

The Structure of Risks in Equilibrium Affine Term Structures of Bond Yields

Anh Le and Kenneth J. Singleton*

November 6, 2012
Preliminary and Incomplete

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) the fundamental sources of risks, including consumption growth, inflation, and yield volatilities are driven by distinct economic shocks; (ii) consumption growth risk alone does not fully account for the predictability of excess returns on bonds; and (iii) 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.

*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. Moreover for default-free debt, and absent strong clientele effects along the yield curve,² yields on bonds of all maturities depend on the same underlying risk factors and, as such, there is a rich cross-section of information about the economic risks underlying the risk premiums demanded by market participants. Bond yields highlight variation in discount rates (as opposed to the cash-flow risks in equity markets) where, arguably, the contributions of consumption risks are most clearly revealed.

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. More concretely, their assumptions about agents' pricing kernel, the conditional distributions of the risk factors, and the market prices of these factor risks give rise to factor representations

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.

²Such segmentation effects have been recently explored by [Greenwood and Vayanos \(2010a\)](#), [Krishnamurthy and Vissing-Jorgensen \(2010\)](#), and [Vayanos and Vila \(2009\)](#), among others. As in most of the empirical research on risk premiums in bond markets, we abstract from supply effects.

of expected excess returns on bonds with known dimension, and the factors driving risk premiums are a subset of those determining the cross-sectional distribution of bond yields. 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.

A premise of most of our econometric analysis is that equilibrium bond yields are affine functions of their underlying risk factors. Most *ETSMs* with *LRR* are constructed so that yields follow an N -factor affine model, with N typically ranging between two and five. For instance, the *LRR* models of explored by [Kojien et al. \(2010\)](#) and [Bansal and Shaliastovich \(2012\)](#) (hereafter B-S) imply four-factor *ETSMs*, two more factors than in the model of [Bansal and Yaron \(2004\)](#). The habit-based *ETSMs* of [Wachter \(2006\)](#) and [Le et al. \(2010\)](#) give rise to two-factor models. With these models 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.³

Central to our empirical analysis of *ETSMs* is a new set of data on yields on US Treasury zero-coupon bonds constructed from daily data on a large cross-section of yields. The Fama-Bliss CRSP and [Gurkanyak, Sack, and Wright \(2007\)](#) (GSW) datasets are the most widely used for empirical analysis of dynamic term structure models (*DTSMs*). The former only contains 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.

2 The Factor Structure of *ETSMs*

Consider an endowment economy, as in most of the extant literature, and let C_t denote real per-capita consumption and $c_t \equiv \log C_t$. We focus on a consumption growth process that encompasses those examined in both the *LRR* and habit literatures:

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

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

where x_t is the *LRR* factor and the consumption-growth and *LRR* shocks (η_{t+1}, e_{t+1}) follow affine processes with conditional variances $Var_{t-1}[\eta_t] = \sigma_{ct}^2$ and $Var_{t-1}[e_t] = \sigma_{xt}^2$. This

³Empirical implementations of reduced-form affine term structure models have typically found that three to four factors explain both the cross-section and time-series properties of bond yields.

Recursive Preferences	LRR	N	Priced Risks	\mathcal{R}	RP	MPR
Bansal, Kiku, and Yaron (2007)	yes	3	$x_t, \sigma_{ct}^2, \bar{\pi}_t$	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
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
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).

specification nests the consumption processes in the models of Bansal, Kiku, and Yaron (2007), Bansal and Shaliastovich (2012), Bollerslev, Tauchen, and Zhou (2009), and Drechsler and Yaron (2011); as well as the habit-based models of Wachter (2006) and Le, Singleton, and Dai (2010), and preference shock model of Bekaert and Engstrom (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) and Bekaert and Engstrom (2010). These shocks may also embody jump components with state-dependent arrival intensities as in Drechsler and Yaron (2011). All of these models are encompassed by our assumption that the shocks follow affine processes (e.g., Duffie, Pan, and Singleton (2000)). Table 1 summarizes the features of these models that are particularly relevant to the goals of our analysis.

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

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

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). Bekaert and Engstrom (2010) study the special case of (3) in which φ_t is an exogenous preference shock, and the conditional variances of (η_t, e_t) differ in “good” and “bad” economic times.

Most of the models with *LRR* are, first and foremost, real business cycle models. 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 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 . Finally, Bollerslev et al. (2009) allow the conditional variance of σ_{ct}^2 (denoted q_t) to be time-varying and to exhibit stochastic volatility. Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) do not specify an inflation process. Bansal and Shaliastovich (2012) depart from these earlier studies and assume that consumption growth Δc_t and the inflation process π_t are conditionally homoskedastic (constant conditional variances). The time-varying quantities for risk are σ_{xt}^2 and the conditional variance of expected inflation, $\sigma_{\pi t}^2$, which leads to a model with $N = 4$.

Common to all of these equilibrium models is the feature that the dimensions (\mathcal{R}) of the expected excess returns or risk premiums (RP) are less than the dimensions of the risk factors underlying bond yields ($N > \mathcal{R}$). In most of the models with LRR , time variation in expected excess returns is induced entirely by variation in the consumption and LRR volatilities that follow first-order Markov processes. Bansal and Shaliastovich (2012) shut down consumption risk, and include time-varying inflation risk. 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 σ_{ct}^2 (so $\mathcal{R} = 1$).

The habit-based models of Wachter (2006) and Le et al. (2010) share the common feature, following Campbell and Cochrane (1999), that innovations in consumption growth and surplus consumption are perfectly correlated and have a time-varying volatility induced by variation in surplus consumption. As such, expected excess returns in these models also lie in a one-dimensional space determined by s_t .

As these models illustrate, the dimensionality \mathcal{R} of expected excess returns in $ETSMs$ is determined by the number of risk factors that drive time-varying volatility. 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 specifications of priced risks and the implied structure of risk premiums. Moreover, as we next discuss, both families of equilibrium models imply restrictions on risk premiums that are robust in the sense that they do not depend on specific values of the parameters governing preferences or on (most features of) the distributions of the state.

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

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

Letting z_t denote the N -dimensional set of time-varying priced risks in an $ETSM$, all but one of the $ETSMs$ 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 .

The set of pricing factors z_t includes the expected growth rate of consumption x_t (the *LRR* factor in those models with this feature) and expected inflation $\bar{\pi}_t = E_t[\pi_{t+1}]$. It also includes the \mathcal{R} volatility factors ζ_t^2 that gives rise to time-varying volatility in bond yields. In models with Epstein-Zin preferences and *LRR*, ζ_t^2 is comprised of one or two of the variances $(\sigma_{ct}^2, \sigma_{xt}^2, \sigma_{\bar{\pi}t}^2)$. In the habit-based models ζ_t^2 is the scalar surplus consumption s_t .

A particular focus of our analysis is the strong implication of these *ETSMs* that expected excess returns can be expressed as

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

where ζ_t^2 is a strict subset of the state z_t comprised of the \mathcal{R} volatility factors. These observations lead us to the following robust implication of these *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 risk⁴ or, more likely, because of a fundamental mis-specification of the structure of preferences of bond investors.⁵

Turning to habit-based models, [Le et al. \(2010\)](#) explore a model in which m_{t+1} has a non-affine structure in order to simultaneously ensure that that zero-coupon yields are exponential-affine functions of z and that their model (approximately) nests prior habit-based term structure models. The expected excess returns $er_t^1(\tau)$, for all τ , are approximately affine functions of surplus consumption 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. Accordingly, we proceed under the assumption that [\(6\)](#) is a reliable approximation for all of the *ETSMs* summarized in [Table 1](#).

⁴[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.

⁵[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*.

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 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 differs from these studies by exploiting the availability of yields on bonds with a cross-section of maturities. Further, and most importantly, we show that this cross-section identifies the time-varying quantities of risk underlying variation in bond-market risk premiums with considerable precision and without facing measurement issues associated with consumption.⁶

To show this we proceed in three steps: First, we express the implications of *RIETSM* for risk premiums in terms of constraints on the loadings on the risk factors in the linear N -factor representation of yields implied by *ETSMs*,

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

where the M factors ζ_t^2 determine the time-varying volatility of the entire set of N factors $z'_t = (x'_t, \zeta_t^2)$. Second, using the fact that, outside of degenerate cases, (7) implies that ζ_t^2 is spanned by contemporaneous bond yields, we express ζ_t^2 in terms of \mathcal{R} portfolios of yields, and argue that the cross-section of bond yields is likely to give very precise estimates of the weights in these portfolios. Finally, we exploit the assumption that ζ_t^2 is an affine variance process to derive maximum likelihood estimates of the unknown weights.

Starting from (4), substituting the factor representation (7) for yields, and representing the conditional means of the affine processes ζ_{t+1}^2 and x_{t+1} as (ignoring constants)⁷

$$E_t[\zeta_{t+1}^2] = \rho \zeta_t^2, \quad \text{and} \quad E_t[x_{t+1}] = K_{1\zeta} \zeta_t^2 + K_{1x} x_t, \quad (8)$$

affine pricing implies that expected excess returns can generically be expressed as

$$er_t^1(n) = (B_n - B_{n-1} K_{1x} - B_1) x_t + (C_n - C_{n-1} \rho - B_{n-1} K_{1\zeta} - C_1) \zeta_t^2. \quad (9)$$

⁶Much 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 \(2009\)](#).

⁷To maintain admissibility, the conditional means of ζ_{t+1}^2 cannot be dependent on the non-volatility factor x_t as its support is typically \mathbb{R}^{N-M} .

Now under *RIETSM*– only ζ_t^2 predicts excess returns– the loadings on x_t must be zero or, equivalently, B_n must satisfy the recursion

$$B_n = B_{n-1}K_{1x} + B_1. \quad (10)$$

Moreover, (10) implies that $B_n = B_{n-h}K_{1x}^h + B_h$ for any horizon h and, therefore, from (9) it follows that the h -period expected excess return of an n -period bond ($er_t^h(n)$) depends only on ζ_t^2 ; that is, ζ_t^2 completely spans $er_t^j(n)$, for all $h > 0$.

Enforcing this constraint leads to the structure of risk premiums implied by *ETSMs* as, under (10), expression (9) simplifies to

$$er_t^1(n) = (C_n - C_{n-1}\rho - B_{n-1}K_{1\zeta} - C_1)\zeta_t^2. \quad (11)$$

Notice that $er_t^1(n)$ does not require (or imply) arbitrage-free pricing (though the key ingredient (6) underlying our derivation of (11) is typically derived in a no-arbitrage setting). Nor was it necessary to parametrically specify the ζ_t^2 process, beyond that it is affine. Further, anticipating our analysis of unspanned risks in Section 6, even if ζ_t^2 is completely unspanned by yields ($C_n \equiv 0$ for all n), (10) and (11) still hold as implications of *RIETSM*.⁸

We next show that this structure on the $er_t^h(n)$ leads to strong identification of the volatility factors ζ_t^2 from the cross-sectional information in the yield curve. For this purpose we suppose that model assessment is based on a collection of J yields y_t ($J > N$), and we let (A, B, C) denote the stacked up loadings on $(1, x_t, \zeta_t^2)$ 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.

An immediate implication of (7) is that, outside degenerate cases discussed in Section 6, ζ_t^2 is fully spanned by bond yields. To express ζ_t^2 in terms of a subset of the yield portfolios \mathcal{P}_t , we partition the loading matrix W into W_x ($(N - M) \times J$) and W_ζ ($M \times J$) and let $\mathcal{P}_{xt} = W_x y_t$ and $\mathcal{P}_{\zeta t} = W_\zeta y_t$. Constructing $W_x y_t$ using (7) and solving for x_t gives

$$x_t = (W_x B)^{-1}(\mathcal{P}_{xt} - W_x(A + C\zeta_t^2)). \quad (12)$$

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

$$y_t = B_{\mathcal{P}}\mathcal{P}_{xt} + \phi_x(A + C\zeta_t^2), \quad (13)$$

where $B_{\mathcal{P}} = B(W_x B)^{-1}$ and $\phi_x = I_J - B_{\mathcal{P}}W_x$. Finally, premultiplying both sides of (13) by W_ζ and solving for ζ_t^2 gives

$$\zeta_t^2 = (W_\zeta \phi_x C)^{-1} \left(\mathcal{P}_{\zeta t} - W_\zeta B_{\mathcal{P}} \mathcal{P}_{xt} - W_\zeta \phi_x A \right). \quad (14)$$

⁸To help build intuition for (11), suppose there is a pricing measure $\tilde{\mathbb{Q}}$ (not necessarily equivalent to the risk-neutral measure since we do not enforce no-arbitrage) that gives rise to (10). Then K_{1x} governs the feedback of x_t under both the \mathbb{P} and $\tilde{\mathbb{Q}}$ measures. By analogy to standard risk-neutral pricing, the identical feedback matrices under \mathbb{P} and $\tilde{\mathbb{Q}}$ imply that x_t does not impact risk premiums.

The assumption of full spanning of ς_t^2 by y_t ensures that the leading matrix $W_\varsigma \phi_x C$ in (14) is invertible. Thus, up to an affine transformation, ς_t^2 is determined by $\mathcal{P}_{\varsigma t} - W_\varsigma B_{\mathcal{P}} \mathcal{P}_{xt}$.

To extract ς_t^2 from y_t , it remains to determine the weights $B_{\mathcal{P}}$. Here we highlight a very convenient representation of $B_{\mathcal{P}}$ in terms of the eigenvalues of K_{1x} . Without loss of generality, we can rotate x so that B_1 is a vector of ones and K_{1x} has the Jordan form. Whence B , and hence $B_{\mathcal{P}}$, are completely determined by the eigenvalues of K_{1x} . 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 $\mathcal{P}_{\varsigma t} - W_\varsigma B_{\mathcal{P}} \mathcal{P}_{xt}$.

Up to this point our derivations do not exploit the fact that ς^2 is a volatility process, beyond the autonomous structure of its conditional mean (8). To proceed with estimation and, most importantly, to ensure that ς_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 ς^2 follows a multivariate autoregressive gamma (*ARG*) process ([Gourieroux and Jasiak \(2006\)](#), [Le, Singleton, and Dai \(2010\)](#)), and the remaining $N - M$ factors x_t are Gaussian conditional on ς_t^2 :

$$\varsigma_{t+1}^2 | \varsigma_t^2 \sim ARG(\rho, c, \nu), \quad (15)$$

$$x_{t+1} - \phi_\varsigma \varsigma_{t+1}^2 | z_t \sim N \left(K_0 + K_{1\varsigma} \varsigma_t^2 + K_{1x} x_t, H_{0x} + \sum_{i=1}^M H_{ix} \varsigma_{it}^2 \right). \quad (16)$$

See [Appendix B](#) for a more detailed construction of the density $f(\varsigma_{t+1}^2 | \varsigma_t^2)$ and the definition of the parameters (ρ, c, ν) .

The ARG distribution is the discrete-time counterpart to the multivariate square-root diffusion ($A_M(M)$ process under \mathbb{P} of [Dai and Singleton \(2000\)](#)). Several recent *ETS*Ms (e.g., [Bollerslev, Tauchen, and Zhou \(2009\)](#) and [Drechsler and Yaron \(2011\)](#)) assume that the innovation in ς^2 is Gaussian, a specification that is logically inconsistent with the non-negativity of conditional variances. In contrast, by adopting an *ARG* process⁹ we ensure that ς_t^2 is an autonomous non-negative process governing the conditional variance of z_{t+1} and that the conditional mean of ς_{t+1}^2 is affine in ς_t^2 ($E_t[\varsigma_{t+1}^2] = \rho \varsigma_t^2 + \nu c$.) The positive semi-definite matrices H_{0x} and H_{ix} ($i = 1, \dots, M$) govern the time-varying volatility of x .

For econometric identification we normalize z so that, without loss of generality, x_t and ς_t^2 are conditionally independent ($\phi_\varsigma = 0$). Additionally, the intercepts K_0 in (16) are normalized to zero; and B_1 , the loadings on x_t for the yield on a one-period bond, are normalized to the row vector of ones. Further, we rotate x so that K_{1x} has the Jordan form; and c is fixed at $\frac{1}{2} \Delta t$. Finally, to prevent ς_t^2 from being absorbed at zero we require that $\nu \geq 1$. The joint density of (x_t, ς_t^2) given by (15) and (16) gives the conditional density of $(\mathcal{P}_{xt}, \mathcal{P}_{\varsigma t})$, after a Jacobian adjustment implied by (12) and (14). This is one of many equivalent normalization schemes that ensure that the extracted ς_t^2 span the $er_t^h(n)$.

⁹The following derivations are easily modified for other choices of non-negative affine distributions for ς_t^2 .

Having characterized the joint density of the N yield portfolios $(\mathcal{P}_{xt}, \mathcal{P}_{\zeta t})$, 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:

$$\mathcal{P}_{et} - W^e B_P \mathcal{P}_{xt} - W^e \phi_x (A + C \zeta_t^2) \sim N(0, \sigma_e^2 I_{J-N}). \quad (17)$$

The parameter set of the model is $(A, C, \rho, \nu, K_{1\zeta}, K_{1x}, H_{0x}, H_{ix}, \sigma_e)$. 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*.

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 $V_t \equiv \mathcal{P}_{\zeta t} - W_\zeta B_P \mathcal{P}_{xt} = \beta \mathcal{P}_t$. 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 ζ_t^2 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. Risk premiums are driven entirely by ζ_t^2 or, equivalently, V_t . 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.

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

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 “*.”

$xr_{t+h}^h(\tau)$ 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 (mean(xr)), the adjusted R^2




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.

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 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.¹⁰ 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

¹⁰<http://www.federalreserve.gov/Pubs/feds/2006/200628/200628abs.html>

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 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

Most empirical studies of arbitrage-free term structure models have assumed that $N \leq 4$ and, as summarized in [Table 1](#), this is also the case of many *ETSMs*. For our empirical analysis we set $N = 4$ and $M = 2$. This means that the set of volatility factors underlying the time variation in expected excess returns will also be $\mathcal{R} = 2$, again consistent with (or more general than) much of the extant literature.

While the information sets generated by z_t and any N linearly independent yield portfolios \mathcal{P}_t are (according to *ETSMs*) 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 are

$$\beta = \begin{bmatrix} 1 & 0 & 10.48 & 58.66 \\ & & (0.243) & (0.078) \\ 0 & 1 & -0.780 & -14.87 \\ & & (0.329) & (1.78) \end{bmatrix}, \quad (18)$$

where robust standard errors are given in parentheses. All four parameters are estimated with considerable precision. This was anticipated owing to the fact that the last 2×2 block

	$H = 2$			$H = 4$			$H = 6$		
	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	0.080***	0.083***	0.877	0.072***	0.071***	0.466	0.056***	0.059***	0.771
PC2	0.031**	0.021	0.571	0.037**	0.035	0.797	0.040***	0.042**	0.914
PC3	0.036***	0.022	0.241	0.039**	0.041*	0.848	0.042**	0.041*	0.637
PC4	0.150***	0.147***	0.620	0.157***	0.157***	0.737	0.126***	0.127***	0.629

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.

of β is fully determined by the Jordan form of K_{1x} which, in turn, is identified primarily from the cross-sectional restrictions in (10). 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 is to extracting V_t is both conceptually and practically very different than those typically pursued in the literature on *ETSMs*.

Does V_t Encompass the Information in \mathcal{P}_t about Yield Volatility?

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

$$E_t(\epsilon_{t+H}^2) = \text{constant} + \beta_{\mathcal{P}} \mathcal{P}_t \quad \text{and} \quad E_t(\epsilon_{t+H}^2) = \text{constant} + \beta_V V_t, \quad (19)$$

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 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 (19) for the forecast errors associated with each of the first four *PCs* comprising \mathcal{P}_t . Across all horizons H and all *PCs*— particularly *PC1*, arguably the most important driver of time varying volatility in yields— V_t captures most of the volatility information linearly spanned by \mathcal{P}_t . For example, for $H = 6$ and *PC1*, the adjusted R^2 from conditioning on \mathcal{P}_t (V_t) is 5.6% (5.9%). To formally evaluate the differences in fits, we conduct χ^2 tests of $H_0 : \beta_{\mathcal{P}} = \beta_V \beta$.¹¹ The probability values (“pvals”) confirm that the small differences in R^2 s are statistically insignificant.

¹¹The probability values of this test reported in Table 3 are robust to the sequential nature of our estimation.

Next, we pursue the alternative approach of constructing the errors in forecasting individual yields using the Blue-Chip financial survey forecasts (BCFF) of bond yields constructed by Wolters Kluwer. We presume that survey forecasts embody at least as much information as the shape of the yield curve (\mathcal{P}_t) and, hence, that the construction of yield-forecast errors ϵ_t from the BCFF data mitigates misspecification owing to omitted information from the first-stage forecasts based on the yield PC s.¹²

We cannot use the survey forecasts directly to construct ϵ_t , because they are forecasts of three-month moving averages of yields.¹³ 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.¹⁴ For each horizon h , $Q_{h-1,t+1}^{y^{(i)}} - Q_{h,t}^{y^{(i)}}$ is the revision of forecaster i within month $t + 1$. After trimming out extreme forecasts, we construct our aggregate squared “innovations” in yields by averaging the squared surprises across all forecasters and summing over all forecast horizons

$$\epsilon_{t+1}^2 = \sum_h \frac{1}{N_{h,t}} \sum_i (Q_{h-1,t+1}^{y^{(i)}} - Q_{h,t}^{y^{(i)}})^2, \quad (20)$$

where $N_{h,t}$ is the number of forecasters for horizon h and we include the horizons $h = 9, 12, 15$, and 18 months. *ETSM*s imply that the $Var_t[\epsilon_{t+h}]$ are affine in V_t , for all $h > 0$. Accordingly, for each yield maturity, we compare the projections

$$E_t(\epsilon_{t+H}^2) = \text{constant} + \beta_{\mathcal{P}} \mathcal{P}_t \quad \text{and} \quad E_t(\epsilon_{t+H}^2) = \text{constant} + \beta_V V_t, \quad (21)$$

again for $H = 2, 4$, and 6.¹⁵ As before, *RIETSM* implies that $\beta_{\mathcal{P}} = \beta_V \beta$.

The adjusted R^2 s for the projections in (21) 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.

Specifically, we account for the use of first-stage estimates of β as well as ϵ_t in estimating the regressions in (19). To account for the serial correlation of errors, we use the Newey-West estimates of the large-sample variance matrix with twelve lags.

¹²The descriptive analysis of Ludvigson and Ng (2010) identifies macro factors that have forecasting power for yields over and above the yields themselves, and Joslin, Pribsch, and Singleton (2011) develop an arbitrage-free term structure model that accommodates such macro forecast factors.

¹³Except 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.

¹⁴The one- and two-quarter forecasts are highly volatile and therefore omitted in our calculations.

¹⁵We 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	0.227***	0.206***	0.617	0.250***	0.219***	0.439	0.266***	0.230***	0.281
1-y	0.221**	0.215***	0.918	0.264***	0.244***	0.721	0.267***	0.242***	0.742
3-y	0.248***	0.243***	0.798	0.294***	0.266***	0.313	0.295***	0.270***	0.214
5-y	0.226***	0.227***	0.950	0.260***	0.249***	0.617	0.261***	0.249***	0.508
7-y	0.271***	0.266***	0.750	0.305***	0.287***	0.333	0.308***	0.287***	0.191
10-y	0.275***	0.269***	0.724	0.311***	0.296***	0.340	0.317***	0.295***	0.205

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.

Does V_t Encompass the Information in \mathcal{P}_t about Risk Premiums?

Next we examine whether V_t also encompasses the information in the yield curve (in the PCs \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 regressing 12-month excess returns on a 10-year bond on \mathcal{P}_t is 32.5%, compared to 28.1% when V_t is used as the predictor. This difference is not statistically significant (pval= 0.148). On the other hand, there is strong evidence that, for shorter maturity bonds, there is predictive information in yields that is not fully captured by V_t . The 3-month and 6-month excess returns on a 1-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, consistent with the implications of *ETSMs*, we find that two quantities of risk extracted from the yield curve using constraints implied by *RIETSM* fully encompass the information in the entire yield curve about interest rate volatility. Furthermore, we reach essentially the same conclusion about 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 .

That there is a statistically significant 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 is not entirely surprising in the light of 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 (2010)). Another influence on the shape of the intermediate segment of the Treasury yield

	$h = 3$ 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	0.159	0.250	0.196	0.081	0.367	0.315	0.379
xr(1)	0.165	0.099	0.004	0.313	0.213	0.003	NA	NA	NA
xr(3)	0.113	0.077	0.134	0.267	0.201	0.089	0.371	0.287	0.122
xr(5)	0.086	0.055	0.131	0.262	0.196	0.060	0.375	0.312	0.275
xr(7)	0.096	0.071	0.167	0.253	0.199	0.073	0.369	0.316	0.271
xr(10)	0.054	0.042	0.424	0.190	0.153	0.124	0.325	0.281	0.148

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.

curve was the hedging activities of mortgage traders (Duarte (2008)). Notwithstanding these considerations, 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 treasury yields.

Is V_t Capturing Inflation or Output Volatility that is Spanned by Bond Yields?

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. Notwithstanding these differences, 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 any of 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*. A distinguishing feature of our analysis is that we link measures of inflation volatility directly to the extracted V_t that, as we have shown, these *ETSMs* identify as the volatility factors that drive excess returns.

Though much of the econometric literature has focused on inflation directly,¹⁶ the monthly CPI inflation data appears quite noisy (choppy) and there are frequent, material revisions. To

¹⁶See, 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	0.283***	0.256***	0.130	0.334***	0.302***	0.066	0.376***	0.313***	0.068
e_{GDP}^2	0.329***	0.276***	0.032	0.403***	0.336***	0.011	0.449***	0.355***	0.002
$e_{CPI}e_{GDP}$	0.031	0.024	0.828	-0.005	0.000	0.190	0.006*	0.009	0.834

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 .

mitigate these measurement issues we construct a measure of inflation uncertainty using 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. Then we project the ϵ_{t+1}^2 in (20) (with Q^π in place of Q^y) onto \mathcal{P}_t and V_t as in (21) 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 $\sigma_{\pi t}^2$

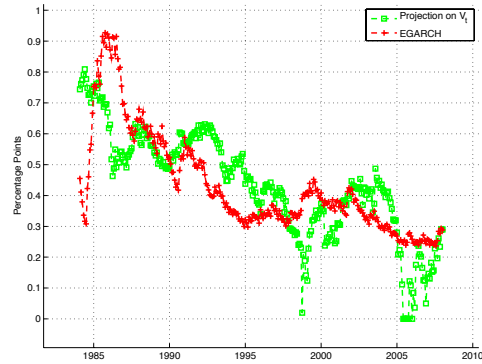
From the first row of Table 6 it is seen that V_t captures almost all of the time variation in the inflation volatility that is spanned by \mathcal{P}_t . For example, for $H = 2$, \mathcal{P}_t and V_t forecast squared average inflation residuals with adjusted R^2 s of 28.3% and 25.6%, respectively, 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, 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) consumption growth is homoskedastic and the source of time-varying risk is stochastic volatility in the surplus consumption ratio. Wachter (2006) and Le, Singleton, and Dai (2010) adopted similar



(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. Negative fitted values were truncated at zero.

specifications of the volatility structure in their *ETSMs*. 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 BCFF 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 (21) 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 ranging from 28% to 45%. Panel B of Figure 2 plots the V -implied volatility estimates together with their EGARCH counterparts. The downward trend in both series is consistent with the “Great Moderation” observed by Stock and Watson (2002) and others.

Although V_t accounts for roughly 80% of the conditional variation in real GDP growth that is spanned by yield PCs . While this is a substantial proportion, the higher predictive power of the PCs \mathcal{P}_t is statistically significant at conventional levels for all three choices of H . This suggests that there information in the yield curve about real economic risks that is not well captured by *ETSMs* with $N = 4$ or fewer state variables of which $\mathcal{R} = 2$ are the time-varying conditional variances of the state. While real GDP growth is far from a perfect substitute for consumption growth, the object of interest in most *ETSMs*, these findings are of interest because GDP growth is a key ingredient in the Federal Reserves setting of monetary policy (e.g., through the “Taylor rule”). These linkages have been explored by,

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2		pval	AdjR2		pval	AdjR2		pval
	4 PCs	V_t		4 PCs	V_t		4 PCs	V_t	
Panel A: Inflation									
3m	0.180	0.171***	0.716	0.215***	0.198***	0.702	0.243***	0.220***	0.647
1-y	0.102***	0.104***	0.720	0.125*	0.113***	0.891	0.149***	0.135***	0.595
3-y	0.098**	0.101***	0.625	0.124***	0.120***	0.848	0.152***	0.147***	0.688
5-y	0.106***	0.109***	0.625	0.129***	0.124***	0.703	0.159***	0.153***	0.485
7-y	0.108**	0.110***	0.332	0.128***	0.124***	0.549	0.152***	0.145***	0.515
10-y	0.117*	0.118**	0.535	0.135**	0.128***	0.641	0.156***	0.148***	0.600
Panel B: Real GDP									
3m	0.056**	0.060**	0.752	0.080*	0.077	0.848	0.098*	0.091	0.715
1-y	0.022*	0.026*	0.939	0.039*	0.045	0.942	0.045**	0.050	0.352
3-y	0.007*	0.011	0.392	0.019	0.025	0.975	0.024**	0.030	0.650
5-y	0.024**	0.028	0.130	0.038*	0.045	0.912	0.043*	0.048	0.874
7-y	0.035*	0.038*	0.073	0.052***	0.058**	0.922	0.056**	0.062*	0.986
10-y	0.036**	0.040	0.682	0.052***	0.057**	0.997	0.057**	0.062*	0.999

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 .

among others, [McCallum \(1994\)](#) and [Gallmeyer, Hollifield, and Zin \(2005\)](#) within *ETSMs* that explicitly incorporate the setting of a short-term rate by a monetary authority.

Pursuing these connections one step further, it is revealing to 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 *ETSMs*. 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.

Very 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 are quite strongly time-varying. Except for one instance, the χ^2 statistics of the joint significance test of the 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 *ESTM*.

In contrast, for the products between real GDP and yield forecast errors, the adjusted R^2 's in Panel B are substantially lower than their counterparts in Panel A, and 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.

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:

$$m_{t+1} - E_t[m_{t+1}] = \epsilon_{R,t+1} + \epsilon_{\pi,t+1}, \quad (22)$$

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

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

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:

$$\begin{aligned} \frac{1}{n}er_t^1(n+1) &\approx Cov_t(m_{t+1}, y_{n,t+1}) \\ &\approx Cov_t(\epsilon_{R,t+1}, y_{n,t+1}) + Cov_t(\epsilon_{\pi,t+1}, y_{n,t+1}). \end{aligned} \quad (24)$$

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 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 an important 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:

$$Cov_t(\epsilon_{R,t+1}, y_{n,t+1}) = Cov_t(\epsilon_{R,t+1}, B_{n,R}\epsilon_{R,t+1} + B_{n,\pi}\epsilon_{\pi,t+1}) \approx B_{n,R}Var_t(\epsilon_{R,t+1}). \quad (25)$$

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. As noted in [Section 1](#), this observation is central to 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 (USV) (e.g., [Collin-Dufresne and Goldstein \(2002\)](#), [Li and Zhao \(2006\)](#), and [Joslin \(2011\)](#)). Since, within *ETSMs* satisfying *RIETSM*, the ζ_t^2 driving excess returns are the sources of time-varying volatility in yields, the presence of USV 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.

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 USV for *ETSMs* we examine whether other conditioning information has predictive content for returns. We can always decompose excess returns $er_t^h(n)$ as

$$\underbrace{ny_{n,t} - hy_{t,h} - (n-h)E[y_{n-h,t+h}|y_t]}_{\text{spanned by } y_t} - \underbrace{(n-h)(E_t[y_{n-h,t+h}] - (n-h)E[y_{n-h,t+h}|y_t])}_{\text{unspanned by } y_t}. \quad (26)$$

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 construct three forecast factors: (i) the yield forecast factor; (ii) the GDP forecast factor; and (iii) the inflation forecast factor. At each point in time, these factors are simply the average forecasts across all forecasters and all forecast horizons. The yield forecast factor is also an average across all yield maturities. Consistent with the construction of our surprise measures, forecasts for the first two quarters are excluded from the averages. Assuming that market professionals condition on more information than y_t , these forecast factors may be informative about the last term in (26). Additionally, we also consider the realized jump means variable (JM) constructed by [Wright and Zhou \(2009\)](#) and extended by [Huang and Shi \(2011\)](#).¹⁷ JM is a measure of the time-varying amplitude of jumps in the US Treasury bond market based on high frequency data on 30-year Treasury bond futures.

¹⁷We thank the authors of these papers for kindly providing us with both the original and the extended data series.

	PC1-4	PC1-4, F		PC1-5	PC1-5, F		PC1-6	PC1-6, F	
	AdjR2	AdjR2	pval	AdjR2	AdjR2	pval	AdjR2	AdjR2	pval
Panel A: 3-month excess returns									
Yield forecast factor	0.088	0.091	0.298	0.121	0.122	0.377	0.156	0.153	0.943
GDP forecast factor	0.088	0.085	0.870	0.121	0.118	0.878	0.156	0.154	0.764
Inflation forecast factor	0.088	0.145	0.003	0.121	0.178	0.001	0.156	0.196	0.005
Mean jump amplitude	0.088	0.151	0.004	0.121	0.194	0.001	0.156	0.225	0.001
Panel B: 12-month excess returns									
Yield forecast factor	0.367	0.375	0.145	0.382	0.388	0.181	0.431	0.430	0.753
GDP forecast factor	0.367	0.370	0.636	0.382	0.384	0.651	0.431	0.436	0.533
Inflation forecast factor	0.367	0.572	0.000	0.382	0.587	0.000	0.431	0.603	0.000
Mean jump amplitude	0.367	0.555	0.000	0.382	0.583	0.000	0.431	0.625	0.000

Table 8: Regressions of cross-maturity average bonds excess returns on the first n yield PCs (PC1- n) and the forecast factors (F).

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 8, where “pval” is the probability value for the t-statistic of the estimated coefficient on the included conditioning variable. Neither the yield-forecast factor nor the GDP-forecast factor shows any significant predictive power in the presence of yield PCs. In contrast, both the inflation-forecast factor and the JM variable show substantial predictive power for risk premiums, much stronger than the yield PCs. Moreover, the same results can be observed across different returns horizons, and independent of the number of PCs (up to six) included in the regressions. Therefore, simply increasing the number of yield-based factors is unlikely to overturn this result.

Given this evidence, and in the light of the strong link between ζ_t^2 and expected excess returns in *ETSMs*, the only way *ETSMs* may rationalize these results is through the presence of USV. In the next subsection, we take up the issue of USV and investigate whether USV within *ETSMs* can indeed rationalize the evidence documented here.

USV in the Context of *ETSMs*

What are the implications of the strong predictive power of the inflation factor (IF) and JM for excess returns for the specification 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:

$$r_Z = \text{rank}([B, C]) < N. \quad (27)$$

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

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2			AdjR2			AdjR2		
	PC1-2	IF	pval	PC1-2	IF	pval	PC1-2	IF	pval
Panel A: Squared yield residuals									
3m	0.226	0.226	0.308	0.253	0.251	0.510	0.266	0.264	0.456
5-y	0.220	0.227	0.238	0.266	0.264	0.662	0.262	0.261	0.394
10-y	0.273	0.284	0.064	0.318	0.318	0.229	0.314	0.316	0.138
Panel B: Macro residuals									
e_{CPI}^2	0.286	0.309	0.118	0.338	0.350	0.244	0.369	0.379	0.261
e_{GDP}^2	0.322	0.320	0.086	0.392	0.392	0.406	0.433	0.431	0.299
$e_{CPI}e_{GDP}$	-0.005	-0.007	0.277	0.002	0.005	0.798	0.001	-0.003	0.871
Panel C: Cross inflation-yield residual products									
3m	0.178	0.182	0.155	0.205	0.207	0.314	0.243	0.241	0.369
5-y	0.100	0.118	0.029	0.128	0.131	0.141	0.165	0.163	0.218
10-y	0.109	0.124	0.068	0.135	0.137	0.349	0.164	0.162	0.533
Panel D: Cross GDP-yield residual products									
3m	0.062	0.059	0.676	0.077	0.073	0.210	0.095	0.095	0.142
5-y	0.032	0.029	0.842	0.042	0.039	0.879	0.046	0.045	0.909
10-y	0.042	0.038	0.830	0.050	0.047	0.733	0.057	0.055	0.829

Table 9: 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).

where the total number of risk factors is $N = 4$ and the number of volatility factors is $\mathcal{R} = 2$, this means that only the two non-volatility factors x_t are spanned by y_t . Further, following the reasoning in [Joslin, Le, and Singleton \(2012\)](#), this implies that the first two yield PCs fully span the information generated by x_t .

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

Now if $r_Z = 1$, then yields at all maturities are perfectly correlated (conditionally and unconditionally). Accordingly, with $N = 4$ in our empirical illustrations, 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. To be conservative 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 for squared residuals. For robustness, we also repeat the same analysis with $r_Z = 3$ (and $r_Z = 4$ to account for models with $N > 4$) and obtain very comparable results (not reported).

We report the projections of yield squared residuals onto the first two PCs and the IF (JM) variable in Panel A of [Table 9](#) ([Table 10](#)). For brevity, we only include three maturities (3 months, 5 years, and 10 years) as the results for other maturities are very similar. Interestingly, despite their power for excess returns, neither the IF nor the JM

	$H = 2$			$H = 4$			$H = 6$		
	AdjR2			AdjR2			AdjR2		
	PC1-2	JM	pval	PC1-2	JM	pval	PC1-2	JM	pval
Panel A: Squared yield residuals									
3m	0.226	0.237	0.223	0.253	0.268	0.068	0.266	0.282	0.086
5-y	0.220	0.217	0.401	0.266	0.266	0.133	0.262	0.262	0.127
10-y	0.273	0.273	0.501	0.318	0.316	0.774	0.314	0.312	0.967
Panel B: Macro residuals									
e_{CPI}^2	0.286	0.306	0.784	0.338	0.345	0.845	0.369	0.371	0.707
e_{GDP}^2	0.322	0.321	0.746	0.392	0.389	0.980	0.433	0.432	0.968
$e_{CPI}e_{GDP}$	-0.005	-0.009	0.643	0.002	-0.001	0.543	0.001	0.000	0.499
Panel C: Cross inflation-yield residual products									
3m	0.178	0.176	0.846	0.205	0.202	0.798	0.243	0.240	0.286
5-y	0.100	0.112	0.375	0.128	0.133	0.729	0.165	0.168	0.877
10-y	0.109	0.123	0.249	0.135	0.143	0.429	0.164	0.170	0.502
Panel D: Cross GDP-yield residual products									
3m	0.062	0.062	0.999	0.077	0.078	0.817	0.095	0.100	0.637
5-y	0.032	0.035	0.355	0.042	0.045	0.540	0.046	0.056	0.649
10-y	0.042	0.042	0.330	0.050	0.051	0.438	0.057	0.064	0.739

Table 10: 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.

variable displays any significant incremental predictive power for yield squared residuals. The adjusted R^2 s by adding the conditioning variables only increase modestly in most cases. All of the probability values (“pval”) for the t-statistics associated with the estimated coefficients of the conditioning variables are larger than 5%, and most of them larger than 10%. Placing this evidence in the context of *ETS*Ms, it suggests that the predictive power of both the IF and the JM variable does not seem to originate from unspanned *yield* volatility. We emphasize that this should not be interpreted as evidence against the existence of unspanned volatility, but rather that unspanned volatility of yields is unlikely behind the predictive power of either the IF or JM for bond excess returns.

To see how such a scenario might be plausible, imagine a scenario that corresponds to our empirical setup ($N = 4$, $\mathcal{R} = 2$, $r_Z = 2$) and assume that one of the non-volatility factors (x_{1t}) and one of the volatility factors (ς_{1t}^2) are unspanned (that is, their respective loadings on the yield pricing equations are uniformly zeros) and the remaining factors (x_{2t} , ς_{2t}^2) fully spanned. Furthermore, assume that the spanned volatility, ς_{2t}^2 , sufficiently modulates the time varying volatility of both spanned variables (x_{2t} , ς_{2t}^2), but both volatility factors are needed to fully capture the predictability of 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 yields PCs without showing any incremental power in predicting yield squared residuals.

Pursuing the above example, noting that the conditional variance of any variable within the *ETS*Ms must be driven by a linear combination of ς_{1t}^2 and ς_{2t}^2 , any economic variables

whose variances are sufficiently distinct from yields volatility must be relatively informative about the unspanned volatility factor ζ_{1t}^2 , and hence the source of predictive power of the IF factor and the JM variable. Guided by this intuition, we next investigate the possibility of unspanned inflation volatility and unspanned GDP volatility as potential sources for the predictive power of the two conditioning variables being considered. Similar analysis is conducted using inflation and GDP squared residuals as well as the product of inflation and GDP residuals constructed from earlier sections. The results are reported in Panel B of [Table 9](#) and [Table 10](#). Very similar to the results for yields volatility, both the IF and JM variables show mostly insignificant ability to predict the squared residuals (and the product of residuals) beyond the first two yield PCs. It is evident that neither unspanned inflation volatility nor unspanned GDP volatility is likely responsible for the predictability documented earlier in this section.

Finally, following an observation from the preceding section that within *ETSMs*, the conditional covariances between yields and inflation as well as the conditional covariances between yields and real GDP must also be linear in ζ_t^2 , we perform similar projections of the cross-products of inflation (GDP) and yield forecast residuals onto PC1-2 and the two conditioning variables being analysed. The results are reported in Panel C and Panel C of [Table 9](#) and [Table 10](#). Again, both the IF and JM variables fail to show any predictive power for these residuals products above the yield PCs.

Taken together, these findings cast doubt on the presumption of many *ETSMs* that variation in expected excess returns is induced entirely by time varying volatility of the state variables. The predictive power of an inflation factor is consistent with evidence reported in [Joslin, Pribsch, and Singleton \(2011\)](#) and [Cieslak and Povala \(2011\)](#). Our analysis goes further by showing that the impact of expected inflation on risk premiums is not operating through its affect on time-varying quantities of risk in the US Treasury bond market. A likely alternative channel is through its affects on the market prices of risk, and the latter are constant or nearly constant (at estimated parameter values) in most extant *ETSMs*.

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 equilibrium term structure models (*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 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 findings suggest that, to be consistent with U.S. macroeconomic and Treasury yield data, affine *ETSMs* should have the features that: (i) the fundamental sources of risks, including consumption growth, inflation, and yield volatilities are driven by distinct economic shocks; (ii) consumption growth risk alone is unlikely to fully account for the predictability of excess returns on bonds; and (iii) inflation risk, and not long-run risks or variation in risk premiums arising from habit-based preferences, is likely to be the dominant risk underlying risk premiums in U.S. Treasury markets.

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,¹⁸ 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 technique¹⁹ 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,²⁰ 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,²¹ 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 3 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 11**. Evidently, regardless

¹⁸For example, the 30-year series were replaced briefly by the 20-year series when the former was discontinued and resumed again.


¹⁹Specifically, we assume that the forward rates are piece-wise constant between any two successive maturities.

²⁰The 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.

²¹They 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 3: 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 $ARG(\rho, c, \nu)$ 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 $c_i > 0$, $\nu_i > 0$, $i = 1, \dots, N$, we define the conditional density of $\zeta_{i,t+1}^2$ given ζ_t^2 as the Poisson mixture of standard gamma distributions:

$$\frac{\zeta_{i,t+1}^2}{c_i} | (P, \zeta_t^2) \sim \text{gamma}(\nu_i + P), \quad \text{where } P | \zeta_t^2 \sim \text{Poisson}(\rho_i \zeta_t^2 / c_i). \quad (28)$$

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 11: 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:

$$f^{\mathbb{Q}}(\zeta_{i,t+1}^2|\zeta_t^2) = \frac{1}{c_i} \sum_{k=0}^{\infty} \left[\frac{\left(\frac{\rho_i \zeta_t^2}{c_i}\right)^k}{k!} e^{-\frac{\rho_i \zeta_t^2}{c_i}} \times \frac{\left(\frac{\zeta_{i,t+1}^2}{c_i}\right)^{\nu_i+k-1} e^{-\frac{\zeta_{i,t+1}^2}{c_i}}}{\Gamma(\nu_i+k)} \right]. \quad (29)$$

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

- 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.
- Bekaert, G. and E. Engstrom (2010). Asset return dynamics under band environment-good environment fundamentals. Technical report, Columbia 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.
- 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. 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 (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.
- 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, L., J. Heaton, and N. Li (2008). Consumption strikes back? measuring long-run risk. *Journal of Political Economy* 116, 260–302.

- 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*, *forthcoming*.
- Joslin, S., M. Priebisch, and K. Singleton (2011). Risk premiums in dynamic term structure models with unspanned macro risks. Technical report, Stanford University.
- Joslin, S., K. Singleton, and H. Zhu (2011). A new perspective on gaussian dynamic term structure models. *Review of Financial Studies* 24, 926–970.
- Koijen, R., H. Lustig, S. Van Nieuwerburgh, and A. Verdelhan (2010). Long-run risk, the wealth consumption ratio, and the temporal pricing of risk. *American Economic Review* 100, 552–556.
- Krishnamurthy, A. and A. Vissing-Jorgensen (2010). The aggregate demand for treasury debt. Technical report, Kellogg School, Northwestern University.
- 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.
- Nelson, C. and A. Siegel (1987). Parsimonious modelling of yield curves. *Journal of Business* 60, 473–489.
- 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 (2002). Has the business cycle changed and why? Technical report, NBER Macroeconomics Annual.
- Stock, J. and M. Watson (2007). Why has u.s. inflation become harder to forecast? *Journal of Money Credit and Banking* 39, 3–33.
- 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.