Do Long-Term Yields Help to Estimate Short-Term Risk Premia?

Andrea Berardi, Roger Brown and Stephen M. Schaefer *

This version: August 5, 2026

Abstract

We identify a new and important property of affine term structure models: at long maturities, the slope of the term structure of risk premia becomes "almost observable". Exploiting this property by fitting models to long-maturity data dramatically improves the precision of estimates of risk premia and expected short rates across the term structure. In a Monte Carlo experiment we show that using rates with maturities up to 25 years, rather than 10 years as is standard in the literature, reduces the standard deviation of estimation errors in risk premia, expected short rates and term premia by over 70%. Applying this approach to US Treasury data, we find that variation in expected short rates accounts for a much larger share of variation in yields (and variation in term premia a correspondingly smaller fraction) than in the widely used estimates reported by the Federal Reserve. Accounting for time-variation in yield volatility is key to successful extraction of the risk premium information in long rates. Our approach therefore departs from much of the empirical literature in two important respects: it incorporates long-maturity yields and allows for time-varying yield volatility, whereas most existing studies ignore long-term yields and rely on constant volatility models.

Keywords: Bond Risk Premia, Long-Term Rates, Yield Volatility, Term Structure
JEL Classification: E43, G12

*Berardi is with Università Ca' Foscari Venezia, Department of Economics, Fondamenta San Giobbe 873, 30121, Venezia, Italy, Email: andrea.berardi@unive.it. Brown E-mail: rbrown@london.edu. Schaefer is with London Business School, Sussex Place, Regent's Park, London NW1 4SA, UK, E-mail: sschaefer@london.edu. Dimitris Papanikolaou was the editor for this article. We are very grateful to an anonymous referee for many helpful suggestions and constructive guidance. We thank Peter Feldhutter, Roberto Gomes Cram, Marco Grotteria, Howard Kung, Ilaria Piatti, Claudio Tebaldi, Jonathan Wright, and seminar participants at the 10th Bundesbank Term Structure Workshop, the 9th University of York Asset Pricing Workshop, and the Collegio Carlo Alberto in Turin for valuable comments.

1 Introduction

Improving the precision with which risk premia can be disentangled from expectations is central to what is arguably the main policy application of models of the term structure. In an address to the Economic Club of New York in 2006, Federal Reserve Chairman Ben Bernanke made the following much-quoted observation about the so-called "Greenspan conundrum":

> "*What does the historically unusual behavior of long-term yields imply for the conduct of monetary policy? The answer, it turns out, depends critically on the source of that behavior. To the extent that the decline in forward rates can be traced to a decline in the term premium, ... the effect is financially stimulative and argues for greater monetary policy restraint, all else being equal. ... However, if the behavior of long-term yields reflects current or prospective economic conditions, the implications for policy may be quite different - indeed, quite the opposite.*"

For policymakers, the ability to identify the source of a change in interest rates as either a change in expectations or a change in the risk (or term) premium is crucial.

Using US Treasury yields with maturities up to 25 years and a stochastic volatility term structure model, we find that expected short rates account for a much larger fraction of the variation in yields and forward rates than in widely used benchmark estimates. In particular, expected short rates account for approximately 70% of the variation in the 10-year yield, compared with around 30% in Adrian et al. (2013) and Christensen and Rudebusch (2012), and about 50% in Kim and Wright (2005) (the "Fed models"). More generally, across the full sample, the decomposition of yield variation differs markedly from standard estimates: expected short rates play a much larger role, while term premia play a correspondingly smaller one. During the recent monetary tightening cycle between March 2022 and October 2023, we estimate that only around 20% of the increase in the 10-year forward rate reflected higher term premia, compared with between 43% and 77% in benchmark models.

Taken together, these results imply markedly different drivers of movements in long-term interest rates, with correspondingly different implications for monetary policy.

Why do our estimates differ so substantially from those of the Fed models? We show that a major part of the answer lies in our use of the information contained in long-maturity forward rates. While the risk-neutral parameters of a term structure model can be obtained from the cross-section of yields, estimates of expectations of future rates and risk premia require the model's parameters under the natural measure. The latter are estimated by combining cross-sectional and time-series information, but here the literature has documented a number of significant problems. These include small-sample bias leading to overestimation of

mean reversion, and imprecision in estimates of the mean of the process generating interest rates (see, for example, Kim and Orphanides (2012) and Bauer et al. (2012)).

Of course, if risk premia − or a known function of risk premia − were observable, this would have the potential to substantially ease the problem of separating the expectations from risk premia. Surprisingly, since risk premia are not generally considered to be observable, this is the central idea in our paper.

We show that, at long maturities, the slope of the term structure of risk premia (TSRP) becomes "almost observable" and that including long-maturity data in estimation, thereby constraining the model-implied slope of the TSRP, significantly improves the precision of estimates of both risk premia and expected short rates.

This near-observability arises because the slope of the forward curve, adjusted by a convexity term that depends only on yield volatility, is equal to the sum of the slope of the TSRP and an expectations term, Δe (defined below), that becomes small at long maturities. The expectations term Δe is related to the slope of the term structure of expected short rates, and its convergence to zero at long maturities is central to our argument. We show that, under conditions routinely imposed in estimation, Δe must converge to zero in all constant volatility models and, intuitively, we expect this result to hold more generally. Since forward rates and yield volatility are observable, this implies that at long maturities the TSRP slope is effectively pinned down by the data.

How does knowledge of the *slope* of the TSRP at long maturities help to estimate the *level* of risk premia at shorter maturities? First, observability of the slope of the TSRP at long maturities, combined with the fact that risk premia must converge to zero at zero maturity, constrains estimates of the entire TSRP. Second, observability of the *time series* of both the slope of the TSRP and the state variables leads to relatively precise estimates of the sensitivity of risk premia to the state variables and, in particular, the market price of risk. We show that it is specifically the greater precision with which the market price of risk is estimated that leads to more precise estimates of the *level* of risk premia. Furthermore, since the risk-neutral mean-reversion matrix is generally estimated relatively accurately, an improvement in the estimated market price of risk will lead to an improvement in the estimate of the mean-reversion matrix under the natural measure, and a corresponding improvement in the estimated expected short rate. Surprisingly, by exploiting the fact that the slope of the forward rate curve at long maturities is little affected by expectations, we obtain more precise estimates of expectations of future rates.

The convexity adjustment is crucial. First, at long maturities convexity is large, typically

larger (in size) than the slope of the forward curve itself.[1] Second, because volatility varies over time, convexity at long maturities is not only large but also highly variable. Accounting for time-varying convexity is therefore important for the successful extraction of the TSRP from long-term forward rates and this necessitates a model with time-varying volatility.[2]

We carry out a series of Monte Carlo experiments to assess the effect of including long-maturity data on the precision of estimates of risk premia and expectations of future short rates. Here we follow Duffee and Stanton (2012) who emphasises the importance of using Monte Carlo to assess the estimation accuracy in term structure models rather than relying on asymptotic properties. For this purpose, we implement a standard stochastic volatility (SV) affine model ($A_1(4)$ in the Dai and Singleton (2000) system) and compare the precision of estimates obtained when the model is fitted to yield data with maturities up to either 25 years or − consistent with most empirical studies − only 10 years. The differences in precision are substantial. Including long-maturity data reduces the median standard deviation of estimation errors in risk premia by over 70%. Surprisingly, when rates up to 25 years rather than 10 years are used in estimation, the relative improvement in the precision of estimates of risk premia and expected short rates at *short* maturities is about the same as, and in some cases even greater than, at long maturities.

This is the first paper to show that the slope of the term structure of risk premia at long maturities is effectively identified from long-term forward rates, yielding substantial gains in the precision of risk premia and expectation estimates.[3] Despite the easy availability of data on long-term yields, the vast majority of empirical studies of the term structure have excluded long-maturity data and limited their attention to maturities of 10 years or less. Only a handful of papers include maturities of 20 or 30 years.[4]

Our results have two significant implications for the specification and estimation of term structure models. The first is that the convention that has developed in the literature of ignoring long-term yield data substantially compromises the precision of estimates of both risk premia and expected short rates because it fails to exploit the near-observability of the term structure of risk premia in long-term forward rates. The second is that, because convexity at long maturities is not only large but significantly time-varying, models must

[1] The only previous study focusing on the impact of volatility on long-term forward rates is Brown and Schaefer (2000), although this topic is also touched on in other studies, such as Sack (2000) and Gurkaynak et al. (2007). The interrelation between volatility, convexity and long-term yields has been examined by Brown and Schaefer (1994) and, more recently, by Giglio and Kelly (2018), Rebonato and Putyatin (2018) and Carr and Wu (2023).

[2] Stochastic volatility term structure models for the estimation of time-varying risk premia include, for example, Cieslak and Povala (2016), Choi et al. (2017), Feldhütter et al. (2018), Berardi and Plazzi (2022) and Leippold and Matthys (2022).

[3] Indeed, we are unaware of any other examples of "observable risk premia".

[4] Appendix A lists one hundred recent empirical studies of the term structure. Of these, only twelve employ data on yields longer than (or equal to) 20 years and of these twelve, only six estimate a structural model, as distinct from curve fitting (Andreasen et al. (2019), Balter et al. (2021) and Pancost (2021) are examples).

accommodate time-variation in yield volatility. Misspecifying volatility in constant volatility affine models – "the basic workhorse in macroeconomics and finance" (Hamilton and Wu (2012)) – is consequential because it prevents effective use of the risk premium information in long-maturity rates, even when these rates are used in estimation.

Stochastic volatility models are often viewed as less successful than constant volatility specifications because they are thought to have difficulty matching the Campbell-Shiller evidence against the expectations hypothesis (see Dai and Singleton (2002) and Joslin and Le (2021)). This view, which appears to be widely accepted,[5] represents a potentially significant challenge to our approach, in which stochastic volatility plays a key role.

In a related paper, Berardi and Schaefer (2026) show that much of the supposed advantage of constant volatility models reflects biases in the estimation of the mean-reversion matrix under the natural measure. Consistent with Bauer et al. (2014), Berardi and Schaefer (2026) find that OLS estimation tends to distort the relative volatility of expectations and risk premia. The apparent success of constant volatility models in matching Campbell-Shiller coefficients should therefore not be interpreted as evidence that they provide more reliable estimates of expectations and risk premia. As we show below, exploiting the information contained in long-maturity forward rates leads to substantially more precise estimates of both quantities and to a markedly different decomposition of long-term yields.

2 From Fama-Bliss to Observable Risk Premia

As discussed extensively in the literature, the main difficulty in estimating term structure models is separating the risk-neutral parameters into those that determine risk premia and those that determine expectations under the physical probability measure $\mathbb{P}$. In particular, the small-sample problems that arise in estimating both the long-term mean and the mean-reversion matrix of the underlying stochastic factors are well documented (see, for example, Bauer et al. (2012), Bauer and Hamilton (2018) and many others).

Perhaps surprisingly, we find that including long-maturity yields leads to more precise estimates of risk premia at *all* maturities and in this section we provide some intuition on why this is the case.

The model used in both the theoretical and empirical analysis is a standard stochastic volatility ATSM belonging to the $A_1(4)$ class in the system of Dai and Singleton (2000). Both the model and the notation we employ are standard and so we relegate details on both to Appendix B.

[5] See, for example, Feldhütter et al. (2018), Gagliardini et al. (2008) and Carpenter et al. (2026).

2.1 Near-Observability of Long-Maturity Δh Improves Estimates of Risk Premia at All Maturities

The "near-observability" of the slope of the TSRP is most easily understood using the well-known decomposition of the forward rate in Fama and Bliss (1987) (FB) and Dai and Singleton (2002) (DS):

$$f_t(\tau) = r_t + h_t(\tau) + c_t(\tau) + e_t(\tau). \tag{1}$$

Here $f_t(\tau)$ is the time-t instantaneous forward rate for maturity τ, r_t is the instantaneous short rate, $h_t(\tau)$ is the risk premium in simple returns, $c_t(\tau)$ is a "convexity term" that depends only on τ and the variance of the τ-period yield, and $e_t(\tau)$ is τ times the expected rate of change in the τ-period yield:

$$h_t(\tau) \equiv \frac{1}{dt}\mathrm{E}_t\left[\frac{dP_t(\tau)}{P_t(\tau)}\right] - r_t, \tag{2}$$

$$c_t(\tau) \equiv -\frac{\tau^2}{2}\frac{1}{dt}\mathrm{Var}_t\left[dY_t(\tau)\right], \tag{3}$$

$$e_t(\tau) \equiv \tau\frac{1}{dt}\mathrm{E}_t\left[dY_t(\tau)\right], \tag{4}$$

where $P_t(\tau)$ and $Y_t(\tau)$ are, respectively, the time-t price and yield of a zero-coupon bond with time-to-maturity τ.

As both FB and DS point out, equation (1) is obtained by taking expectations of a relation that holds by definition. Equation (1) therefore holds for any arbitrage-free term structure model.

In both FB and DS, the risk premium is defined in terms of log-returns and, as a result, the risk premium, $h_t(\tau)$, that depends on investor preferences, is conflated with the convexity term, $c_t(\tau)$, that arises simply as a result of non-zero yield volatility and the non-linear relation between price and yield. Since the economic nature of these two quantities is quite different and because a change in convexity says nothing about the risk premium that investors require to hold bonds, we keep the risk premium and convexity terms separate and use the FB/DS decomposition in the form given in equation (1).[6]

Taking the difference between equation (1) evaluated at maturities τ_2 and τ_1, $\tau_2 > \tau_1$, and using the notation $\Delta\ell \equiv \ell_t(\tau_2) - \ell_t(\tau_1), \ell = f, c, h, e,$ – where for brevity we suppress

[6]The analysis in both FB and DS is conducted in discrete time. The corresponding continuous-time analogue is $Lh_t(\tau) = f_t(\tau) - r_t - e_t(\tau)$. Here, $Lh_t(\tau) \equiv \frac{1}{dt}\mathrm{E}_t\left[d\log\left(P_t(\tau)\right)\right] - r_t$ denotes the time-t risk premium in log returns on a τ-period bond. In our diffusion setting, this is related to the risk premium in simple returns, $h_t(\tau)$, through $Lh_t(\tau) = h_t(\tau) + c_t(\tau)$.

both the time index and the maturity argument – we have:

$$\underbrace{\Delta f - \Delta c}_{\text{“observable”}} = \underbrace{\Delta h + \Delta e}_{\text{“model dependent”}} \quad . \tag{5}$$

The variables on the LHS of equation (5), Δf and Δc, are “observable” since both can be obtained directly from price data and without the need for a model. Thus, any model that fits both observed forward rates and yield volatility must give the same value for $\Delta h + \Delta e$, and in this sense, $\Delta h + \Delta e$ is also observable. The task for a term structure model is then to separate these “model dependent” terms.

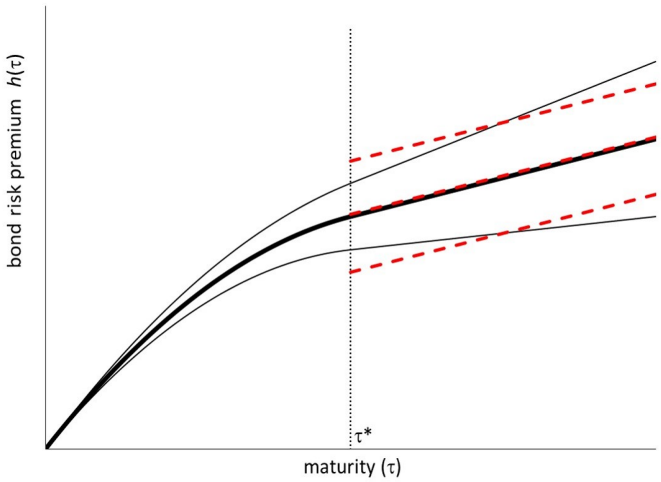
We find empirically that Δe is small at long maturities and, since $\Delta f - \Delta c$ is observable, Δh – the slope of the TSRP – becomes “nearly observable”:

$$\underbrace{\Delta f - \Delta c}_{\text{“observable”}} = \Delta h + \Delta e \approx \underbrace{\Delta h}_{\text{“nearly observable”}} \quad . \tag{6}$$

The near-observability of the slope of the TSRP at long maturities leads to more precise estimates of risk premia at all, including shorter, maturities. To understand *why*, it is useful to recall that, although long-maturity rates provide information about the slope of the TSRP at long maturities, risk premia must converge to zero as maturity approaches zero, thereby constraining the entire TSRP. **Figure 1** illustrates this intuition.

Figure 1
Term Structure of Risk Premia

This figure shows the term structure of risk premia for three possible sets of parameters. The red dashed lines represent the “almost observable” slope of h obtained from long-term rates, $\Delta h \approx \Delta f - \Delta c$.



Suppose that a model is fitted to "short-maturity" rates, i.e., up to the maturity τ^*. The three black lines represent the term structure of model values of $h(\tau)$ for three candidate sets of parameters; the corresponding estimates of $h(\tau)$ are different as a result of estimation error in the parameters. The effect of now including data on long-maturity rates is to "pin down" Δh, the slope of the term structure of long-term risk premia. Assuming that equation (6) holds (i.e., that $\Delta e \to 0$ for long maturities), each of the three red dashed lines has a slope equal to this "observable" value of $(\Delta f - \Delta c)$, but of the three candidate term structures of $h(\tau)$, only the middle one implies a value of Δh for maturities longer than τ^* that is consistent with the observed value.

Thus, including long rates in the estimation constrains the estimated model parameters to be consistent with the "observable" slope of the TSRP and this improves the precision of risk premium estimates, not only at "long" maturities (those longer than τ^* in **Figure 1**), but also at short maturities (those shorter than τ^*). Indeed, in our empirical results, even though risk premia converge to zero at zero maturity, we find that the *relative* improvement in the precision of estimated risk premia is at least as large at short maturities as it is at long maturities.

This result depends crucially on whether Δe is indeed small at long maturities and whether, as a result, Δh is "observable".

2.2 Why Δe Is Small at Long Maturities

Why, and under what conditions, should we expect Δe to be small at long maturities? From equation (4),

$$\Delta e \equiv e_t(\tau_2) - e_t(\tau_1) = \tau_2 \frac{1}{dt} \mathrm{E}_t[dY_t(\tau_2)] - \tau_1 \frac{1}{dt} \mathrm{E}_t[dY_t(\tau_1)]. \tag{7}$$

Defining $\overline{f}_t(\tau_1, \tau_2)$ as the average forward rate at time t between maturities τ_1 and τ_2, i.e.,

$$\overline{f}_t(\tau_1, \tau_2) \equiv \frac{1}{\Delta\tau}[\tau_2 Y_t(\tau_2) - \tau_1 Y_t(\tau_1)], \tag{8}$$

where $\Delta\tau \equiv \tau_2 - \tau_1$, we have:

$$\Delta e = \Delta\tau \frac{1}{dt} \mathrm{E}_t[d\overline{f}_t(\tau_1, \tau_2)]. \tag{9}$$

In other words, Δe is $\Delta\tau$ times the expected change in the average forward rate between maturities τ_1 and τ_2 and, at sufficiently long maturities, we should expect this to be small. First, it is intuitively reasonable that there exist values of τ_1 and τ_2 – however distant – for which the arrival rate of information about the expected change in the $\overline{f}_t(\tau_1, \tau_2)$ is minimal.

Second, under the conditions described in Martin and Ross (2019), as maturity tends to infinity, both forward rates and yields converge to a constant, implying that the expected change in $\overline{f}_t(\tau_1, \tau_2)$ also converges to zero.[7]

We next examine the properties of Δe across the three main classes of ATSMs: those with (i) constant volatility and constant risk premia, (ii) constant volatility and time-varying risk premia, and (iii) stochastic volatility and time-varying risk premia.

2.2.1 Δe with Constant Volatility and Constant Risk Premia

In all affine models, Δe can be written as:

$$\Delta e = \Delta B'(\tau) K \left(\Theta - X_t\right), \tag{10}$$

where X_t is the time-t vector of state variables, K is the mean-reversion matrix under the physical measure $\mathbb{P}$, and Θ is the long-run mean vector in the process for X, also under $\mathbb{P}$. The vector $B(\tau)$ measures the sensitivity of bond prices to X and depends on the mean-reversion matrix under the risk-neutral measure $\mathbb{Q}$, given by $\widetilde{K} = K + \Sigma\Lambda_1$, where Σ is the volatility matrix in the process for X and Λ_1 is the matrix relating the market price of risk to X.[8]

For constant volatility ATSMs with constant risk premia, $\Lambda_1 = 0$ and therefore $\widetilde{K} = K$. In this case, the time-t expected short rate at horizon $t + \tau$ is given by:

$$\mathrm{E}_t[r_{t+\tau}] = r_t + e_t(\tau), \tag{11}$$

so that Δe coincides with the *slope* of the term structure of expected short rates:

$$\Delta e = \Delta \mathrm{E}[r] \equiv \mathrm{E}_t[r_{t+\tau_2}] - \mathrm{E}_t[r_{t+\tau_1}]. \tag{12}$$

Hence, if the term structure of expected short rates becomes flat at long horizons, Δe is small.

2.2.2 Δe with Constant Volatility and Time-Varying Risk Premia

For constant volatility ATSMs with state-dependent risk premia, because Λ_1 is non-zero, $\widetilde{K}$ is no longer equal to K and, as a result, Δe is no longer equal to the slope of the term

[7] See also Dybvig et al. (1996).

[8] See Appendix B for the process of X under the risk-neutral and physical measures (equations (B.1) and (B.3)), the bond price equation (B.5), and the market price of risk specification (equation (B.22)).

structure of expected short rates. However, it is straightforward to show (see Appendix C) that if all eigenvalues of $\widetilde{K}$ are positive (a necessary condition for the state variables under the $\mathbb{Q}$ measure to be stationary) and $\widetilde{K}^{-1}K$ is finite, then Δe converges to zero at long maturities.[9]

With time-varying risk premia, a flat term structure of long-horizon expected short rates is no longer sufficient for Δe to remain small at long maturities because there is a difference between Δe and $\Delta \mathrm{E}[r]$ equal to Δd ($\Delta d \equiv d_t(\tau_2) - d_t(\tau_1)$). We refer to the variable $d_t(\tau)$ as the "duration adjustment" term because, as we show in Appendix B.2, it is related to the impact of time-varying risk premia on the duration term $B(\tau)$.[10] If the conditions on the eigenvalues of $\widetilde{K}$ and $\widetilde{K}^{-1}K$ given above are met, Δe nonetheless becomes small at long maturities, but its two components, $\Delta E[r]$ and Δd, may converge at different rates.

2.2.3 Δe with Stochastic Volatility and Time-Varying Risk Premia

For most ATSMs with stochastic volatility there is no closed form solution for $B(\tau)$.[11] In this case, the conditions under which Δe becomes small for large τ cannot be derived analytically and the behavior of Δe at long maturities remains an empirical issue. This includes the stochastic volatility model used in our empirical analysis.

2.2.4 The Relative Variability of Δh and Δe

In each of these three cases, and even if Δe becomes small at long maturities, "near-observability" of Δh is only useful if its variability is sufficiently higher than the variability of Δe. This is an empirical issue: even if risk premia are time-varying, the variability of Δh may be too small for $(\Delta f - \Delta c)$ to be a good approximation for Δh.

In our empirical analysis, however, this is not the case. In our estimates, the ratio of the variance of Δh to that of Δe is 0.12 between 1 and 10 years, but rises to 14.3 between 15 and 25 years – more than 100 times larger. Thus the variation in $(\Delta f - \Delta c)$ at maturities up to 10 years is driven primarily by changes in expectations, while between 15 and 25 years, mainly by changes in risk premia.

[9] Appendix C also considers the case in which one eigenvalue is positive but close to zero.

[10] See equations (B.14–B.16) on the derivation and interpretation of the duration adjustment term $d_t(\tau)$.

[11] Exceptions include Cox et al. (1985) and Vasicek and Fong (1982).

2.3 The Importance of Time-Varying Volatility

Constant volatility (Gaussian) ATSMs have become widely used in macro-finance. Do the benefits of including long-maturity yields in estimation also apply to constant volatility models? Specifically, can the information on risk premia contained in long-term forward rates be successfully extracted using a model with constant volatility?

From the FB/DS relation, $\Delta f - \Delta c = \Delta h + \Delta e$, and so "pinning down" Δh at long maturities requires not only that Δe is small, but that the model accurately estimates the convexity-adjusted slope of the forward curve, $(\Delta f - \Delta c)$. In a world with constant yield volatility, this simply requires that the model provides a good fit to both the slope of the forward curve and the (time-invariant) term structure of yield volatility. In a constant volatility model, Δc for a given maturity interval is a constant, while, especially at long maturities, time-series variation in *actual* yield volatility leads to substantial variation in Δc. Thus, in a constant volatility model, the error in estimating Δc feeds directly into the estimate of Δh as $(\Delta f - \Delta c)$.

To illustrate the materiality of time-varying yield volatility for the approximation $\Delta h \approx \Delta f - \Delta c$, we compute "model-free" estimates of Δf and Δc.

Monthly estimates of Δf are obtained from STRIP prices.[12] As noted earlier (equation (3)), $c_t(\tau)$ depends only on yield variance. To estimate Δc, we consider two proxies for $\text{Var}_t[dY_t(\tau)]$: a regression-based estimate obtained by regressing realized variance in month $t+1$ on realized variance in month t and the squared MOVE index at the end of month t,[13] and a simpler estimate based directly on the squared MOVE index, scaled so that the implied average yield variance matches the time-series variance of the 20-year zero-coupon rate.

Figure 2 shows Δf together with these two estimates of Δc for the 15- to 25-year maturity interval. The estimates are very similar, and variation in Δc explains a substantial part of the variation in Δf. Both estimates of Δc exhibit significant variation over time. The regression estimate has a mean of -133 bps and a standard deviation of 49 bps. The estimate using only the MOVE index has a mean of -148 bps and a standard deviation of 97 bps. To put this in context, the slope of the forward curve, Δf, has a mean of -76 bps and a standard deviation of 72 bps. Thus, the variability of Δc at long maturities is of the same order of magnitude as Δf. It follows that a model with constant yield volatility (and therefore constant Δc) that fits the forward curve well cannot also provide a good fit to $(\Delta f - \Delta c)$ and therefore cannot fit Δh.

[12] For a description of the data and the reasons we use STRIPS, see Appendix D.1.

[13] This is the proxy we use in estimation. Its construction is described in Appendix D.1.

Of course, it is logically possible that, while failing to fit convexity, a constant volatility model might nonetheless produce good estimates of risk premia but in this case, it cannot also fit the forward curve. Overall, it seems that constant volatility models are flexible enough to provide a good fit to yields and, in some cases, a very good fit. For example, Adrian et al. (2013) use a five-factor ATSM with constant volatility and report yield pricing errors across maturities from one to ten years that do not exceed 0.4 bps on average and have a standard deviation that is below one basis point (Adrian et al. (2013), p. 116). If, because of their flexibility, ATSMs with constant volatility provide a good fit to long-term forward rates, then, as just explained, at long maturities it follows that either risk premia (Δh) or expectations (Δe), or both are mis-estimated.

3 Estimation

3.1 Data and Estimation Method

The model is estimated using end-of-month observations of zero-coupon yields and forward rates obtained from STRIPS prices with maturities ranging from 2 to 25 years. We also fit to yield variance for long-term maturities of 10, 15, 20, and 25 years. The estimation period runs from January 1992 through December 2023. Appendix D.1 describes the data.

Model parameters are estimated by maximum likelihood using an approximate Kalman filter to infer the unobserved state variables. Details of the estimation method are provided in Appendix D.2.

Two features of our estimation procedure differ from standard empirical practice. First, we use STRIPS data rather than zero-coupon yield estimates derived from coupon-bearing bonds, in order to avoid dependence on interpolation schemes such as splines.[14] Appendix F.1 nevertheless shows that our main results are robust to using the yield curve estimates of Gurkaynak et al. (2007) or Liu and Wu (2021).

Second, unlike most of the literature, we estimate the model using rates with maturities up to 25 years (rather than 10 years) and fit to forward rates as well as yields. We view both choices as reflecting the weighting assigned across maturities. Since zero-coupon bond prices, yields, and forward rates contain the same information, the choice among them is fundamentally a question of weighting, governed by assumptions on fitting errors.

The literature offers little justification for assigning zero weight to maturities beyond 10

[14] Cochrane and Piazzesi (2008) express the same concern regarding the yield curve estimates of Gurkaynak et al. (2007), which rely on interpolation, versus Fama-Bliss data (available from CRSP), which do not. The Fama-Bliss data have a maximum maturity of 5 years, which does not meet our needs.

years or for preferring yields over forward rates. Our results suggest these choices matter. In particular, including long-term rates substantially improves the precision of estimates of both risk premia and expected short rates. Moreover, from the FB/DS decomposition, it is the τ-maturity forward rate at time t, rather than the corresponding zero-coupon yield, that directly contains information about the time-t risk premium on a τ-maturity zero-coupon bond.

While we do not claim to provide an optimal weighting scheme, our Monte Carlo experiments examine how alternative choices regarding the inclusion of long-term rates, yields, forward rates, and yield volatility affect estimation. The results suggest that these choices have a significant impact on the estimates.[15]

3.2 In-Sample Evidence on Long-Maturity Risk Premia

Appendix D.3 reports the full estimation output, including parameter estimates and goodness-of-fit measures. The fit to yields is comparable to ATSMs with a similar number of state variables (e.g., Cieslak and Povala (2016)), with better performance at long than short maturities, and annualised yield volatility is matched with an average error of around 3 bps.

We now focus on forward rate decompositions and the associated long-term risk premia estimates. Panel A of **Figure 3** breaks down Δf into Δe, Δc, and Δh for the maturity interval 15 to 25 years.

The estimates of Δe are extremely small: across the whole 32-year sample period, the mean value is -0.3 bps and the standard deviation is 13.5 bps.[16]

Since Δe is very small and $\Delta f = \Delta e + \Delta c + \Delta h$, the implied value of Δh in Panel A is well approximated by $(\Delta f - \Delta c)$. Although Δf is negative for most of the sample, after adjusting for convexity Δh is mostly positive, implying that the term structure of estimated risk premia at long maturities is almost always upward sloping.

As the model provides a good fit to differences in both forward rates and convexity, model estimates of differences in average risk premia closely match the corresponding "model-free" values obtained from the approximation $\Delta h \approx \Delta f - \Delta c$, where Δf is calculated from actual forward rates and Δc from actual yield variance, i.e., by computing the τ-maturity convexity as $c_t(\tau) = -\frac{1}{2}\tau^2 \frac{1}{dt}\text{Var}_t[dY_t(\tau)]$.

Panel B of **Figure 3** compares model and "model-free" estimates of the average term structures of long-term forward rate spreads, differences in risk premia, and differences in

[15] For details, see Appendix F.2.

[16] If we calculate a $\pm$ two-standard-deviation confidence interval through Monte Carlo simulations, we observe that the estimated values are never statistically different from zero.

convexity, holding the maturity difference constant at 10 years. Model estimates of Δf, Δc and Δh match almost perfectly their "model-free" counterparts.[17]

Average forward rate spreads and convexity differences are both negative and downward sloping. Since convexity differences are larger than forward rate spreads, the term structure of average risk premia is increasing and concave.

Panel C of **Figure 3** shows estimates of Δe for the maturity intervals 1 to 10 years, 5 to 15 years, and 15 to 25 years. The variability of Δe declines sharply with maturity. At short maturities, $(\Delta f - \Delta c)$ is strongly affected by changes in expectations, whereas at long maturities Δe becomes small and $(\Delta f - \Delta c)$ provides a good proxy for Δh.

These results highlight the usefulness of long-maturity yields in separating expectations from risk premia. In contrast, most of the literature employs yield data with maturities of 10 years or less, where Δe – and therefore the slope of the forward curve – remains strongly influenced by expectations. The fact that the relative variability of expectations and risk premia varies across the term structure has been noted previously (e.g., Cieslak and Povala (2016), Figure 3). What has not been recognised are the implications of this fact for the estimation of term structure models.

3.3 Decomposition of Risk Premia

We now turn to the in-sample estimates of the *level* of bond risk premia. The average term structure of estimated bond risk premia is upward sloping, from 80 bps at 5 years to 146 bps at 10 years and 230 bps at 20 years. The curve is concave, with the slope declining from 10 bps per year between 10 and 15 years to 4 bps per year between 20 and 25 years.

To understand the determinants of risk premia, we decompose them into four components corresponding to the contributions of the variance factor v and the three latent factors x_1, x_2, and x_3.[18] The sensitivities of risk premia to the state variables vary with maturity: for v, it is positive and almost linear in maturity, while for x_3 it is negative and convex, reaching a minimum at around nine years. Sensitivities to x_1 and x_2 are much smaller, particularly for maturities up to 10 years.[19]

Figure 4 shows a decomposition of the estimated risk premia for $\tau = 2, 5, 10, 20$ years into their four components. For a given maturity, the red line shows the component attributable

[17] If we consider the time series of differences in risk premia Δh for the maturity pairs $20-10, \ldots, 25-15$ years, the correlation between monthly changes in the estimated values and monthly changes in the corresponding "model-free" values ranges from 0.77 to 0.81.

[18] See equation (B.24) in Appendix B.3 for details.

[19] For brevity, we do not report this figure. The sensitivities are calculated using normalized changes in the state variables, making their relative magnitudes directly comparable.

to the corresponding state variable, while the black line shows total risk premia. The first observation is that the state variables x_1 and x_2 – correlated with the level and slope of the yield curve, respectively – account for only a very small proportion of the variability of risk premia at any maturity.

Because the impact of v on risk premia is roughly proportional to maturity, its contribution is negligible at short maturities but substantial at long maturities. Risk premia at short maturities are largely determined by x_3, while at longer maturities both v and x_3 are important. Since v measures the level of volatility, higher volatility raises risk premia across the term structure roughly in proportion to maturity. By contrast, x_3, which is correlated with the slope of the volatility term structure, lowers risk premia when long-term volatility increases relative to short-term volatility.

4 Monte Carlo Simulations

Section 2 proposed two conjectures. The first was that, at long maturities, the convexity-adjusted slope of the forward curve would be little affected by expectations and, by virtue of the FB/DS decomposition, would therefore provide a good approximation to the slope of the TSRP. The second was that, if the first conjecture were true, including information on long-term rates could improve the precision of estimates of both risk premia and expected short rates at all maturities. **Figure 3** showed that the first conjecture is extremely well supported by the in-sample estimates, and we now turn to the second question: whether including long-maturity rates leads to improved precision in estimates of risk premia and expected short rates.

4.1 The Information in Long-Term Yields

To test this conjecture, we run a Monte Carlo experiment. We use the in-sample parameter estimates from our SV model to generate 500 paths of the state variables, each of length 384 months (the same length as our sample). For each trial, and again using the in-sample parameters, we compute the term structure of yields, forward rates, and yield volatilities. Using the simulated yields etc. as "data", we then estimate the model twice: once using maturities up to 25 years and a second time using maturities only up to 10 years.

For each trial, we construct two series of 384 "estimation errors" in (i) the risk premium and (ii) the expected short rate. The estimation error is defined as the difference between the value obtained using the parameters estimated from the simulated data (the "estimated

value") and the value obtained using the in-sample parameters (the "true value").

It might be expected that, since the approximation $\Delta h \approx \Delta f - \Delta c$ works well only at long maturities, any increase in the accuracy of estimates of risk premia and expected short rates would be similarly limited to long maturities. However, as also discussed in Section 2, if the near-observability of the TSRP at long maturities leads to improved estimates of the market prices of risk, this improves the precision of estimates of risk premia and expected short rates at both short and long maturities. In fact, we find that the relative improvement is actually marginally *greater* at short maturities.

Table 1 provides summary statistics on estimation errors in risk premia and expected short rates from the Monte Carlo trials. In panel A, estimates use simulated data with maturities up to 25 years (model SV_{25}) and, in Panel B, up to 10 years (model SV_{10}). In Panel A, results are reported for maturities between two years and 25 years. In Panel B, since data with maturities longer than 10 years are excluded, we report results only for maturities of 10 years or less. The reported statistic in each case is calculated across the 500 trials.

In Panel A, columns (2−4) report summary statistics in basis points for errors in estimated risk premia: the mean error, the median standard deviation of the estimation error within each trial, and the standard deviation (across trials) of the within-trial mean error, respectively. The results in column (2) show no evidence of bias, with the mean error in the estimated risk premium for each maturity less than 1 bp. Column (3) reports the median standard deviation of the estimation error, ranging from 4.1 bps at a maturity of two years to 20.4 bps at 25 years. For the 10-year maturity, the median standard deviation is 11.6 bps. To put this in context, the in-sample estimates of the 10-year risk premium have a mean value of 146 bps and an interquartile range of [88,206] bps, so the standard deviation of the errors in the 10-year risk premium is around 8% of its average level. Column (4) reports the standard deviation of the mean error, which is of the same order of magnitude as the median standard deviation of the error around the mean.

Columns (5−7) contain the corresponding results for the expected short rate, and as before there is no evidence of bias (column (5)). The median standard deviation of the estimation error in the expected short rate is smaller than for the risk premia, and it is notable that the standard deviation of the mean error is actually larger than the median within-trial standard deviation.

Panel B of **Table 1** now addresses a central question in this study: how the precision of these estimates changes when data on yields longer than 10 years are excluded. The mean error in both risk premia (column (2)) and the expected short rate (column (5)), although somewhat larger than in Panel A, remains relatively small.

The key results are in columns (3−4) for risk premia and (6−7) for expected short rates, where we report the *ratio* of each statistic to the corresponding statistic in Panel A. For example, the value of 3.48 for the 10-year risk premium in column (3) means that the median standard deviation of the estimation error using data up to 10 years is 3.48 times the value reported in Panel A, i.e., 40.4 bps versus 11.6 bps. The results for the median standard deviation are striking: none of these ratios is below three, implying that excluding data with maturities longer than 10 years at least triples the standard deviation of the estimation error. It is particularly striking that the largest ratios are for the shortest maturities. For the two-year maturity, excluding long-maturity data increases the standard deviation of the estimation error in the risk premium by a factor of 3.81. Excluding data for maturities between 10 and 25 years therefore almost quadruples the size of the estimation error for the two-year risk premium.

Column (6) reports the corresponding results for the expected short rate. While Panel A shows that the median standard deviation is larger for risk premia than for expected short rates, Panel B shows that the improvement in precision for estimates of both risk premia and expected short rates − as measured by these ratios − is very similar.

Columns (4) and (7) report the corresponding results for the standard deviation of the mean errors. In both cases, although the ratios are above one − implying that including long maturity data leads to higher precision − they are substantially smaller than those for the median standard deviation. In other words, the inclusion of long-maturity rates in estimation has a much larger impact on the standard deviation of the time series of errors than on the standard error of their mean. We return to this finding in Section 4.5 below.

More generally, the precision gains documented here do not arise solely from the inclusion of long-maturity yields. As shown in Appendix F.2, fitting the model to forward rates as well as yields, and incorporating yield volatility in the estimation, are also important for successfully extracting the risk premium information contained in the term structure.

4.2 A Regression Interpretation

Section 2.1 describes the benefit of including long-maturity yields in estimation in terms of constraining the model estimates of risk premia to be consistent with the "almost observable" slope of the TSRP. We now take this argument a further step to pin down more precisely how information in long rates increases the precision of risk premium estimates across the whole of the term structure.

The near-observability of the TSRP at long maturities − shown to hold well in our

empirical estimates in Section 3.2 – implies that, for two long maturities τ_1 and τ_2 and using the functional form for Δh,[20] we have:

$$\Delta f - \Delta c \approx \Delta h = -\Delta B' \, \Sigma(\Lambda_0 + \Lambda_1 X_t), \tag{13}$$

where, as before, $\Delta \ell \equiv \ell_t(\tau_2) - \ell_t(\tau_1)$, $\ell = f, c, h, e$, and $\Delta B \equiv B(\tau_2) - B(\tau_1)$.

Equation (13) can be interpreted as a regression of $\Delta f - \Delta c$ on the state vector X_t. For long maturities, Δe, the "residual" in the regression, is small, so the coefficient on X_t, $-\Delta B' \, \Sigma\Lambda_1$, can be estimated relatively precisely. For short maturities, however, as shown in Panel C of **Figure 3**, the variability of Δe is much larger, making the estimated loading on X_t a less reliable measure of $-\Delta B'\Sigma\Lambda_1$.

4.3 Error Decomposition

The previous section suggested that the higher precision in estimates of $\hat{h}_t(\tau)$ when long rates are included in estimation results from a more precise estimate of $B(\tau)\Sigma\Lambda_1$, the sensitivity of the risk premium to the state vector. We now decompose the errors in $\hat{h}_t(\tau)$ to determine which of its determinants is mainly responsible for the improved precision.

In our simulations we find that both X_t and Σ are recovered with minimal error, both when yields up to 10 years are used in estimation and when yields up to 25 years are used. In our decomposition we therefore ignore the estimation error in X_t and focus on the components of the error in $\hat{h}_t$ that result from estimation errors in $B(\tau)$ and $\Sigma\Lambda_1$.[21] We denote these by:

$$\varepsilon_B \equiv \hat{B}'(\tau) - B'(\tau) \qquad \text{and} \qquad \varepsilon_{\Sigma\Lambda_1} \equiv \hat{\Sigma}\hat{\Lambda}_1 - \Sigma\Lambda_1.$$

Panel A of **Table 2** gives the results for the SV_{25} model and Panel B for the SV_{10} model. Column (2) shows the mean error in the estimated time invariant component of $\hat{h}_t$ and the remaining columns the results for the time-varying component. Columns (3) and (4) give the mean and median standard deviation of errors related to ε_B, columns (5) and (6) the corresponding values for $\varepsilon_{\Sigma\Lambda_1}$ and columns (7) and (8) for the product of the errors in $B(\tau)$ and $\Sigma\Lambda_1$. Columns (9) and (10) contain the results for the total error.

For both SV_{25} and SV_{10}, the errors in $\hat{h}_t$ that result from both ε_B and the cross-product $\varepsilon_B\varepsilon_{\Sigma\Lambda_1}$ are very small (a standard deviation of two bps or less). However, the errors that derive from $\varepsilon_{\Sigma\Lambda_1}$ are both much larger and quite different for the two models, with the

[20] See equation (B.23).

[21] Details of the decomposition are given in Appendix E.

standard deviation of the errors for SV_{10} over three times that for SV_{25}.[22] For example, for $\tau = 10$, the standard deviation of the $\varepsilon_{\Sigma\Lambda_1}$ is 11.74 bps for SV_{25} but more than three times higher at 40.03 bps for SV_{10}. Thus the lower variability of the errors associated with $\varepsilon_{\Sigma\Lambda_1}$ explains almost all the reduction in the estimation error of $\hat{h}_t$. Since the estimation error in Σ is negligible, the improvement in the precision of $\hat{h}_t$ is almost entirely attributable to a more precise estimate of Λ_1.

In summary, the reason that the inclusion of long-maturity data improves estimates of risk premia is that it leads to a more precise estimate of the market price of risk matrix Λ_1.

4.4 The Relation between the Precision of Estimates of Risk Premia and Expected Short Rates

As already mentioned, we find that estimates of $B(\tau)$, Σ, and X_t are very precise, whether or not long rates are used in estimation. The fact that the estimation error in $B(\tau)$ is small means that $\widetilde{K} = K + \Sigma\Lambda_1$ is also well determined. If so, an improvement in the estimate of Λ_1 automatically leads to a more precise estimate of K. Since K determines the variation in the expected short rate, an improvement in the precision of risk premia should be automatically accompanied by an improvement in the precision of the expected short rate, and this is precisely what we find in the Monte Carlo experiment reported in **Table 1**.

The relation between the precision of estimates of risk premia and expected short rates is itself interesting. The small-sample problems that plague estimation of the mean-reversion matrix K in term structure models are well known (e.g., Bauer et al. (2012)). We find that including long rates in estimation improves not only estimates of risk premia but also the precision of estimates of K. The reason that long rates improve the precision of estimates of risk premia is that the slope of the term structure of forward rates is little affected by expectations. It is therefore somewhat ironic − but true − that including rates that are little affected by expectations leads to more precise estimates of expected short rates.

4.5 Why Long Rates Improve Mean Estimates Less

Although including long rates substantially improves the precision of estimates of time-varying risk premia and expected short rates, **Table 1** shows that the effect on the precision of their sample means is much smaller. Under two mild conditions – that the model fits yields closely and that the factor loadings and latent state variables are estimated with negligible

[22] For SV_{25}, the mean error attributable to $\varepsilon_{\Sigma\Lambda_1}$ is close to zero. For SV_{10}, the mean error is negative but is largely offset by an error of similar magnitude and opposite sign in the constant term.

error – the mean estimation error in risk premia depends primarily on errors in the estimated residual dynamics of the latent factors and is therefore approximately invariant to whether long rates are included in estimation. As shown analytically in Berardi and Schaefer (2026), this explains why including long rates dramatically reduces the variance of estimation errors while having little effect on bias or on the precision of estimated mean risk premia. The same intuition applies to mean expected short rates, although less precisely because of the presence of the duration adjustment term ($d_t(\tau)$) in the difference between expected short rates and expected yield changes.

5 New Estimates of the Term Premium

Using Monte Carlo simulations, the previous section showed that, in the presence of stochastic volatility, incorporating long-maturity yields substantially improves the identification of both risk premia and expected short rates. We now turn to the data.

The main empirical finding of this section is that the improved identification delivered by the stochastic volatility (SV) model leads to a markedly different decomposition of long-term yields than that implied by the benchmark models used within the Federal Reserve (Fed). Whereas standard estimates attribute most movements in long-term yields to variation in term premia, the SV model assigns a much larger role to changes in expected short rates. The differences are economically large and, as we show below, imply substantially different interpretations of recent episodes of Treasury market behaviour.

We first compare estimates from the SV model with those from the widely used Fed term premium models. We then examine the period of quantitative tightening between March 2022 and October 2023 as a case study illustrating the practical implications of the different decompositions.

5.1 Comparison with Fed Models

How much do estimates of risk premia and term premia change when stochastic volatility and long-maturity yields are incorporated into estimation?

To answer this question, we compare estimates from our SV model, estimated using yields up to 25 years, with estimates from the three benchmark constant volatility (CV) Fed models: Kim and Wright (2005) (KW), Adrian et al. (2013) (ACM), and Christensen and Rudebusch (2012) (CR).[23]

[23]The estimates of term premia using the KW, ACM, and CR models are produced by staff in the Federal Reserve System

Although term premium estimates from the SV model are positively correlated with those from the Fed models, there are major differences in the decomposition of long-term yields and forward rates.

Since the term premium includes convexity,[24] it might be imagined that estimates of term premia − as distinct from risk premia − would depend less on the inclusion of stochastic volatility and that, as a result, estimates of the term premium obtained with the SV model would be similar to those from a CV model. We show that this is not the case and term premia in the SV model have much lower variability than in ACM and CR, although similar variability to KW. Further, while there is at least some similarity between *term premia* in the SV and Fed models, this is not true, as we show, for *risk premia*.

Panel A of **Table 3** reports summary statistics for estimates of the 10-year term premium ($TP \equiv TP_t(10)$) and average expected short rate ($ER \equiv ER_t(10)$). The values in column (2) show that the mean of TP in the SV model is 55 bps, close to CR (64 bps), but below KW (76 bps), and well below ACM (103 bps). The annualised standard deviation of monthly changes (column (3)) is 44 bps in SV, similar to KW (47 bps), but about half the variability in ACM (82 bps) and one-third of that in CR (124 bps).

Panel A of **Figure 5** plots the time series of TP estimates and, here, we note two points in particular. First, although there are significant differences between the TP estimates in the three Fed models over the 32 years in the sample, they all display a marked downward trend that mirrors the downward trend in rates over the period. In the SV model, TP estimates have no downward trend for the first 20 years, while after the GFC, TP estimates are lower in all four models. Second, despite differences in level and, especially, volatility, the figure suggests that the different estimates of TP are nonetheless quite highly correlated, and this is indeed the case. In levels, the SV estimates have a correlation of 0.72 with KW, 0.88 with ACM and 0.81 with CR.[25] Interestingly, the TP estimates in each of the models, including the SV model, share a similar level of correlation with the MOVE index (between 0.5 and 0.6). Correlation with the MOVE index is noted by ACM in connection with their own TP estimates, but it turns out that it is also a feature of the KW, CR and SV estimates. The fact that term premia in the SV model are correlated with the MOVE index is easily understood since the model has time-varying volatility and risk premia that depend strongly on yield variability. For a CV model, however, this feature is more difficult to rationalise.

and are updated on a daily basis at: https://www.federalreserve.gov/pubs/feds/2005/200533/200533abs.html for KW; https://libertystreeteconomics.newyorkfed.org/2014/05/treasury-term-premia-1961-present/ for ACM; and https://www.frbsf.org/research-and-insights/data-and-indicators/treasury-yield-premiums/ for CR.

[24] See equations (B.17−B.21).

[25] The monthly changes in TP are less correlated but not dramatically so. The correlation of monthly changes in SV estimates is 0.63 with KW, 0.75 with ACM, and 0.66 with CR.

Panel B of **Figure 5** shows that the time series of ER follow a broadly similar path in the four models, though levels and, particularly, volatilities differ. As reported in panel A of **Table 3** (column (4)), the average value of ER in SV is 362 bps compared with 340 bps for KW, 311 bps for ACM and 284 for CR. The volatility of ER (column (5)) is 70 bps in SV, lower in KW (48 bps) and ACM (63 bps) and higher in CR (87 bps).

Because all models provide a good fit to observed yields, differences in the estimated variability of term premia imply different interpretations of yield movements. Columns (6) and (7) in panel A show the results of the variance decomposition of the 10-year yield.

In the Fed models, the term premium accounts for between 50% (KW) and 74% (CR) of the variability of changes in the 10-year yield. This is consistent with the common view in the literature that fluctuations in long-term yields are driven primarily by changes in compensation for interest-rate risk (see also Jotikasthira et al. (2015) and Moench (2019)). In the SV model, however, the contribution of the term premium falls to only 32%, implying that expected short rates account for the majority of yield variation.

Thus, incorporating stochastic volatility and long-maturity yields leads to a fundamentally different decomposition of long-term interest rates. Rather than viewing yield fluctuations primarily as changes in risk compensation, the SV model suggests that revisions in expectations about the future path of short-term rates play the dominant role.

The differences become even more pronounced if we shift our attention from yields to forward rates. Panels C and D of **Figure 5** show the time series of the 10-year forward term premium ($FTP \equiv FTP_t(10)$) and expected short rate ($\mathrm{E}[r] \equiv \mathrm{E}_t[r_{t+10}]$), respectively. Panel B of **Table 3** reports summary statistics and variance decompositions for FTP and $\mathrm{E}[r]$.

The expected short rate $\mathrm{E}[r]$ is *less* volatile in the Fed models than in SV, with the annualised volatility of monthly changes 39 bps for KW and ACM, 50 bps for CR, and 68 bps for SV (column (5)). Conversely, FTP is significantly *more* volatile in the Fed models, with volatilities of 63 (KW), 114 (ACM), and 135 bps (CR), compared with 47 bps for SV (column (3)).

The variance decomposition, reported in columns (6) and (7), provides perhaps the clearest illustration of the discrepancy between the SV and Fed models. Under the SV specification, $\mathrm{E}[r]$ accounts for 64% of the variability of monthly changes in the 10-year forward rate, while FTP accounts for only 36%. In contrast, ACM and CR attribute more than 90% of forward-rate variation to changes in the forward term premium and less than 10% to changes in expectations.

Finally, although the FTP estimates in our SV model and the Fed models do not appear to be dramatically different, there are very significant differences in the estimates of the

underlying *risk premium*. Panel E of **Figure 5** compares the 10-year risk premium ($h(10)$) in SV and ACM.[26] The two series are positively correlated (0.72 in levels and 0.35 in monthly changes), but differ substantially. The average level of the risk premium in ACM is about twice that in SV (305 vs. 146 bps) but its volatility is *nine times* higher (850 vs. 91 bps).

Taken together, our results suggest that conventional term premium estimates may substantially overstate the contribution of risk compensation to Treasury yield fluctuations.

Using a Monte Carlo experiment similar to Section 4, Appendix F.3 shows that our SV model provides a significant improvement in the precision of term premium estimates with respect to a standard CV model.

5.2 The 2022-2023 Quantitative Tightening

The so-called "Greenspan conundrum", described by Ben Bernanke and cited in the Introduction, asks whether the "historically unusual behavior of long-term yields" observed in 2004-2005 reflects changes in the term premium or in expectations of future short-term interest rates.

More recently, unusual yield dynamics re-emerged during the post-pandemic normalization of Fed policy. In November 2021, the Fed began tapering its quantitative easing program and, by March 2022, net asset purchases had ended, effectively halting the expansion of the Fed's balance sheet. Quantitative tightening was formally announced in May 2022 and continued for about eighteen months. At the November 2023 FOMC meeting, the Fed kept the policy rate unchanged and communicated a cautious, data-dependent stance, contributing to a decline in long-term yields as expectations of further tightening eased. Between November 2021 and October 2023, the yield curve shifted up and became sharply flatter: the 2-year yield rose from 50 to 500 bps, the 10-year yield from 150 to 500 bps, and the 10-year forward rate from 200 to 550 bps.

This episode again prompts Bernanke's question, namely whether the strong and persistent increase in long-term yields primarily reflects changes in expectations or in the term premium. In this respect, it provides a natural setting to compare the SV and the Fed models. We focus on the period from March 2022 to October 2023, when the Fed balance sheet expansion had effectively ceased.

Figure 6 reports the time series of the components of the 10-year forward rate: the forward term premium (FTP) in Panel A and the expected short rate ($\mathrm{E}[r]$) in Panel B, for

[26] In their paper, ACM provide a detailed explanation of the estimation procedure. Applying that procedure to GSW data, we closely replicate their results and recover the parameter estimates for K, Σ, Λ_0 and Λ_1, which are needed to calculate the risk premium.

the SV model and Fed models (KW, ACM, and CR). All series are normalized to zero in March 2022.

In the SV model, the increase in the forward rate is driven primarily by expectations, with only a limited role for the term premium. Column (2) of **Table 4** shows that the increase of 62 bps in FTP over the period represents only 21% of the total increase in the fitted value of $f_t(10)$. Thus the corresponding increase in E[r] accounts for 79% of the total increase (column (4)). Columns (3) and (5) report the annualised standard deviation of monthly changes and show that for the SV model the volatility of changes in FTP is approximately half that of E[r] (41 vs. 84 bps).

By contrast, in the Fed models FTP explains a much higher fraction of the increase: 77% in ACM, 56% in CR, and 43% in KW. Overall, FTP volatility is much higher in the Fed models over this period – about twice as high as SV for KW and more than three times as high for ACM – while E[r] volatility is roughly half that in SV. This is seen clearly in the variance decomposition in columns (6) and (7) of **Table 4**. In the SV model, FTP accounts for 30% of forward rate variability, compared with 65% (KW), 76% (CR), and 89% (ACM).

The contrast between the two sets of estimates is economically significant. Under the benchmark Fed models, the rise in long-term forward rates during the tightening cycle is interpreted primarily as an increase in term premia. Under the SV model, the same episode is interpreted largely as a reassessment of the future path of short-term interest rates. The policy implications of these interpretations are very different and, to paraphrase Bernanke, depending on how expectations and risk premia are identified, "the implications for policy may be quite different – indeed, quite the opposite".

6 Conclusion

At long maturities, the slope of the term structure of risk premia (TSRP) becomes almost observable, and approximately equal to the slope of the forward curve, net of a convexity adjustment that depends solely on yield volatility. The convexity adjustment, again at long maturities, is both large and significantly time-varying and so an accurate estimate of the slope of the TSRP requires a model with time-varying volatility.

Using a stochastic volatility (SV) model, we find that including long-term rates in estimation dramatically improves the precision of estimates of both risk premia and expected short rates. A Monte Carlo experiment shows that when yield data with maturities of up to 25 years, rather than 10 years are used in estimation, the median standard error of both expected short rates and risk premia falls by approximately 70%. Significantly, the improve-

ment in precision from employing long-maturity data is just as large at short maturities (e.g., two years) as at long maturities (e.g., 10 years).

Estimating our SV model using US Treasury yields with maturities up to 25 years, we find estimates of expected short rates and term premia that differ markedly from those reported for the Fed models of Kim and Wright (2005), Adrian et al. (2013), and Christensen and Rudebusch (2012). Specifically, we find that expected short rates explain a much larger fraction of yield changes, and term premia a much smaller fraction.

Employing long-maturity data in estimation and a model with time-varying volatility represents a significant departure from prevailing empirical practice. Most studies ignore yields with maturities beyond 10 years and the models used in the macro-finance literature and by the monetary authorities generally have constant volatility. And while models that incorporate time-varying volatility have been quite extensively studied in the literature, the critical role that time-varying volatility plays in extracting risk premia from long-term rates has not previously been recognized. Indeed, when a constant volatility (CV) model is fitted to data generated by our SV model, there is little or no benefit from using long-maturity data. This is because CV models cannot account for time-varying convexity, and this leads to inaccurate estimates of the slope of the TSRP.

For policymakers and others for whom accurate estimation of risk premia and expected short rates is a priority, the information on risk premia that is contained in long-term rates can provide significant benefits. Models that incorporate time-varying volatility and are estimated on data that includes long-maturity yields offer the potential for substantial improvements in the accuracy of estimates of risk premia, term premia and expected short rates.

Appendix

A List of Term Structure Papers

Table A.1 reports a list of one hundred empirical term structure papers from which we compute the percentage of empirical studies on the term structure using yields with maturities of at least 20 years.

B The Model

This appendix introduces the standard notation used in the ATSM literature (Duffie and Kan (1996), Dai and Singleton (2000), and Duffee (2002)), derives the decomposition of forward rates and yields into their components, and presents the model used in our empirical analysis.

B.1 Some Standard Notation

The term structure is determined by an n-vector of state variables, X_t, with dynamics, under the risk-neutral measure $\mathbb{Q}$, described by:

$$dX_t = \widetilde{K}\left(\widetilde{\Theta} - X_t\right)dt + \Sigma\sqrt{S_t}d\widetilde{W}_t, \tag{B.1}$$

where $\widetilde{K}$ and Σ are $n \times n$ matrices, $\widetilde{\Theta}$ is an n-vector, and $\widetilde{W}_t$ is a n-vector of independent Brownian motions. S_t is a diagonal $n \times n$ matrix with the i^{th} diagonal element given by:

$$[S_t]_{ii} = \alpha_i + \beta_i' X_t, \tag{B.2}$$

where α_i is a scalar, β_i is a n-vector and the values in equation (B.2) are nonnegative.

The market price of risk, denoted by the n-vector Λ_t, relates $\widetilde{W}_t$ to W_t, the n-vector of independent Brownian motions under the natural measure $\mathbb{P}$, $dW_t = d\widetilde{W}_t - \Lambda_t dt$, and determines the process of X_t under the natural measure:

$$dX_t = K\left(\Theta - X_t\right)dt + \Sigma\sqrt{S_t}dW_t, \tag{B.3}$$

where $K\left(\Theta - X_t\right) = \widetilde{K}\left(\widetilde{\Theta} - X_t\right) + \Sigma\sqrt{S_t}\Lambda_t$. The short rate in all ATSMs can be written as:

$$r_t = \delta_0 + \delta' X_t, \tag{B.4}$$

where δ_0 is a scalar and δ is an $n \times 1$ vector. The time-t price of a zero-coupon bond with face value one and time-to-maturity τ, $P_t(\tau)$, has the exponentially affine closed-form solution:

$$P_t(\tau) = \exp[A(\tau) - B'(\tau)X_t], \tag{B.5}$$

where the coefficients $A(\tau)$ and $B(\tau)$ depend on the underlying model parameters and solve a system of ordinary differential equations (see Dai and Singleton (2000)). The term structure of zero-coupon yields is therefore affine in the state vector and the time-t yield on the τ period bond, $Y_t(\tau)$, is given by:

$$Y_t(\tau) = -\frac{1}{\tau}\left[A(\tau) - B(\tau)X_t\right]. \tag{B.6}$$

Unless all the elements β_i in equation (B.2) are zero, the diffusion term in equation (B.1) is a function of X_t, which implies that yield volatilities are time-varying. In particular, the time-t instantaneous variance of the τ-maturity yield, $V_t(\tau)$, is given by:

$$V_t(\tau) = \frac{1}{\tau^2}B'(\tau)\left(\Sigma\, S_t \Sigma'\right) B(\tau). \tag{B.7}$$

Since the model is time homogeneous, the time-t instantaneous forward rate for date $t+\tau$, $f_t(\tau) = \frac{1}{P_t(\tau)}\frac{\partial P_t(\tau)}{\partial t}$, is given by:

$$f_t(\tau) = r_t - B'(\tau)\Sigma\sqrt{S_t}\Lambda_t - \frac{1}{2}B'(\tau)\left(\Sigma\, S_t \Sigma'\right) B(\tau) + B'(\tau)K(\Theta - X_t). \tag{B.8}$$

B.2 Forward Rate Decomposition

The decomposition of the forward rate in equation (B.8) is exactly the FB/DS decomposition in equation (1), $f_t(\tau) = r_t + h_t(\tau) + c_t(\tau) + e_t(\tau)$, where:

$$h_t(\tau) = -B'(\tau)\Sigma\sqrt{S_t}\Lambda_t, \tag{B.9}$$

$$c_t(\tau) = -\frac{1}{2}B'(\tau)\left(\Sigma S_t \Sigma'\right) B(\tau), \tag{B.10}$$

$$e_t(\tau) = B'(\tau)K(\Theta - X_t). \tag{B.11}$$

An alternative representation expresses the forward rate in terms of the time-t expected short rate at time $t+\tau$ under the $\mathbb{P}$ measure:

$$\mathrm{E}_t\left[r_{t+\tau}\right] = r_t + B_G'(\tau)K(\Theta - X_t), \tag{B.12}$$

where $B_G(\tau)$ (with subscript G denoting the Gaussian benchmark), $B'_G(\tau) \equiv \delta' K^{-1}(\mathrm{I}-e^{-K\tau})$, is the factor loading in a Gaussian affine model with the same mean-reversion matrix K and long-run mean Θ, but a constant market price of risk (for example, as in Vasicek (1977) and Langetieg (1980)). Using equations (B.11) and (B.12), we may rewrite the FB/DS decomposition of the forward rate as:

$$f_t(\tau) = \mathrm{E}_t\left[r_{t+\tau}\right] + h_t(\tau) + c_t(\tau) + d_t(\tau), \tag{B.13}$$

where

$$d_t(\tau) = \left[B'(\tau) - B'_G(\tau)\right] K(\Theta - X_t). \tag{B.14}$$

We define $d_t(\tau)$ as the "duration adjustment" term, since the difference $(B'(\tau) - B'_G(\tau))$ represents the combined impact on $B(\tau)$ (the "duration" coefficient) of stochastic volatility and the essentially affine specification. This term goes to zero if $S_t = \mathrm{I}$ *and* Λ_t is a constant in equation (B.3).

From equations (B.11) and (B.14) and the definition of $e_t(\tau)$ in equation (4), we may also write $d_t(\tau)$ as:

$$d_t(\tau) \equiv \tau \frac{1}{dt}\left(\mathrm{E}_t\left[dY_t(\tau)\right] - \mathrm{E}_t\left[dY_t^G(\tau)\right]\right) = e_t(\tau) - e_t^G(\tau), \tag{B.15}$$

where $e_t^G(\tau)$ is the time-t expected change in the τ-period yield $Y_t^G(\tau)$ in a Gaussian model that, as described above, shares the same values of K and Θ. Thus $d_t(\tau)$ measures the impact of stochastic volatility and time-varying risk premia on the expected change in the τ-period yield. Dai and Singleton (2002) (in their equation [10]) show that $d_t(\tau)$ can also be expressed as the sum of the expected one-period changes in future log excess returns over the τ-maturity horizon:

$$d_t(\tau) \equiv \sum_{i=1}^{\tau-1} \mathrm{E}_t\left[h^*_{t+i}(\tau - i) - h^*_{t+i-1}(\tau - i)\right], \tag{B.16}$$

where, using our notation, $h^*_t(\tau) = h_t(\tau) + c_t(\tau)$. In the case of a model with constant volatility and constant risk premia, all these expected changes are zero.

The difference between the forward rate at maturity τ and the corresponding expected short rate gives the τ-maturity "forward term premium", $FTP_t(\tau)$ (see, e.g., Dai and Singleton (2002) and Kim and Wright (2005)):

$$FTP_t(\tau) = f_t(\tau) - \mathrm{E}_t\left[r_{t+\tau}\right]. \tag{B.17}$$

The term-average of $FTP_t(\tau)$ is the τ-maturity "term premium", $TP_t(\tau)$:

$$TP_t(\tau) = \frac{1}{\tau}\int_0^\tau FTP_t(s)ds = Y_t(\tau) - ER_t(\tau), \tag{B.18}$$

where $ER_t(\tau)$ is the τ-maturity average expected short rate:

$$ER_t(\tau) = \frac{1}{\tau}\int_0^\tau \mathrm{E}_t\left[r_{t+s}\right] ds. \tag{B.19}$$

Although the term premium is conventionally viewed as a risk premium,[27] there is a "wedge" between the term premium and the risk premium. Using equation (B.13), the forward term premium $FTP_t(\tau)$ can be written as:

$$FTP_t(\tau) = f_t(\tau) - \mathrm{E}_t\left[r_{t+\tau}\right] = h_t(\tau) + c_t(\tau) + d_t(\tau), \tag{B.20}$$

and so includes not only the risk premium, $h_t(\tau)$, but also convexity, $c_t(\tau)$, and the duration adjustment term, $d_t(\tau)$. The term premium therefore includes the term-average of all three terms:

$$TP_t(\tau) = Y_t(\tau) - ER_t(\tau) = TARP_t(\tau) + CX_t(\tau) + DA_t(\tau), \tag{B.21}$$

where $TARP_t(\tau) = \frac{1}{\tau}\int_0^\tau h_t(s)ds$ is the τ-maturity term-average risk premium, $CX_t(\tau) = \frac{1}{\tau}\int_0^\tau c_t(s)ds$ the average convexity term, and $DA_t(\tau) = \frac{1}{\tau}\int_0^\tau d_t(s)ds$ the average duration adjustment term.

B.3 Model Specification

We now specify the model used in the empirical analysis.

We implement a stochastic volatility model that belongs to the $A_1(4)$ class in the Dai and Singleton (2000) system. Specifically, we assume that the dynamics of the term structure are determined by three zero-mean latent factors x_1, x_2 and x_3, and, crucially, by a (latent) variance factor, v, which drives the conditional volatility of the other three state variables. In addition to being jointly related in the diffusion component through v, the factors are potentially mutually dependent via their conditional means.

We collect the latent factors in the vector $X = (v\ x_1\ x_2\ x_3)'$ and in equation (B.1) assume that S_t is a diagonal (4×4) matrix with all elements equal to v_t. We also assume that the functional form of the market price of risk is given by an "essentially affine" specification

[27] Cohen et al. (2018) express what appears to be a widely held view about the way that conventionally defined term premia should be understood: "This term premium is normally thought of as the extra return (a risk premium) that investors demand to compensate them for the risk associated with a long-term bond, particularly for long maturities."

(see Duffee (2002) and Duarte (2004)) of this type:

$$\Lambda_t = \sqrt{S_t^-}\,(\Lambda_0 + \Lambda_1 X_t)\,, \tag{B.22}$$

where S_t^- denotes the inverse of S_t, Λ_0 is a (4×1) vector and Λ_1 a (4×4) full matrix. This specification of the market price of risk implies that the risk-neutral mean-reversion coefficient and the long-term mean in equation (B.1) can be written as $\widetilde{K} = K + \Sigma\Lambda_1$ and $\widetilde{\Theta} = \widetilde{K}^{-1}(K\Theta - \Sigma\Lambda_0)$, respectively.

The essentially affine form allows risk premia to vary over time and change sign. The expression for the τ-maturity risk premium (equation (B.9)) thus becomes:

$$h_t(\tau) = -B'(\tau)\Sigma\,(\Lambda_0 + \Lambda_1 X_t)\,. \tag{B.23}$$

Equation (B.23) implies that risk premia are affine in the state variables and can therefore be decomposed into the four additive components:

$$h_t(\tau) = \sum_{i=1}^{4} h_{t,i}(\tau) = \sum_{i=1}^{4} \Gamma_{0,i}(\tau) + \sum_{i=1}^{4} \Gamma_{1,i}(\tau) X_{t,i}, \tag{B.24}$$

where $X_{t,i}, i = 1, \ldots, 4$, with $X_1 = v$, $X_2 = x_1$, $X_3 = x_2$, $X_4 = x_3$, $\Gamma_{0,i}(\tau) \equiv -B_i'(\tau)\Sigma_{ii}\Lambda_{0,i}$, and

$$\Gamma_{1,i}(\tau) \equiv \sum_{j=1}^{4} \left(-B_j'(\tau)\Sigma_{jj}\Lambda_{1,ji}\right),$$

where $\Gamma_{1,i}(\tau)$ measures the sensitivity of risk premia to state variable i at maturity τ.

Finally, we impose on the model the admissibility constraint that the off-diagonal terms in the first row of matrix K are zero and the following identification constraints (see Dai and Singleton (2000, 2002) and Duffee (2002)): (i) $\theta_{x_1} = \theta_{x_2} = \theta_{x_3} = 0$, in vector Θ of equation (B.3); (ii) all elements in the first row of matrix Λ_1 in equation (B.22) equal to zero; (iii) $\delta_v = \delta_{x_3} = 0$ and $\delta_{x_1} = \delta_{x_2} = 1$ in vector δ of equation (B.4), which means that the short rate is given by a constant (δ_0) plus the sum of the latent variables x_1 and x_2.

C Δe in Constant Volatility Models

In all ATSMs, Δe is given by:

$$\Delta e = \Delta B' K\,(\Theta - X_t) = \Delta B' \widetilde{K}\left(\widetilde{K}^{-1}K\right)(\Theta - X_t)\,,$$

where $\Delta B \equiv B(\tau_2) - B(\tau_1)$.

If the model has constant volatility, $B^{'}(\tau)$ is given by:

$$B^{'}(\tau) = \delta^{'} \widetilde{K}^{-1} \left[\mathrm{I} - e^{-\widetilde{K}\tau} \right].$$

(If the model also has constant risk premia, then $\widetilde{K} = K$). Writing $\widetilde{K}$ in terms of its eigenvalues, Ψ, and eigenvectors, U:

$$\widetilde{K} = U \Psi U^{-1},$$

and therefore:

$$B^{'}(\tau)\widetilde{K} = \delta^{'} \left[\mathrm{I} - U e^{-\Psi\tau} U^{-1} \right],$$

and

$$\begin{aligned} \Delta B^{'} \widetilde{K} &= \delta' U \left[e^{-\Psi\tau_1} - e^{-\Psi\tau_2} \right] U^{-1} \\ &= \delta' U \left(e^{-\Psi\tau_2} \left[e^{-\Psi(\tau_1 - \tau_2)} - \mathrm{I} \right] \right) U^{-1}. \end{aligned}$$

If all the eigenvalues are positive, i.e., $\Psi > 0$, then, holding $(\tau_2 - \tau_1)$ fixed, $\Delta B^{'} \widetilde{K}$ becomes small as $(\tau_1, \tau_2) \to \infty$ since $(e^{-\Psi\tau_1}, e^{-\Psi\tau_2}) \to 0$. If one or more eigenvalues are close to zero (but non-negative), then $\Delta B^{'} \widetilde{K}$ is also small, since $\left[e^{-\Psi(\tau_1 - \tau_2)} - \mathrm{I} \right]$ is small. Finally, if $\widetilde{K}^{-1} K$ is finite, then Δe is also small.

D Model Estimates

D.1 Data

Our empirical analysis is based on US Treasury STRIP prices (STRIP is an acronym for the Separate Trading of Registered Interest and Principal), which are obtained from Street Software and Bloomberg.

The interest and principal payments made by US Treasury notes and bonds are traded and registered separately and STRIP prices are direct observations of the discount factors that define the zero-coupon yield curve. Street Software's data is daily covering the period 1 January 1992 to 31 July 2010, while Bloomberg data (also daily) run from 1 August 2010 to 31 December 2023. Daily prices are available on each STRIP, resulting in around 160 observations of discount factors on a typical day. We use only coupon STRIPS and run a spline interpolation through the observed yields to obtain estimates of the yield curve at

annual maturities ranging between 1 and 30 years.

Estimates of realized yield variance for a given end-month are computed as the annualized variance of daily changes in yields during that month. We build an estimator for the τ-maturity yield variance at time t by regressing the τ-maturity realized variance at time $t+1$ on the realized variance at time t and the time-t implied variance calculated from the MOVE index.

The fitted values of the regression are our estimates of the time series of the τ-maturity yield variance. In the paper, references to estimates of volatility or variance derived directly from the data (as distinct from model estimates) indicate these fitted values. In the estimation, we use these estimates of yield variance for maturities of 10, 15, 20 and 25 years.

D.2 State-Space Representation and Estimation

Maximum-likelihood estimation of the model using the Kalman filter requires specification of the measurement and transition equations.

The measurement system consists of 52 equations. It includes yields and forward rates with maturities between 2 and 25 years, as well as yield variance at maturities of 10, 15, 20, and 25 years.[28] All variables in this system are expressed as affine functions of the latent factors X_t, with a vector of observation errors ϵ_t:

$$M_t = \mu_0(\phi) + \mu(\phi)\, X_t + \epsilon_t,$$

where M_t includes the 52 observable series at time t. The vector μ_0 and matrix μ are determined by the model equations and depend on the full parameter vector ϕ. The errors in ϵ_t are assumed to be normally distributed and homoskedastic.

The state dynamics are given by the discrete-time (monthly) counterpart of the continuous-time model in equation (B.3). The discrete-time representation is obtained from the solution to the stochastic differential equations governing the state variables, as in Christoffersen et al. (2014). The conditionally affine nature of the process implies that X_t follows a Markov vector autoregressive (VAR) process:

$$X_{t+1} = \eta_0(\phi) + \eta(\phi)\, X_t + u_{t+1},$$

where the vector η_0 and the matrix η also depend on the full set of model parameters ϕ. An important feature of the model is the presence of time-varying conditional volatility that

[28] We also estimate the model using a limited set of observable yields and forward rates (for example, including only values for the maturities 2, 5, 10, 15, 20, and 25 years), but the estimation results do not differ significantly.

is driven by the state variables. This implies that the shock u_{t+1} is heteroskedastic, with a conditional variance that is affine in the state vector X_t.

The parameters of the state-space model are estimated by maximum likelihood, using an approximate Kalman filter to infer the unobserved state variables. Although the state vector follows affine dynamics, its conditional distribution is non-Gaussian. Consequently, we approximate the conditional transition density using analytical expressions for the first two conditional moments of the state vector (see Duan and Simonato (1999), Duffee and Stanton (2012), and Christoffersen et al. (2014)).

D.3 Estimation Output

Panel A of **Table A.2** reports statistics on the goodness of fit. Errors are defined as the difference between the actual and model values of yields, forward rates and yield variances. The average standard deviation of errors, across maturities from 2 to 25 years, is 6.0 bps for yields and 24.9 bps for forward rates. The model fits medium- and long-term yields more accurately than short-term yields and, if we exclude the maturities below 5 years, the average standard deviation of yield errors falls to 4.6 bps. The average difference between monthly ex-post yield variance and our ex-ante estimate, for the maturities 10, 15, 20 and 25 years, is zero and the standard deviation of the differences is on average 0.06 bps. These errors imply an average standard deviation of the differences in yield volatilities of 2.8 bps.

Panel B of **Table A.2** reports the estimated model parameters. The latent factor x_1 is highly persistent and strongly correlated with the first principal component of yields, i.e., the level of the yield curve, with a correlation coefficient of 0.96. The factor x_2 has a mean-reversion coefficient of approximately 0.21 and is also strongly correlated (correlation coefficient 0.94) with the second principal component of yields, i.e., the slope of the yield curve.

The variance factor v exhibits relatively weak mean reversion, with a coefficient of about 0.16, and is highly correlated (correlation coefficient 0.99) with the first principal component of yield variances. The latent factor x_3 also displays weak mean reversion and appears to represent a second volatility-related factor, as it exhibits the highest correlation among the four state variables (approximately 0.30) with the second principal component of yield variances.

E Decomposition of Errors in Risk Premia Simulations

The risk premium for maturity τ is:

$$h_t(\tau) \equiv -B'(\tau)\Sigma\Lambda_0 - B'(\tau)\Sigma\Lambda_1 X_t.$$

The corresponding estimated value (in each trial) can be written as:

$$\hat{h}_t(\tau) = -\hat{B}'(\tau)\hat{\Sigma}\hat{\Lambda}_0 - \hat{B}'(\tau)\hat{\Sigma}\hat{\Lambda}_1 X_t.$$

As discussed in Section 4.3, the estimation error in X_t is extremely small and so, in computing the estimated value of h_t, we do not distinguish between X_t and $\hat{X}_t$. The error in $h_t(\tau)$ is therefore:

$$\begin{aligned}
\varepsilon_t(\tau) &= \hat{h}_t(\tau) - h_t(\tau) \\
&= -\big(\hat{B}'(\tau)\hat{\Sigma}\hat{\Lambda}_0 + \hat{B}'(\tau)\hat{\Sigma}\hat{\Lambda}_1 X_t\big) + \big(B'(\tau)\Sigma\Lambda_0 + B'(\tau)\Sigma\Lambda_1 X_t\big) \\
&= -\big(\varepsilon_B\Sigma\Lambda_0 + B'(\tau)\varepsilon_{\Sigma\Lambda_0} + \varepsilon_B\varepsilon_{\Sigma\Lambda_0}\big) - \big(\varepsilon_B\Sigma\Lambda_1 + B'(\tau)\varepsilon_{\Sigma\Lambda_1} + \varepsilon_B\varepsilon_{\Sigma\Lambda_1}\big)X_t, \qquad \text{(E.25)}
\end{aligned}$$

where:

$$\begin{aligned}
\varepsilon_B &\equiv \hat{B}'(\tau) - B'(\tau), \\
\varepsilon_{\Sigma\Lambda_0} &\equiv \hat{\Sigma}\hat{\Lambda}_0 - \Sigma\Lambda_0, \\
\varepsilon_{\Sigma\Lambda_1} &\equiv \hat{\Sigma}\hat{\Lambda}_1 - \Sigma\Lambda_1.
\end{aligned}$$

F Robustness Checks

F.1 Different Data Sets

Much of the empirical literature on the term structure employs the yield estimates produced by Gurkaynak et al. (2007) while, for reasons explained in Section 3.1, we rely on STRIPS prices. To test whether the main empirical results depend on our use of STRIPS prices, we re-estimate our SV model using yields, forward rates and yield volatilities from two alternative data sets: Gurkaynak et al. (2007) (GSW) and Liu and Wu (2021) (LW).

In both cases, we obtain a goodness of fit similar to that reported in **Table A.2**, with the average (across maturities) standard deviation of estimation errors in yields, forward rates and yield volatility equal, respectively, to 5.1 bps, 26.1 bps and 3.6 bps for GSW and 7.9

bps, 25.7 bps and 3.4 bps for LW (compared with 6.0 bps, 24.9 bps and 2.8 bps for STRIPS data).

Panel A of **Figure A.1** reports estimates of Δf and Δe for the maturity interval 15 to 25 years obtained using GSW and LW data. These estimates are similar to those in Panel A of **Figure 3** which use STRIPS data. Critically, the estimates obtained using both the GSW and LW data imply values of Δe that are small and thus consistent with the near-observability of Δh implied by the approximation given in equation (6). The two time series, which appear almost identical, have both a mean around 1 bp and a standard deviation of about 14 bps (vs. -0.1 bps and 13.4 bps for the estimates with STRIPS).

The GSW and LW estimates of risk premia, h, are also similar to those obtained with STRIPS data. Panel B of **Figure A.1** plots the 10-year risk premium estimated using GSW and LW data against the corresponding values estimated with STRIPS. The average size of the risk premium is 119 bps for GSW and 123 bps for LW, compared with 146 bps for STRIPS. The R^2 obtained by regressing the GSW risk premium against the STRIPS risk premium is 0.90, while in the case of the LW risk premium is 0.97.

Overall, the estimates using GSW and LW data confirm that our main results are not sensitive to the use of STRIPS prices.

F.2 Including Forward Rates and Yield Volatility in Estimation

As discussed in Section 3.1, our estimation method differs from standard practice in a number of respects; not only in the use of long-maturity data but also in the fact that we fit to forward rates and yield volatility as well as zero-coupon yields. Since forward rates contain the same information as yields, it may not be immediately obvious what benefit there is in including forward rates along with yields.

The issue is one of weighting. Although the term structure of zero-coupon yields up to 25 years contains the same information as the corresponding term structure of forward rates, the weight attached to the one-year forward rate at 25 years in the 25-year zero-coupon yield is only 4%. Since our "expectations-free" estimate of Δh is obtained from long-term forward rates, it makes sense that we would wish to put significant weight on these in estimation.

In this appendix we again use Monte Carlo to examine the effect on the precision of our estimates of excluding some of these data in estimation. The experiment follows the procedure used in **Table 1** except here we use maturities of up to 25 years in each case.

The results are given in **Table A.3**. This has three panels: in Panel A, we fit to zero-coupon yields and yield volatility, but exclude forward rates; in Panel B, we fit to forward

rates and yield volatility, but exclude zero-coupon yields and in Panel C, we fit to forward rates and zero-coupon yields, but exclude yield volatility. The key results for the precision of risk premium estimates are given in columns (3) and (4) of each panel for, respectively, the median standard deviation of the estimation error and the standard deviation of the mean error. The corresponding results for the expected short rate are given in columns (6) and (7) and the statistics are expressed in the form of the ratio to the corresponding statistic in Panel A of **Table 1**.

For the median standard deviation of estimation errors in risk premia (column (3)), all the ratios in Panels A, B and C are above one. In other words, failing to fit to any of forward rates, yields or yield volatility increases estimation error in the risk premium. The reduction in precision is significant: excluding forward rates increases the median standard deviation by between 17% and 60%; excluding yields by between 7% and 43% and excluding yield volatility by between 75% and 153%. Significantly, the reduction in precision is greatest when we exclude yield volatility and smallest when we exclude zero-coupon yields − the variable most commonly used in fitting models in the literature. Comparing Panel C of **Table A.3** and Panel B of **Table 1** shows that failing to fit to volatility removes much of the substantial benefit from including long-term rates.

The results for the standard deviation of the mean error (column (4)) follow the same pattern except that the ratios are closer to one. Once again, eliminating yields and fitting only to forward rates and yield volatility (Panel B) has the smallest impact on estimation errors − indeed at longer maturities it seems to have almost no effect − while failing to fit to either forward rates or yield volatility has a substantial effect.

The results for errors in the expected short rate (columns (6) and (7)) follow a similar pattern, the main exception being Panel C where, for long maturities, failing to fit to yield volatility has a much smaller effect on the precision of both the median standard deviation of errors and the standard deviation of the mean error than for estimates of the risk premium.

We do not claim that the particular combination of variables that we use in fitting the model is optimal, any more than it can be claimed that the standard practice of excluding rates longer than 10 years and fitting only to zero-coupon yields is "optimal". However, the results strongly suggest that fitting to yields alone significantly compromises the precision of estimates of both risk premia and expected short rates and that, at least for our SV model, fitting to yield volatility is critical.

F.3 Improved Precision in Term Premium Estimates

A key motivation for using a stochastic volatility (SV) model is to capture time-variation in convexity so that the slope of the term structure of risk premia at long maturities can be accurately separated from the slope of the forward curve. Our simulations in Section 4 show that this leads to a substantial improvement in the precision of estimates of *risk* premia at all maturities. The *term* premium, on the other hand, combines the risk premium with convexity and so estimating the term premium does not, apparently, require separate estimation of convexity. Perhaps, therefore, in estimating the term premium there is little benefit in using a SV model rather than a constant volatility (CV) model, such as the Fed benchmark models. This appendix addresses this issue directly by comparing, in a Monte Carlo setting, the precision of term premium and expected short rate estimates obtained from the SV model and from an otherwise comparable CV specification.

The CV model we use here is a standard four-factor model ($A_0(4)$, in Dai and Singleton (2000) notation) estimated using the exactly same technique and data that we employ for our SV model. Specifically, we fit the model to both yields and forward rates for maturities from two to 25 years and also to the estimated variance of yields at maturities of 10, 15, 20 and 25 years. In this model, yield volatility for a given maturity is constant and so "fitting to variance" simply means fitting to average yield variance. We assume that the state variables are latent factors and, as for our model, use a maximum likelihood Kalman filter methodology to obtain estimates of the factors and the model parameters.

In sample, the model fits yields and forward rates well, with an average (across maturities) standard deviation of errors of 4.6 bps and 21.4 bps, respectively, while the standard deviation of the errors obtained by fitting average yield volatility ranges between 16 bps (25-year) and 21 bps (10-year). Thus the estimation errors for yields and forward rates are comparable to those obtained for the SV model (see Panel A of **Table A.2**), while the CV model, of course, cannot fit the time series of volatility.

In **Table A.4** we report the results of a Monte Carlo experiment evaluating the relative precision of the SV and CV models in estimating term premia and average expected short rates. The calculations follow the procedure used in Section 4 and, for each trial, we calculate the estimation errors in term premia (TP) and average expected short rates (ER) as the differences between the estimated and "true" values (i.e., those simulated using the in-sample parameters) and compute the mean and the standard deviation of these errors.

In Panel A, columns (2−4) report respectively, the mean error, the median value of the within-trial standard deviation of the errors and the standard deviation (across trials) of the within-trial mean error for TP. The results in column (2) show no evidence of bias,

with the mean error in the estimated term premium for all maturities less than 1 bp. The median value of the standard deviation of the estimation error in column (3) ranges from 2.9 to 5.8 bps. Column (4) shows that, as in the case of risk premia, the standard deviation of the mean error is approximately of the same order of magnitude as the median standard deviation of the error around the mean.

Columns (5−7) contain the corresponding results for ER. As before, there is no evidence of bias (column (5)). The median standard deviation of the estimation error (column (6)) and the standard deviation of the mean error (column (7)) are similar to those obtained for TP.

Panel B of **Table A.4** now compares the precision of the SV and CV models when both are estimated using yield data up to a maturity of 25 years.[29] As in Panel A, mean errors (columns (2) and (5)) are shown in basis points, while the remaining statistics are shown as ratios to the corresponding statistic in Panel A, i.e., relative to the result using the SV model. Columns (3) and (6) show that the level of imprecision in the CV model relative to the SV model is large for both TP and ER, as the median standard deviation of errors in the CV model estimates is around five times larger than for the SV model.

Columns (4) and (7) report the corresponding results for the standard deviation of the mean errors. Since the ratios are above one, the SV model has higher precision, but the improvement is smaller than for the median standard deviation. Therefore, just as we found for risk premia, the inclusion of stochastic volatility has a much greater impact on the size of errors around the mean than on the error in the mean itself.

Overall, and despite the fact that the term premium "includes convexity", the benefits of including stochastic volatility for estimating term premia are very similar to those reported for risk premia.

[29] The estimating of the CV model with yield data only up to 10 years provides very similar results. Indeed, since the CV model cannot capture the time-variation in convexity and thus the slope of the TSRP, the precision of CV model estimates using long-term yields (up to 25 years) shows no significant improvement with respect to the standard case of yields up to 10 years.

References

Adrian, T., R. Crump, and E. Moench. 2013. Pricing the term structure with linear regressions. *Journal of Financial Economics* 110:110–138.

Andreasen, M. M., J. H. Christensen, and G. D. Rudebusch. 2019. Term structure analysis with big data: one-step estimation using bond prices. *Journal of Econometrics* 212:26–46.

Balter, A. G., A. Pelsser, and P. C. Schotman. 2021. What does a term structure model imply about very long-term interest rates? *Journal of Empirical Finance* 62:202–219.

Bauer, M., G. Rudebusch, and J. Wu. 2014. Term premia and inflation uncertainty: Empirical evidence from an international panel dataset: Comment. *American Economic Review* 104:323–37.

Bauer, M. D., and J. D. Hamilton. 2018. Robust bond risk premia. *Review of Financial Studies* 31:399–448.

Bauer, M. D., G. D. Rudebusch, and J. C. Wu. 2012. Correcting estimation bias in dynamic term structure models. *Journal of Business & Economic Statistics* 30:454–467.

Berardi, A., and A. Plazzi. 2022. Dissecting the yield curve: the international evidence. *Journal of Banking & Finance* 134:106286.

Berardi, A., and S. Schaefer. 2026. Does the slope of the yield curve actually predict Treasury returns? London Business School working paper.

Brown, R., and S. Schaefer. 1994. Interest rate volatility and the shape of the term structure. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences* 347:563–576.

Brown, R., and S. Schaefer. 2000. Why long term forward interest rates (almost) always slope downwards. London Business School working paper.

Carpenter, J. N., F. Lu, and R. F. Whitelaw. 2026. Government bond risk and return in the US and China. *Journal of Financial Economics* 176:104224.

Carr, P., and L. Wu. 2023. Decomposing long bond returns: a decentralized theory. *Review of Finance* 27:997–1026.

Choi, H., P. Mueller, and A. Vedolin. 2017. Bond variance risk premiums. *Review of Finance* 21:987–1022.

Christensen, J. H., and G. D. Rudebusch. 2012. The response of interest rates to US and UK quantitative easing. *Economic Journal* 122:F385–F414.

Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui. 2014. Nonlinear Kalman filtering in affine term structure models. *Management Science* 60:2248–2268.

Cieslak, A., and P. Povala. 2016. Information in the term structure of yield curve volatility. *Journal of Finance* 71:1393–1436.

Cochrane, J., and M. Piazzesi. 2008. Decomposing the yield curve. University of Chicago Booth working paper.

Cohen, B., P. Hördahl, and F. Xia. 2018. Term premia: models and some stylised facts. BIS Quarterly Review, September.

Cox, J., J. Ingersoll Jr, and S. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53:385–407.

Dai, Q., and K. Singleton. 2000. Specification analysis of affine term structure models. *Journal of Finance* 55:1943–1978.

Dai, Q., and K. Singleton. 2002. Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of Financial Economics* 63:415–441.

Duan, J.-C., and J.-G. Simonato. 1999. Estimating and testing exponential-affine term structure models by Kalman filter. *Review of Quantitative Finance and Accounting* 13:111–135.

Duarte, J. 2004. Evaluating an alternative risk preference in affine term structure models. *Review of Financial Studies* 17:379–404.

Duffee, G. 2002. Term premia and interest rate forecasts in affine models. *Journal of Finance* 57:405–443.

Duffee, G., and R. Stanton. 2012. Estimation of dynamic term structure models. *Quarterly Journal of Finance* 2:1–51.

Duffie, D., and R. Kan. 1996. A yield-factor model of interest rates. *Mathematical Finance* 6:379–406.

Dybvig, P. H., J. E. Ingersoll Jr, and S. A. Ross. 1996. Long forward and zero-coupon rates can never fall. *Journal of Business* pp. 1–25.

Fama, E., and R. Bliss. 1987. The information in long-maturity forward rates. *American Economic Review* 77:680–692.

Feldhütter, P., C. Heyerdahl-Larsen, and P. Illeditsch. 2018. Risk premia and volatilities in a nonlinear term structure model. *Review of Finance* 22:337–380.

Gagliardini, P., P. Porchia, and F. Trojani. 2008. Ambiguity aversion and the term structure of interest rates. *Review of Financial Studies* 22:4157–4188.

Giglio, S., and B. Kelly. 2018. Excess volatility: beyond discount rates. *Quarterly Journal of Economics* 133:71–127.

Gurkaynak, R., B. Sack, and J. Wright. 2007. The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54:2291–2304.

Hamilton, J. D., and J. C. Wu. 2012. Identification and estimation of Gaussian affine term structure models. *Journal of Econometrics* 168:315–331.

Joslin, S., and A. Le. 2021. Interest rate volatility and no-arbitrage affine term structure models. *Management Science* 67:7391–7416.

Jotikasthira, C., A. Le, and C. Lundblad. 2015. Why do term structures in different currencies co-move? *Journal of Financial Economics* 115:58–83.

Kim, D., and J. Wright. 2005. An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. FEDS working paper 2005-33.

Kim, D. H., and A. Orphanides. 2012. Term structure estimation with survey data on interest rate forecasts. *Journal of Financial and Quantitative Analysis* 47:241–272.

Langetieg, T. C. 1980. A multivariate model of the term structure. *Journal of Finance* 35:71–97.

Leippold, M., and F. Matthys. 2022. Economic policy uncertainty and the yield curve. *Review of Finance* 26:751–797.

Liu, Y., and J. C. Wu. 2021. Reconstructing the yield curve. *Journal of Financial Economics* 142:1395–1425.

Martin, I. W., and S. A. Ross. 2019. Notes on the yield curve. *Journal of Financial Economics* 134:689–702.

Moench, E. 2019. The term structures of global yields. BIS Papers 102.

Pancost, A. N. 2021. Zero-coupon yields and the cross-section of bond prices. *Review of Asset Pricing Studies* 11:209–268.

Rebonato, R., and V. Putyatin. 2018. The value of convexity: A theoretical and empirical investigation. *Quantitative Finance* 18:11–30.

Sack, B. 2000. Using Treasury STRIPS to measure the yield curve. FEDS working paper 2000-42.

Vasicek, O. 1977. An equilibrium characterization of the term structure. *Journal of Financial Economics* 5:177–188.

Vasicek, O. A., and H. G. Fong. 1982. Term structure modeling using exponential splines. *Journal of Finance* 37:339–348.

Table 1
Information in Long-Maturity Yields

This table reports the results of a Monte Carlo experiment addressing the additional information contained in long-maturity yields for the estimation of risk premia and expected short rates. We use the in-sample parameter estimates from our stochastic volatility (SV) model to simulate 500 paths for the four state variables with the same length as our data, i.e., 384 months, and, for each month and trial in the simulation, compute the term structure of yields, forward rates and yield volatilities. Then, for each trial, we fit the SV model first using yields, forward rates and yield volatilities up to 25 years (SV_{25}) and then using yields, forward rates and yield volatilities only up to 10 years (SV_{10}). For each trial, we calculate the estimation errors in risk premia and expected short rates as the differences between the estimated and "true" values (i.e., those simulated using the in-sample parameters) and compute the mean and the standard deviation of these errors. Panel A reports, for SV_{25}, (i) the average, over trials, of the mean errors; (ii) the median, over trials, of the standard deviation of the errors; and (iii) the standard deviation, over trials, of the mean error for estimates of risk premia and expected short rates with maturities 2, 5, 10, 15, 20, and 25 years. Panel B reports the same statistics obtained from SV_{10}, i.e., by fitting the SV model to the simulated data using maturities up only to 10 years, with the statistics (ii) and (iii) expressed as a ratio to the corresponding statistics in Panel A (i.e., SV_{25}). The values in Panel A and in columns (2) and (5) of Panel B are in basis points.

Panel A: SV model with maturities up to 25 years (SV_{25})

	Risk premia			Expected short rates		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors (3)	St. dev. of mean errors (4)	Average of mean errors (5)	Median of st. dev. of errors (6)	St. dev. of mean errors (7)
2	0.37	4.12	5.27	-0.17	4.03	4.81
5	0.56	7.93	9.47	-0.40	6.12	7.46
10	0.38	11.60	11.76	-0.41	6.71	7.94
15	0.04	14.71	11.87	-0.21	6.73	9.06
20	-0.25	17.36	12.16	0.03	7.06	11.60
25	-0.45	20.39	14.80	0.24	7.23	14.70

Panel B: SV model with maturities up to 10 years (SV_{10})

	Risk premia			Expected short rates		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors Ratio wrt SV_{25} (3)	St. dev. of mean errors Ratio wrt SV_{25} (4)	Average of mean errors (5)	Median of st. dev. of errors Ratio wrt SV_{25} (6)	St. dev. of mean errors Ratio wrt SV_{25} (7)
2	0.72	3.81	1.17	-4.93	3.74	1.15
5	-0.42	3.74	1.16	-10.18	3.96	1.11
10	-5.95	3.48	1.13	-16.26	3.89	1.32

Table 2
Decomposition of Errors in Bond Risk Premia

This table contains the decomposition of the errors in estimated risk premia obtained by fitting models to simulated data. In particular, we use the in-sample parameter estimates from our stochastic volatility (SV) model to simulate 500 paths for the four state variables with the same length as our data, i.e., 384 months, and, for each month and trial in the simulation, compute the term structure of yields, forward rates and yield volatilities. Then, for each trial, we fit the SV model first using yields, forward rates and yield volatilities up to 25 years (SV_{25}) and then using yields, forward rates and yield volatilities only up to 10 years (SV_{10}). For each trial, we calculate the errors in risk premia as the differences between the estimated and "true" values (i.e., those simulated using the in-sample parameters), and decompose these errors using equation (E.25): error from constant term, $-\left(\varepsilon_B \Sigma\Lambda_0 + B'(\tau)\varepsilon_{\Sigma\Lambda_0} + \varepsilon_B \varepsilon_{\Sigma\Lambda_0}\right)$ (column (2)); error from $B(\tau)$, $-\varepsilon_B \Sigma\Lambda_1 X_t$ (columns (3–4)); error from $\Sigma\Lambda_1$, $-B'(\tau)\varepsilon_{\Sigma\Lambda_1} X_t$ (columns (5–6)); cross-product error, $-\varepsilon_B \varepsilon_{\Sigma\Lambda_1} X_t$ (columns (7–8)). Columns (9–10) report the values in columns (2–3) of **Table 1**. For each error component, the table reports (i) the average across simulation trials of the mean estimation error ("Mean"), and (ii) the median across trials of the standard deviation of the estimation error ("St. dev."). Panel A shows the decomposition for SV_{25} and Panel B for SV_{10}. All values are expressed in basis points.

Panel A: Decomposition of errors in risk premia simulations for SV_{25}

	Error from constant	Error from $B(\tau)$		Error from $\Sigma\Lambda_1$		Cross-product error		Total error	
Maturity (yrs) (1)	Mean (2)	Mean (3)	St. dev. (4)	Mean (5)	St. dev. (6)	Mean (7)	St. dev. (8)	Mean (9)	St. dev. (10)
2	-1.11	-0.03	0.63	1.56	4.10	-0.05	0.08	0.37	4.12
5	-1.50	-0.07	1.19	2.25	8.09	-0.11	0.16	0.56	7.93
10	-0.35	-0.18	1.62	1.05	11.74	-0.14	0.21	0.38	11.60
15	1.81	-0.27	1.80	-1.39	14.33	-0.12	0.22	0.04	14.71
20	4.40	-0.32	1.84	-4.23	17.45	-0.10	0.23	-0.25	17.36
25	7.14	-0.32	1.78	-7.17	20.28	-0.10	0.22	-0.45	20.39

Panel B: Decomposition of errors in risk premia simulations for SV_{10}

	Error from constant	Error from $B(\tau)$		Error from $\Sigma\Lambda_1$		Cross-product error		Total error	
Maturity (yrs) (1)	Mean (2)	Mean (3)	St. dev. (4)	Mean (5)	St. dev. (6)	Mean (7)	St. dev. (8)	Mean (9)	St. dev. (10)
2	35.15	0.04	0.49	-35.05	15.59	0.58	0.28	0.72	15.69
5	68.64	0.03	0.89	-70.26	29.48	1.17	0.51	-0.42	29.68
10	81.53	0.00	1.45	-89.56	40.03	2.09	0.80	-5.95	40.33

Table 3
Comparison of Term Premia in SV and Fed Models

Panel A of this table reports summary statistics on the 10-year term premium and average expected short rate estimated by our stochastic volatility model (SV) and the Fed models from Kim and Wright (2005) (KW), Adrian et al. (2013) (ACM), and Christensen and Rudebusch (2012) (CR). Column (2) shows the mean and column (3) the standard deviation of estimated 10-year term premia, where mean values are computed from the variables in level, while annualised standard deviations from monthly changes in the variables. Columns (4) and (5) show mean and standard deviation for estimated 10-year average expected short rates. The values are expressed in basis points. Columns (6) and (7) report, for each model, the variance decomposition of the 10-year yield into the term premium and the average expected short rate. These values are in percentage terms. Panel B reports the same statistics for the decomposition of the 10-year forward rate. The sample period is from January 1992 to December 2023, with the exception of CR for which data are available from January 1998.

Panel A: Term premium and average expected short rate

	Term premium (TP)		Av. expected short rate (ER)		% of variance explained	
Model	Mean	St. dev.	Mean	St. dev.	TP	ER
(1)	(2)	(3)	(4)	(5)	(6)	(7)
SV	55	44	362	70	32	68
KW	76	47	340	48	50	50
ACM	103	82	311	63	66	34
CR	64	124	284	87	74	26

Panel B: Forward term premium and expected short rate

	Forward term premium (FTP)		Expected short rate (E[r])		% of variance explained	
Model	Mean	St. dev.	Mean	St. dev.	FTP	E[r]
(1)	(2)	(3)	(4)	(5)	(6)	(7)
SV	84	47	429	68	36	64
KW	150	63	385	39	62	38
ACM	163	114	344	39	94	6
CR	103	135	320	50	92	8

Table 4
Decomposition for the 2022-2023 Quantitative Tightening

This table contains summary statistics for the decomposition of the 10-year forward rate into the forward term premium (FTP) and the expected short rate (E[r]) estimated by our stochastic volatility model (SV) and the Fed models from Kim and Wright (2005) (KW), Adrian et al. (2013) (ACM), and Christensen and Rudebusch (2012) (CR) for the March 2022 to October 2023 Fed Quantitative Tightening. Column (2) reports, for each model, the percentage of the total change in the forward rate over the period explained by the term premium component and column (3) the annualised standard deviation of monthly changes in the estimated forward term premium (expressed in bps). Columns (4–5) contain the same statistics for the expected short rate, while columns (6–7) show the percentage of the variance of monthly changes in the 10-year forward rate explained by its term premium and expected short rate components, respectively.

	Forward term premium (FTP)		Expected short rate (E[r])			
Model	Share of forward rate change	Standard deviation	Share of forward rate change	Standard deviation	% of variance explained FTP	E[r]
(1)	(2)	(3)	(4)	(5)	(6)	(7)
SV	21	41	79	84	30	70
KW	43	78	57	44	65	35
ACM	77	138	23	37	89	11
CR	56	105	44	49	73	27

Figure 2
Time Variation in Convexity

This figure reports estimates of Δf and Δc, calculated as the difference between the 25-year and 15-year maturity forward rates and convexity terms, respectively. The values of Δc reported in the figure include both those obtained directly using yields calculated from the STRIPS data ($\Delta c = -(1/2)(25^2 \times \text{Var}[dY(25)] - 15^2 \times \text{Var}[dY(15)])$) and those obtained using the MOVE index rescaled to the average level of the 20-year yield variance ($\Delta c_{MOVE} = -(1/2)(25^2 - 15^2) \times (MOVE_{rescaled})^2$). The sample period is January 1992 to December 2023.

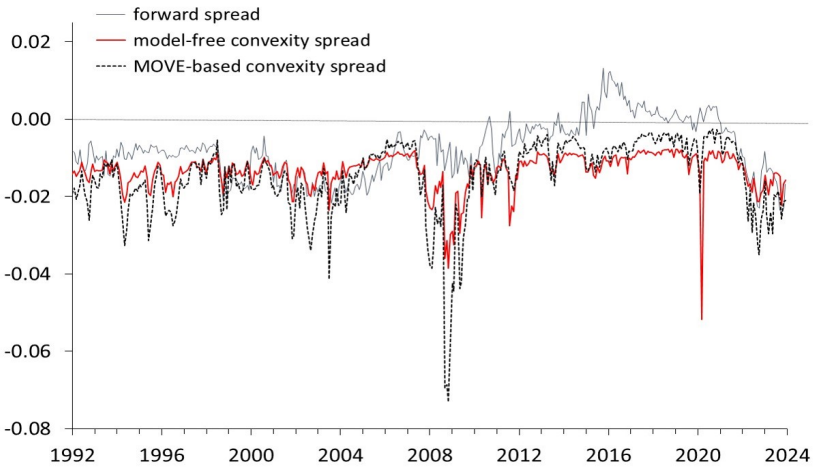
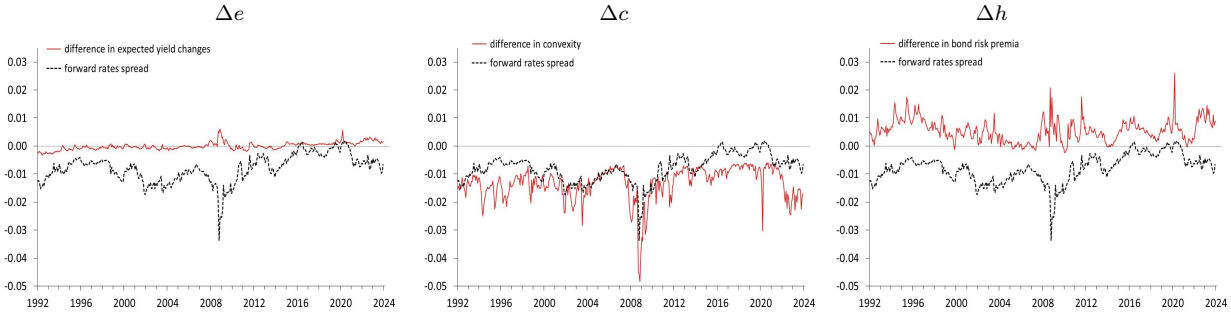


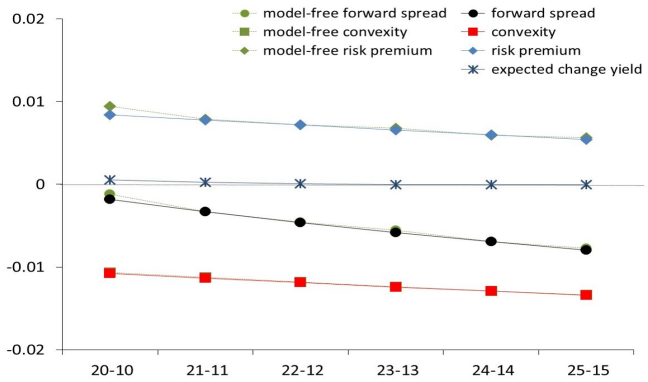
Figure 3
Forward Spread and Its Components

This figure shows estimates of Δf, calculated as the spread between the 25- and 15-year maturity instantaneous forward rates, and its components. Panel A shows the decomposition of the model values of Δf into its three components: the difference in expected yield changes (Δe), convexity (Δc), and bond risk premia (Δh). Panel B reports the term structure of the average values of the forward rate spread (Δf), difference in risk premia (Δh) and difference in convexity (Δc) for a range of maturities and a constant difference in maturity of 10 years. The model values are compared with "model-free" values calculated using measurable proxies for forward rates and convexity terms. Panel C reports the model estimates of Δe for different maturity intervals: (i) 10 years minus 1 year; (ii) 15 years minus 5 years; (iii) 25 years minus 15 years. The sample period is from January 1992 to December 2023.

Panel A: Forward spread components



Panel B: Decomposition of term structure of average forward rate spreads



Panel C: Δe for different maturity intervals

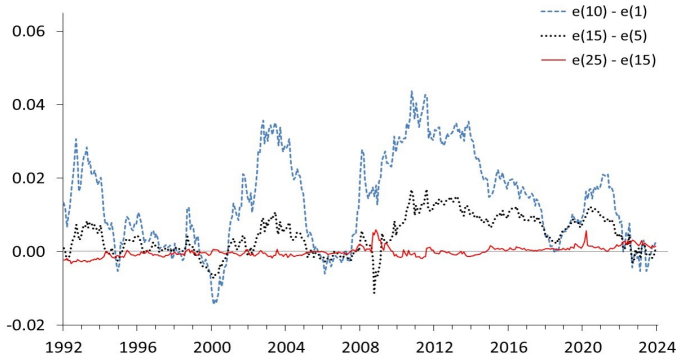
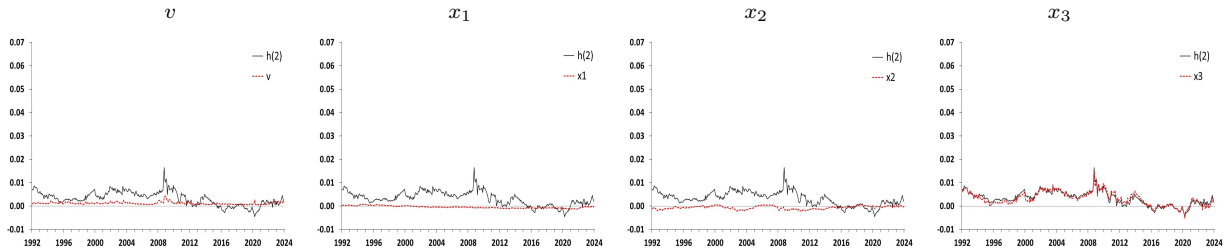


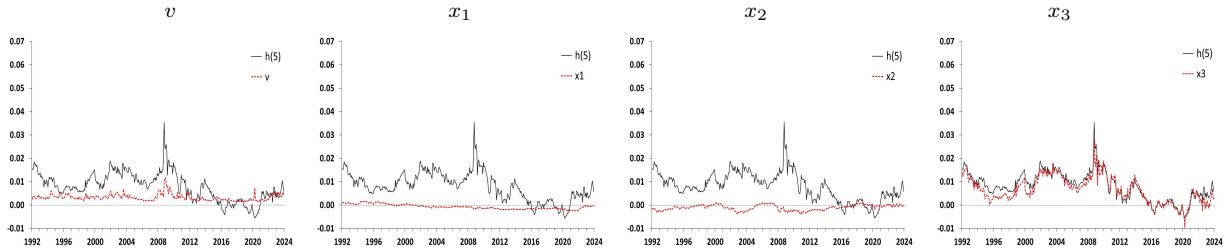
Figure 4
Decomposition of Bond Risk Premia

This figure shows the decomposition of bond risk premia into the components related to the four state variables, i.e., the variance factor v, the level factor x_1, the slope factor x_2, and the slope of volatility factor x_3, over the period January 1992 to December 2023. Panel A reports the decomposition for the 2-year risk premium, while Panel B, Panel C and Panel D show the decomposition for the 5-year, 10-year and 20-year risk premium, respectively.

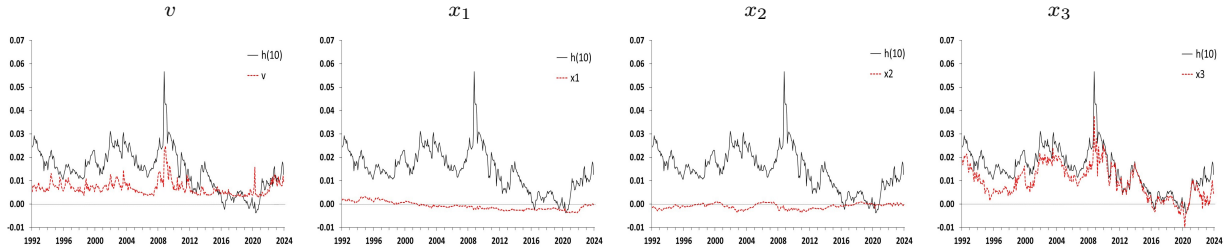
Panel A: 2-year risk premium



Panel B: 5-year risk premium



Panel C: 10-year risk premium



Panel D: 20-year risk premium

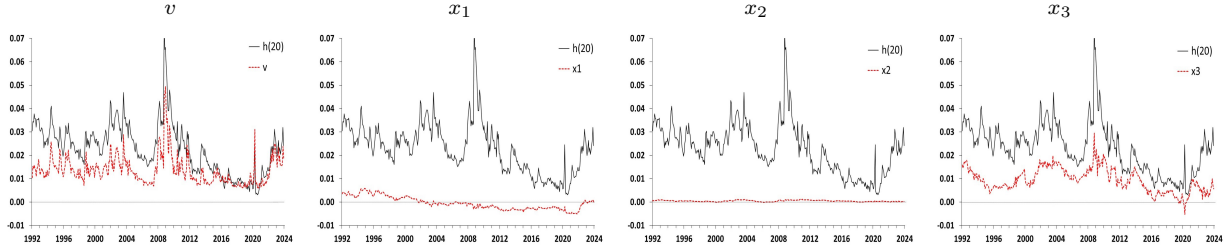
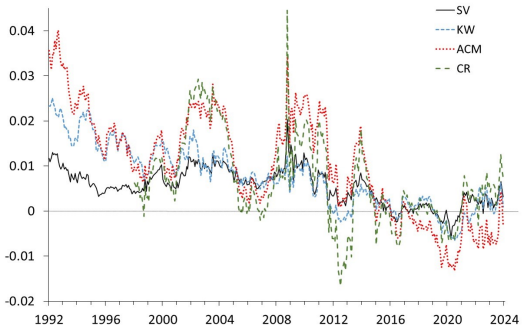


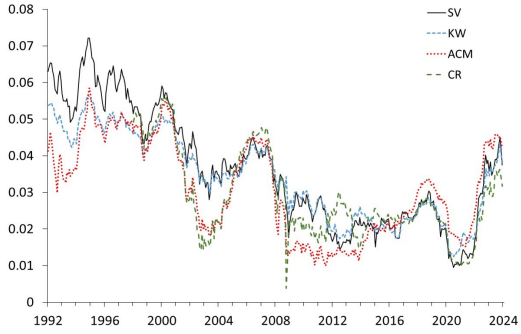
Figure 5
Estimates of Term Premia and Expected Rates

Panel A of this figure shows the time series of the 10-year term premium (calculated as the difference between the 10-year yield and average expected short rate) for our stochastic volatility model (SV) and the Fed models: Kim and Wright (2005) (KW), Adrian et al. (2013) (ACM), and Christensen and Rudebusch (2012) (CR). Panel B reports the series for the 10-year average expected short rate. Panels C and D report estimates for the 10-year forward term premium and expected short rate, respectively, along with the 10-year forward rate. Panel E compares the 10-year risk premium estimated by SV with the equivalent series implied by ACM estimates. The sample period is from January 1992 to December 2023, with the exception of CR for which data are available from January 1998.

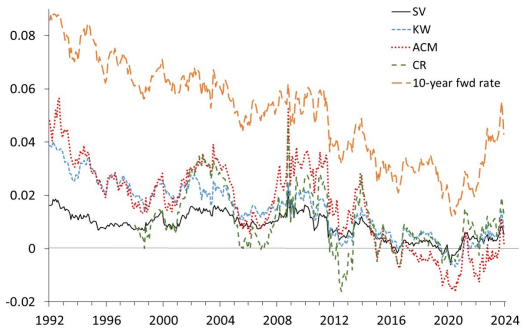
Panel A: Term premium



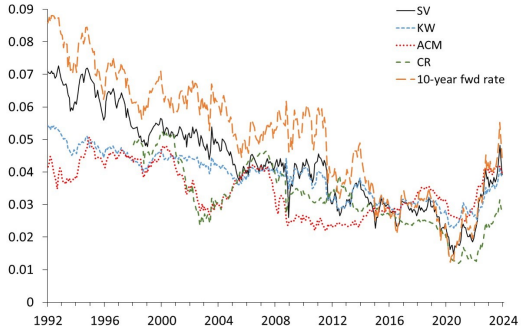
Panel B: Average expected short rate



Panel C: Forward term premium



Panel D: Expected short rate



Panel E: Risk premium

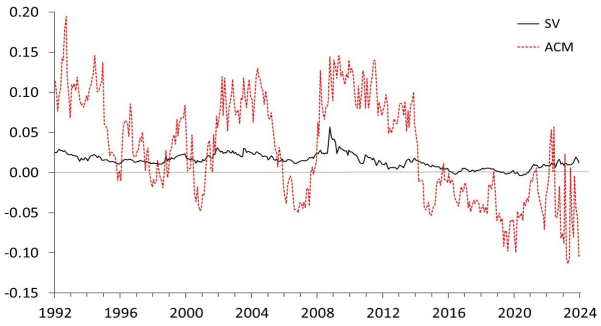
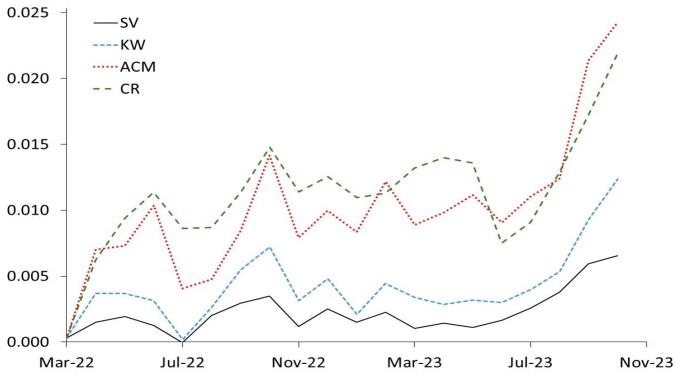


Figure 6
The 2022-2023 Quantitative Tightening

This figure shows the time series of the 10-year forward term premium (Panel A) and expected short rate (Panel B) for our stochastic volatility model (SV) and the Fed models (Kim and Wright (2005) (KW), Adrian et al. (2013) (ACM), and Christensen and Rudebusch (2012) (CR)) over the Fed quantitative tightening period from March 2022 to October 2023. The values are the cumulative changes from March 2022.

Panel A: Forward term premium



Panel B: Expected short rate

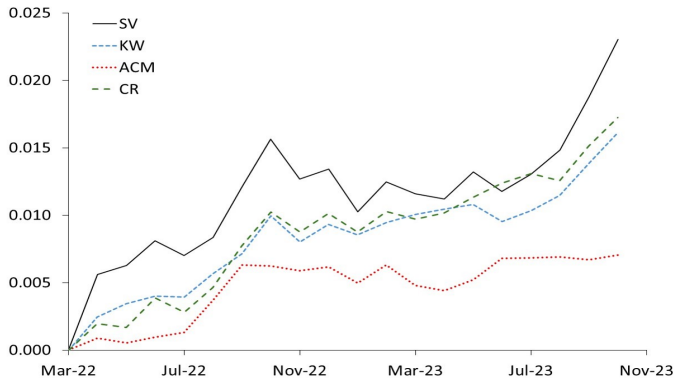


Table A.1
List of Empirical Term Structure Papers

This table reports a list of one hundred empirical studies on the term structure for the period from 2000 to 2025.

Abrahams, M., T. Adrian, R. Crump, E. Moench, and R. Yu. Journal of Monetary Economics 2016
Adrian, T., R. Crump, and E. Moench. Journal of Financial Economics 2013
Almeida, C., J. Graveline, and S. Joslin. Journal of Econometrics 2011
Andersen, T., and L. Benzoni. Journal of Finance 2010
Andreasen, M., J. Christensen, and G. Rudebusch. Journal of Econometrics 2019
Andreasen, M., J. Christensen, and S. Riddell. Review of Finance 2021
Ang, A., and M. Piazzesi. Journal of Econometrics 2003
Ang, A., G. Bekaert, and M. Wei. Journal of Finance 2008
Ang, A., M. Piazzesi, and M. Wei. Journal of Econometrics 2006
Argyropoulos, E., and E. Tzavalis. Journal of Empirical Finance 2015
Aruoba, B. Journal of Business and Economic Statistics 2020
Backus, D., S. Foresi, A. Mozumdar, and L. Wu. Journal of Financial Economics 2001
Balter, A., A Pelsser, and P. Schotman. Journal of Empirical Finance 2021
Bauer, M. Journal of Business & Economic Statistics 2018
Bauer, M., and J. Hamilton. Review of Financial Studies 2018
Bauer, M., G. Rudebusch, and J. Wu. Journal of Business & Economic Statistics 2012
Bekaert, G., E. Engstrom, and A. Ermolov. Journal of Financial Economics 2021
Bekaert, G., and A. Ermolov. Journal of Financial and Quantitative Analysis 2023
Bikbov, R., and M. Chernov. Journal of Econometrics 2011
Bikbov, R., and M. Chernov. Journal of Econometrics 2013
Bikbov, R., and M. Chernov. Management Science 2009
Bolder, D., G. Johnson, and A. Metzler. BoC working paper 2004
Buraschi, A., and A. Jiltsov. Journal of Financial Economics 2005
Byrne, J., S. Cao, and D. Korobilis. Journal of Empirical Finance 2017
Campbell, J., and L. Viceira. American Economic Review 2001
Candelon, B., and R. Moura. Journal of Financial Econometrics 2024
Carr, P., and L. Wu. Review of Finance 2023
Cheridito, P., D. Filipović, and R. Kimmel. Journal of Financial Economics 2007
Chernov, M., and P. Mueller. Journal of Financial Economics 2012
Choi, A., and K. Kang. Journal of Banking and Finance 2023
Christensen, J., and G. Rudebusch. Journal of Financial Econometrics 2015
Christensen, J., and G. Rudebusch. Review of Economics and Statistics 2019
Christensen, J., F. Diebold, and G. Rudebusch. Journal of Econometrics 2011
Christensen, J., J. Lopez, and G. Rudebusch. Journal of Money, Credit and Banking 2010
Christiansen, C. Journal of Banking and Finance 2005
Christiansen, C., and J. Lund. Working paper 2005
Cieslak, A., and P. Povala. Journal of Finance 2016
Cieslak, A., and P. Povala. Review of Financial Studies 2015
Cochrane, J., and M. Piazzesi. American Economic Review 2005
Cochrane, J., and M. Piazzesi. Working paper 2008
Cohen, B., P. Hordahl, and F. Xia. BIS working paper 2018
Collin Dufresne, P., and R. Goldstein. Journal of Finance 2002
Coroneo, L., D. Giannone, and M. Modugno. Journal of Business & Economic Statistics 2016
Creal, D., and J. Wu. Journal of Econometrics 2015
Dahlquist, M., and H. Hasseltoft. Journal of International Economics 2013
Dai, Q., and K. Singleton. Journal of Finance 2000
Dai, Q., and K. Singleton. Journal of Financial Economics 2002
D'Amico, S., D. Kim, and M. Wei. Journal of Financial and Quantitative Analysis 2018
Dewachter, H., and M. Lyrio. Journal of Money, Credit and Banking 2006
Dewachter, H., L. Iania, and M. Lyrio. Journal of Applied Econometrics 2012

Dewachter, H., M. Lyrio, and K. Maes. Journal of Applied Econometrics 2006
Diebold, F. and C. Li. Journal of Econometrics 2006
Diebold, F., G. Rudebusch and B. Aruoba. Journal of Econometrics 2006
Doh, T. Journal of Economic Dynamics and Control 2011
Doshi, H., K. Jacobs, and R. Liu. Journal of Financial Economics 2024
Duarte, J. Review of Financial Studies 2004
Duffee, G. Journal of Finance 2002
Duffee, G. Review of Financial Studies 2011
Duffee, G., and R. Stanton. Quarterly Journal of Finance 2012
Feldhutter, P. Quarterly Journal of Finance 2016
Feldhutter, P., C. Heyerdahl Larsen, and P. Illeditsch. Review of Finance 2018
Golinski, A., and P. Spencer. Journal of Financial Economics 2017
Gomez-Cram, R., and A. Yaron. Review of Financial Studies 2021
Grishchenko, O., and J. Huang. Journal of Fixed Income 2013
Gurkaynak, R., B. Sack, and J. Wright. Journal of Monetary Economics 2007
Hansen, A. Journal of Financial Econometrics 2023
Haubrich, J., G. Pennacchi, and P. Ritchken. Review of Financial Studies 2012
Hordahl, P., and O. Tristani. International Journal of Central Banking 2014
Imakubo, K., and J. Nakajima. BoJ working paper 2015
Jacobs, K., and L. Karoui. Journal of Financial Economics 2009
Joslin, S. Management Science 2017
Joslin, S., and A. Le. Management Science 2021
Joslin, S., and Y. Konchitchki. Journal of Financial Economics 2018
Joslin, S., K. Singleton, and H. Zhu. Review of Financial Studies 2011
Joslin, S., M. Priebsch, and K. Singleton. Journal of Finance 2014
Jotikasthira, C., A. Le, and C. Lundblad. Journal of Financial Economics 2015
Joyce, M., P. Lildholdt, and S. Sorensen. Journal of Banking and Finance 2010
Joyce, M., I. Kaminska, and P. Lildholdt. Review of Finance 2012
Kaminska, I., Z. Liu, J. Relleen, and E. Vangelista. Journal of Banking and Finance 2018
Kim, D., and A. Orphanides. Journal of Financial and Quantitative Analysis 2012
Kim, D., and J. Wright. FED working paper 2005
Kim, D., and K. Singleton. Journal of Econometrics 2012
Koijen, R., T. Nijman, and B. Werker. Review of Financial Studies 2009
Kopp, E., and P. Williams. IMF working paper 2018
Leippold, M., and F. Matthys. Review of Finance 2022
Lemke, W., and A. Vladu. ECB working paper 2017
Li, H. Journal of Financial Econometrics 2025
Liu, Y., and J. Wu. Journal of Financial Economics 2021
Malik, S., and A. Meldrum. Journal of Banking and Finance 2016
Moench, E., and S. Soofi-Siavash. Journal of Financial Economics 2022
Monfort, A., F. Pegoraro, J. Renne, and G. Roussellet. Journal of Econometrics 2017
Opschoor, D., and M. van der Wel. Journal of Business & Economic Statistics 2025
Pancost, A. Review of Asset Pricing Studies 2021
Rebonato, R., and V. Putyatin. Quantitative Finance 2018
Rudebusch, G., and T. Wu. Economic Journal 2008
Sack, B. FED working paper 2000
Wright, J. American Economic Review 2011
Wright, J. Economic Journal 2012
Wu, J., and D. Xia. Journal of Applied Econometrics 2020
Yin, X., and G. Yang. Journal of Empirical Finance 2024

Table A.2
Estimation Output

Panel A of the table reports the standard deviation of fitting errors in yields, forward rates and yield variances for the SV model estimated using data with maturities up to 25 years (SV_{25}). All values are expressed in basis points. Panel B contains the maximum likelihood estimates of the model parameters. The coefficients are ordered as $[v; x_1; x_2; x_3]$. Bootstrapped standard errors are calculated from Monte Carlo simulations: (*) indicates statistical significance at least at the 5% level.

Panel A: Fitting errors

Maturity (years)	2 - 5	6 - 9	10	11 - 14	15	16 - 19	20	21 - 24	25
St. dev. of errors on yields	13.3	5.5	5.6	5.2	3.8	3.6	4.7	4.5	2.3
St. dev. of errors on forward rates	23.9	22.7	16.9	19.7	26.2	32.9	30.4	22.6	36.9
St. dev. of errors on yield variance			0.09		0.05		0.05		0.06

Panel B: Estimated parameters

K

	$\kappa_{\cdot 1}$	$\kappa_{\cdot 2}$	$\kappa_{\cdot 3}$	$\kappa_{\cdot 4}$
$\kappa_{1\cdot}$	0.1585(∗)	0	0	0
$\kappa_{2\cdot}$	0.0049(∗)	0.0310(∗)	0.0976(∗)	0.0089(∗)
$\kappa_{3\cdot}$	−0.0024(∗)	0.0412(∗)	0.2086(∗)	0.0081(−)
$\kappa_{4\cdot}$	−0.0138(∗)	0.0642(−)	0.0063(∗)	0.0841(∗)

Θ

θ_1	0.3595(∗)
θ_2	0
θ_3	0
θ_4	0

Λ_1

	$\lambda_{\cdot 1}$	$\lambda_{\cdot 2}$	$\lambda_{\cdot 3}$	$\lambda_{\cdot 4}$
$\lambda_{1\cdot}$	0	0	0	0
$\lambda_{2\cdot}$	−0.1461(∗)	−0.5557(∗)	−0.0430(∗)	0.1167(∗)
$\lambda_{3\cdot}$	0.0137(−)	−0.2125(−)	−1.4403(∗)	0.9752(∗)
$\lambda_{4\cdot}$	0.0657(∗)	−0.0512(−)	0.0133(∗)	−0.3521(∗)

Λ_0

λ_{01}	−0.0079(∗)
λ_{02}	−0.0180(∗)
λ_{03}	0.0148(∗)
λ_{04}	0.0210(−)

Σ

	$\sigma_{\cdot 1}$	$\sigma_{\cdot 2}$	$\sigma_{\cdot 3}$	$\sigma_{\cdot 4}$
$\sigma_{1\cdot}$	0.0651(∗)	0	0	0
$\sigma_{2\cdot}$	0	0.0167(∗)	0	0
$\sigma_{3\cdot}$	0	0	0.0249(∗)	0
$\sigma_{4\cdot}$	0	0	0	0.1059(∗)

δ

δ_1	0
δ_2	1
δ_3	1
δ_4	0

δ_0
0.0644(∗)

Table A.3
Information in Forward Rates and Stochastic Volatility

This table reports the results of a Monte Carlo experiment addressing the additional information contained in forward rates and stochastic volatility for the estimation of risk premia and expected short rates. We generate simulated series for state variables, yields, forward rates, yield volatility, risk premia and expected short rates and calculate estimation errors as described in **Table 1**. The SV model is fitted using data up to 25 years but with different subsets of variables included in the measurement equations of the Kalman filter. In particular, Panel A shows the results when we include data on yields and yield volatilities, but exclude data on forward rates, Panel B when we include data on forward rates and yield volatilities but exclude data on yields, and, finally, Panel C when we include data on yields and forward rates but exclude data on yield volatilities. For each case, we compute the same statistics as contained in **Table 1**, with the statistics in columns (3-4) and (6-7) expressed as a ratio to the corresponding statistics in Panel A of **Table 1** (i.e., SV_{25}). The values in columns (2) and (5) are in basis points.

Panel A: SV model with maturities up to 25 years and no data on forward rates

	Risk premia			Expected short rates		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors Ratio wrt SV_{25} (3)	St. dev. of mean errors Ratio wrt SV_{25} (4)	Average of mean errors (5)	Median of st. dev. of errors Ratio wrt SV_{25} (6)	St. dev. of mean errors Ratio wrt SV_{25} (7)
2	-2.9	1.60	1.08	3.1	1.72	1.16
5	-2.5	1.61	1.10	2.8	1.72	1.21
10	2.2	1.54	1.21	-1.1	1.57	1.33
15	6.7	1.37	1.38	-4.5	1.35	1.34
20	9.2	1.25	1.55	-6.9	1.21	1.23
25	9.7	1.17	1.55	-8.3	1.20	1.15

Panel B: SV model with maturities up to 25 years and no data on yields

	Risk premia			Expected short rates		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors Ratio wrt SV_{25} (3)	St. dev. of mean errors Ratio wrt SV_{25} (4)	Average of mean errors (5)	Median of st. dev. of errors Ratio wrt SV_{25} (6)	St. dev. of mean errors Ratio wrt SV_{25} (7)
2	-3.4	1.43	1.07	3.2	1.45	1.05
5	-3.2	1.37	1.02	0.2	1.59	0.99
10	-0.1	1.22	0.98	-2.7	1.28	1.04
15	0.8	1.08	0.98	-4.2	1.08	1.06
20	-0.8	1.07	1.02	-5.3	0.95	1.00
25	-3.8	1.07	1.07	-6.2	1.02	0.96

Panel C: SV model with maturities up to 25 years and no data on yield volatility

	Risk premia			Expected short rates		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors Ratio wrt SV_{25} (3)	St. dev. of mean errors Ratio wrt SV_{25} (4)	Average of mean errors (5)	Median of st. dev. of errors Ratio wrt SV_{25} (6)	St. dev. of mean errors Ratio wrt SV_{25} (7)
2	1.8	2.53	1.13	-3.7	2.46	1.21
5	4.3	2.36	1.10	-8.0	2.31	1.23
10	4.9	1.86	1.14	-10.6	1.61	1.41
15	1.0	1.75	1.47	-11.7	0.96	1.49
20	-6.2	2.01	2.08	-12.6	0.87	1.41
25	-15.9	2.32	2.47	-13.3	1.15	1.33

Table A.4
Precision in Term Premia Estimation

This table reports the results of a Monte Carlo experiment addressing the additional information contained in long-maturity yields and stochastic volatility for the estimation of term premia and average expected short rates. We generate simulated series for state variables, yields, forward rates, yield volatility, term premia and average expected short rates and calculate estimation errors as described in **Table 1**. Panel A reports, for SV, (i) the average, over trials, of the mean errors; (ii) the median, over trials, of the standard deviation of the errors; and (iii) the standard deviation, over trials, of the mean error for estimates of term premia and average expected short rates with maturities 2, 5, 10, 15, 20, and 25 years. Panel B reports the same statistics obtained from CV, with the statistics (ii) and (iii) expressed as a ratio to the corresponding statistics in Panel A (i.e., SV). The values in Panel A and in columns (2) and (5) of Panel B are in basis points.

Panel A: SV model

	Term premium			Average expected short rate		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors (3)	St. dev. of mean errors (4)	Average of mean errors (5)	Median of st. dev. of errors (6)	St. dev. of mean errors (7)
2	0.3	2.9	3.7	-0.1	3.4	4.0
5	0.5	4.3	5.4	-0.2	4.6	5.6
10	0.5	5.4	6.4	-0.3	5.6	6.5
15	0.4	5.6	6.7	-0.3	5.7	6.8
20	0.3	5.7	7.1	-0.3	5.8	7.1
25	0.3	5.8	7.8	-0.2	5.9	7.8

Panel B: CV model

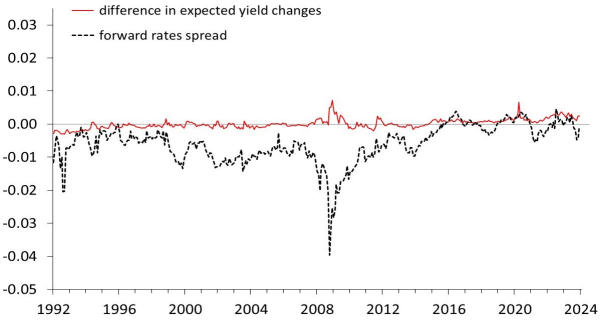
	Term premium			Average expected short rate		
Maturity (years) (1)	Average of mean errors (2)	Median of st. dev. of errors Ratio wrt SV (3)	St. dev. of mean errors Ratio wrt SV (4)	Average of mean errors (5)	Median of st. dev. of errors Ratio wrt SV (6)	St. dev. of mean errors Ratio wrt SV (7)
2	1.9	5.30	1.35	-4.3	4.56	1.21
5	0.5	5.51	1.32	-0.3	5.16	1.30
10	-3.0	5.35	1.29	2.9	5.07	1.26
15	-4.7	5.33	1.21	3.6	5.15	1.19
20	-2.7	4.86	1.11	2.9	4.80	1.11
25	-0.8	4.55	1.02	1.3	4.53	1.01

Figure A.1
Estimates with Alternative Datasets

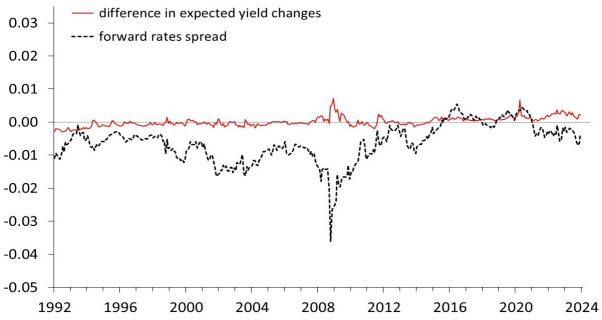
This figure reports estimates of some key variables obtained estimating our stochastic volatility (SV) models using data from Gurkaynak et al. (2007) (GSW) and Liu and Wu (2021) (LW). Panel A shows the forward spread between the 25- and 15-year maturity (Δf) along with the difference in expected yield changes for the same maturities (Δe). Panel B plots the 10-year risk premium estimated using GSW or LW data (vertical axis) vs. the same series obtained estimating the SV model with STRIPS (horizontal axis). The sample period is January 1992 to December 2023.

Panel A: Forward rate spread and difference in expected yield changes (Δe)

GSW

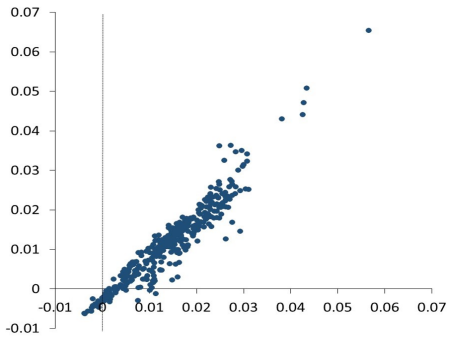


LW



Panel B: Plot of 10-year risk premium estimated with GSW or LWU vs. estimates with STRIPS

GSW



LW

