

The Structure of Risks in Equilibrium Affine Models of Bond Yields *

Anh Le and Kenneth J. Singleton[†]

April 1, 2013

Abstract

Many equilibrium term structure models (*ETSMs*) in which the state of the economy follows an affine process imply that the variation in expected excess returns on bond portfolio positions is fully spanned by the conditional variances of the state variables. We show that these two assumptions alone— an affine state process with conditional variances that span expected excess returns— are sufficient to econometrically identify the factors determining risk premiums in these *ETSMs* from data on the term structure of bond yields. Using this result we derive maximum likelihood estimates of the conditional variances of the state— the “quantities of risk”— and evaluate the goodness-of-fit of a large family of affine *ETSMs*. These assessments of fit are fully robust to the values of the parameters governing preferences and the evolution of the state, and to whether or not the economy is arbitrage free. Our findings suggest that, to be consistent with U.S. macroeconomic and Treasury yield data, affine *ETSMs* should have the features that: (i) inflation risk, and not long-run risks or variation in risk premiums arising from habit-based preferences, is a significant (and perhaps the dominant) risk underlying risk premiums in U.S. Treasury markets; and (ii) risks that are unspanned by bond yields have substantial explanatory power for risk premiums consistent with time-varying market prices of risks.

*We thank Christian Heyerdahl-Larsen (EFA discussant), Greg Duffee, Ivan Shaliastovich (AFA discussant), and seminar participants at the 2012 EFA meetings, 2013 AFA meetings, Bank of Canada, Imperial College London, Kenan-Flagler Business School, Warwick Business School, and University of Wisconsin-Madison for helpful comments.

[†]Le: Kenan-Flagler Business School, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 (email: anh_le@unc.edu); Singleton: Graduate School of Business, Stanford University, Stanford, CA 94305, and National Bureau of Economic Research (email: kenneths@stanford.edu).

1 Introduction

Equilibrium affine dynamic term structure models (*ETSMs*) imply that time-variation in expected excess returns on nominal risk-free bonds is driven by changes in the conditional variances of the models' real economic and inflation risks.¹ Within the long-run risks (*LRR*) models of [Bansal, Kiku, and Yaron \(2007\)](#) and [Bansal and Shaliastovich \(2012\)](#), risk premiums are exact affine functions of the time-varying quantities of priced risks. Risk premiums in the habit-based *ETSMs* of [Wachter \(2006\)](#) and [Le, Singleton, and Dai \(2010\)](#) are well approximated by affine functions of the variance of surplus consumption.

We argue in this paper that the term structure of bond yields is particularly revealing about the structure of time-varying risks in these models. The availability of a broad spectrum of maturities provides a market-based parsing of the effects of short- and long-lived risk factors on excess returns.² In fact we show that, *knowing only* the properties of *ETSMs* that (i) the state z_t follows an affine process and (ii) the conditional variances of z_t span expected excess returns, risk premiums are fully identified (can be extracted) from the cross-section of yields on default-free bonds. These extracted variances ζ_t^2 span the time-varying volatilities of long-run risks or surplus consumption and the volatility of inflation in any *ETSMs* that are nested within the presumed affine structure of z_t . Using this result, we compute maximum likelihood estimates of the time-series process ζ_t^2 and assess the goodness-of-fit of a large family of *ETSMs*. These assessments are valid regardless of the true values of the parameters governing preferences and the evolution of the state.

Our approach to model assessment blends the structure of the consumption risks embodied in typical models with habit formation or *LRR* with the focus on factor structures and market prices of risk in reduced-form, affine term structure models ([Dai and Singleton \(2003\)](#), [Piazzesi \(2010\)](#)). We highlight two robust features of recent *ETSMs*: First, they imply that the only sources of variation in expected excess returns on bonds are the time-varying volatilities of the risk factors underlying consumption growth and *LRR* or surplus consumption. Second *ETSMs*, as typically formulated, imply full spanning of the quantities of risk ζ_t^2 by bond yields. Together, these two features of *ETSMs* lead to expressions for ζ_t^2 in terms of the cross-section of bond yields that can be estimated with a high degree of precision.

The *ML* estimates of ζ^2 explain a striking high percentage of the variation in risk premiums in the US Treasury market. Indeed, consistent with most *ETSMs*, ζ^2 largely subsumes the forecast power of the term structure for excess bond returns. Moreover, variation in ζ_t^2 shows strong co-movement with the conditional covariances between yields and inflation, whereas there is much less co-movement with the conditional covariances between yields and real GDP growth, a proxy for real growth. This suggests that inflation risk is a significant

¹This follows, either exactly or to a good approximation, from the assumptions that the state of the economy follows an affine process and that agents' marginal rate of substitution in an exponential-affine form. See, for instance, [Eraker and Shaliastovich \(2008\)](#) and [Eraker \(2008\)](#) for discussions of equilibrium affine pricing models based on recursive preferences and long-run risks (*LRR*), and [Le, Singleton, and Dai \(2010\)](#) for a nonlinear model with habit formation that admits affine pricing.

²We focus on default-free debt and abstract from the segmentation effects explored by [Greenwood and Vayanos \(2010a\)](#), [Krishnamurthy and Vissing-Jorgensen \(2012\)](#), and [Vayanos and Vila \(2009\)](#), among others.

(and perhaps the dominant) risk underlying risk premiums in U.S. Treasury markets. These findings are robust to the structure of agents' preferences and, as such, they strengthen the complementary findings in B-S for a model with long-run risks.

Viewed within the broader family of affine dynamic term structure models, these results may seem surprising. The affine models with stochastic volatility examined in [Dai and Singleton \(2002\)](#) typically fail to produce risk premiums that replicate the failure of the expectations theory of the term structure. [Ahn, Dittmar, and Gallant \(2002\)](#) and [Collin-Dufresne, Goldstein, and Jones \(2009\)](#) conclude that the fitted volatilities of yields from their affine models bear limited resemblance to the time-varying volatilities estimated from semi-parametric time-series models. The novelty of our analysis is that we show how to extract the volatility factors, that according to *ETSMs* underlie variation in risk premiums, very precisely from the cross-section of bond yields. Not only does our ζ^2 subsume the information in the entire yield curve about time-varying yield volatility, it also spans virtually all of the information in yields about risk premiums.

While encouraging, this evidence represents only a partial vindication of affine *ETSMs* as they imply that no incremental information beyond the current yield curve should be useful for forecasting excess returns. To explore this strong proposition we focus on variables identified in the literature as having predictive power for excess returns after controlling for the information in bond yields.³ Not only are we able to identify conditioning information with substantial incremental predictive power for excess returns, but this information has virtually no forecast power for ζ^2 . Thus, on both fronts, this evidence challenges the structure of many *ETSMs*. Instead, the evidence suggests that there are time-varying market prices of risks induced by macro variables that are unspanned by bond yields.⁴

Central to our empirical analysis of *ETSMs* is a new dataset of yields on US Treasury zero-coupon bonds constructed from daily data on a large cross-section of coupon yields. The Fama-Bliss CRSP and [Gurkanyak, Sack, and Wright \(2007\)](#) (GSW) datasets are the most widely used. The former only includes maturities out to five years (a limitation in our view for studying risk premiums in equilibrium models), while a limitation of the latter is that the authors construct smoothed fitted yields using an extended [Nelson and Siegel \(1987\)](#) model. Using the extensive CRSP database on yields on individual Treasury coupon bonds, and applying the same filter to remove bonds that are illiquid or have embedded options, and the same Fama-Bliss bootstrap method as CRSP (see [Bliss \(1997\)](#)), we construct a consistent set of zero-coupon bond yields with maturities out to ten years over a sample period from 1972 through 2007. As we document subsequently, there is substantial predictive power of long-dated forward rates for excess returns relative to what is found with the smoothed GSW data. This feature of our data will play a key role in the subsequent assessment of the nature of the economic forces underlying risk premiums.

The remainder of this paper is organized as follows. [Section 2](#) provides further motivation

³For example, see [Joslin, Priebsch, and Singleton \(2013\)](#), [Cieslak and Povala \(2011\)](#), [Wright and Zhou \(2009\)](#), and [Huang and Shi \(2011\)](#) among many others.

⁴This evidence also runs counter to a large class of Gaussian macro-finance term structure models that imply that macro risk factors are spanned by the current yield curve. See [Joslin, Le, and Singleton \(2012\)](#) and the references therein.

of our analysis within the factor structures of a large family of *ETSMs*. Robust estimation of the volatility factors underlying variation in excess returns is explored in [Section 3](#). In [Section 4](#) we describe the construction of your yield data. The relationships between our extracted volatility factors, excess returns on bonds, and yield volatilities are explored in [Section 5](#). Finally, the economic importance and statistical significance for forecasts of excess returns of factors that are unspanned by bond yields are explored in [Section 6](#).

2 The Factor Structure of *ETSMs*

As in most of the extant literature we consider an endowment economy with $c_t \equiv \log C_t$, the logarithm of real per-capita consumption C_t , following the process

$$\Delta c_{t+1} = \mu_g + x_t + \eta_{t+1} \tag{1}$$

$$x_{t+1} = \rho x_t + e_{t+1}, \tag{2}$$

where x_t is the *LRR* factor and the mean-zero shocks (η_{t+1}, e_{t+1}) follow affine processes with conditional variances $Var_t[\eta_{t+1}] = \sigma_{ct}^2$ and $Var_t[e_{t+1}] = \sigma_{xt}^2$ (see, e.g., [Duffie, Pan, and Singleton \(2000\)](#)). This specification nests the consumption processes in the models of [Bansal, Kiku, and Yaron \(2007\)](#), [Bollerslev, Tauchen, and Zhou \(2009\)](#), [Drechsler and Yaron \(2011\)](#), and [Bansal and Shaliastovich \(2012\)](#); as well as the habit-based models of [Wachter \(2006\)](#) and [Le, Singleton, and Dai \(2010\)](#). The conditional variances of the shocks (η_{t+1}, e_{t+1}) may have a multi-dimensional factor structure, as for instance in [Bollerslev, Tauchen, and Zhou \(2009\)](#). These shocks may also embody jump components with state-dependent arrival intensities as in [Drechsler and Yaron \(2011\)](#). [Table 1](#) summarizes the features of these models that are particularly relevant to the goals of our analysis. N denotes the dimension of the set of priced risks z_t , and \mathcal{R} is the implied dimension of the factor-space for excess returns on bonds.

The logarithm of the kernel for pricing nominal bonds typically takes the form

$$m_{t+1} = \gamma_0 \log \delta - \gamma_1 \Delta c_{t+1} - \gamma_2 \varphi_{t+1} - \pi_{t+1}, \tag{3}$$

where π_{t+1} is the log of the inflation rate, $\log(P_{t+1}/P_t)$. In the models of habit formation with external habit level H_t , φ_{t+1} is the growth rate of the consumption surplus ratio $s_t = \log[(C_t - H_t)/C_t]$ and $\varphi_{t+1} = (s_{t+1} - s_t)$. [Wachter \(2006\)](#) and [Le et al. \(2010\)](#) examine models with two priced risks ($N = 2$), s_t and π_t , and with Δc_{t+1} conditionally perfectly correlated with s_t as in [Campbell and Cochrane \(1999\)](#).

In models with *LRR*, the real pricing kernel is given by [\(3\)](#) without the term $-\pi_{t+1}$ and $\varphi_{t+1} = r_{c,t+1}$, the one-period return on a claim to aggregate consumption flows. The models of [Bansal and Yaron \(2004\)](#) and [Bansal, Kiku, and Yaron \(2007\)](#) have three priced risks ($N = 3$) under the assumption that $\sigma_{ct}^2 = \sigma_{xt}^2$. [Drechsler and Yaron \(2011\)](#) include a stochastic drift to σ_{ct}^2 (denoted q_t in [Table 1](#)) and affine jumps in both σ_{ct}^2 and x_t giving $N = 6$. [Bollerslev et al. \(2009\)](#) allow for a time-varying conditional variance of σ_{ct}^2 (denoted q_t). Finally, [Bansal and Shaliastovich \(2012\)](#) assume that $(\Delta c_t, \pi_t)$ is a conditionally homoskedastic process, and

Recursive Preferences	LRR	N	Priced (z_t)	\mathcal{R}	RP	MPR
Bansal, Kiku, and Yaron (2007)	yes	3	$x_t, \sigma_{ct}^2, \bar{\pi}_t$	1	σ_{ct}^2	const
Bollerslev, Tauchen, and Zhou (2009)	no	3	$\sigma_{ct}^2, q_t, \bar{\pi}_t$	2	σ_{ct}^2, q_t	const
Drechsler and Yaron (2011)	yes	6	$x_t, \sigma_{ct}^2, q_t, \pi_t$ Jump Risks	1	σ_{ct}^2	const
Bansal and Shaliastovich (2012)	yes	4	$x_t, \sigma_{xt}^2, \sigma_{\bar{\pi}t}^2, \bar{\pi}_t$	2	$\sigma_{xt}^2, \sigma_{\bar{\pi}t}^2$	const
Habit Formation	LRR	N	Priced Risks	\mathcal{R}	RP	MPR
Wachter (2006)	no	2	$s_t, \bar{\pi}_t$	1	s_t	t-vary
Le, Singleton, and Dai (2010)	no	2	$s_t, \bar{\pi}_t$	1	s_t	t-vary

Table 1: Features of *ETSMS*. The column “LRR” indicates whether the model has *LRR*; N is the number of priced risks which are indicated in column four; \mathcal{R} is the model-implied dimension of expected excess returns on nominal bonds, and the factors driving these risk premiums are indicated in column six; and the column “MPR” indicates whether the market prices of risk are constant or time-varying (t-vary).

their sources of time-varying quantities of risk are σ_{xt}^2 and the conditional variance $\sigma_{\bar{\pi}t}^2$ of expected inflation ($\bar{\pi}_t = E_t[\pi_{t+1}]$), so $z'_t = (x_t, \sigma_{xt}^2, \sigma_{\bar{\pi}t}^2, \bar{\pi}_t)$ and $N = 4$.

The expected excess return over a horizon h on a τ -period zero-coupon bond with yield to maturity y_t^τ is

$$er_t^h(\tau) = -(\tau - h)E_t[y_{t+h}^{\tau-h}] + \tau y_t^\tau - y_t^h. \quad (4)$$

All but one of the *ETSMS*s summarized in Table 1 imply that y_t^τ is an affine function of z_t :

$$y_t^\tau = A(\tau) + B(\tau) \cdot z_t, \quad \forall \tau \geq 0. \quad (5)$$

The exception is Wachter’s model which does not admit affine pricing but, to a good approximation, Le et al. (2010)’s affine pricing model nests her model (see below). Accordingly, we take (5) as given in our subsequent analysis. Then, together, (4) and (5) imply that $er_t^h(\tau)$ is an affine function of z_t .

Common to all of these equilibrium models are the features that the dimensions (\mathcal{R}) of the expected excess returns or risk premiums are less than the dimensions of the risk factors underlying bond yields ($N > \mathcal{R}$), and time variation in risk premiums is induced entirely by variation in the quantities of risks listed in column *RP* of Table 1:

$$er_t^h(\tau) = \mathcal{A}(\tau, h) + \mathcal{B}(\tau, h) \cdot \zeta_t^2, \quad (6)$$

where ζ_t^2 is a strict subset of the state z_t comprised of the \mathcal{R} volatility factors. In the habit-based models of Wachter (2006) and Le et al. (2010) expected excess returns lie in a one-dimensional space determined by $\zeta_t^2 = s_t$. Similarly, among the *LRR* models, Drechsler and Yaron (2011) have the largest number of priced risks, but the structure of their model is such that expected excess returns on all bonds of all maturities are proportional to $\zeta_t^2 = \sigma_{ct}^2$.

These observations lead us to the following robust implication of many *ETSMs*:

RIETSM: *The dimensionality \mathcal{R} of the expected excess returns $er_t^h(\tau)$ is common for all horizons h and all bond maturities τ . Moreover, the set of risk factors ζ_t^2 underlying variation in the $er_t^h(\tau)$ is a subvector of the state z_t determining bond yields, and ζ_t^2 is the sole source of time-varying volatilities in bond yields. That is, the ζ_t^2 listed under “RP” in [Table 1](#) span the time-varying volatilities of bond yields.*

At the heart of *RIETSM* for *ETSMs* other than habit-based models is the assumption that the market prices of risk Λ for the risk factors z_t are state-independent (the weights on $t + 1$ variables in [\(3\)](#) are constants). If information other than ζ_t^2 is incrementally useful for forecasting excess returns, then either these *ETSMs* have omitted time-varying quantities of risks that are relevant in bond markets or the market prices Λ_t of the risks ζ_t^2 are time varying. Holding \mathcal{R} fixed, state-dependence of Λ_t could arise because the linearizations inherent in affine *ETSMs* leave out empirically relevant dimensions of risk⁵ or because of a fundamental mis-specification of the structure of preferences of bond investors.⁶

In habit-based models surplus consumption s_t is a source of a time-varying market price of consumption risk. In the model of [Le et al. \(2010\)](#), for instance, m_{t+1} has a non-affine structure and expected excess returns $er_t^1(\tau)$, for all τ , are approximately affine functions of s_t and $\sqrt{s_t}$. The accuracy of [\(6\)](#), now viewed as a linearization of the model-implied nonlinear interplay between the time-varying quantity and market price of s_t risk, is parameter-value dependent. However, when valued at their maximum likelihood estimates, virtually all of the variation in $er_t^1(\tau)$ is induced by s_t , so the affine approximation is very accurate. Since their model effectively nests Wachter’s model, we proceed under the assumption that [\(6\)](#) is a reliable approximation for all of the *ETSMs* summarized in [Table 1](#).

3 Robust Evaluation of the Constraint that Expected Excess Returns are Spanned by Volatility Factors

RIETSM is a powerful implication of *ETSMs*. When combined with the assumption that z_t follows an affine process, it allows us to extract a set of \mathcal{R} risk factors from the term structure of bond yields that fully span expected excess returns. Moreover, it leads to tests of goodness-of-fit of affine *ETSMs* that are robust to values of the parameters governing agents preferences and the distributions of the non-volatility factors impacting bond yields.

That the risk factors in *ETSMs* can be extracted from market returns has long figured prominently in the literature on pricing equity portfolios. For instance, [Bansal, Kiku, and Yaron \(2007\)](#), [Constantinides and Ghosh \(2009\)](#), and [Marakani \(2009\)](#) explore the fits of

⁵[Bansal and Yaron \(2004\)](#) and [Bansal and Shaliastovich \(2012\)](#), among others, argue that the linearizations within their *LRR* models are inconsequential for their empirical analyses.

⁶[Bonomo et al. \(2010\)](#), for instance, argue that replacing Kreps-Porteous preferences by preferences exhibiting disappointment aversion resolves some of the weaknesses of the [Bansal and Yaron \(2004\)](#) model with regard to matching the predictive power of dividend yields for consumption growth and excess returns on stocks. The pricing kernel implied by their model implicitly exhibits time-varying market prices of *LRR*.

two-factor models with *LRR* using the price-dividend ratio for the aggregate stock market and the nominal risk-free rate to extract a *LRR* factor from asset returns and consumption data. Our complementary analysis shows that the tight linkages in the cross-section among expected excess returns on bonds of different maturities identifies the time-varying quantities of risk underlying variation in risk premiums with a high degree of precision and without facing measurement issues associated with consumption.⁷

For this purpose we suppose that model assessment is based on a collection of J yields y_t ($J > N$), and we partition the state z_t into the $N - \mathcal{R}$ non-volatility factors ν_t and the \mathcal{R} factors ζ_t^2 governing the time-varying volatility in z_t . *ETSMs* imply that the yield on an n -period zero-coupon bond has the affine representation

$$y_{nt} = (A_n + B_n\nu_t + C_n\zeta_t^2)/n, \quad (7)$$

and we let (A, B, C) denote the stacked up loadings on $(1, \nu_t, \zeta_t^2)$ from (7) for y_t . Additionally, we assume that there exists an $N \times J$ full-rank weight matrix W such that the N portfolios of yields $\mathcal{P}_t = Wy_t$ are measured without error. We provide an in depth justification for this approach as part of our empirical analysis in [Section 5](#).

To show that the structure of *RIESTM* leads to the precise extraction of ζ_t^2 from y_t we proceed in three steps. First, the counterparts of (7) for the N portfolios \mathcal{P}_t are inverted to express ζ_t^2 as an affine function \mathcal{P}_t . The spanning of z_t by \mathcal{P} is typically assumed in dynamic term structure models ([Joslin, Le, and Singleton \(2012\)](#)); it is not specific to *ETSMs*. Second, we show that *ETSMs*, through their property *RIESTM*, imply a strong restriction on how the B_n 's are linked across maturities n . Finally, we argue that (and provide intuition for why) the weights that determine ζ_t^2 as a linear combination of \mathcal{P}_t can be estimated very precisely from the cross-section of the yield curve. These derivations are subsequently used to compute maximum likelihood estimates of the unknown weights for ζ_t^2 .

Step 1: an immediate implication of (7) is that, outside degenerate cases discussed in [Section 6](#), ζ_t^2 is fully spanned by bond yields. Partition the loading matrix W determining the yield portfolios \mathcal{P} into W_ν ($(N - \mathcal{R}) \times J$) and W_ζ ($\mathcal{R} \times J$) and let $\mathcal{P}_{\nu t} = W_\nu y_t$ and $\mathcal{P}_{\zeta t} = W_\zeta y_t$. Constructing $W_\nu y_t$ using (7) and solving for ν_t gives

$$\nu_t = (W_\nu B)^{-1}(\mathcal{P}_{\nu t} - W_\nu(A + C\zeta_t^2)). \quad (8)$$

Substituting back into (7) leads to an expression for y_t in terms of $\mathcal{P}_{\nu t}$ and ζ_t^2 :

$$y_t = B_{\mathcal{P}}\mathcal{P}_{\nu t} + \phi_\nu(A + C\zeta_t^2), \quad (9)$$

where $B_{\mathcal{P}} = B(W_\nu B)^{-1}$ and $\phi_\nu = I_J - B_{\mathcal{P}}W_\nu$. Finally, premultiplying both sides of (9) by W_ζ and solving for ζ_t^2 gives

$$\zeta_t^2 = (W_\zeta \phi_\nu C)^{-1} \left(\mathcal{P}_{\zeta t} - W_\zeta B_{\mathcal{P}} \mathcal{P}_{\nu t} - W_\zeta \phi_\nu A \right). \quad (10)$$

⁷Much of the debate about the goodness-of-fit of *LRR* models to macroeconomic data has focused on measurement issues associated with consumption. See, for instance, the discussions in [Bansal, Kiku, and Yaron \(2007\)](#) and [Beeler and Campbell \(2009\)](#).

The assumption of full spanning of ς_t^2 by y_t ensures that the leading matrix $W_\varsigma \phi_\nu C$ in (10) is invertible. Thus, up to an affine transformation, ς_t^2 is determined by $\mathcal{P}_{\varsigma t} - W_\varsigma B_{\mathcal{P}} \mathcal{P}_{\nu t}$. Since W_ς is known, it remains to determine the weights $B_{\mathcal{P}}$ in order to extract (a nonsingular rotation of) ς_t^2 from y_t .

In Step 2 we exploit the restrictions on the B_n (and hence B) implied by *RIESTM*. Let

$$E_t[\varsigma_{t+1}^2] = \rho \varsigma_t^2 \quad \text{and} \quad E_t[\nu_{t+1}] = K_{1\varsigma} \varsigma_t^2 + K_{1\nu} \nu_t \quad (11)$$

represent the conditional means of ς_{t+1}^2 and ν_{t+1} (ignoring constants) under the data generating process.⁸ Substituting the factor representation (7) for yields into (4) leads to the following generic (under affine pricing) expression for expected excess returns:

$$er_t^1(n) = (B_n - B_{n-1} K_{1\nu} - B_1) \nu_t + (C_n - C_{n-1} \rho - B_{n-1} K_{1\varsigma} - C_1) \varsigma_t^2. \quad (12)$$

Now when only ς_t^2 predicts excess returns, as under *RIETSM*, the loadings on ν_t must be zero or, equivalently, B_n must satisfy the recursion

$$B_n = B_{n-1} K_{1\nu} + B_1. \quad (13)$$

Moreover, (13) implies that $B_n = B_{n-h} K_{1\nu}^h + B_h$ for any horizon h and, therefore, from (12) it follows that the h -period expected excess return of an n -period bond ($er_t^h(n)$) depends only on ς_t^2 ; that is, ς_t^2 completely spans $er_t^h(n)$, for all $h > 0$.

With this constraint enforced, expression (12) simplifies to

$$er_t^1(n) = (C_n - C_{n-1} \rho - B_{n-1} K_{1\varsigma} - C_1) \varsigma_t^2. \quad (14)$$

Notice that $er_t^1(n)$ does not require (or imply) arbitrage-free pricing (though the key ingredient (6) underlying our derivation of (14) is typically derived in a no-arbitrage setting). Nor was it necessary in deriving (14) to specify the ς_t^2 process, beyond that it is affine. Further, anticipating our analysis of unspanned risks in Section 6, even if ς_t^2 is completely unspanned by yields ($C_n \equiv 0$ for all n), (13) and (14) still hold as implications of *RIETSM*.

The restrictions on (13) can be exploited to extract precise estimates of the volatility factors ς_t^2 from the cross-sectional information in the yield curve. Here we highlight a very convenient representation of $B_{\mathcal{P}}$ in terms of the eigenvalues of $K_{1\nu}$. Without loss of generality, we can rotate ν so that B_1 is a vector of ones and $K_{1\nu}$ has the Jordan form. Whence B (and therefore $B_{\mathcal{P}}$), are completely determined by the eigenvalues of $K_{1\nu}$. In the context of risk-neutral pricing, these eigenvalues are the persistence parameters $\lambda^{\mathbb{Q}}$ that [Joslin, Singleton, and Zhu \(2011\)](#) show are estimable with considerable precision from the cross-section of bond yields. Similarly, we anticipate (and subsequently confirm) that our identification strategy will reveal the time variation in and predictive content of ς_t^2 very precisely through $\mathcal{P}_{\varsigma t} - W_\varsigma B_{\mathcal{P}} \mathcal{P}_{\nu t}$.

To help build intuition for this claim, recall that in arbitrage-free affine term structure models the loadings in (7) are governed by the parameters of the pricing (risk-neutral)

⁸To maintain admissibility, the conditional means of ς_{t+1}^2 cannot be dependent on the non-volatility factor ν_t as its support is typically $\mathbb{R}^{N-\mathcal{R}}$.

feedback matrix $K_{1z}^{\mathbb{Q}}$ and they satisfy recursions of the same form as (13).⁹ Moreover, $K_{1z}^{\mathbb{Q}}$ is typically estimated very precisely from the cross-section of bond yields. Now suppose in our setting that there is a pricing measure \mathbb{Q} (not necessarily equivalent to the risk-neutral measure since we do not enforce no-arbitrage) that gives rise to (13). Then $K_{1\nu}$ governs the feedback of ν_t under both the historical and \mathbb{Q} measures. This implicit assumption of *ETMSs* ensures that ν_t does not impact risk premiums (the market price of ν risk is constant) and (heuristically) that estimates of $K_{1\nu}$ inherit the precision of estimates of $K_{1\nu}^{\mathbb{Q}}$.

Up to this point our derivations do not exploit the fact that ζ^2 is a volatility process, beyond the autonomous structure of its conditional mean (11). To proceed with estimation and, most importantly, to ensure that ζ_t^2 is interpretable as the conditional volatility of z_t , we adopt a parametric affine model for the conditional distribution of z . Specifically, we assume that ζ^2 follows a multivariate autoregressive gamma (*ARG*) process (Gourieroux and Jasiak (2006), Le, Singleton, and Dai (2010)), and the remaining $N - \mathcal{R}$ factors ν_t are Gaussian conditional on ζ_t^2 :

$$\zeta_{t+1}^2 | \zeta_t^2 \sim ARG(\rho, c, \gamma), \quad (15)$$

$$\nu_{t+1} - \phi_\zeta \zeta_{t+1}^2 | z_t \sim N \left(K_0 + K_{1\zeta} \zeta_t^2 + K_{1\nu} \nu_t, H_{0\nu} + \sum_{i=1}^M H_{i\nu} \zeta_{it}^2 \right). \quad (16)$$

See Appendix B for a more detailed construction of the density $f(\zeta_{t+1}^2 | \zeta_t^2)$ and the definition of the parameters (ρ, c, γ) .

The ARG distribution is the discrete-time counterpart to the multivariate square-root diffusion (the $A_{\mathcal{R}}(\mathcal{R})$ process of Dai and Singleton (2000)). Several recent *ETSMs* (e.g., Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011)) assume that the innovation in ζ^2 is Gaussian, a specification that is logically inconsistent with the non-negativity of conditional variances. In contrast, by adopting an ARG process¹⁰ we ensure that ζ_t^2 is an autonomous non-negative process governing the conditional variance of z_{t+1} and that the conditional mean of ζ_{t+1}^2 is affine in ζ_t^2 ($E_t[\zeta_{t+1}^2] = \rho \zeta_t^2 + \gamma c$.) The positive semi-definite matrices $H_{0\nu}$ and $H_{i\nu}$ ($i = 1, \dots, M$) govern the time-varying volatility of ν .

For econometric identification we normalize z so that, without loss of generality, ν_t and ζ_t^2 are conditionally independent ($\phi_\zeta = 0$). Additionally, the intercepts K_0 in (16) are normalized to zero; and B_1 , the loadings on ν_t for the yield on a one-period bond, are normalized to the row vector of ones. Further, we rotate ν so that $K_{1\nu}$ has the Jordan form; and c is fixed at $\frac{1}{2}\Delta t$. Finally, to prevent ζ_t^2 from being absorbed at zero we require that $\gamma \geq 1$. The joint density of (ν_t, ζ_t^2) given by (15) and (16) gives the conditional density of $(\mathcal{P}_{\nu t}, \mathcal{P}_{\zeta t})$, after a Jacobian adjustment implied by (8) and (10). This is one of many equivalent normalization schemes that ensure that the extracted ζ_t^2 span the $er_t^h(n)$.

⁹More precisely, $K_{1z}^{\mathbb{Q}}$ is such that: $E_t^{\mathbb{Q}}[z_{t+1}] = K_{1z}^{\mathbb{Q}} z_t$, ignoring constants. Due to admissibility requirements in standard no-arbitrage models, the volatility factors ζ_{t+1}^2 cannot be forecastable by the non-volatility factors ν_t under the \mathbb{Q} and \mathbb{P} measures. It immediately follows that the recursions for the loadings on the non-volatility factors (B_n) cannot be dependent on the volatility factors' loadings (C_{n-1}).

¹⁰The following derivations are easily modified for other choices of non-negative affine distributions for ζ_t^2 .

Having characterized the joint density of the N yield portfolios $(\mathcal{P}_{\nu t}, \mathcal{P}_{\zeta t})$, the construction of the likelihood function for all J yields y_t is completed by including an additional $J - N$ yield portfolios $\mathcal{P}_{et} = W^e y_t$, with W^e linearly independent of W , that are measured with additive errors. For simplicity we assume that these portfolios are observed with *i.i.d.* errors with zero means and common variance:

$$\mathcal{P}_{et} - W^e B_P \mathcal{P}_{\nu t} - W^e \phi_\nu (A + C \zeta_t^2) \sim N(0, \sigma_\epsilon^2 I_{J-N}). \quad (17)$$

The parameter set is $(A, C, \rho, \gamma, K_{1\zeta}, K_{1\nu}, H_{0\nu}, H_{i\nu}, \sigma_\epsilon)$. Quasi Maximum Likelihood (QML) estimates of these parameters will serve as inputs into goodness-of-fit tests of affine *ETSMs* that are robust to specification of the remaining economic structure of these *ETSMs*.¹¹

Model assessment can be based on any set of \mathcal{R} volatility factors that are a basis for the \mathcal{R} -dimensional space of admissible volatilities ζ_t^2 . When the volatility factors ζ_t^2 are fully spanned by bond yields we can, without loss of generality, base our empirical analysis on the spanning vector

$$V_t \equiv \mathcal{P}_{\zeta t} - W_\zeta B_P \mathcal{P}_{xt} = \beta \mathcal{P}_t. \quad (18)$$

Since $\mathcal{P}_{\zeta t}$ corresponds to the first M entries of \mathcal{P}_t , the first $M \times M$ block of the $M \times N$ matrix β is the identity matrix. The fitted V_t we obtain are not literally the volatility factors in say an *ETSM* with *LRR*, as the structural volatility factors are not identified under the minimal structure we have imposed on the yield curve. Rather, *ETSMs* imply that V_t is a non-singular rotation of ζ_t^2 and, hence, that projections of yields or excess returns onto our fitted V_t must be identical to those onto ζ_t^2 .

Equipped with *ML* estimates of the parameters β and the resulting fitted V_t , we proceed to examine the following implications of *ETSMs*. First, we examine whether V_t fully captures the information about volatility spanned by yields. That is, we investigate whether yields have incremental forecasting power for the conditional variances of bond yields beyond the information in V_t . The answer should be no, since *ETSMs* imply that yield volatilities are fully determined by V_t . Second, and similarly, after conditioning on V_t , information in the yield curve should have no predictive content for realized excess returns. Third, all of the *ETSMs* we have discussed also imply that the conditional variances of inflation and consumption growth are fully determined by ζ_t^2 . As such, after conditioning on V_t , information in the yield curve should have no predictive content for the conditional volatilities of these macro variables.

We recognize, and indeed emphasize, that this strategy for extracting ζ_t^2 presumes that the time-varying volatilities of bond yields are spanned by the yields themselves, y_t . Outside of knife-edge cases, this is a direct implication of *ETSMs* and therefore is an assumption that we maintain in our initial econometric analysis of risk premiums. However, as noted in the introduction, the literature on reduced-form affine term structure models casts doubt on the validity of this assumption. There is substantial evidence of unspanned (by yields) volatility in bond markets (e.g., [Collin-Dufresne, Goldstein, and Jones \(2009\)](#)). Additionally,

¹¹In our setting the volatility process is likely to be tightly determined by the cross-section of bond yields; essentially, the likelihood function is searching for the two \mathbb{Q} eigenvalues that describe volatility. Thus, our QML estimates are likely to be robust to a wide variety of affine specifications of the volatility process.

using high-frequency data to estimate time-varying second moment and jump components of bond yields, [Wright and Zhou \(2009\)](#) and [Andersen and Benzoni \(2010\)](#), among others, provide evidence that there is information about bond excess returns embodied in yields that is not well captured by monthly data on the shape of the yield curve. In [Section 5](#) we maintain the spanning assumption enforced by *ETSMS* to explore what is learned about risk premiums under *RIETSM*. We then relax this assumption in [Section 6](#) where we explore the implications of unspanned yield volatility for risk premiums.

4 The Information in Long-Maturity Yields

Prior to embarking on this empirical analysis, we briefly discuss the importance of using long-maturity bond yields in assessing the dynamic properties of risk premiums in Treasury markets. We raise this issue, because our prior is that the choice of splines underlying the construction of the zero-coupon bond yields typically used in the analysis of *ETSMS* may well matter for the properties of the fitted risk premiums. [Cochrane and Piazzesi \(2005\)](#), [Duffee \(2011\)](#), and [Bansal and Shaliastovich \(2012\)](#), among others, chose to focus on maturities out to five years when studying risk premiums.

To shed light on whether longer maturity yields contain non-redundant information about risk premiums, we estimate the unconstrained linear projections of realized excess returns $xr_{t+h}^h(\tau)$ onto yield-curve information at date t . We set h to six or twelve months and consider three choices of conditioning information: (i) yields on maturities from one to five years ($J = 5$); (ii) yields on maturities from one to five years plus the seven-, eight-, and ten-years ($J = 8$); and (iii) yields on maturities from one to ten years ($J = 10$). Of interest is the incremental predictive power of the latter cases ($J = 8$ and $J = 10$) relative to when only yields up to five years are used in prediction.

Using the CRSP treasury bond data and similar algorithms as described by [Fama and Bliss \(1987\)](#), we construct a consistent set of “Fama-Bliss” zero yields out to ten years to maturity through to the end of 2007 (the UFB dataset). Our sample period starts in January, 1984 after the abandonment of the monetary policy experiment between 1979 and 1982. To avoid the extreme market conditions of the ongoing crisis, we end our sample in December, 2007. BIC ([Schwarz \(1978\)](#)) scores are used to select the preferred forecasting model among these three specifications.

The adjusted R^2 's from the projections are reported in [Table 2](#), along with the probability values (*pvals*) of the chi-square tests of joint significance of the yields with maturities beyond five years. From these results it is clear that long maturity yields contain substantial extra predictive power over and above the first five yields. For example, for the annual holding period and the cross-sectional average of the excess returns ($\text{mean}(xr)$), the adjusted R^2 increases by 4% (6%) to 42% (44%) by including three (five) longer-maturity yields. Moreover, for both this average and the individual excess returns, the BIC selection criterion always chooses information sets that include the long-maturity yields.

We explored the robustness of these findings in two ways. First, we re-estimated the projections with yields back to March, 1972. Earlier data was discarded owing to the relative

Panel A: Six-Month Holding Period								
	R(1-5)		R(1-5,7,8,10)			R(1-10)		
	Adj R^2	BIC	Adj R^2	pval	BIC	Adj R^2	pval	BIC
mean(xr)	0.26	-10.51	0.29	0.01	-10.62	0.31	0.00	-10.63*
xr(3)	0.28	-13.56	0.32	0.00	-13.74*	0.32	0.00	-13.66
xr(5)	0.25	-10.69	0.30	0.00	-10.88*	0.31	0.00	-10.83
xr(7)	0.25	-9.04	0.27	0.04	-9.09*	0.29	0.00	-9.09
xr(10)	0.20	-7.12	0.26	0.02	-7.23	0.29	0.00	-7.31*

Panel B: Twelve-Month Holding Period								
	R(1-5)		R(1-5,7,8,10)			R(1-10)		
	Adj R^2	BIC	Adj R^2	pval	BIC	Adj R^2	pval	BIC
mean(xr)	0.38	-8.55	0.42	0.02	-8.79	0.44	0.00	-8.81*
xr(3)	0.35	-12.25	0.40	0.02	-12.60*	0.40	0.04	-12.54
xr(5)	0.36	-9.21	0.42	0.00	-9.58*	0.43	0.00	-9.55
xr(7)	0.38	-7.60	0.42	0.05	-7.82	0.43	0.04	-7.82*
xr(10)	0.36	-5.69	0.41	0.02	-5.89	0.43	0.00	-5.95*

Table 2: Adjusted R^2 (Adj R^2) from regressing six- (Panel A) and twelve-month (Panel B) excess returns $xr(n)$ of bonds with n years to maturity on yields with 1-5 years to maturity (R(1-5)), 1-10 years to maturity (R(1-10)), and 1-5, 7,8,10 years to maturity (R(1-5,7,8,10)). Yields are extracted from the UFB dataset. pval’s are for the joint significance tests of the loadings on longer-than-5-year maturities. The models chosen by the BIC scores (divided by 100) are indicated by an “*.”

sparseness of long-maturity bonds. Using this longer data set there is even stronger evidence that long-maturity yields have predictive content for excess returns in Treasury markets.

Second, we re-estimated the projections using the GSW dataset from the Federal Reserve’s website.¹² For the long sample, the long-term GSW yields embodied much less incremental forecasting power than our constructed UFB data. For the shorter sample period the patterns were comparable, but with GSW data the BIC criterion selected the specification $J = 5$ for both the mean(xr) and the excess returns on the seven- through ten-year bonds.

Based on this evidence we conclude that long-dated yields are informative about risk premiums and that the prior literature may have overlooked this information owing to the highly smoothed forward rates implicit in spline used to construct the GSW data. The weights on forward rates in the projections with our UFB data are displayed in [Figure 1](#). Clearly visible is the “tent-shape” pattern of loadings for the first five forward rates documented by [Cochrane and Piazzesi \(2005\)](#). The pattern of these loadings is essentially unchanged when all ten forward rates are included as predictors. Interestingly, the loadings on the long-maturity forwards for years six through ten form an inverted “tent-shape” pattern which is also robust to whether the first five forward rates are included or not.

With this evidence in mind, we proceed to evaluate the goodness-of-fit of the robust

¹²<http://www.federalreserve.gov/Pubs/feds/2006/200628/200628abs.html>

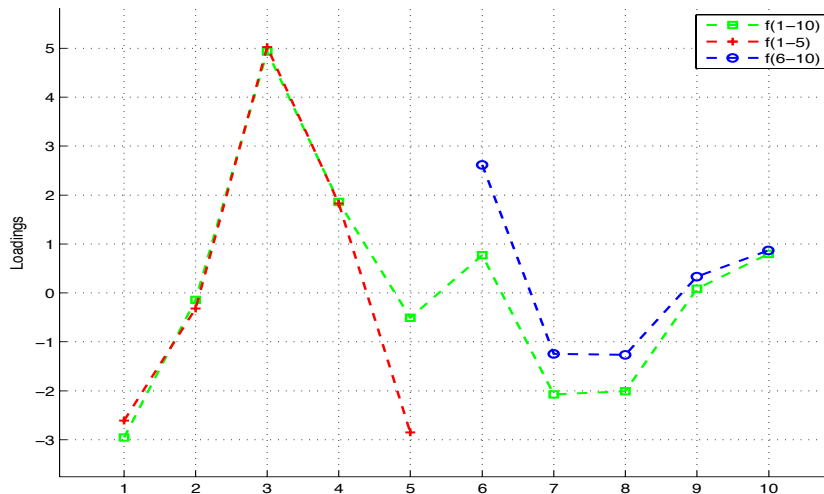


Figure 1: Loadings from the projections of $mean(xr_{t+h})$ with $h = 12$ (annual holding period) onto the first five one year forward rates (f(1-5)), the first ten one-year forward rates (f(1-10)), and the set of forward rates from 5 years to 6 years, 6 years to 7 years, ... 9 years to 10 years (f(6-10)). The UFB data are used to construct forward rates and excess returns.

features *RIETSM* of *ETSMs* using the historical bond yields $R(1 - 5, 7, 8, 10)$ constructed from the CRSP data.

5 Volatility Factors and Expected Excess Returns

With the models summarized in Table 1 in mind, we explore the nature of risk premiums in models with $N = 4$ risk factors governed by an affine process with $\mathcal{R} = 2$ sources of time-varying conditional variances. Expanding the number of such risks will add “dimension” (and thereby flexibility) to risk premiums. Yet, since these models have been put forth as successful representations of observed variation in expected excess returns, we proceed to evaluate their goodness-of-fit taking as given the assumed dimensions of the priced risks and the implied structure of risk premiums. Moreover, as documented in Duffee (2011) and Joslin, Singleton, and Zhu (2011), models of the term structure with $N > 4$ are potentially over-parameterized and this can severely distort the model-implied risk premiums.

5.1 Extracting the Volatility Factors From the Current Yield Curve

While the information sets generated by z_t and any N linearly independent yield portfolios \mathcal{P}_t are (according to *ETMSs*) theoretically identical, there is the issue in practice of accurate measurement of yields (pricing of bonds). For instance, even though in theory z_t is spanned by y_t , the sample projections of realized excess returns onto the information set generated by the

observed yields y_t^o are in general consistent estimators of their true theoretical counterparts only when y_t^o is priced (nearly) perfectly by the *ETSMs*. It is now standard practice to accommodate measurement errors on all bond yields and to using filtering in estimation of macro-finance *DTSMs*. The errors $y_t^o - y_t$ can be large in macro-finance *DTSMs*, especially when N is small (Joslin, Le, and Singleton (2012)).

Fortunately for our purposes the diversification that comes from using “portfolios” of yields, and in particular from setting \mathcal{P} to the first N *PCs* of bond yields, substantially mitigates these measurement issues. Joslin et al. (2012) show that the Kalman filter estimates of the model-implied low-order *PCs* in reduced-form *DTSMs* are (nearly) identical to their observed counterparts, even when the model-implied pricing errors $y_t^o - y_t$ become large. With this evidence in mind, we exploit the theoretical equivalence of the information sets generated by z_t and \mathcal{P}_t and conduct our analysis using the *PCs* \mathcal{P}_t . Proceeding under the assumption of no measurement errors on \mathcal{P}_t amounts to holding *ETSMs* to the same (high) standards of fit as for reduced-form arbitrage-free models.

Setting $N = 4$, the joint distribution of \mathcal{P}_t is determined by the joint density (15) - (16) along with the relevant Jacobian, as described in Section 3. The remaining $J - 4$ *PCs* of the J yields y_t , \mathcal{P}_{et} , are assumed to be priced up to additive errors according to (17). With these distributions in hand, and after imposing normalizations, we estimate the parameters by quasi-maximum likelihood (*QML*). The resulting estimates of the free parameters in the 2×4 matrix β that determines the volatility-spanning vector V_t in (18) are

$$\beta = \begin{bmatrix} 1 & 0 & 10.48 & 58.66 \\ & & (1.599) & (0.145) \\ 0 & 1 & -0.780 & -14.87 \\ & & (0.688) & (1.430) \end{bmatrix}, \quad (19)$$

where robust standard errors are given in parentheses. With one exception, all parameters are estimated with considerable precision. This was anticipated owing to the fact that the last 2×2 block of β is fully determined by the Jordan form of $K_{1\nu}$ which, in turn, is identified primarily from the cross-sectional restrictions in (13). We stress that our ability to exploit cross-sectional information is key to this precision. By way of contrast, the loadings from the time-series projection of squared residuals (from the projection of $PC1_{t+1}$ onto \mathcal{P}_t) onto \mathcal{P}_t are estimated much less precisely. Thus, our approach to extracting V_t is both conceptually and practically very different than those typically pursued in the literature on *ETSMs*.

Since V_t is one among an infinity of spanning vectors for the volatilities in any affine model with $N = 4$ and $\mathcal{R} = 2$ that imposes *RIETSM*, there is not a natural economic interpretation of the components of V_t . More informative are the correlations between V_t and the *PCs* of bond yields. Projecting each of the first four *PCs* of bond yields onto V_t gives the adjusted R^2 s of (0.72, 0.57, 0.05, 0.64). It follows that the information in V_t is highly correlated with the first, second, and fourth *PCs* of yields.

The high correlation between V_t and *PC4* is intriguing as *PC4* shows substantial correlation with the forward rate factor constructed by Cochrane and Piazzesi (2005) and also with real economic growth (Joslin et al. (2013)). Both of these variables are known to have strong

	$H = 2$			$H = 4$			$H = 6$		
	Adj. R^2			Adj. R^2			Adj. R^2		
	4 PCs	V_t	pval	4 PCs	V_t	pval	4 PCs	V_t	pval
PC1	0.188***	0.184***	0.750	0.302***	0.288***	0.532	0.366***	0.353***	0.644
PC2	0.073*	0.047	0.653	0.130	0.096	0.675	0.154*	0.143	0.912
PC3	0.088***	0.051	0.142	0.186**	0.141	0.323	0.242**	0.200*	0.376
PC4	0.268***	0.255***	0.612	0.388**	0.379***	0.798	0.442**	0.429***	0.654

Table 3: Comparison of projections of squared forecast errors, obtained from a $VAR(1)$ model of \mathcal{P}_{t+1} , onto \mathcal{P}_t versus V_t . “pval” is the robust probability value for the chi-square test of the null hypothesis $\beta_{\mathcal{P}} = \beta_V \beta$. The superscripts (*, **, ***) denote the level of significance (at 10%, 5%, 1%, respectively) of the χ^2 test of joint significance of the slope coefficients.

predictive power for excess bond returns. We turn next to an investigation of whether V_t encompasses the information in \mathcal{P}_t about yield volatility and excess returns.

5.2 Does V_t Encompass the Information in \mathcal{P}_t About Risk?

When ζ_t^2 is spanned by y_t , the volatility factors V_t span the conditional variances of the state z_t . To assess the empirical support for this implication of *RIETSM*, we let ϵ_t denote the error in forecasting \mathcal{P}_t based on information at time $t - 1$ and we examine the projections

$$E_t \left(\sum_{j=1}^H \epsilon_{t+j}^2 \right) = \text{constant} + \beta_{\mathcal{P}} \mathcal{P}_t \quad \text{and} \quad E_t \left(\sum_{j=1}^H \epsilon_{t+j}^2 \right) = \text{constant} + \beta_V V_t, \quad (20)$$

for $H = 2, 4, 6$ (in months). According to the risk-structure of *ETSMs*, \mathcal{P}_t and V_t should have the same explanatory power for the squared forecast errors. Additionally, given that $V_t = \beta \mathcal{P}_t$, *ETSMs* imply the constraint $H_0 : \beta_{\mathcal{P}} = \beta_V \beta$.

As a first approach to testing these constraints we use the models’ assumption that \mathcal{P}_t follows a first-order vector-autoregression ($VAR(1)$) under \mathbb{P} to obtain consistent estimates of the forecasting errors ϵ_t . The squared fitted residuals are then projected onto \mathcal{P}_t or V_t . In [Table 3](#) we report the adjusted R^2 statistics of the two sets of projections in (20) for the forecast errors associated with each of the first four *PCs* comprising \mathcal{P}_t . Across all horizons H and all *PCs*, V_t has virtually the same forecasting power as \mathcal{P}_t . For example, for $H = 2$ and *PC1*, the adjusted R^2 from conditioning on \mathcal{P}_t (V_t) is 18.8% (18.4%). To formally evaluate the differences in fits, we conduct χ^2 tests of $H_0 : \beta_{\mathcal{P}} = \beta_V \beta$.¹³ The probability values (“pvals”) confirm that the small differences in R^2 s are statistically insignificant.

For comparison, the errors in forecasting individual yields are also constructed using the Blue-Chip financial survey forecasts (BCFF) of bond yields constructed by Wolters Kluwer.

¹³The probability values of this test reported in [Table 3](#) are adjusted to account for the use of first-stage estimates of β as well as ϵ_t in estimating the regressions in (20). To account for the serial correlation of errors, we use the Newey-West estimates of the large-sample variance matrix with twelve lags.

Survey forecasts should embody at least as much information as the shape of the yield curve (\mathcal{P}_t) and, hence, the construction of yield-forecast errors ϵ_t from the BCFF data mitigates misspecification owing to omitted information in projections onto yield PC s alone.¹⁴

We cannot use the survey forecasts directly to construct ϵ_t , because they are forecasts of three-month moving averages of yields.¹⁵ Let

$$Q_{h,t}^{y(i)} = E_t^{(i)} [y_{t+h} + y_{t+h+1} + y_{t+h+2}]$$

denote the h -month ahead forecast of average yields formed by the i^{th} forecaster, and $Q_{h,t}^y = \text{median}(Q_{h,t}^{y,i})$ be the corresponding cross-sectional median forecast. We use the median as opposed to the mean forecast to minimize possible effects of outliers. For each horizon h , $Q_{h-1,t+1}^y - Q_{h,t}^y$ is the revision due to new information arriving within month $t+1$. We construct our aggregate “innovations” in yields by summing over all forecast horizons

$$\epsilon_{t+1} = \sum_h (Q_{h-1,t+1}^y - Q_{h,t}^y), \quad (21)$$

where we include the horizons $h = 9, 12, 15$, and 18 months. *ETSMs* imply that the $\text{Var}_t[\epsilon_{t+h}]$ are affine in V_t , for all $h > 0$. Accordingly, for each yield maturity, we compare the projections

$$E_t \left(\sum_{j=1}^H \epsilon_{t+j}^2 \right) = \text{constant} + \beta_{\mathcal{P}} \mathcal{P}_t \quad \text{and} \quad E_t \left(\sum_{j=1}^H \epsilon_{t+j}^2 \right) = \text{constant} + \beta_V V_t, \quad (22)$$

again for $H = 2, 4$, and 6.¹⁶ As before, *RIETSM* implies that $\beta_{\mathcal{P}} = \beta_V \beta$.

The adjusted R^2 s for the projections in (22) are reported in Table 4 for maturities $n = 3$ months, 1, 3, 5, 7, and 10 years. Across these maturities and all choices of H , the second moments conditioned on V_t and \mathcal{P}_t are again very similar and, in fact, in many cases the projections based on V_t have larger adjusted R^2 s. It is therefore not surprising that the null hypothesis that V_t captures all of the forecasting power of \mathcal{P}_t typically cannot be rejected.

Next we examine whether V_t also encompasses the information in the yield curve (in the PC s \mathcal{P}_t) about expected excess returns in Treasury markets. Comparisons of the projections of realized excess returns onto V_t and \mathcal{P}_t are displayed in Table 5 for holding periods of lengths $h = 3, 6$, and 12 months. Overall, V_t captures a substantial portion of the predictive content of \mathcal{P}_t , particularly for longer maturity bonds. For example, the adjusted R^2 from

¹⁴The descriptive analysis of Ludvigson and Ng (2010) identifies macro factors that have forecasting power for yields over and above the yields themselves, and Joslin, Priebisch, and Singleton (2013) develop an arbitrage-free term structure model that accommodates such macro forecast factors.

¹⁵Except for the three-month and six-month maturities, the BCFF forecasts are for averages of par yields. See Appendix A for details of the construction of zero yield forecasts. Additionally, the BCFF forecasts are over *calendar* quarters. We follow the interpolation approach of Chun (2010) to build forecasts for non-calendar quarters. The same interpolation technique is used to construct forecasts for horizons not provided by the BCFF newsletter.

¹⁶We omit $H = 1$, because the BCFF surveys are conducted over a two-day period somewhere between the 20th and 26th of a given month t and so ϵ_{t+1} is not, strictly speaking, a surprise relative to the information set at the end of month t .

	$H = 2$			$H = 4$			$H = 6$		
	Adj R^2			Adj R^2			Adj R^2		
	4 PCs	V_t	pval	4 PCs	V_t	pval	4 PCs	V_t	pval
3m	0.059**	0.052***	0.828	0.132***	0.116***	0.940	0.201***	0.170***	0.631
1-y	0.096***	0.093***	0.969	0.216***	0.205***	0.809	0.328***	0.295***	0.622
3-y	0.100**	0.104***	0.900	0.262***	0.254***	0.418	0.400***	0.380***	0.382
5-y	0.102**	0.106***	0.999	0.249***	0.250***	0.821	0.388***	0.383***	0.704
7-y	0.119	0.123*	0.997	0.264***	0.260***	0.822	0.371***	0.359***	0.612
10-y	0.125	0.129	0.998	0.248***	0.245***	0.900	0.326***	0.319***	0.890

Table 4: Regressions of squared residuals constructed from BCFF yield forecasts on PCs and V_t . The superscripts (*, **, ***) denote the level of significance (at 10%, 5%, 1%, respectively) of the χ^2 test of joint significance of the regression slopes.

	$h = 3$ months			$h = 6$ months			$h = 12$ months		
	Adj. R^2			Adj. R^2			Adj. R^2		
	4 PCs	V_t	pval	4 PCs	V_t	pval	4 PCs	V_t	pval
mean(xr)	0.088	0.063	0.198	0.250	0.196	0.145	0.367	0.315	0.256
xr(1)	0.165	0.099	0.003	0.313	0.213	0.004			
xr(3)	0.113	0.077	0.140	0.267	0.201	0.192	0.371	0.287	0.597
xr(5)	0.086	0.055	0.082	0.262	0.196	0.127	0.375	0.312	0.431
xr(7)	0.096	0.071	0.225	0.253	0.199	0.133	0.369	0.316	0.205
xr(10)	0.054	0.042	0.374	0.190	0.153	0.176	0.325	0.281	0.175

Table 5: Projections of bonds excess returns onto \mathcal{P}_t and V_t for holding periods of length 3, 6, and 12 months on bonds of maturities 1, 3, 5, 7, and 10 years. $mean(xr)$ is the cross-sectional average of the excess returns.

regressing twelve-month excess returns on a ten-year bond on \mathcal{P}_t is 32.5%, compared to 28.1% when V_t is used as the predictor, and the difference is not significant (pval= 0.175). For shorter maturity bonds there is predictive information in yields that is not fully captured by V_t . The three-month and six-month excess returns on a one-year zero are predicted by \mathcal{P}_t with adjusted R^2 s of 16.5% and 31.3%, compared to 9.9% and 21.3% by V_t . For both cases, the pvals of the difference tests are smaller than 1%, indicating rejection of the constraint $\beta_{\mathcal{P}} = \beta_V \beta$ at conventional significance levels.

Summarizing, we find that there exists a strictly positive volatility process extracted from the yield curve under the constraint *RIETSM* that fully encompasses the information in the entire yield curve about interest rate volatility. This result is an implication of *ETSMs* and it shows that reduced-form affine models are fully capable of replicating yield volatilities conditioned on y_t .

Equally notable is our finding that V_t largely spans the information in bond yields about

expected excess returns (risk premiums). We stress that these findings are not built in through the construction of V_t . Moreover, the estimates of $\beta_{\mathcal{P}}$ and β_V are jointly significant for almost all choices of maturity and H . Therefore the similar forecasting power of V_t and \mathcal{P}_t is not a manifestation of lack of predictive power of V_t . The differentiating feature of our construction of V , relative to the extraction of V from fitted arbitrage-free, reduced-form affine models, is that we are able to exploit the information in the cross-section of the yield curve under the constraint that V is the sole driver of expected excess bond returns.

RIETSM is a strong restriction and we are not surprised that it is not fully satisfied at all maturities. In particular, there is a component of risk premiums on one-year bonds that varies with the shape of the yield curve and is not fully spanned by V_t . This is anticipated by recent studies on liquidity factors in Treasury markets. Bond supplies, foreign demands, and clienteles have been shown to affect the shape of the U.S. Treasury curve (e.g., Greenwood and Vayanos (2010b) and Krishnamurthy and Vissing-Jorgensen (2012)). Another influence on the shape of the intermediate segment of the Treasury yield curve was the hedging activities of mortgage traders (Duarte (2008)). Nevertheless, it is notable how much of the variation in risk premiums is captured by the two-dimensional V_t extracted from the volatility structure of US Treasury yields.

5.3 Is V_t Capturing Inflation or Output Volatility?

Of equal interest are the connections between our volatility factors V_t and macroeconomic risks, in particular consumption and inflation risks. As we discussed in Section 2, the literature on equilibrium pricing models has adopted a wide variety of distinct (non-nested) specifications of conditional variances of these macro risks. All of these *ETSMs* imply that the time-varying conditional variances of both $(\Delta c_t, \pi_t)$ and their expected values $(x_t, \bar{\pi}_t)$ are linearly spanned by bond yields. Moreover, as summarized by *RIETSM*, they imply that our V_t derived from yields must be as powerful as \mathcal{P}_t in predicting their conditional variances. In fact, given the affine structure of *ETSMs*, we know in addition that the conditional covariances between $(\Delta c_t, \pi_t)$ or $(x_t, \bar{\pi}_t)$ and y_t are linear in V_t . This can be explored by examining whether \mathcal{P}_t and V_t have equal forecasting power for the products of innovations in y_t and these macro factors.

Consider first the role of inflation volatility. Piazzesi and Schneider (2007), Rudebusch and Wu (2007), Doh (2011), and Wright (2011) argue, within the context of affine term structure models, that a decline in inflation uncertainty was partially responsible for the decline in term premiums during the past twenty years. Doh (2011) and Bansal and Shaliastovich (2012) in particular focus on *ETSMs* with *LRR*. In contrast to these studies, we link measures of inflation volatility directly to the extracted V_t , the volatility factors that *ETSMs* identify as the determinants of excess returns.

Though much of the econometric literature has focused on inflation directly,¹⁷ the monthly CPI inflation data appears quite noisy (choppy) and there are frequent, material revisions. To mitigate these measurement issues we construct a measure of inflation uncertainty using

¹⁷See, for examples, the analyses of time-varying inflation volatility in Engle (1982) and Stock and Watson (2007).

	$H = 2$			$H = 4$			$H = 6$		
	Adj. R^2			Adj. R^2			Adj. R^2		
	4 PCs	V_t	pval	4 PCs	V_t	pval	4 PCs	V_t	pval
e_{CPI}^2	0.108	0.106	0.960	0.296**	0.283***	0.886	0.420***	0.393***	0.637
e_{GDP}^2	0.087**	0.074**	0.471	0.249***	0.189***	0.035	0.376***	0.271***	0.001
$e_{CPI}e_{GDP}$	0.037	0.025	0.999	0.049	0.034	0.699	0.097	0.078	0.421

Table 6: Regressions of squared inflation residuals (e_{CPI}^2), squared real GDP residuals (e_{GDP}^2), and the product of inflation and real GDP residuals ($e_{CPI}e_{GDP}$) on \mathcal{P}_t and V_t .

the Blue Chip Economic Indicator (BCEI) survey forecasts. Analogously to our treatment of yield forecasts, we let

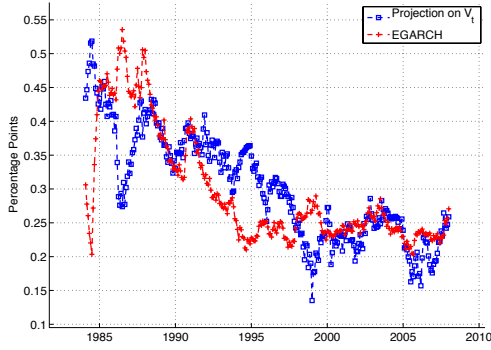
$$Q_{h,t}^{\pi(i)} = E_t^{(i)} [\pi_{t+h} + \pi_{t+h+1} + \pi_{t+h+2}],$$

as survey forecasts of inflation are for quarterly horizons. Similarly, $Q_{h,t}^{\pi}$ denotes the h -month median inflation forecast. Then we project the ϵ_{t+1}^2 in (21) (with Q^{π} in place of Q^y) onto \mathcal{P}_t and V_t as in (22) and compare the fits of the two sets of regressions. This approach amounts to assessing whether the volatility factors V_t span the time-varying volatility of expected inflation $\sigma_{\pi t}^2$.

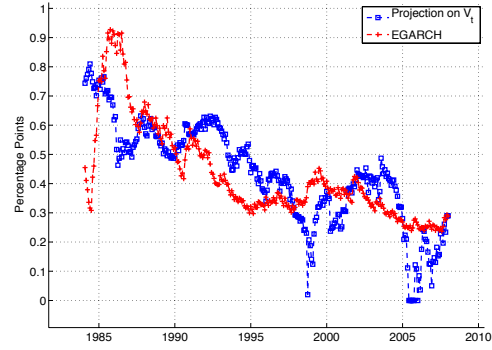
From the first row of Table 6 it is seen that V_t captures almost all of the time variation in the inflation volatility that is spanned by \mathcal{P}_t . For example, for $H = 4$, \mathcal{P}_t and V_t forecast squared average inflation residuals with adjusted R^2 s of 29.6% and 28.3%, respectively, and the difference is not statistically significant. Panel (a) of Figure 2 displays the fitted volatilities of inflation conditional on V_t . Consistent with the view that inflation risk declined in the late 1980's and 1990's, there is a persistent decline in fitted volatility over this portion of our sample. Interestingly, while the downward trend is also visible in the univariate EGARCH-fitted volatility,¹⁸ the latter declines much more rapidly in the late 1980's and then the two measures catch up with each other a decade later. During the later portion of our sample the EGARCH estimates seem to stabilize whereas the V -implied volatilities are relatively more choppy.

Under *RIETSM* our extracted V_t encompasses the sources of both time-varying inflation and real economic risks. In the model of Bansal, Kiku, and Yaron (2007) the $\mathcal{R} = 2$ dimensions of V_t are both required to span consumption risk. On the other hand, under the specification in Bansal and Shaliastovich (2012) V_t encompasses both inflation ($\sigma_{\pi t}^2$) and *LRR* (σ_{xt}^2) risks and, as such, its dimension exceeds that of the one-dimensional quantity of consumption risk. Given the limited structure we have imposed in extracting V_t , we can not distinguish between these non-nested formulations of aggregate risks. Nevertheless, the implication of *RIETSM* that V_t and \mathcal{P} constitute the same information about real economic risks is testable using V_t . In the habit-based model of Campbell and Cochrane (1999), Wachter (2006), and

¹⁸We construct the EGARCH volatility series using the same series of residuals ϵ_t computed as described in the text.



(a) Inflation Volatility



(b) GDP Volatility

Figure 2: Fitted conditional volatilities of inflation (Panel (a)) and GDP growth (Panel (b)) from projections of squared forecast errors onto the yield-curve extracted volatilities V_t . For comparison, each panel includes the fitted values from an EGARCH(1,1) model of volatility estimated using the same series of residuals. Negative fitted values were truncated at zero.

Le, Singleton, and Dai (2010)) the source of time-varying risk is stochastic volatility in the surplus consumption ratio. Again, the current yield curve should be informative about this real economic risk.

In an attempt to shed some light on connection between V_t and real economic risks, while at the same time avoiding the well known and severe measurement problems with consumption data, we examine the BCFE survey forecasts of real GDP growth and associated real GDP forecast errors (constructed using the same approach used for constructing inflation forecast errors). The squared GDP forecast errors are projected onto \mathcal{P}_t and V_t as in (22) and the fits of these two sets of regressions are compared. Most notable about the second row of Table 6 is the finding that both \mathcal{P}_t and V_t have substantial predictive content for the volatility of real GDP growth, with adjusted R^2 s ranging from 7% to 38%. Panel B of Figure 2 plots the V -implied volatility estimates together with their EGARCH counterparts. The downward trend in both series is symptomatic of the “Great Moderation” observed by Stock and Watson (2002) and others.

V_t accounts for roughly 80% of the conditional variation in real GDP growth that is spanned by yield PC s. While this is a substantial proportion, the higher predictive power of the PC s \mathcal{P}_t is statistically significant at conventional levels for $H = 4$ and $H = 6$. This suggests that, except for the very short horizons, there is information in the yield curve about real economic risks that is not well captured by ETS Ms with $N = 4$ (or fewer) risk factors and $\mathcal{R} = 2$ time-varying quantities of risk. While real GDP growth is an imperfect surrogate for consumption growth, the object of interest in most ETS Ms, these findings are of interest both because GDP growth is correlated with Δc_t and it is a key ingredient in the Federal Reserve’s setting of monetary policy (e.g., through the “Taylor rule”). These linkages have

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2			AdjR2			AdjR2		
	4 PCs	V_t	pval	4 PCs	V_t	pval	4 PCs	V_t	pval
Panel A: Inflation									
3-m	0.043	0.046	0.750	0.119***	0.110***	0.629	0.244***	0.210***	0.464
1-y	0.030	0.033	0.842	0.096***	0.092***	0.792	0.212***	0.193***	0.393
3-y	0.020	0.024	0.830	0.073***	0.072***	0.777	0.184***	0.166***	0.458
5-y	0.030	0.036	0.981	0.106***	0.104***	0.625	0.236***	0.214***	0.458
7-y	0.035	0.040	0.989	0.125***	0.119***	0.670	0.258***	0.226***	0.427
10-y	0.026	0.030	0.985	0.101***	0.097***	0.432	0.225***	0.195***	0.254
Panel B: Real GDP									
3-m	0.042***	0.034**	0.465	0.107***	0.096**	0.903	0.174	0.164	0.952
1-y	0.025	0.023	0.578	0.073	0.066	0.907	0.114	0.109	0.915
3-y	0.041	0.030	0.747	0.109	0.085	0.653	0.162	0.133	0.583
5-y	0.054***	0.040*	0.244	0.132*	0.102	0.454	0.206	0.165	0.361
7-y	0.051***	0.039**	0.543	0.126*	0.107	0.502	0.202**	0.175	0.419
10-y	0.068***	0.056*	0.312	0.165***	0.144	0.254	0.256***	0.223	0.176

Table 7: Regressions of the products of CPI inflation forecast residuals and yield forecast residuals (Panel A) and the products of real GDP forecast residuals and yield forecast residuals (Panel B) on \mathcal{P}_t and V_t .

been explored by, among others, [McCallum \(1994\)](#) and [Gallmeyer, Hollifield, and Zin \(2005\)](#) within *ETS*Ms that include a monetary authority’s rule for setting a short-term rate.

Next we examine whether the fitted volatilities (based on V_t) of inflation and GDP growth have predictive content for risk premiums in Treasury markets, a central thesis of *ETS*Ms. When we project the 3-, 6-, and 12-month average excess returns onto the fitted volatility of GDP we get R^2 ’s of 6%, 16%, and 25%, respectively. Using the fitted volatility of inflation as a predictor gives comparable (slightly lower) adjusted R^2 ’s. Thus, fitted volatility from the real side of the economy roughly matches the explanatory power of inflation risk for excess returns in bond markets. This finding supports the inclusion of both σ_{xt}^2 and $\sigma_{\pi t}^2$ as quantities of risk in [Bansal and Shaliastovich \(2012\)](#).

Additional insights into the effects of real and inflation shocks on risk premiums come from inspection of the links between V_t and the conditional covariances between yields, real GDP, and inflation. In the prototypical *LRR* model, V_t fully determines the conditional covariance between output growth and inflation. From the last row of [Table 6](#) it is seen that neither \mathcal{P}_t nor V_t has predictive power for the product of the inflation and GDP residuals. In the light of our earlier findings that V_t and \mathcal{P}_t have strong predictive power for the conditional volatilities of these variables, the covariance results suggest that there is a canceling effect that renders the covariance largely time-invariant, at least for our conditioning set.

Quite different patterns emerge from examining the conditional covariances between yields and inflation and yields and GDP ([Table 7](#), Panels A and B, respectively). Across all yield

maturities, the conditional covariances between inflation and yields at horizons $H = 4$ and $H = 6$ are quite strongly time-varying. For these horizons, the χ^2 statistics of the joint test of zero regression coefficients are all significant at the 1% level. Moreover, the predictive content of \mathcal{P}_t is well captured by V_t , consistent with the structure of many *ESTMs*. The lack of predictability at the short horizon $H = 2$ is similar to the patterns observed in (the first row of) [Table 6](#) and, to some extent, in (the last two rows of) [Table 4](#).

In contrast, for the products between real GDP and yield forecast errors, the adjusted R^2 's in Panel B are, with a few exceptions, lower than their counterparts in Panel A, and more importantly the corresponding χ^2 test statistics are mostly insignificant. Evidently V_t is substantially more informative about the conditional covariances between yields and inflation than about the covariances between yields and real GDP growth.

To help with the interpretation of these results, suppose that the innovation to the nominal pricing kernel decomposes into a real and a nominal part:

$$m_{t+1} - E_t[m_{t+1}] = \epsilon_{R,t+1} + \epsilon_{\pi,t+1}, \quad (23)$$

and that the innovation to the n -period bond yield takes the form:

$$y_{n,t+1} - E_t[y_{n,t+1}] = B_{n,R}\epsilon_{R,t+1} + B_{n,\pi}\epsilon_{\pi,t+1}. \quad (24)$$

The inflation shock $\epsilon_{\pi,t+1}$ is driven by innovations to inflation and inflation expectations, while the real shock $\epsilon_{R,t+1}$ is governed by innovations to the consumption process. For conditionally Gaussian shocks and ignoring (largely negligible) Jensen terms we can write:

$$\begin{aligned} \frac{1}{n}er_t^1(n+1) &\approx Cov_t(m_{t+1}, y_{n,t+1}) \\ &\approx Cov_t(\epsilon_{R,t+1}, y_{n,t+1}) + Cov_t(\epsilon_{\pi,t+1}, y_{n,t+1}). \end{aligned} \quad (25)$$

Thus, in this simplified setting, bond-return predictability arises through the real channel, via the time variation in $Cov_t(\epsilon_{R,t+1}, y_{n,t+1})$ or through the inflation channel, $Cov_t(\epsilon_{\pi,t+1}, y_{n,t+1})$. An implication of *ETSMS* is that V_t drives the time variation of both of the right-hand side terms. Thus, up to this approximation, the results in [Table 7](#) indicate that inflation risks represent the dominant source of variation in bond risk premiums.

Why does the inflation channel seem relatively more important? This is despite the similarities between the conditional variances of real GDP and CPI observed so far: both display similar patterns (in [Figure 2](#)), both are strongly predictable by V_t (as seen in [Table 6](#)), and both can forecast future bonds' excess returns reasonably well. To the extent that shocks to real GDP largely capture the time variation in $\epsilon_{R,t+1}$, the weak predictability of V_t for the conditional covariances between shocks to real GDP and CPI, documented in the last row of [Table 6](#), suggests that the time variation in $Cov_t(\epsilon_{R,t+1}, \epsilon_{\pi,t+1})$ implied by *ETSMS* is inconsequential. As a result, ignoring constants, we have:

$$Cov_t(\epsilon_{R,t+1}, y_{n,t+1}) = Cov_t(\epsilon_{R,t+1}, B_{n,R}\epsilon_{R,t+1} + B_{n,\pi}\epsilon_{\pi,t+1}) \approx B_{n,R}Var_t(\epsilon_{R,t+1}). \quad (26)$$

The relative importance of the real channel depends not only on the time variation in the conditional variances of the real shocks but also on how important the real shock is as part

of the shocks to nominal bond yields. If $B_{n,R}$ is very close to zero (or very small relative to $B_{n,\pi}$), so that nominal shocks are more important for nominal bond yields, then inflation risks will be the more dominant determinant of predictability in excess returns. Of course, to the extent that consumption shocks are different from GDP shocks, a more decisive verification of this requires better quality data on aggregate consumption than is currently available.

6 “Unspanned” Risks and Excess Returns on Bonds

A maintained assumption in virtually the entire literature on affine *ETSMs* is that the state of the economy is spanned by returns on bonds and certain equity claims. This observation is central to the empirical strategy used in many prior studies of *LRR* and equity returns. However, for the bond market, there is an extensive literature documenting the existence of unspanned stochastic volatility of bond yields (USYV) (e.g., [Collin-Dufresne and Goldstein \(2002\)](#), [Li and Zhao \(2006\)](#), and [Joslin \(2011\)](#)). More broadly, we are interested in the possibility of unspanned quantities of risk (USQR): time varying volatilities that are not spanned by bond yields and that have predictive content for excess bond returns. The potential presence of USQR raises the important possibility that *ETSMs* have mis-specified the conditional distribution of y_t in a way that may have distorted estimated risk premiums for bond markets. This section explores this possibility for US Treasury bond markets.

6.1 Further Evidence on the Predictability of Excess Returns

Up to this point we have focused on the predictability of excess returns based on yields (y_t), because *ETSMs* imply that ζ_t^2 (and hence the $er_t^h(n)$) are spanned by y_t . Prior to exploring the implications of USQR for *ETSMs* we examine whether other conditioning information has predictive content for returns. We can always decompose excess returns $er_t^h(n)$ as

$$\underbrace{ny_{n,t} - hy_{t,h} - (n-h)E[y_{n-h,t+h}|y_t]}_{\text{spanned by } y_t} - \underbrace{(n-h)(E_t[y_{n-h,t+h}] - E[y_{n-h,t+h}|y_t])}_{\text{unspanned by } y_t}. \quad (27)$$

So additional information will be useful for forecasting excess returns if it forecasts the component of future yields that is unspanned by y_t .

With this in mind, we approach the potential presence of USQR from two econometric perspectives. First, we focus on the special case of USYV and use relatively higher frequency (than monthly) data to construct semi-nonparametric estimates of the volatilities of US Treasury yields and ask whether these estimates have incremental forecast power for excess returns relative V_t . This approach is motivated by the possibility that high-frequency data embodies more information about yield volatility than that captured by V_t extracted from the cross-section of monthly data on yields through our parametric affine model of y_t .

The data we constructed in [Section 4](#) is available daily. Using this daily time-series we compute measures of realized variance over the horizons of one, three, and six months. Time-series models for these realized variances are then estimated by one of the following two approaches: *MIDAS* estimates using equation (4.2) in [Andreou, Ghysels, and Kourtellos](#)

	R^2 s for PC s				R^2 s for $\bar{x}r_t(6m)$			R^2 s for $\bar{x}r_t(12m)$		
	$PC1$	$PC2$	$PC3$	$PC4$	V	V+H	pval(H)	V	V+H	pval(H)
MIDAS 1M	0.157	0.199	0.153	0.390	0.196	0.228	0.031	0.315	0.326	0.719
MIDAS 3M	0.197	0.222	0.177	0.453	0.196	0.231	0.034	0.315	0.329	0.736
MIDAS 6M	0.215	0.213	0.194	0.515	0.196	0.243	0.024	0.315	0.336	0.767
EGARCH 1M	0.116	0.133	0.092	0.295	0.196	0.252	0.002	0.315	0.338	0.319
EGARCH 3M	0.116	0.134	0.094	0.297	0.196	0.253	0.001	0.315	0.339	0.327
EGARCH 6M	0.116	0.134	0.096	0.300	0.196	0.253	0.003	0.315	0.338	0.340

Table 8: The adjusted R^2 s from projecting high frequency measures of the volatilities of yields (H) onto V , and projecting excess returns onto V and H . $\bar{x}r_t(T)$ denotes the average excess return over the holding period of length T across bonds with maturities of one through ten years at annual intervals.

(2010), and *EGARCH* estimates using the framework in Nelson (1991).¹⁹ The first column of Table 8 indicates the horizon (in months) over which the realized variances that serve as inputs into the forecasting equations are computed. Estimation is with daily data and then the fitted variances (H) are sampled at a monthly frequency for estimation of the forecasts of excess returns. The columns with the headings PCj , $j = 1, \dots, 4$, present the adjusted R^2 s from projections of the realized variances of the first four PC s of bond yields onto V_t . For the first three PC s these R^2 are quite low indicating that V_t bears little resemblance to the alternative measures of variance based on higher frequency data. There is a higher correlation of the variance of $PC4$ with V_t , but still it is well below 50% for most cases.

Given these differences, a question of interest is whether H has incremental forecasting power relative to V_t for excess bond returns. For the holding periods of six and twelve months, we average the realized excess returns for the bonds with original maturities of one through ten years (excluding the one year for the twelve-month holding period), and project these returns onto V alone and onto (V, H) where H includes the forecasted realized volatilities for the first four PC s. The high-frequency H add a small, though statistically significant, amount of predictive power for six-month returns, and the same is true for three-month holding periods (not shown). On the other hand, for the one-year holding period the H show no incremental explanatory power relative to V . Overall, while there is some useful information about yield risk premiums in the high-frequency data, our extracted V appears to subsume the vast majority of the information in cash bond prices.

As a second approach to examining USQR we focus on broader information sets: (i) the yield forecast factor (YF); (ii) the GDP forecast factor (GF); and (iii) the inflation forecast factor (IF). At each point in time, these factors are simply the median forecasts across all

¹⁹For the *MIDAS* estimates of $PC1$ for example, the regressors used are lagged realized variances of $PC1$ for the following horizons: one day, two weeks, one month, three months, and six months. For the *EGARCH* estimates, we first fit the daily time series of $PC1$ through an $AR(1)$ process to extract a daily series of residuals which are then used to estimate an $EGARCH(1,1)$ model. We then compound the estimated model to compute variance forecasts at desired horizons.

Forecast factor	PC1-4	PC1-4, F		PC1-5	PC1-5, F		PC1-6	PC1-6, F	
	R^2	R^2	pval	R^2	R^2	pval	R^2	R^2	pval
Panel A: 3-month excess returns									
YF	0.088	0.091	0.298	0.121	0.122	0.377	0.156	0.153	0.943
GF	0.088	0.085	0.870	0.121	0.118	0.878	0.156	0.154	0.764
IF	0.088	0.158	0.000	0.121	0.192	0.000	0.156	0.209	0.001
JM	0.088	0.151	0.004	0.121	0.194	0.001	0.156	0.225	0.001
(IF, JM)	0.088	0.194	0.006 0.047	0.121	0.237	0.003 0.016	0.156	0.255	0.011 0.013
Panel B: 12-month excess returns									
YF	0.367	0.375	0.145	0.382	0.388	0.181	0.431	0.430	0.753
GF	0.367	0.370	0.636	0.382	0.384	0.651	0.431	0.436	0.533
IF	0.367	0.584	0.000	0.382	0.601	0.000	0.431	0.616	0.000
JM	0.367	0.555	0.000	0.382	0.583	0.000	0.431	0.625	0.000
(IF, JM)	0.367	0.695	0.000 0.000	0.382	0.739	0.000 0.000	0.431	0.737	0.000 0.000

Table 9: Adjusted R^2 s from regressions of cross-maturity average bonds excess returns on the first n yield PCs (PC1- n) and the forecast factors F. Specifications YF, GF, IF, JM correspond to regressions in which F is the yield forecast factor, GDP forecast factor, inflation forecast factor, and the jump amplitude factor, respectively. “pval” is the probability value for the t-statistic of the estimated coefficient on F. (IF,JM) refers to regressions in which excess returns are projected onto PC1- n and both the inflation forecast factor and the jump amplitude factor with the corresponding probability values provided in the “pval” column.

Blue-Chip professional forecasters and all forecast horizons, excluding (as before) the first two quarters. The yield forecast factor is also a median across all yield maturities. Assuming that market professionals condition on more information than y_t , these forecast factors may be informative about the last term in (27). We proceed by projecting cross-maturity average excess returns on the first n yield PCs for $n = 4, 5, 6$, augmenting the projection with each of the four conditioning variables. The adjusted R^2 s are reported in Table 9, where “pval” is the probability value for the t-statistic of the estimated coefficient on the included conditioning variable. Neither YF nor GF show any significant predictive power in the presence of yield PCs. In contrast, IF shows substantial predictive power for risk premiums, much stronger than the yield PCs. Moreover, the results are similar across horizons, and as the number of included PCs (up to six) is increased. Therefore, a further increase in the number of yield-based factors is unlikely to overturn these results.

Finally, considerable attention has recently been given to modeling risk premiums (e.g., [Bollerslev and Todorov \(2011\)](#)) and to forecasting volatility (e.g., [Hansen \(2011\)](#)) using ultra high-frequency data on cash and derivative prices. In principle, these high-frequency constructs applied to US Treasury data should embody at least as much information about risk as V , since they are computed non-parametrically using intra-day data from derivatives

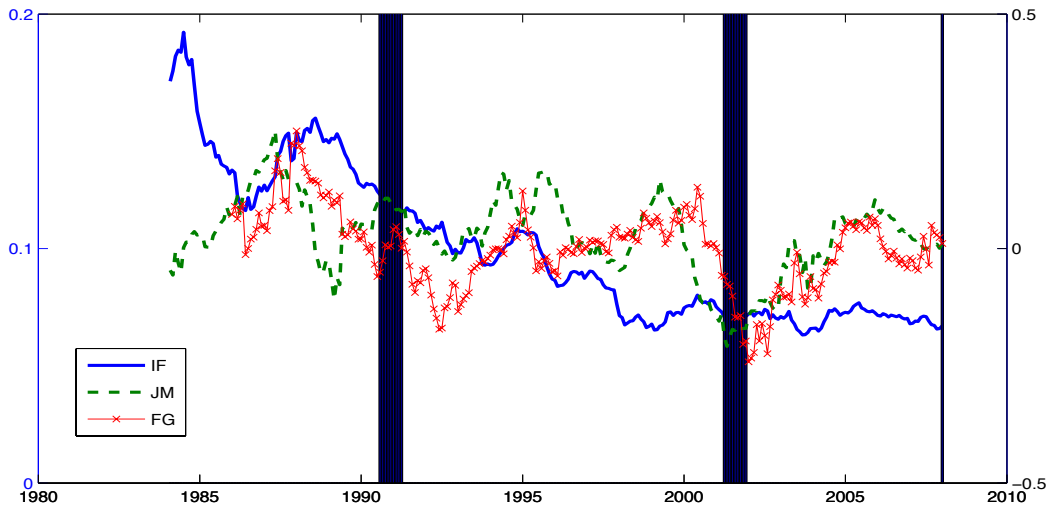


Figure 3: Time series of the inflation forecast factor (IF) and the mean jump amplitude factor (JM), compared to the liquidity factor constructed by Fontaine and Garcia (2012) (FG). Shaded areas indicate NBER recessions.

markets. To explore this possibility, in our case as a means of revealing USQR, we consider the realized jump means (JM) constructed by Wright and Zhou (2009) and extended by Huang and Shi (2011), and the measure of realized volatility (RV) constructed by Mueller, Vedolin, and Yen (2011).²⁰ Both series are computed from intra-day data on thirty-year Treasury bond futures. JM is a measure of the time-varying amplitude of jumps in the US Treasury bond market, and RV is a measure of the historical realized volatility over a day.

The row “JM” in Table 9 reveals that JM has substantial incremental forecasting power for excess returns relative to the yield PCs, confirming a similar finding by Wright and Zhou (2009) using different yield data. On the other hand, we find that inclusion of RV as a conditioning variable has a relatively small impact on fitted excess returns. RV has no incremental explanatory power relative to our extracted volatilities V for excess returns over six-month holding periods. For a one-year holding period the adjusted R^2 increases from 32% to 34%. Thus, once again we find that high-frequency measures of volatility reveal limited new information about risk premium in Treasury markets relative to V .

Focusing on the two variables with substantial incremental forecasting power, JM and IF, we wondered whether they represent separate sources of predictability. When both JM and IF are included in the projections (see the last rows of each panel of Table 9, specification (IF,JM)) there is a sizable increase in the adjusted R^2 relative to when they appear separately. For example, when the holding period is twelve months, the R^2 statistics increase by more than 10% regardless of the number of yield PCs included. The reason for the complementary forecast powers of JM and IF is visually clear in Figure 3 where it is seen that they operate

²⁰We thank the authors of these papers for providing us with their data series.

at different frequencies. Whereas IF is trend-like, JM shows more cyclical variation. We found that JM correlates rather strongly, and particularly so since the mid 1990's, with the market liquidity factor that Fontaine and Garcia (2012) constructed from Treasury on- and off-the-run spreads. This suggests that JM is capturing the effects of financial frictions or flight-to-liquidity phenomena as documented in Greenwood and Vayanos (2010b) and Krishnamurthy and Vissing-Jorgensen (2012).

These findings represent a major challenge for those *ETSMs* that imply that the volatility factors V fully span all relevant information about expected excess returns in bond markets. Indeed it seems that to rationalize our findings within this class of affine *ETSMs*, it is necessary to introduce quantities of risk that are unspanned by bond yields. Since V , by construction, is spanned by y_t , it would then be possible for both risk premiums to be fully determined by ζ_t^2 and yet for other conditioning variables beyond our constructed V to have predictive power for bond excess returns.²¹

6.2 USQR in the Context of *ETSMs*

To develop the implications of unspanned risks within *ETSMs*, consider again the stacked yield pricing equations (7). The presence of unspanned risks is equivalent to the rank of the loading matrix $[B, C]$ being less than the number of risk factors:

$$r_Z = \text{rank}([B, C]) < N. \quad (28)$$

At one extreme, suppose that ζ_t^2 is fully *unspanned* by bond yields. In this case the loadings C are zero for all maturities, as y_t cannot depend directly on ζ_t^2 . For our empirical implementation with $(N = 4, \mathcal{R} = 2)$, this means that only the two non-volatility factors ν_t are spanned by y_t and that $r_Z = 2$. Further, following the reasoning in Joslin, Le, and Singleton (2012), this implies that the first two yield PCs fully span the information generated by ν_t .

More generally, when $r_Z < N$, it can be shown that there are $N - r_Z$ linear combinations of $(\nu_t', \zeta_t^{2'})'$ that cannot be inverted from bond yields. The information set generated by the r_Z linear combinations of $(\nu_t', \zeta_t^{2'})'$ that *can* be inverted from bond yields should be well approximated by the first r_Z yield PCs, particularly for small r_Z .

Now if $r_Z = 1$ (yields depend on a single linear combination of the state $z_t' = (\nu_t', \zeta_t^{2'})'$), then yields at all maturities are perfectly correlated, conditionally and unconditionally. Whence, for our choice of $N = 4$, realistic values of r_Z are 2 and 3. The larger is r_Z the richer is the information set generated by the spanned components, and the more difficult it is to rationalize our prior evidence on the predictability of excess returns. Therefore, we set $r_Z = 2$ in our subsequent analysis and use the first two yield PCs to control for the spanned components. This way, we give IF and JM the best chances to exhibit their predictive power

²¹An alternative rationalization could be that a higher dimensional ζ_t^2 is needed. However, our choice of $(N = 4, \mathcal{R} = 2)$ encompasses the majority of *ETSMs*, as illustrated in Table 1. Furthermore, given the spanning properties of these models, we have found that the coefficients of the projections of squared yield-forecast errors onto y_t are well approximated a matrix of rank two, which implies that a larger \mathcal{R} is not likely to mitigate this challenge.

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2			AdjR2			AdjR2		
	PC1-2	IF	pval	PC1-2	IF	pval	PC1-2	IF	pval
Panel A: Squared yield residuals									
3m	0.048	0.045	0.845	0.125	0.123	0.768	0.194	0.197	0.944
5-y	0.098	0.095	0.996	0.246	0.250	0.477	0.363	0.377	0.337
10-y	0.128	0.125	0.863	0.253	0.250	0.953	0.323	0.321	0.908
Panel B: Macro residuals									
e_{CPI}^2	0.110	0.111	0.626	0.295	0.302	0.564	0.418	0.422	0.595
e_{GDP}^2	0.093	0.091	0.855	0.247	0.244	0.899	0.364	0.362	0.926
$e_{CPI}e_{GDP}$	0.002	-0.001	0.506	0.033	0.035	0.386	0.083	0.090	0.239
Panel C: Cross inflation-yield residual products									
3m	0.047	0.051	0.206	0.119	0.120	0.733	0.225	0.223	0.871
5-y	0.036	0.039	0.278	0.113	0.111	0.734	0.235	0.233	0.706
10-y	0.034	0.035	0.464	0.104	0.101	0.679	0.211	0.209	0.937
Panel D: Cross GDP-yield residual products									
3m	0.034	0.030	0.099	0.103	0.100	0.007	0.166	0.163	0.111
5-y	0.056	0.055	0.629	0.126	0.126	0.310	0.183	0.187	0.606
10-y	0.076	0.074	0.959	0.168	0.166	0.433	0.244	0.242	0.515

Table 10: Regressions of various squared residuals and residuals products constructed from BCFF yield forecasts on on the first two yield PCs and the inflation-forecast factor (IF).

for squared residuals. Though not reported here, we repeated the same analysis with $r_Z = 3$ (and $r_Z = 4$ to account for models with $N > 4$) and obtained comparable results.

The first question we address is whether the correlation between (IF, JM) and expected excess bond returns is through their correlations with unspanned yield volatility (USYV). The projections of yield squared residuals onto the first two PCs and IF (JM) are displayed in Panel A of Table 10 (Table 11) for three maturities. Interestingly, despite their predictive power for excess returns, neither IF nor JM has significant incremental predictive power for yield squared residuals at conventional significance levels. Evidently, the predictive power of (IF, JM) does not originate from USYV; there may be USYV, but its origins are other than time-variation in (IF, JM) .

To see how this might arise, consider an environment with $(N = 4, \mathcal{R} = 2, r_Z = 2)$ and suppose that: (a) one of the non-volatility factors (ν_{1t}) and one of the volatility factors (ς_{1t}^2) are unspanned and hence do not enter the affine pricing expression (7), for all maturities n ; (b) the remaining factors $(\nu_{2t}, \varsigma_{2t}^2)$ are fully spanned by y_t ; (c) the time varying volatility of $(\nu_{2t}, \varsigma_{2t}^2)$ is driven entirely by ς_{2t}^2 ; and (d) both volatility factors have predictive content for bonds excess returns. In such a setting, if IF and JM are informative about ς_{1t}^2 , they would be able to forecast excess returns beyond the yield PCs without showing any incremental power in predicting the conditional variance of y_t .

Pursuing this scenario, within the structure of *ETSMs*, there remains the question of which unspanned risk is being captured by the quantity ς_{1t}^2 . Panels B of Table 10 and

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2			AdjR2			AdjR2		
	PC1-2	JM	pval	PC1-2	JM	pval	PC1-2	JM	pval
Panel A: Squared yield residuals									
3m	0.048	0.075	0.269	0.125	0.191	0.057	0.194	0.302	0.133
5-y	0.098	0.096	0.825	0.246	0.249	0.115	0.363	0.370	0.101
10-y	0.128	0.125	0.662	0.253	0.250	0.666	0.323	0.320	0.471
Panel B: Macro residuals									
e_{CPI}^2	0.110	0.107	0.782	0.295	0.294	0.748	0.418	0.416	0.898
e_{GDP}^2	0.093	0.099	0.745	0.247	0.259	0.789	0.364	0.383	0.881
$e_{CPI}e_{GDP}$	0.002	-0.001	0.093	0.033	0.030	0.331	0.083	0.082	0.894
Panel C: Cross inflation-yield residual products									
3m	0.047	0.044	0.886	0.119	0.117	0.643	0.225	0.223	0.996
5-y	0.036	0.035	0.946	0.113	0.114	0.783	0.235	0.236	0.910
10-y	0.034	0.039	0.632	0.104	0.115	0.640	0.211	0.225	0.648
Panel D: Cross GDP-yield residual products									
3m	0.034	0.036	0.274	0.103	0.113	0.377	0.166	0.193	0.467
5-y	0.056	0.061	0.645	0.126	0.145	0.526	0.183	0.220	0.461
10-y	0.076	0.080	0.330	0.168	0.187	0.274	0.244	0.279	0.210

Table 11: Regressions of various squared residuals and residuals products constructed from BCFF yield forecasts on on the first two yield PCs and the jump means variable.

Table 11 address the possibility that it captures the influence of unspanned volatility in inflation or real output growth. In fact, we see that neither JM nor IF has predictive power for the conditional second moments of these macro variables. Panels C and D examine the conditional covariances of yields with inflation and GDP growth, and again (IF, JM) fail to show any predictive power for these conditional second moments beyond the yield PCs.

7 Concluding Remarks

In this paper, we set out to explore in depth the nature of risk premiums in US Treasury bond markets over the past thirty years through the lens of investors' pricing kernels as parameterized in studies of preference-based $ETSMs$. Many prominent $ETSMs$ in which the state of the economy z_t follows an affine process imply that the variation in expected excess returns on bond portfolio positions is fully spanned by the set of conditional variances ζ_t^2 of z_t . We show that these two assumptions alone—spanning of expected excess returns by the variances ζ_t^2 of affine processes z_t —are sufficient to econometrically identify the quantities of risk that span risk premiums from the term structure of bond yields. Using this result we derive maximum likelihood estimates of ζ_t^2 and evaluate the goodness-of-fit of the family of affine $ETSMs$ that imply this tight link between premiums and quantities of risk. These assessments are fully robust to the values of the parameters governing preferences and the

evolution of the state z_t , and to whether or not the economy is arbitrage free.

Our main findings are four-fold: (i) time variation in two quantities of risk (yield volatilities) fully subsume the information in the yield curve about both the conditional variances of yields and excess bond returns; (ii) inflation risk is a significant (and perhaps a dominant macroeconomic) risk underlying risk premiums in U.S. Treasury markets; (iii) there is a second, distinct “flight to quality” risk reflected by the impact of yield-jump amplitudes on excess returns; and (iv) there is a significant unspanned (by bond yields) component of risk premiums in the US Treasury market that is not accommodated by extant *ETSMs* and that is not explained by unspanned volatility in yields, inflation, or output growth.

Taken together, these findings cast doubt on the presumptions of many affine term structure models. At one extreme, finding (i) runs counter to the presumption in the large class of Gaussian models that time variation in risk premiums arises entirely through variation in the market prices of risks. The quantities of risk V_t span the information in the yield curve upon which many of these Gaussian models have relied on in explaining excess bond returns. This is consistent with the implications of recent *ETSMs* that have emphasized the importance of volatility movement on asset prices (e.g., [Bansal, Kiku, Shaliastovich, and Yaron \(2012\)](#)). At the other extreme, our finding challenge the presumption in many *ETSMs* that time-varying quantities of risk are sufficient to capture the state-dependence of risk premiums in the Treasury market. Both an inflation factor and a measure of the amplitude of jumps in bond yields have substantial explanatory power for excess bond returns. To reconcile this evidence with *ETSMs*, one could rely on unspanned (by yields) quantities of risk. However, in exploring this possibility, we were able to rule out (for our data and sample period) the possibilities of these factors affecting expected excess returns through unspanned volatilities in bond yields, inflation, or output growth.

Two plausible, alternative explanations of our findings are time-varying market prices of risk (as documented in studies of Gaussian models) and flight-to-quality or specialness of Treasury bonds. The predictive power of the inflation factor is consistent with [Joslin, Priebisch, and Singleton \(2013\)](#) though their Gaussian model channels the effects of inflation on risk premiums entirely through a time-varying market price of risk. Our analysis suggests, when modeling bond yields, it may be important to extend *ETSMs* to allow for both time-varying quantities and market prices of risks. As for the second possible explanation, neither *ETSMs* nor reduced-form affine models of Treasury yields have typically accommodated financial frictions/liquidity considerations as explicit sources of variation in risk premiums. However, the contribution of jump amplitudes to explained variation in risk premiums is large, plays a distinct role relative to inflation risk, and its effects are present for holding periods of at least one year. Thus, accommodating a “liquidity” factor when studying the pricing of Treasury bonds within *ETSMs* may well be essential for obtaining a reliable assessment of how macroeconomic risks impact risk premiums in this market.

A Construction of Zero Yield Forecasts

In this section, we give details of the construction of the zero yield forecasts used in the paper. We first show how to interpolate the raw forecasts (of yields over calendar quarters) obtain forecasts for non-calendar quarters. Next, we show how to construct zero yield forecasts from forecasts of par yields. To fix notation, let's denote the n -period zero yields and n -period par yields by $y_{n,t}$ and $\tilde{y}_{n,t}$, respectively.

A.1 Non-Calendar Quarter Forecasts

Recall that from the Blue Chip surveys, we obtain forecasts of average par yields over calendar quarters. For example, the one- and two-quarter forecasts as of the end of December, January, and February are all:

$$\tilde{F}_{1,t} = E_t[\tilde{y}_{120,Jan} + \tilde{y}_{120,Feb} + \tilde{y}_{120,Mar}] \quad \text{and} \quad \tilde{F}_{2,t} = E_t[\tilde{y}_{120,Apr} + \tilde{y}_{120,May} + \tilde{y}_{120,Jun}].$$

Obviously, the forecast horizons are different from one month to another depending on which month of the quarter at which the forecasts are formed. To equate the forecast horizons throughout the sample, we follow the approach of [Chun \(2010\)](#) and interpolate the raw forecasts such that the one-quarter forecasts are always for the average par yields of the first three months from $t + 1$ to $t + 3$, the two-quarter forecasts are for the average par yields from $t + 4$ to $t + 6$:

$$\tilde{Q}_{1,t} = E_t[\tilde{y}_{120,t+1} + \tilde{y}_{120,t+2} + \tilde{y}_{120,t+3}], \quad \tilde{Q}_{2,t} = E_t[\tilde{y}_{120,t+4} + \tilde{y}_{120,t+5} + \tilde{y}_{120,t+6}],$$

and so on. Specifically, for the first month of each quarter (Jan, Apr, Jul, Oct), we compute the q -quarter forecasts as:

$$\tilde{Q}_{q,t} = \frac{2}{3}\tilde{F}_{q,t} + \frac{1}{3}\tilde{F}_{q+1,t}.$$

Likewise, for the second month of each quarter (Feb, May, Aug, Nov), we compute the q -quarter forecasts as:

$$\tilde{Q}_{q,t} = \frac{1}{3}\tilde{F}_{q,t} + \frac{2}{3}\tilde{F}_{q+1,t}.$$

For the third month of each quarter, we leave the raw forecasts intact. The same treatment is applied to bonds of all maturities.

A.2 Zero Yield Forecasts

For each day t and each forecast horizon q , we have forecasts for par yields of various maturities. The set of maturities has changed from time to time. For the first four years of the sample (1984-1987), there are only three maturities included: 3-month, 3-year, and 30-year. For the remaining twenty years of the sample, forecasts of eight different maturities are reported each

month. Except for some brief changes,²² this set includes: 3-month, 6-month, 1-year, 2-year, 5-year, 7-year, 10-year, and 30-year.

To fix to a set of maturities over time, and more importantly, to convert forecasts of par yields into forecasts of corresponding zero yields, we simply apply the standard Fama-Bliss bootstrap technique²³ to the set of par yield forecasts on each day t and each forecast horizon q . In so doing, we treat the par yield forecasts as if they were actual observed par yields. Additionally, in the bootstrapping calculations, we ignore the fact that the forecasts are for averages over three successive months. The output of the bootstrap, for example for the 10-year maturity and one-quarter forecast horizon, is interpreted as:

$$Q_{1,t} = E_t[y_{120,t+1} + y_{120,t+2} + y_{120,t+3}].$$

In applying the bootstrapping techniques to the average forecasts directly,²⁴ the resulting zero yield forecasts suffer from the Jensen effect. However, we now show that this effect is economically very negligible.

To prove that the Jensen effect is very minimal, we perform the following exercise. We fit yields data over our sample period to three prominent term structure models: $A_0(4)$, $F_1(4)$, and $F_2(4)$. The first model is the standard gaussian affine with no-arbitrage. The next two models have one and two stochastic volatility factors, respectively but without no-arbitrage imposed. To be consistent with the models considered in this paper, we choose four factors in each model. Next, we use each model to generate par yield forecasts for six different maturities,²⁵ averaged over three-month windows, corresponding exactly to the forecasts reported in the BCFF newsletters. To each set of forecasts, we apply the bootstrapping technique described above to obtain an approximate set of zero yield forecasts averaged over three-month windows. Finally, we compare these bootstrapped forecasts to the true forecasts of zero yields generated by each model.

Figure 4 plots the bootstrapped forecasts and the true forecasts generated by the $F_2(4)$ model for the 1-year and 10-year bonds:

$$E_t[y_{12,t+16} + y_{12,t+17} + y_{12,t+18}] \text{ and } E_t[y_{120,t+16} + y_{120,t+17} + y_{120,t+18}].$$

Note that these forecasts are for six quarters out. Visually, the bootstrapped and the true forecasts are essentially identical. Very similar pattern obtains if we use the other two models, any forecast horizon or any bond maturity.

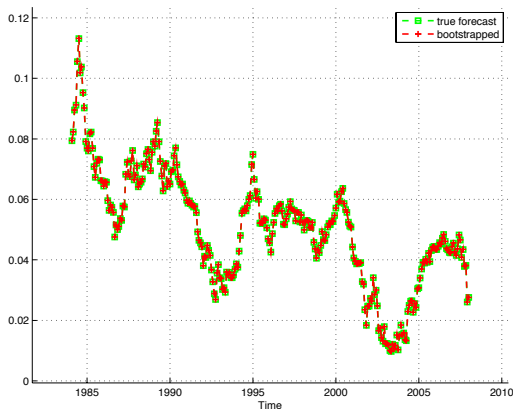
To show this more concretely, we regress the approximate forecasts on the true forecasts and report the loading, the adjusted R-squared statistics in Table 12. Evidently, regardless

²²For example, the 30-year series were replaced briefly by the 20-year series when the former was discontinued and resumed again.

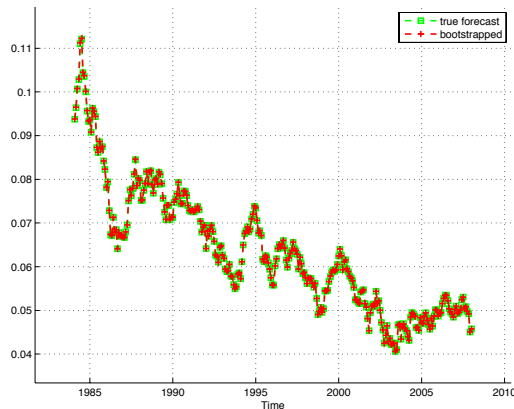
²³Specifically, we assume that the forward rates are piece-wise constant between any two successive maturities.

²⁴The more proper approach is to apply the bootstrapping technique to each future realization of the par yield curve and take average over all resulting zero curves. But such a method is not feasible since we do not observe the dynamics underlying each forecast.

²⁵They are 6-month, 1-year, 3-year, 5-year, 7-year, and 10-year.



(a) 1-yr bond, $F_2(4)$



(b) 10-year bond, $F_2(4)$

Figure 4: Comparison of bootstrapped forecasts versus true forecasts of average zero-yields averaged over the sixth quarter.

of the forecast horizons (three quarters or six quarters out), bond maturity (1-year, 3-year, 5-year, or 10-year), or the true underlying models ($A_0(4)$, $F_1(4)$, or $F_2(4)$), the adjusted R-squared is always perfect and the loading is essentially one. The average differences as well as the RMSEs between the bootstrapped and true forecasts are, for all cases, less than or equal to three basis points.

B $ARG(\rho, c, \gamma)$ Processes

Following [Le, Singleton, and Dai \(2010\)](#) we assume that, conditional on ζ_t^2 , the components of ζ_{t+1}^2 are independent. To specify the conditional distribution of ζ_{t+1}^2 , we let ϱ be an $N \times N$ matrix with elements satisfying

$$0 < \varrho_{ii} < 1, \quad \varrho_{ij} \leq 0, \quad 1 \leq i, j \leq N.$$

Furthermore, for each $1 \leq i \leq N$, we let ρ_i be the i^{th} row of the $N \times N$ non-singular matrix $\rho = (I_{N \times N} - \varrho)$. Then, for constants $c_i > 0$, $\gamma_i > 0$, $i = 1, \dots, N$, we define the conditional density of $\zeta_{i,t+1}^2$ given ζ_t^2 as the Poisson mixture of standard gamma distributions:

$$\frac{\zeta_{i,t+1}^2}{c_i} | (P, \zeta_t^2) \sim \text{gamma}(\gamma_i + P), \quad \text{where } P | \zeta_t^2 \sim \text{Poisson}(\rho_i \zeta_t^2 / c_i). \quad (29)$$

Here, the random variable $P \in (0, 1, 2, \dots)$ is drawn from a Poisson distribution with intensity modulated by the current realization of the state vector ζ_t^2 .

		Three-quarter forecasts				Six-quarter forecasts			
		Loading	Adj R2	Ave. Diff.	RMSE	Loading	Adj R2	Ave. Diff.	RMSE
$A_0(4)$	1-yr bond	1.000	1.000	-2.436	2.436	1.000	1.000	-0.999	0.999
	3-yr bond	1.000	1.000	-2.446	2.463	1.000	1.000	-0.251	0.317
	5-yr bond	1.000	1.000	-2.426	2.439	1.000	1.000	-0.206	0.268
	10-yr bond	1.001	1.000	-2.388	2.401	1.001	1.000	-0.576	0.608
$F_1(4)$	1-yr bond	0.999	1.000	2.018	2.027	1.001	1.000	-0.752	0.799
	3-yr bond	1.000	1.000	2.931	2.943	1.001	1.000	-0.380	0.417
	5-yr bond	1.000	1.000	3.034	3.042	1.001	1.000	-0.584	0.624
	10-yr bond	0.999	1.000	2.778	2.796	1.001	1.000	-1.119	1.144
$F_2(4)$	1-yr bond	1.000	1.000	-0.425	0.446	1.000	1.000	-0.274	0.292
	3-yr bond	1.001	1.000	-0.515	0.559	1.001	1.000	0.274	0.327
	5-yr bond	1.001	1.000	-0.812	0.866	1.001	1.000	0.174	0.299
	10-yr bond	1.002	1.000	-1.175	1.250	1.002	1.000	-0.227	0.408

Table 12: Regression of bootstrapped forecasts on true forecasts. Average differences (Ave. Diff.) and root mean squared errors (RMSE) are in basis points.

The conditional density function of $\zeta_{i,t+1}^2$ takes the form:

$$f^{\mathbb{Q}}(\zeta_{i,t+1}^2 | \zeta_t^2) = \frac{1}{c_i} \sum_{k=0}^{\infty} \left[\frac{\left(\frac{\rho_i \zeta_t^2}{c_i}\right)^k}{k!} e^{-\frac{\rho_i \zeta_t^2}{c_i}} \times \frac{\left(\frac{\zeta_{i,t+1}^2}{c_i}\right)^{\gamma_i+k-1} e^{-\frac{\zeta_{i,t+1}^2}{c_i}}}{\Gamma(\gamma_i+k)} \right]. \quad (30)$$

Using conditional independence, the distribution of ζ_{t+1}^2 , conditional on ζ_t^2 , is given by $f(\zeta_{t+1}^2 | \zeta_t^2) = \prod_{i=1}^N f(\zeta_{i,t+1}^2 | \zeta_t^2)$. When the off-diagonal elements of the $N \times N$ matrix ρ are non-zero, the autoregressive gamma processes $\{\zeta_i^2\}$ are (unconditionally) correlated. Thus, even in the case of correlated ζ_{it}^2 , the conditional density of $\zeta_{i,t+1}^2$ is known in closed form.

References

- Ahn, D., R. Dittmar, and A. Gallant (2002). Quadratic term structure models: Theory and evidence. *Review of Financial Studies* 15, 243–288.
- Andersen, T. and L. Benzoni (2010). Do bonds span volatility risk in the u.s. treasury market? a specification test for affine term structure models. *Journal of Finance* LXV, 603–653.
- Andreou, E., E. Ghysels, and A. Kourtellos (2010). Forecasting with mixed-frequency data. Technical report, University of North Carolina.
- Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2012). Volatility, the macroeconomy and asset prices. Technical report, NBER Working Paper.
- Bansal, R., D. Kiku, and A. Yaron (2007). Risks for the long run: Estimation and inference. Technical report, Duke University.
- Bansal, R. and I. Shaliastovich (2012). A long-run risks explanation of predictability puzzles in bond and currency markets. Technical report, Duke University.
- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Beeler, J. and J. Campbell (2009). The long-run risks model and aggregate asset prices: An empirical assessment. Technical report, Harvard University.
- Bliss, R. R. (1997). Testing term structure estimation methods. *Advances in Futures and Options Research* 9, 197–231.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T. and V. Todorov (2011). Tails, fear, and risk premia. *Journal of Finance* LXVI, 2165–2211.
- Bonomo, M., R. Garcia, N. Meddahi, and R. Tedongap (2010). Disappointment aversion, long-run risks and aggregate asset prices. Technical report, forthcoming, *Review of Financial Studies*.
- Campbell, J. and J. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205–251.
- Chun, A. (2010). Expectations, bond yields and monetary policy. Technical report, forthcoming, *Review of Financial Studies*.
- Cieslak, A. and P. Povala (2011). Understanding bond risk premia. Technical report, University of Lugano.

- Cochrane, J. and M. Piazzesi (2005). Bond risk premia. *American Economic Review* 95, 138–160.
- Collin-Dufresne, P., R. Goldstein, and C. Jones (2009). Can interest rate volatility be extracted from the cross section of bond yields? *Journal of Financial Economics* 94, 47–66.
- Collin-Dufresne, P. and R. S. Goldstein (2002). Do bonds span the fixed income markets? theory and evidence for ‘unspanned’ stochastic volatility. *Journal of Finance* 57, 1685–1730.
- Constantinides, G. and A. Ghosh (2009). Asset pricing tests with long run risks in consumption growth. Technical report, Graduate School of Business, University of Chicago.
- 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). Expectations puzzles, time-varying risk premia, and affine models of the term structure. *Journal of Financial Economics* 63, 415–441.
- Dai, Q. and K. Singleton (2003). Term structure dynamics in theory and reality. *Review of Financial Studies* 16, 631–678.
- Doh, T. (2011). Long run risks in the term structure of interest rates: Estimation. Technical report, Federal Reserve Bank of Kansas City, RWP 08-11.
- Drechsler, I. and A. Yaron (2011). What’s vol got to do with it. *Review of Financial Studies* 24.
- Duarte, J. (2008). The causal effect of mortgage refinancing on interest-rate volatility: Empirical evidence and theoretical implications. *Review of Financial Studies* 21, 1689–1731.
- Duffee, G. (2011). Information in (and not in) the term structure. *Review of Financial Studies* 24, 2895–2934.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of u.k. inflation. *Econometrica* 50, 987–1008.
- Eraker, B. (2008). Affine general equilibrium models. *Management Science* 54, 2068–2080.
- Eraker, B. and I. Shaliastovich (2008). An equilibrium guide to designing affine asset pricing models. *Mathematical Finance* 18, 519–543.
- Fama, E. F. and R. R. Bliss (1987). The information in long-maturity forward rates. *American Economic Review* 77(4), 680–692.

- Fontaine, J.-S. and R. Garcia (2012). Bond liquidity premia. *Review of Financial Studies* 25, 1207–1254.
- Gallmeyer, M., B. Hollifield, and S. Zin (2005). Taylor rules, mcallum rules and the term structure of interest rates. *Journal of Monetary Economics* 52, 921–950.
- Gourieroux, C. and J. Jasiak (2006). Autoregressive gamma processes. *Journal of Forecasting* 25, 129–152.
- Greenwood, R. and D. Vayanos (2010a). Bond supply and excess bond returns. Technical report, LSE.
- Greenwood, R. and D. Vayanos (2010b). Price pressure in the government bond market. *American Economic Review P&P*, 585–590.
- Gurkanyak, R., B. Sack, and J. Wright (2007). The U.S. treasury yield curve: 1961 to the present. *Journal of Monetary Economics* 54, 2291–2304.
- Hansen, P. (2011). Forecasting volatility using high frequency data. Technical report, forthcoming, *Handbook on Economic Forecasting*.
- Huang, J. and Z. Shi (2011). Determinants of bond risk premia. Technical report, Penn State University.
- Joslin, S. (2011). Can unspanned stochastic volatility models explain the cross section of bond volatilities? Technical report, USC.
- Joslin, S., A. Le, and K. Singleton (2012). Why gaussian macro-finance term structure models are (nearly) unconstrained factor-vars. *Journal of Financial Economics*.
- Joslin, S., M. Pribsch, and K. Singleton (2013). Risk premiums in dynamic term structure models with unspanned macro risks. Technical report, forthcoming, *Journal of Finance*.
- Joslin, S., K. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term structure models. *Review of Financial Studies* 24, 926–970.
- Krishnamurthy, A. and A. Vissing-Jorgensen (2012). The aggregate demand for treasury debt. Technical report, forthcoming, *Journal of Political Economy*.
- Le, A., K. Singleton, and J. Dai (2010). Discrete-time affine^q term structure models with generalized market prices of risk. *Review of Financial Studies* 23, 2184–2227.
- Li, H. and F. Zhao (2006). Unspanned stochastic volatility: Evidence from hedging interest rate derivatives. *Journal of Finance* 61, 341–378.
- Ludvigson, S. and S. Ng (2010). Macro factors in bond risk premia. *Review of Financial Studies* 22, 5027–5067.

- Marakani, S. (2009). Long run consumption risks: Are they there? Technical report, Northwestern University.
- McCallum, B. T. (1994). Monetary policy and the term structure of interest rates. NBER Working Paper No. 4938.
- Mueller, P., A. Vedolin, and Y. Yen (2011). Bond variance risk premia. Technical report, London School of Economics.
- Nelson, C. and A. Siegel (1987). Parsimonious modelling of yield curves. *Journal of Business* 60, 473–489.
- Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59, 347–370.
- Piazzesi, M. (2010). Affine term structure models. In Y. Ait-Sahalia and L. Hansen (Eds.), *Handbook of Financial Econometrics*, Chapter 12, pp. 691–766. Elsevier B.V.
- Piazzesi, M. and M. Schneider (2007). Equilibrium yield curves. In D. Acemoglu, K. Rogoff, and M. Woodford (Eds.), *NBER Macroeconomics Annual*. Cambridge, MA: MIT Press.
- Rudebusch, G. and T. Wu (2007). Accounting for a shift in term structure behavior with no-arbitrage and macro-finance models. *Journal of Money Credit and Banking* 39, 395 – 422.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- Stock, J. and M. Watson (2007). Why has u.s. inflation become harder to forecast? *Journal of Money Credit and Banking* 39, 3–33.
- Stock, J. H. and M. W. Watson (2002). Has the business cycle changed and why? *NBER Macroeconomics Annual* 17, 160–218.
- Vayanos, D. and J. Vila (2009). A preferred-habitat model of the term structure of interest rates. Technical report, London School of Economics.
- Wachter, J. A. (2006). A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79(2), 365–399.
- Wright, J. (2011). Term premiums and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review* 101, 1514–34.
- Wright, J. and H. Zhou (2009). Bond risk premia and realized jump risk. *Journal of Banking and Finance* 33, 2333–2345.