measuring the slowly evolving trend in U.S. inflation with professional forecasts, Australian National University CAMA 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.
- 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, in *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.
- 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.
- Swanson, Eric T., 2016, A macroeconomic model of equities and real, nominal, and defaultable debt, Working paper, University of California, Irvine.
- Swanson, Eric T., 2018, Risk aversion, risk premia, and the labor margin with generalized recursive preferences, *Review of Economic Dynamics* 28, 290–321.

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

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.