Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks*

Scott Joslin† Marcel Priebsch‡ Kenneth J. Singleton§

First Draft: April, 2009
This Draft: October 7, 2010

Abstract

This paper quantifies how variation in real economic activity and inflation in the U.S. influenced the market prices of level, slope, and curvature risks in U.S. Treasury markets. To accomplish this we develop a novel arbitrage-free $DTSM$ in which macroeconomic risks– in particular, real output and inflation risks– impact bond investment decisions separately from information about the shape of the yield curve. Estimates of our preferred macro-$DTSM$ over the twenty-three year period from 1985 through 2007 reveal that *unspanned* macro risks explained a substantial proportion of the variation in forward terms premiums. Unspanned macro risks accounted for nearly 90% of the conditional variation in short-dated forward term premiums, with unspanned real economic growth being the key driving factor. Over horizons beyond three years, these effects were entirely attributable to unspanned inflation. Using our model, we also reassess some of Chairman Bernanke's remarks on the interplay between term premiums, the shape of the yield curve, and macroeconomic activity.

---

*We are grateful for feedback from seminar participants at MIT, Stanford University, the University of Chicago, the Federal Reserve Board and Federal Reserve Bank of San Francisco, the International Monetary Fund, the Western Finance Association (San Diego) and for comments from Greg Duffee, Patrick Gagliardini, Imen Ghattassi, Monika Piazzesi, Oreste Tristani, and Jonathan Wright. An earlier version of this paper circulated under the title "Risk Premium Accounting in Macro-Dynamic Term Structure Models."

†MIT Sloan School of Management, `sjoslin@mit.edu`

‡Department of Economics, Stanford University, `priebsch@stanford.edu`

§Graduate School of Business, Stanford University, and NBER, `kenneths@stanford.edu`

1 Introduction

The cross-correlations of bond yields are well described by a low-dimensional factor model in the sense that the first three principal components (PCs) of bond yields – "level," "slope," and "curvature" – explain well over 95% of their variation (e.g., Litterman and Scheinkman (1991)). Very similar three-factor representations emerge from arbitrage-free, dynamic term structure models ($DTSM$s), at least for a wide range of maturities.[1] Yet in spite of the central role of level, slope, and curvature factors in both dynamic modeling and investment strategy, little is known about how macroeconomic shocks affect the market prices of these risks.

This paper quantifies how variation in real economic activity and inflation in the U.S. influenced the market prices of level, slope, and curvature risks in U.S. Treasury markets over period from 1985 through 2007. To accomplish this we develop a novel arbitrage-free $DTSM$ in which macroeconomic risks– in particular, real output and inflation risks– impact bond investment decisions separately from information about the shape of the yield curve. This is consistent with the descriptive evidence in Cooper and Priestley (2008) and Ludvigson and Ng (2009) that the state of the macroeconomy has predictive content for excess returns over and above the PCs of bond yields.[2] By accommodating macro risks that are *theoretically unspanned* by model-implied bond yields we also allow for the possibility that macro risks are distinct priced risks from yield-curve risks. This feature of our model cannot be replicated by $DTSM$s with macro variables that are unambiguously reflected in the yield curve (enter directly as risk factors), as such formulations imply that the macro factors are spanned by the model-implied PCs of bond yields.[3]

We find that *unspanned* macro risks explain a substantial proportion of the variation in forward terms premiums in U.S. Treasury markets. The effects are largest for forward loans initiated in the near future and they decline monotonically with the inception date of a loan out to about three years. Unspanned macro risks account for

[1] See, for instance, Dai and Singleton (2000) and Duffee (2002). Dai and Singleton (2002) and Piazzesi (2005) find that the addition of a fourth factor helps in capturing variation at the very short end of the yield curve owing (in part) to institutional features of the money markets. Our subsequent analysis is easily extended to accommodate a wider span of maturities and additional priced yield-curve risks.

[2] Complementary supporting evidence comes from $DTSM$s fit to yields alone where it has been found that the fourth and fifth PCs of bond yields forecast excess returns, but they contribute little to explaining the cross-sectional distribution of yields (e.g., Cochrane and Piazzesi (2005) and Duffee (2009a)). These higher-order PCs are correlated with macro information.

[3] Studies that enforce theoretical spanning include Ang and Piazzesi (2003), Ang, Dong, and Piazzesi (2007), Rudebusch and Wu (2008), Ravenna and Seppala (2007a), Smith and Taylor (2009), Bikbov and Chernov (2010), and Chernov and Mueller (2009).

nearly 90% of the conditional variation in short-dated forward term premiums and this effect asymptotes to about 20% for three-year horizons and beyond. Unspanned real economic growth is the key factor at short horizons while these effects are entirely attributable to unspanned inflation over horizons longer than three years. Using our model, we also reassess some of Chairman Bernanke's remarks on the interplay between term premiums, the shape of the yield curve, and macroeconomic activity.

Forward term premiums were large during the depths of recessions, consistent with a counter-cyclical pattern. But during much of the 1990's forward term premiums were procyclical with economic growth. In addition, the precipitous decline in long-dated forward term premiums during the period of Greenspan's conundrum is largely explained by economic forces that were orthogonal to our measures of economic growth and inflation, raising the possibility that economic mechanisms other than those captured in standard neo-Keynesian models of bond yields were important determinants of risk premiums in Treasury markets. Finally, shocks to long-dated forward term premiums affect real economic growth virtually entirely through their effects on its unspanned component. All of these results on unspanned macro risks and risk premiums are absent (by construction) from extant models that enforce theoretical spanning of M_t by bond yields.

Like the large empirical literature on $DTSM$s that precedes us, in implementing our model, we face the practical problem of having a large number of free parameters. To achieve parsimony researchers have arbitrarily set some parameters to zero or set those parameters to zero that have insignificant individual t-statistics based on a first-round analysis of a more flexible $DTSM$.[4] We propose a more systematic approach that uses likelihood-based model selection criteria to search over 2^{19} models for the "best" parsimonious parameterization of the risk premiums on exposures to the level, slope, and curvature risks. By design, our model selection exercise also addresses the near-cointegration of bond yields and macro factors, and the well-documented small-sample bias in estimated risk premiums in unconstrained Markov models.

While the literature on $DTSM$s is vast, we are unaware of any prior research that explores the relationship between unspanned macro shocks and risk premiums in bond markets within arbitrage-free pricing models. Independently, Duffee (2009a) proposes a latent factor (yields-only) model for accommodating unspanned risks in bond markets. However he does not explore the econometric identification of such a model, nor does he empirically implement a $DTSM$ with unspanned risks. We

[4]The first strategy is prevalent in the literature on macro-$DTSM$s, because of the large numbers of parameters governing market prices of risk and the $\mathbb{Q}$ distribution of the pricing factors in high-dimensional factor models; see, for example, Ang, Dong, and Piazzesi (2007). Dai and Singleton (2000) and Bikbov and Chernov (2010) are examples of papers that follow the second strategy.

formally derive a canonical form for Gaussian $DTSM$s with unspanned information that affects expected excess returns, and provide a convenient normalization that ensures econometric identification. Moreover, as we illustrate, the global optimum of the associated likelihood function is achieved extremely quickly. Wright (2009) and Barillas (2010) use our framework to explore the effects of inflation uncertainty on bond market risk premiums using international data, and optimal bond portfolio choice in the presence of macro-dependent market prices of risk, respectively.

2 A Canonical Gaussian Macro-$DTSM$ with *Unspanned* Macro Risks

Figure 1 plots the zero-coupon yield curves at the end of February, 1988 and February, 1991 (see Section 2.3 for details on our data.) Although the yield curves are similar, these dates represent distinctly different stages of the business cycle. February 1988 marked a period of high growth (in the top 10% of our sample based on our measure of economic growth), while February, 1991 marked a period of very low growth (the second lowest growth during our sample). The differences between the corresponding, one-year ahead realizations of the yield curves (lines with markers) seem large if in fact agents base forecasts solely on information about the current shape of the yield curve. On the other hand, if the state of the macroeconomy incrementally drives premiums that investors demand for bearing yield-curve risks, then large opposing changes from the current curve become more plausible. Indeed, the directions of the changes in Figure 1 are consistent with a model in which the level of the yield curve rises and the slope falls when real economic growth is high, and vice-versa when it is low. With these observations in mind, we proceed to develop a macro-$DTSM$ that attempts to capture these effects of macro information on risk premiums in bond markets over and above the PCs of bond yields.

Consider a standard (discrete-time) R-factor *Gaussian DTSM* in which the R-dimensional vector of pricing factors $\mathcal{P}_t$ and the short-rate r_t satisfy:

$Ar\mathbb{Q}:$ $r_t = \rho_0 + \rho_{\mathcal{P}} \cdot \mathcal{P}_t$, for scalar ρ_0 and R-vector $\rho_{\mathcal{P}}$; and

$A\mathcal{P}\mathbb{Q}:$ there is a pricing measure $\mathbb{Q}$ under which $\mathcal{P}_t$ follows the process

$$\mathcal{P}_t = K^{\mathbb{Q}}_{0\mathcal{P}} + K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}\mathcal{P}_{t-1} + \sqrt{\Sigma_{\mathcal{P}\mathcal{P}}}\epsilon^{\mathbb{Q}}_{\mathcal{P}t}, \tag{1}$$

where $\epsilon^{\mathbb{Q}}_{\mathcal{P}t} \sim N(0, I_R)$, the price of a τ-period zero coupon bond is $E^{\mathbb{Q}}_t[e^{-\sum_{s=0}^{\tau-1} r_{t+s}}]$, and for any price process, P_t, $\{e^{-\sum_{s=0}^{t-1} r_{t+s}} P_t\}$ is a $\mathbb{Q}$-martingale.

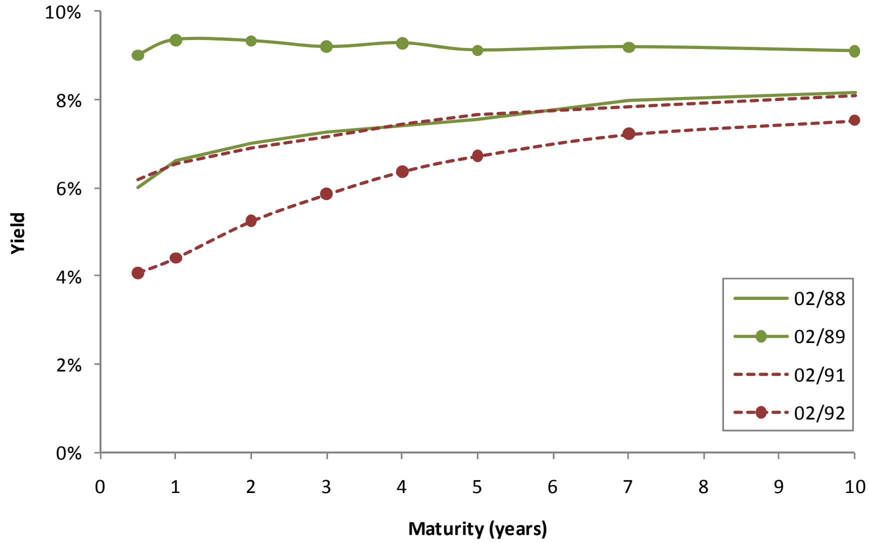

Figure 1: Yield Curves at the end of February, 1988 and February, 1991. The lines with markers show the corresponding yields curves one-year ahead.

Assumptions $Ar\mathbb{Q}$ and $A\mathcal{P}\mathbb{Q}$ ensure affine pricing so yields on zero-coupon bonds are affine functions of the pricing factors $\mathcal{P}$ as in Duffie and Kan (1996). The m-year zero coupon yield takes the form

$$y_{t,m} = A_m + B_m \cdot \mathcal{P}_t, \tag{2}$$

where (A_m, B_m) satisfy well-known Riccati difference equations (see Appendix A for a summary). We also suppose that $\mathcal{P}_t$ is chosen to be of minimal dimension to fully explain the cross-section of bond yields.[5]

Together, $Ar\mathbb{Q}$ and $A\mathcal{P}\mathbb{Q}$ give rise to a bond pricing model that subsumes most of the extant R-factor, Gaussian $DTSM$s in the literature. Regardless of the mix of latent, macroeconomic, or yield variables that enter the R-vector $\mathcal{P}$ of risk factors, the model-implied bond yields (2) are identical. Put differently, the pricing implications of canonical latent-factor $DTSM$s and macro-finance models in which macro variables enter as pricing factors are equivalent. What typically differentiates these $DTSM$s are the assumptions made about historical distribution of risks in the underlying macroeconomy.

[5]We take up the issue of measurement errors on bond yields in Sections 2.2 and 2.3.

Our interest is in how macro risk affects the risk premiums associated the changing shape of the yield curve. Accordingly, following Joslin, Singleton, and Zhu (2010) (hereafter JSZ), we rotate $\mathcal{P}$ so that it becomes the first R PCs of bond yields.[6] Then, letting M_t denote the set of macro variables of interest, we assume that the reduced-form of the N-vector $Z_t' = (\mathcal{P}_t', M_t')$ implied by the macroeconomy is:

$AZ\mathbb{P}$: Under the historical distribution $\mathbb{P}$, Z_t follows the VAR

$$\begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} = \begin{bmatrix} K_{0\mathcal{P}}^{\mathbb{P}} \\ K_{0M}^{\mathbb{P}} \end{bmatrix} + \begin{bmatrix} K_{\mathcal{P}\mathcal{P}}^{\mathbb{P}} & K_{\mathcal{P}M}^{\mathbb{P}} \\ K_{M\mathcal{P}}^{\mathbb{P}} & K_{MM}^{\mathbb{P}} \end{bmatrix} \begin{bmatrix} \mathcal{P}_{t-1} \\ M_{t-1} \end{bmatrix} + \sqrt{\Sigma_Z}\epsilon_{Zt}^{\mathbb{P}}, \tag{3}$$

where $\epsilon_{Zt}^{\mathbb{P}} \sim N(0, I_N)$, the $N \times N$ matrix Σ_Z is nonsingular, and $\Sigma_{\mathcal{P}\mathcal{P}}$ is the upper $R \times R$ block of Σ_Z.

Underlying $AZ\mathbb{P}$ is the presumption that there are at least N aggregate risks impinging on bond yields, that these risks are reflected in the innovation vector $\epsilon_Z^{\mathbb{P}}$, and that the resulting (reduced-form) data-generating process for Z_t is (3). Assumption $AZ\mathbb{P}$ allows for general feedback between $\mathcal{P}_t$ and M_t.[7] Additionally, it implies that M_t is *unspanned* by $\mathcal{P}_t$: knowledge of the shape the yield curve is not, by itself, sufficient to describe the effects of the aggregate shocks embodied in M_t on bond risk premiums.[8]

Notationally, we let $UMA_0^R(N)$ denote the family of Gaussian $DTSM$s with R pricing factors and $N - R$ unspanned macro conditioning variables, analogous to the notation of Dai and Singleton (2000).

2.1 A Canonical Form for $UMA_0^R(N)$

Having constructed a macro-$DTSM$ with unspanned macro risks, we now show that our formulation is canonical in that all N-factor, Gaussian $DTSM$s with unspanned macro risks are observationally equivalent to a $DTSM$ with our proposed structure.

[6] An equivalent pricing model is obtained with $\mathcal{P}_t$ equal to any R linearly independent combinations of y_t. For instance, once could select R distinct zero-coupon yields.

[7] In this respect, (3) is very similar to the descriptive six-factor model studied by Diebold, Rudebusch, and Aruoba (2006). As in their analysis,we emphasize the joint determination of the macro and yield variables (potential two-way feedback). We add the structure of a no-arbitrage pricing model so that it is possible to explore the properties of risk premiums in bond markets.

[8] More precisely, $\sigma(M_{it}) \subsetneq \sigma(\mathcal{B}_t \cup M_t^{(-i)})$, where $\mathcal{B}_t$ be the information in fixed income security prices and $M_t^{(-i)}$ is the vector of macro variables excluding the ith variable M_{it}. Here we use the notation $\sigma(\cdot)$ for a σ-field or information set. We define the information in fixed income prices at time t to be the σ-field generated by the prices of the payoffs $g(r_{t+t_1}, r_{t+t_2}, \ldots, r_{t+t_n})$; that is, by $\{P(Z_t) = E_t^{\mathbb{Q}}[g(r_{t+t_1}, r_{t+t_2}, \ldots, r_{t+t_n})] : g \in C_0\}$.

Moreover, there is a convenient, minimal set of normalizations for identification of our model with unspanned macro risks that facilitates finding the global optimum to the associated likelihood function.

Let y_t denote the J-vector of bond yields ($J > N$) to be used in assessing the fit of a macro-$DTSM$. The rotation that sets $\mathcal{P}$ to the first R PCs of these yields has the property that the parameters governing $(\rho_0, \rho_{\mathcal{P}}, K^{\mathbb{Q}}_{0\mathcal{P}}, K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}})$– those governing the pricing distribution for bond yields– are fully determined by the parameter set $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_{\infty})$, where $\lambda^{\mathbb{Q}}$ denotes the R-vector of ordered non-zero $\mathbb{Q}$-eigenvalues of $K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}}$,[9] and $r^{\mathbb{Q}}_{\infty}$ denotes the long-run mean of the short rate under $\mathbb{Q}$. The matrix $\Sigma_{\mathcal{P}\mathcal{P}}$ depends, of course, on the portfolio of yields comprising $\mathcal{P}$. $(\lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_{\infty})$ are rotation-invariant (that is, independent of the choice of pricing factors) and, hence, are economically interpretable parameters. Appendix B gives the explicit construction of $(\rho_0, \rho_{\mathcal{P}}, K^{\mathbb{Q}}_{0\mathcal{P}}, K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}})$ from $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_{\infty})$ for our choice of PCs as pricing factors.

The following proposition, proved in Appendix B, characterizes our canonical form for the family $UMA^R_0(N)$:[10]

Proposition 1 *Every canonical model for the family $UMA^R_0(N)$ is observationally equivalent to the following canonical Gaussian macro-DTSM: the state vector is $Z'_t = (\mathcal{P}'_t, M'_t)$, where $\mathcal{P}_t$ are the first R principal components of y_t; $r_t = \rho_0 + \rho_{\mathcal{P}} \cdot \mathcal{P}_t$; the $\mathbb{P}$ representation of Z_t is given by (3); and the $\mathbb{Q}$ representation of $\mathcal{P}$ is given by (1). Moreover, there is a unique mapping between the parameters governing $(\rho_0, \rho_{\mathcal{P}}, K^{\mathbb{Q}}_{0\mathcal{P}}, K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}})$ and the parameter set $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_{\infty})$. Thus, the parameter vector for the full model is $(K^{\mathbb{P}}_0, K^{\mathbb{P}}_1, \Sigma_Z, \lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_{\infty})$.*

2.2 What Do We Mean by Unspanned Macro Risks?

Both $\mathcal{P}_t$ and M_t subsume information about the primitive shocks impinging on the macroeconomy. Under $AZ\mathbb{P}$, the residual OM_t in the linear projection

$$M_t = \gamma_0 + \gamma_1' \mathcal{P}_t + OM_t \tag{4}$$

is also informative about these shocks. That is, the unspanned OM_t – the macro risks that cannot be replicated by portfolios of yields y_t or yield-PCs $\mathcal{P}_t$– embodies priced economic risks, so investors' pricing kernel cannot be represented in terms of

[9] As in JSZ, we can accommodate repeated and complex eigenvalues. As they show, a minor modification allows us to consider zero eigenvalues in the canonical form.

[10] This proposition is easily extended to the case where the first-order Markov specification in (3) is relaxed to allow Z_t to follow a higher-order VAR under $\mathbb{P}$. As we will see, this extension is not needed for our dataset and sample period.

$\mathcal{P}_t$ alone. Therefore, OM_t may materially affect risk premiums and it potentially has predictive power for future bond yields.

In contrast, following Ang and Piazzesi (2003), the vast majority of reduced-form macro-$DTSM$s adopt variants of the assumptions $Ar\mathbb{Q}$ and $A\mathcal{P}\mathbb{Q}$, include M_t in the set of R risk factors, say $\mathcal{P}_t^M$. Then, in place of $AZ\mathbb{P}$, it is assumed that the R-vector $\mathcal{P}_t^M$ follows a Gaussian VAR under $\mathbb{P}$. This formulation is observationally equivalent to one in which the risk factors are $\mathcal{P}_t$ (the first R PCs of bond yields), $\mathcal{P}_t$ follows a VAR under $\mathbb{P}$, and M_t is spanned by model-implied bond yields:[11]

$$M_t = \gamma_0 + \gamma_1' \mathcal{P}_t. \tag{5}$$

Economic environments that maintain (5) have the property that all aggregate risk impinging on the future shape of the yield curve can be fully summarized by the yield PCs $\mathcal{P}_t$. In particular, (5) implies that the past history of M_t is irrelevant for forecasting both future M and y, once one has conditioned on $\mathcal{P}_t$; and that the premiums that bond investors demand for bearing output growth and inflation risks (and any other factors in M_t) can be expressed as functions of $\mathcal{P}_t$ alone.

We stress that whether or not a macro-$DTSM$ embodies the spanning property (5) is wholly independent of the issue of errors in measuring either bond yields or macro factors. As typically parameterized in the literature, measurement errors are independent of economic agents' decision problems and, hence, of the economic mechanisms that determine bond prices.

Interestingly, the framework of Kim and Wright (2005), the model cited by Chairman Bernanke when discussing the impact of the macro economy on bond market risk premiums, breaks the perfect spanning condition (5) in a different way. Kim and Wright assume that M_t is inflation, and they arrive at their version of (4) by assuming that *expected* inflation is spanned by the pricing factors in the bond market. They additionally assume that $\mathcal{P}$ follows an autonomous Gaussian process under $\mathbb{Q}$ so their model and ours imply exactly the same bond prices. However, the $\mathbb{P}$-distribution of Z_t implied by their assumptions is:

$$\begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} = \begin{bmatrix} K_{0\mathcal{P}}^{\mathbb{P}} \\ \gamma_0 \end{bmatrix} + \begin{bmatrix} K_{\mathcal{P}\mathcal{P}}^{\mathbb{P}} & 0 \\ \gamma_1' K_{\mathcal{P}\mathcal{P}}^{\mathbb{P}} & 0 \end{bmatrix} \begin{bmatrix} \mathcal{P}_{t-1} \\ M_{t-1} \end{bmatrix} + \sqrt{\Sigma_Z} \begin{bmatrix} \epsilon_{\mathcal{P}t}^{\mathbb{P}} \\ \eta_t \end{bmatrix}, \tag{6}$$

where $\eta_t = (\nu_t + \gamma_1' \sqrt{\Sigma_{\mathcal{P}\mathcal{P}}} \epsilon_{\mathcal{P}t}^{\mathbb{P}})$. Thus, the Kim-Wright formulation leads to a constrained special case of (3) (Assumption $AZ\mathbb{P}$) under which the history of M_t has no

[11] See Joslin, Le, and Singleton (2010) for a proof of this equivalence. It holds unless there exists eigenvector of $K_{1\mathcal{P}}^{\mathbb{Q}}$ which is orthogonal to ρ_1 but not orthogonal to the column span of γ_1 (Joslin (2006)), which violates our assumption the $\mathcal{P}_t$ is minimal. Also, though we focus on first-order Markov processes under both $\mathbb{P}$ and $\mathbb{Q}$, an analogous spanning condition holds in extant models that allow for higher-order lags in the VAR representations of the risk factors.

forecasting power for futures values of M or $\mathcal{P}$, once one conditions on the history of $\mathcal{P}$. As we will see, the zero restrictions in (6) are strongly rejected in our data.

2.3 How Large Are Unspanned Macro Risks?

Some insight into the relative empirical plausibility of the relationships (4) and (5) comes from linear projections of observed values of the candidate macro factors M_t^o onto the PCs of observed bond yields, $\mathcal{P}_t^o$. The results will depend on the choice of M_t and the presumed number of pricing factors ($dim(\mathcal{P}_t) = R$).

For our subsequent empirical analysis we include measures of real economic activity (GRO) and inflation (INF) in M_t. INF is measured as the first PC of the monthly log difference of seasonally adjusted CPI (all items) and the personal consumption deflator.[12] This choice reflects the emphasis given to these aggregate price indices by the Federal Reserve in setting monetary policy. GRO is measured by the Chicago Fed National Activity Index, a measure of current real economic conditions.[13] We make the parsimonious choice of $M_t' = (GRO_t, INF_t)$ to encompass the two macro risks– measures of real economic growth and of inflation– that have received the most attention in prior studies. Ang et al. (2006) and Jardet, Monfort, and Pegoraro (2010) focus on models with $R = 3$ with output growth being the sole macro risk. Similarly, Kim and Wright (2005) explore $DTSM$s with $R = 3$ and expected inflation being the sole macro risk. Bikbov and Chernov (2010) and Chernov and Mueller (2009) examine models with $R = 4$ and with both output and inflation as macro risks. All of these prior studies enforce versions of the spanning condition (5).

Following the literature, we tie the choice of R to the cross-sectional factor structure of yields on bonds over the range of maturities being studied. Our sample extends from January, 1985 through December, 2007. There is substantial evidence that the Federal Reserve changed its policy rule during the early 1980's, following

[12] Price level indices evidence large transitory shocks, one source of which may be measurement error. To filter out this noise, we construct an exponentially decaying weighted average of past inflation, in the spirit of a hidden components model whereby true inflation follows an AR(1) process and observed inflation is equal to true inflation plus an *i.i.d.* measurement error. Wachter (2006) and Kim (2008) apply similar filters.

[13] The Chicago Federal Reserve constructs the CFNAI from economic indicators from the categories: production and income (23 series), employment and hours (24 series), personal consumption and housing (15 series), and sales, orders, and inventories (23 series). The data is inflation adjusted. The methodology used is similar to that used by Stock and Watson (1999) to construct their index of real economic acvitivity, and it is also related to the PCs of economic activity used by Ludvigson and Ng (2009) to forecast excess returns in bond markets. We also set GRO to an exponentially weighted average of past values of the CFNAI, in the spirit of the Chicago Fed's use of three-month moving averages of CFNAI when assessing the state of real economic activity.

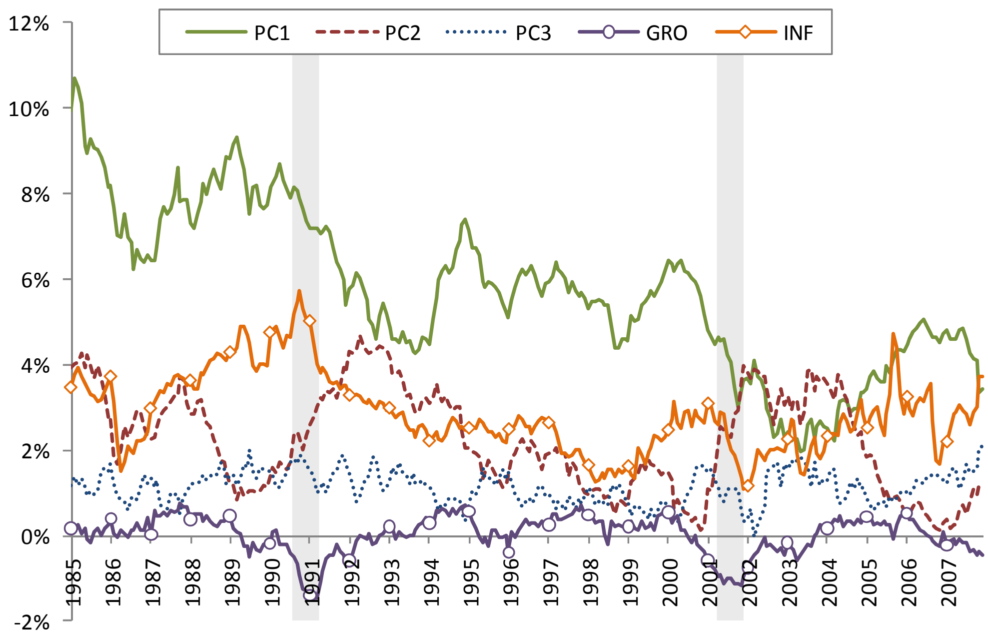

Figure 2: **Term Structure and Macro Variables** This figure plots the time series of $(PC1, PC2, PC3)$ of US Treasury-implied zero yields and smoothed GRO and INF. The vertical bars mark NBER recessions.

a significant policy experiment (Clarida, Gali, and Gertler (2000), Taylor (1999), and Woodford (2003)). Our starting date is well after the implementation of new operating procedures, and covers the Greenspan and early Bernanke regimes. The U.S. Treasury nominal zero-coupon bond yields comprising y_t have maturities six months and one through five, seven, and ten years $(J = 8)$.[14] Over 99% of the variation in these bond yields is explained by their first three PCs, so we set $R = 3$ and (invoking Proposition 1) normalize $\mathcal{P}_t$ to be the first three PCs of bond yields. The components of the state $Z_t' = (\mathcal{P}_t', GRO_t, INF_t)$ are displayed in Figure 2.[15]

[14] The zero curves for U.S. Treasury series were provided to us by Anh Lee of the University of North Carolina. He constructed these zero curves using the same bond selection criteria as in the Fama-Bliss data used in many previous studies. Importantly, we are using a consistent series out to ten years to maturity and throughout our sample period.

[15] Letting $\ell_{j,i}$ denote the loading on yield i in the construction of PCj, the PCs have been rescaled so that (1) $\sum_{i=1}^{8} \ell_{1,i}/8 = 1$, (2) $\ell_{2,10y} - \ell_{2,6m} = 1$, and (3) $\ell_{3,10y} - 2\ell_{3,2y} + \ell_{3,6m} = 1$. This puts all the PCs on similar scales. We then convert INF and GRO to an annual scale. Now all variables take on values in $[-5\%, 10\%]$.

For our data and sample period, the projection of GRO onto the first three PCs of yields gives an (adjusted) R^2 of 16%. Thus over 80% of the variation in GRO arises from risks distinct from $(PC1, PC2, PC3)$. Adding $PC4$ and $PC5$ as regressors only raises the R^2 for GRO to 34%. The comparable R^2s for INF are 55% $(PC1 - PC3)$ and 57% $(PC1 - PC5)$. Additionally, when projecting changes in GRO and INF onto changes in $(PC1-PC5)$ the R^2s are even smaller at 2.8% and 4.8%, respectively. We are confident that any errors in measuring bond prices are inconsequential for these projections, since the fitted PCs are virtually identical whether one assumes that these PCs are priced perfectly or are measured with errors within our $DTSM$s. We follow the extant literature on macro-$DTSM$s and assume that $M_t = M_t^o$, but even if one acknowledges some error in measuring M_t it does not seem plausible that such errors could explain the very low R^2 for say GRO.

Equally importantly, not only is OM_t large, but it has considerable predictive power for excess returns, over and above $\mathcal{P}_t$. For instance, for one-year holding period returns on two-year Treasury bonds in our data, the adjusted R^2 from the projection of excess returns onto $\mathcal{P}_t$ was 0.22, while onto $\{\mathcal{P}_t, GRO_t, INF_t\}$ it was 0.49. Consistent with this evidence, our canonical macro-$DTSM$ accommodates large unspanned components of GRO and INF that embody economically meaningful risks that impact expected excess returns on bonds separately from the yield PCs.

2.4 What Do Macro-$DTSM$s Reveal About the Relationship Between Unspanned-Macro and Yield-Curve Risks?

The no-arbitrage structure of a canonical macro-$DTSM$ reveals the market prices of the yield-curve risks $\mathcal{P}_t$ and their dependence on both spanned and unspanned macro risks, and allows maximum likelihood (ML) estimation of the parameters governing the $\mathbb{Q}$ distribution of $\mathcal{P}_t$. The $\mathbb{P}$ distribution of $(\mathcal{P}_t, M_t)$, on the other hand, can be studied within the structure of a macro-$DTSM$ or using an unconstrained VAR. Of interest then is whether a macro-$DTSM$ can shed new light on the effects of OM_t on risk premiums, beyond what is learned from a VAR. To help motivate our subsequent empirical analysis, this section briefly addresses these issues.

We start from the premise that the pricing kernel for the payoff space of all portfolios of zero-coupon bonds with weights that lie in the information set generated by $\{Z_t\}$ is fully characterized by the canonical form of Proposition 1.[16] The market price of the $\mathcal{P}_t$ risks is constructed from the drift $\mu_{\mathcal{P}}^{\mathbb{P}}(Z_t)$ of $\mathcal{P}_t$ under $\mathbb{P}$ (obtained

[16]That a pricing kernel exists for this payoff space under value additivity and other standard regularity conditions is established in Hansen and Richard (1987).

from (3)) and the drift $\mu^{\mathbb{Q}}_{\mathcal{P}}(\mathcal{P}_t)$ of $\mathcal{P}_t$ under $\mathbb{Q}$ (obtained from (1)) as follows:

$$\Lambda_{\mathcal{P}}(Z_t) = \Sigma_{\mathcal{P}\mathcal{P}}^{-1/2}\left(\mu^{\mathbb{P}}_{\mathcal{P}}(Z_t) - \mu^{\mathbb{Q}}_{\mathcal{P}}(\mathcal{P}_t)\right). \tag{7}$$

Using (7) we will quantitatively explore the effects of shocks to unspanned macro risks on expected excess returns and term premiums. Additionally, the market prices of the spanned macro risks are also identified in our model as they are computable linear combinations of the the market prices of the PC risks, $\Lambda_{\mathcal{P}}(Z_t)$.

On the other hand, the market prices of the *total*– spanned plus unspanned– macro risks are not econometrically identified, because nominal bond prices are not sensitive to the risk premiums that investors demand for bearing the unspanned macro risks.[17] This is a generic feature of all reduced-form macro-$DTSM$s constructed to price nominal zero-coupon bonds.

Thus, what we potentially gain from the construction and estimation of models in $UMA_0^R(N)$ are quantitative assessments of the effects of unspanned macro risks on $\Lambda_{\mathcal{P}}(Z_t)$, and more reliable empirical characterizations of the market prices of spanned macro risks. To the extent that unspanned macro risks are quantitatively important, models that assume that M_t can be replicated by a linear combination of $\mathcal{P}_t$ have misspecified both the market prices of spanned and unspanned macro risks.

Our macro-$DTSM$ with pricing factors $\mathcal{P}$ ($N = 3$) and state $Z_t' = (\mathcal{P}_t', M_t')$ ($N = 5$) is nested within a canonical macro-$DTSM$ in which the pricing factors are the state Z_t ($N = R = 5$). We do not explore this five-factor pricing model, because including M_t as pricing factors implies that M_t is spanned by the first five PCs of bond yields. We have already seen that this spanning condition is not supported by our data. Moreover, we know from JSZ and Duffee (2009b) that over-specifying the dimension of $\mathcal{P}$ is not innocuous. They found that estimated five-factor yield-only models of Treasury yields imply wholly implausible Sharpe ratios for certain portfolios of bonds. Estimation of a macro-$DTSM$ with Z_t as pricing factors ($R = 5$) would likely lead to similarly unreliable assessments of bond portfolio risk.[18]

Finally, can a macro-$DTSM$ yield insights about the $\mathbb{P}$ distribution of Z_t that cannot otherwise be discerned from studying a VAR model for Z_t? The answer is

[17]The market prices of unspanned inflation risk are potentially identified from TIPS yields, as in D'Amico, Kim, and Wei (2008) and Campbell, Sunderam, and Viceira (2009). However, the introduction of TIPS raises new issues related to illiquidity and data availability, so we follow most of the extant literature and focus on bond yields alone.

[18]Recall that a macro-$DTSM$ with pricing factors Z_t is observationally equivalent to one with the first five PCs as pricing factors and the spanning condition (5) enforced. Why the small-sample ML estimates of $DTSM$s with over-parameterized pricing distributions (extraneous pricing factors) give unreliable assessments of the risks of bond portfolios is an interesting question that we hope to take up in future research.

revealed by the form of the likelihood function for our canonical model. In specifying our model we adopt a first-order $\mathbb{P}-VAR$ model for $Z_t' = (\mathcal{P}_t', M_t')$, because all three of the Akaike (1973) (AIC), Hannan and Quinn (1979) (HQIC), and Schwarz (1978) Bayesian information (SBIC) criteria point to a first-order process among VARs with one through twelve lags (to capture possible seasonal and annual effects).[19] Suppose, in addition, that Z_t is measured without error and that $PC^{e\prime} \equiv (PC4, PC5, \ldots, PC8)$ is priced with *i.i.d.* $N(0, \Sigma_e)$ errors. Then the conditional density of (Z_t, PC_t^e) is:

$$\begin{aligned}\ell(Z_t, PC_t^e | Z_{t-1}; \Theta) &= \ell(PC_t^e | Z_t, Z_{t-1}; \Theta) \times \ell(Z_t | Z_{t-1}; \Theta) \\ &= \ell(PC_t^e | Z_t, Z_{t-1}; \lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}}, L_Z, L_e) \times \ell(Z_t | Z_{t-1}; K_Z^{\mathbb{P}}, K_0^{\mathbb{P}}, L_Z),\end{aligned} \tag{8}$$

where L_Z and L_e are the Cholesky factorizations of Σ_Z and Σ_e, respectively.

Absent any restrictions on Θ, the conditional density $\ell(Z_t | Z_{t-1}; \Theta)$ depends on $(K_0^{\mathbb{P}}, K_Z^{\mathbb{P}}, L_Z)$, but not on $(\lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}})$; whereas the density of PC_t^e depends only on the risk-neutral parameters and Σ_e. This separation of the parameter space implies, analogously to the results in JSZ, that ML estimates of the conditional mean (optimal forecasts) of Z_t are identical to those obtained from an unconstrained factor-VAR.[20] Furthermore, based on the analysis in Joslin, Le, and Singleton (2010), ML estimates of the entire $\mathbb{P}$ distribution of Z_t are likely to be nearly invariant to the imposition of no-arbitrage restrictions within any canonical model for the family $UMA_0^R(N)$.

What breaks this irrelevance of no-arbitrage restrictions in our subsequent empirical work are the restrictions across the $\mathbb{P}$ and $\mathbb{Q}$ distributions of Z_t that we are led to impose by our comprehensive model-selection exercise. Specifically, as we next explain, the likelihood criterion exploits the precision with which the cross-sectional distribution of bond yields pins down the parameters of the pricing distribution of Z_t to gain precision in estimating its historical distribution. The resulting macro-$DTSM$ gives a different and, we believe, more reliable characterization of the risk characteristics of bond portfolios than what is implied by an unconstrained factor-VAR.

[19] Our order-selection results are compatible with the higher-order VARs adopted by Ang, Dong, and Piazzesi (2007) and Jardet, Monfort, and Pegoraro (2010) for their macro-$DTSM$s with spanned macro factors. They assume that $N = R = 3$ and, thus, their state vector omits one or more yield PCs or macro variables relative to our Z_t.

[20] More precisely, the unconstrained factor-VAR that is the natural alternative to macro-$DTSM$ has the state Z_t following the Gaussian VAR of assumption $AZ\mathbb{P}$ and the PCs of bond yields satisfying the observation equation

$$\mathcal{P}_t^e = G + H\mathcal{P}_t + e_t, \quad e_t \sim N(0, \Sigma_e),$$

for conformable matrices G and H.

3 Model Selection and Risk Premium Accounting

Our canonical model for $UMA_0^R(N)$ has forty-five parameters governing the $\mathbb{P}$ distribution of Z (those comprising $K_0^{\mathbb{P}}$, $K_Z^{\mathbb{P}}$, and L_Z). There are four additional parameters governing the $\mathbb{Q}$ distribution of Z ($r_\infty^{\mathbb{Q}}$ and $\lambda^{\mathbb{Q}}$). Faced with such a large number of free parameters, standard practice has been to estimate a maximally flexible $DTSM$, set to zero many of the parameters in $(K_0^{\mathbb{P}}, K_Z^{\mathbb{P}})$ and $(K_{0\mathcal{P}}^{\mathbb{Q}}, K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}})$ that are statistically insignificant at a conventional significance level, and then to analyze the constrained model. The approach to model selection taken in our analysis is both more focused and more systematic in that we use formal model selection criteria to pick our preferred parsimonious model of risk premiums.

In comparing models we assume that the first three PCs of yields are priced perfectly by our $DTSM$. This, or an analogous assumption for individual yields, is common in the literature (e.g., Dai and Singleton (2000) and Collin-Dufresne, Goldstein, and Jones (2008)). Others (e.g., Ang, Dong, and Piazzesi (2007), Duffee (2009a)) have assumed instead that all of the yield PCs are measured with errors. For unconstrained versions of our canonical macro-$DTSM$ we found that virtually all of the quantitative properties of our estimated models were robust to inclusion of measurement errors on the PCs.[21] For this reason, and to reduce the computational burden of model selection, we henceforth focus on the model in which the observed and model-implied $\mathcal{P}_t$ coincide.

3.1 Constraining the Market Prices of PC Risks

The scaled market prices of $\mathcal{P}$ risk, $\Sigma_{\mathcal{P}\mathcal{P}}^{1/2}\Lambda_{\mathcal{P}}(Z_t)$, depend on the fifteen parameters of the matrix $\mathcal{K} \equiv (K_{\mathcal{P}Z}^{\mathbb{P}} - [K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}\ 0_{3\times 2}])$ governing state-dependence, where $K_{\mathcal{P}Z}^{\mathbb{P}}$ is the first three rows of $K_Z^{\mathbb{P}}$, and the three intercept terms. We address two distinct aspects of model specification with our selection exercise. First, we seek the *best* set of zero restrictions on these eighteen parameters governing $\Sigma_{\mathcal{P}\mathcal{P}}^{1/2}\Lambda_{\mathcal{P}}(Z_t)$, where we trade off fit against the costs of over-parameterization. Our selection strategy is not implementable outside of a $DTSM$ as both $K_Z^{\mathbb{P}}$ and $K_Z^{\mathbb{Q}}$ must be estimable.

We show in Appendix C that, to a first-order approximation, our constraints on $\mathcal{K}$ can be interpreted directly as constraints on the expected excess returns to pure exposures to the $\mathcal{P}$ risks. That is, the constraints on the first row of $\mathcal{K}$ can be interpreted as constraints on excess returns on the portfolio whose value changes

[21] Similarly, JSZ and Joslin, Le, and Singleton (2010) find that adding measurement errors on the PC-based pricing factors has negligible effects on the goodness-of-fit of Gaussian yield-only $DTSM$s and macro-$DTSM$s with spanned macro risks, respectively.

(locally) one-to-one with changes in $PC1$, but whose value is unresponsive to changes in $PC2$ or $PC3$. Similar interpretations apply to the second and third rows of $\mathcal{K}$ for $PC2$ and $PC3$, respectively. By examining the behavior of the expected excess returns on these PC-mimicking portfolios, $xPCj_t$ ($j = 1, 2, 3$), we gain a new and informative perspective on the nature of priced risks in Treasury markets. This economic interpretation of the parameter constraints highlights a benefit of our canonical form; no such model-free interpretation is possible for similar constraints in a latent factor model. Moreover, by considering constraints on $\mathcal{K}$ rather that $\Lambda_{\mathcal{P}}$, the interpretation of our constraints does not rest on an arbitrary factorization of the innovation covariance matrix.

Second, in applying these selection criteria we are mindful of the near unit-root behavior of $\mathcal{P}$ under both $\mathbb{P}$ and $\mathbb{Q}$. There is substantial evidence that bond yields are nearly cointegrated (e.g., Giese (2008), Jardet, Monfort, and Pegoraro (2010)). We also find that $PC1$, $PC2$, and INF exhibit behavior consistent with a near cointegrating relationship, whereas $PC3$ and GRO appear stationary. We do not believe that $(PC1, PC2, INF)$ literally embody unit-root components. At the same time it is desirable to enforce a high degree of persistence under $\mathbb{P}$, since ML estimators of drift parameters are known to be biased in small samples. This bias tends to be proportionately larger the closer a process is to a unit root process (Phillips and Yu (2005), Tang and Chen (2009)).

Moreover, when $K_Z^{\mathbb{P}}$ is estimated from a VAR, its largest eigenvalue tends to be sufficiently below unity to imply that expected future interest rates out ten years or longer are virtually constant (see below). This is inconsistent with surveys on interest rate forecasts (Kim and Orphanides (2005)),[22] and leads to the attribution of too much of the variation in forward rates to variation in risk premiums.

To address this persistence bias we exploit two robust features of $DTSM$s: the largest eigenvalue of $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$ tends to be close to unity, and the cross-section of bond yields precisely identifies the parameters of the $\mathbb{Q}$ distribution (in our case, $r^{\mathbb{Q}}$ and $\lambda^{\mathbb{Q}}$). Any zero restrictions in $\mathcal{K}$ called for by our model selection criteria effectively pull $K_Z^{\mathbb{P}}$ closer to $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$, so the former may inherit more of the high degree of persistence inherent in the latter matrix. In addition, we call upon our model selection criteria to evaluate whether setting the largest eigenvalues of the feedback matrices $K_Z^{\mathbb{P}}$ and $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$ equal to each other improves the fit of our macro-$DTSM$. Through both channels we are effectively examining whether the high degree of precision with which the cross-section of yields pins down $\lambda^{\mathbb{Q}}$ is reliably informative about the degree of persistence in the data-generating process for Z_t. Again, this exploration is not

[22]Similar considerations motivated Cochrane and Piazzesi (2008), among others, to enforce even more persistent unit-root behavior under $\mathbb{P}$ in their models.

possible absent the structure of a $DTSM$.

3.2 Selecting Among 2^{19} Parameterizations

Since there are eighteen free parameters governing $\Sigma_{\mathcal{PP}}^{1/2}\Lambda_{\mathcal{P}}(Z_t)$, there are 2^{18} possible configurations of $DTSM$s with some of risk-premium parameters set to zero. We examine each of these models with and without the eigenvalue constraint across $K_Z^{\mathbb{P}}$ and $K_{\mathcal{PP}}^{\mathbb{Q}}$, so a total of 2^{19} specifications. Though 2^{19} is large, the rapid convergence to the global optimum of the likelihood function obtained using our normalization scheme makes it feasible to undertake this search using formal model selection criteria. For each of the 2^{19} specifications examined, we compute full-information ML estimates of the parameters and then evaluate the AIC, HQIC, and SBIC information criteria.[23] The criteria HQIC and SBIC are consistent (i.e., asymptotically they select the correct configuration of zero restrictions on $\mathcal{K}$), while the AIC criterion may asymptotically over fit (have too few zero restrictions) with positive probability.[24]

The resulting frontiers of maximal values of the log-likelihood function achieved for each choice of the number of zero restrictions on $\mathcal{K}$, with and without the eigenvalue ("EV") restriction imposed, are displayed in Figure 3. The tangent points of the HQIC criterion is at twelve restrictions: eleven zero restrictions on $\mathcal{K}$ and the EV constraint. SBIC calls for exactly the same set of restrictions as the HQIC criterion plus a zero on the loading for $PC2$ in the market price of $PC1$ (level) risk, though the values of SBIC for the cases of twelve and thirteen restrictions are very close. Notably, the consistent HQIC and SBIC criteria call for enforcing near-cointegration through the EV constraint. The AIC criterion's selection of a much less parsimonious model, including relaxation of the EV constraint, suggests that in our application it is selecting too few restrictions. Indeed, when we apply the AIC criterion to select among all 2^{18} models with the EV constraint enforced, it selects the model with precisely the same zero restrictions chosen by the HQIC criterion.

Based on these results, we proceed to investigate in more depth the macro-$DTSM$ that enforces the EV constraint along with the common set of eleven restrictions on the market prices of the yield-curve risks.

---

[23] Bauer (2010) proposes a complementary approach to model selection based on the posterior odds ratio from Bayesian analysis. Owing to the computational complexity of his approach, an intermediate step is inserted to narrow down the set of models to be compared. Additionally, Bauer considers a standard $A_0(3)$ $DTSM$, so there is no macro conditioning information used in estimation or in specifying his market prices of risks, and he does not address near-cointegration.

[24] These properties apply both when the true process is stationary and when it contains unit roots, as is discussed in Lutkepohl (2005), especially Propositions 4.2 and 8.1.

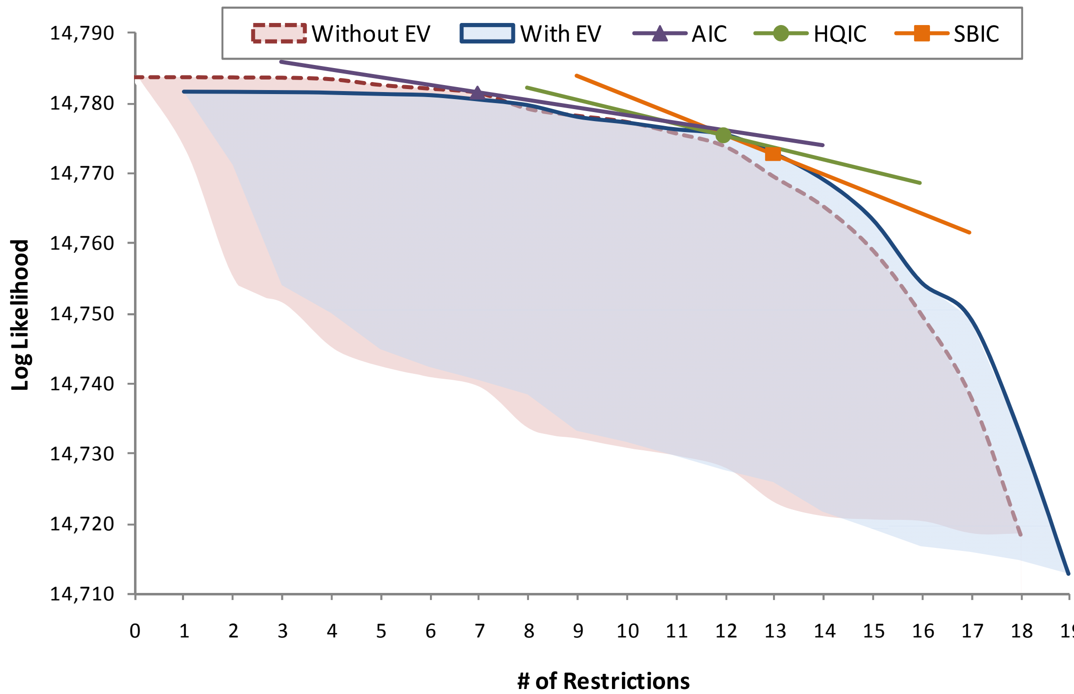

Figure 3: Frontiers of values of the log-likelihood function over 2^{19} specifications of our macro-$DTSM$, along with the tangency points for the AIC, HQIC, and SBIC information criteria. "EV" refers to the eigenvalue constraint across $K_Z^{\mathbb{P}}$ and $K_{\mathcal{PP}}^{\mathbb{Q}}$.

3.3 Risk Premium Accounting: Model Comparison

To highlight the properties of the selected model we compare the properties of four Gaussian $DTSM$s with $(R = 3, N = 5)$: the unconstrained canonical model ($\mathcal{CM}$); model $\mathcal{CM}$ with the constraint EV that sets the largest eigenvalue of $K_Z^{\mathbb{P}}$ equal to the corresponding eigenvalue of $K_{\mathcal{PP}}^{\mathbb{Q}}$ ($\mathcal{CM}_E$); model $\mathcal{CM}$ with the eleven zero restrictions on $\mathcal{K}$ ($\mathcal{CM}^0$); and our preferred model that imposes both the EV and the eleven zero restrictions on $\mathcal{K}$ ($\mathcal{CM}_E^0$).

Maximum likelihood estimates of the parameters governing the $\mathbb{Q}$ distribution of Z_t from model $\mathcal{CM}_E^0$ are displayed in the first column of Table 1.[25] The estimates for the other three models are virtually indistinguishable from these estimates, typically differing in the fourth decimal place. This says that the parameters of the $\mathbb{Q}$ distribution are determined largely by the cross-sectional restrictions on bond yields,

[25]Throughout our analysis asymptotic standard errors are computed by numerical approximation to the Hessian and using the delta method.

Param	$\mathcal{CM}_E^0$	Param	$\mathcal{CM}$	$\mathcal{CM}^0$	$\mathcal{CM}_E$	$\mathcal{CM}_E^0$
$r_\infty^{\mathbb{Q}}$	0.0925	$\lvert\lambda_1^{\mathbb{P}}\rvert$	0.9735	0.9904	0.9971	0.9970
	(0.0058)		(0.0113)	(0.0081)	(0.0005)	(0.0005)
$\lambda_1^{\mathbb{Q}}$	0.9970	$\lvert\lambda_2^{\mathbb{P}}\rvert$	0.9532	0.9525	0.9515	0.9514
	(0.0005)		(0.0170)	(0.0107)	(0.0185)	(0.0107)
$\lambda_2^{\mathbb{Q}}$	0.9649	$\lvert\lambda_3^{\mathbb{P}}\rvert$	0.9532	0.9525	0.9515	0.9514
	(0.0026)		(0.0170)	(0.0107)	(0.0185)	(0.0107)
$\lambda_3^{\mathbb{Q}}$	0.8875	$\lvert\lambda_4^{\mathbb{P}}\rvert$	0.9305	0.8803	0.9312	0.8808
	(0.0119)		(0.0329)	(0.0146)	(0.0342)	(0.0143)
		$\lvert\lambda_5^{\mathbb{P}}\rvert$	0.7593	0.8803	0.7600	0.8808
			(0.0409)	(0.0146)	(0.0410)	(0.0143)

Table 1: ML estimates of the $\mathbb{Q}$ parameters for model $\mathcal{CM}_E^0$, and of the moduli of the eigenvalues of $K_Z^{\mathbb{P}}$ for models $\mathcal{CM}$, $\mathcal{CM}^0$, $\mathcal{CM}_E$, and $\mathcal{CM}_E^0$. Asymptotic standard errors are given in parentheses.

and not by their time-series properties under the $\mathbb{P}$ distribution. Models $\mathcal{CM}_E$ and $\mathcal{CM}_E^0$ exploit this precision to restrict the degree of persistence of Z_t under $\mathbb{P}$.

There are two basic features of a $DTSM$ that determine the $\mathbb{P}$ distribution of bond yields: the model-implied loadings of the yields onto the pricing factors $\mathcal{P}_t$ in (2), and the $\mathbb{P}$ distribution of $\mathcal{P}_t$. The loadings are fully determined by the $\mathbb{Q}$ parameters $(r_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, \Sigma_{\mathcal{PP}})$ (Appendix A). We have just seen that the parameters $(r_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$ are nearly identical across models and, as it turns out, so are the ML estimates of $\Sigma_{\mathcal{PP}}$. Consequently, the loadings (A_m, B_m) are also (essentially) indistinguishable across the four models examined.

Thus, any differences in the model-implied risk premiums must be attributable to different estimates of the feedback matrix $K_Z^{\mathbb{P}}$. The eigenvalues of $K_Z^{\mathbb{P}}$ are displayed in the remaining columns Table 1.[26] The largest $\mathbb{P}$-eigenvalue in the canonical model $\mathcal{CM}$ is smaller than in the constrained models. The relatively small value implies that expected future short-term rates beyond ten years are (nearly) constant in model $\mathcal{CM}$. Equivalently, it implies, counterfactually, that virtually all of the variation in long-dated forward rates arises from variation in risk premiums.

Comparison of model $\mathcal{CM}^0$ to models $(\mathcal{CM}_E, \mathcal{CM}_E^0)$ sheds light on how the model-selection criteria use the relative precision in estimating $K_{\mathcal{PP}}^{\mathbb{Q}}$ to select constraints

[26] The fact that there are pairs of equal moduli in all three models means that there are complex roots in $K_Z^{\mathbb{P}}$. The complex parts were small in absolute value.

on $K_Z^{\mathbb{P}}$. Enforcing the eleven zero restrictions in model $\mathcal{CM}^0$ increases the largest eigenvalue of $K_Z^{\mathbb{P}}$ from 0.973 to 0.990, and thus closes most of the gap between models $\mathcal{CM}$ and $\mathcal{CM}_E^0$. Not only does model $\mathcal{CM}^0$ lie on the optimal "Without EV" frontier in Figure 3, it implies that Z_t is sufficiently persistent under $\mathbb{P}$ for long-dated forecasts of the short rate to display considerable time variation. A further increase in the largest eigenvalue comes with the additional EV constraint in model $\mathcal{CM}_E^0$, as called for by the criteria (HQIC, SBIC).

Estimates for models $\mathcal{CM}_E$ and $\mathcal{CM}_E^0$ of the matrix $\mathcal{K}$ and the intercept vector $(K_{0\mathcal{P}}^{\mathbb{P}} - K_{0\mathcal{P}}^{\mathbb{Q}})$ governing one-month expected excess returns on PC risk exposures are displayed in Table 2. Recall that the ordering of the variables in Z_t, corresponding to the columns in Table 2, is $(PC1, PC2, PC3, GRO, INF)$. Focusing first on model $\mathcal{CM}_E^0$ (Table 2b), we see that the first and second rows of $\mathcal{K}$ have (statistically significant) non-zero entries, while the last row is chosen to be zero by our model selection criteria, as are all three entries of $(K_{0\mathcal{P}}^{\mathbb{P}} - K_{0\mathcal{P}}^{\mathbb{Q}})$. It follows that exposures to $PC1$ and $PC2$ risks are priced, but exposure to $PC3$ risk is not priced, at the one month horizon and during our sample period.[27] The finding that both level and slope risks are priced differs from Cochrane and Piazzesi (2008) (only level risk is priced), but agrees with Duffee (2009b). Both of these analyses omit macro information which we find assists in identifying the prices of these yield-curve risks.

Specifically, the expected excess returns $xPC1_t$ and $xPC2_t$ both depend in statistically significant ways on $PC1$ and GRO. Expected excess returns on $PC1$ risk in model $\mathcal{CM}_E^0$ are also influenced by $PC2$ (slope) and inflation. Note that the effect of INF on $xPC1$ in model $\mathcal{CM}_E^0$ is larger and more precisely estimated than its counterpart in model $\mathcal{CM}_E$.

The positive signs on GRO and INF imply that risk premiums on $PC1$ exposures are pro-cyclical (positively correlated with GRO and INF). This can be seen graphically in Figure 4a for models $\mathcal{CM}_E$ and $\mathcal{CM}_E^0$, where the shaded areas represent the NBER-designated recessions. Exposures to $PC1$ lose money when rates fall, which is when investors holding long bond positions make money. This explains the predominantly negative expected excess returns on the level-mimicking portfolio in Figure 4a. There is broad agreement across these two models about the (annualized) expected excess return on a level-mimicking portfolio, $xPC1_t$. These excess returns take on their largest (absolute) values during the 1990 and 2001 recessions.

[27] An alternative approach to analyzing risk premiums within a macro-$DTSM$ would have been to adapt the methods in JSZ to enforce the constraint that risk premiums lie in a two-dimensional space. This approach would have let us proceed without taking a stand on which of the risks $\mathcal{P}$ are priced. However, this two-dimensional restriction is a much weaker restriction on $\mathcal{K}$ than that of the *best* model chosen by our selection criteria.

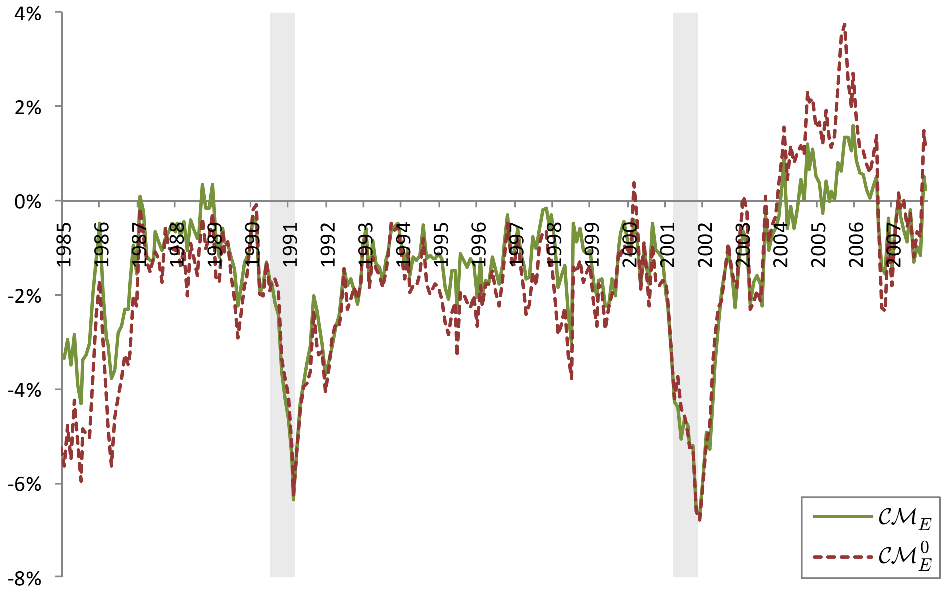

(a) Excess Return on Level-Mimicking Portfolio

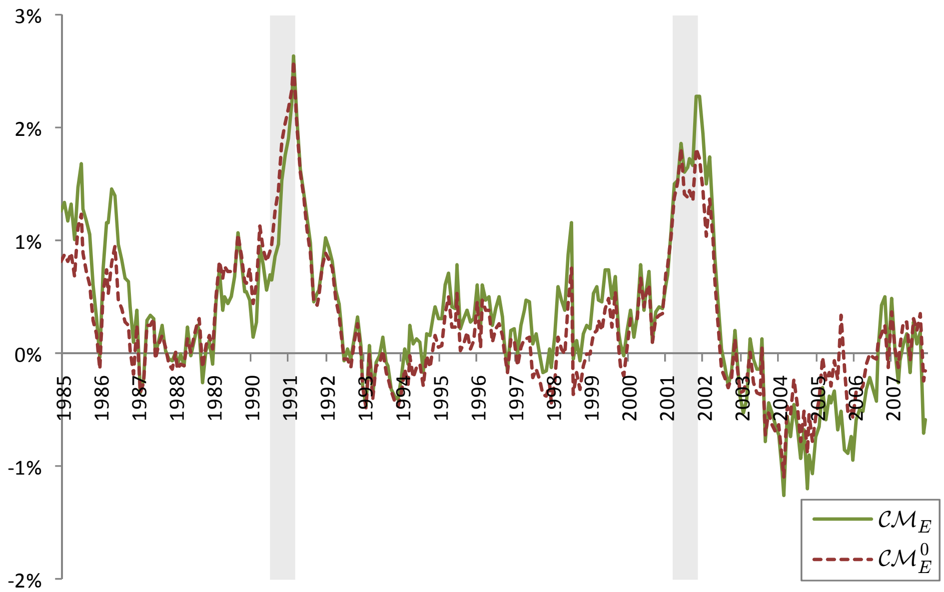

(b) Excess Return on Slope-Mimicking Portfolio

Figure 4: Expected excess returns on the level- and the slope-mimicking portfolios implied by models $\mathcal{CM}_E$ and $\mathcal{CM}_E^0$.

$\mathcal{P}$	const	$PC1$	$PC2$	$PC3$	GRO	INF
$PC1$	-0.0007	-0.0231	-0.0207	0.0525	0.1518	0.0394
	(0.0006)	(0.0189)	(0.0180)	(0.0578)	(0.0445)	(0.0373)
$PC2$	0.0002	0.0373	0.0064	-0.0712	-0.1782	-0.0416
	(0.0008)	(0.0120)	(0.0171)	(0.0501)	(0.0377)	(0.0309)
$PC3$	0.0012	-0.0265	-0.0053	-0.0984	-0.0092	0.0492
	(0.0007)	(0.0132)	(0.0111)	(0.0433)	(0.0329)	(0.0223)

(a) Risk premium parameters for model $\mathcal{CM}_E$

$\mathcal{P}$	const	$PC1$	$PC2$	$PC3$	GRO	INF
$PC1$	0	-0.0481	-0.0248	0	0.1560	0.0856
		(0.0098)	(0.0107)		(0.0364)	(0.0195)
$PC2$	0	0.0224	0	-0.0803	-0.1641	0
		(0.0056)		(0.0280)	(0.0336)	
$PC3$	0	0	0	0	0	0

(b) Risk premium parameters for model $\mathcal{CM}_E^0$

Table 2: ML estimates of the parameters $\mathcal{K}$ governing expected excess returns on the PC mimicking portfolios. Standard errors are given in parentheses. Zeros from our model selection criteria appear in the $\mathcal{CM}_E^0$ estimates in the lower panel.

The negative sign on GRO in the second row of $\mathcal{K}$ implies that premiums on exposure to slope risk are counter-cyclical. It is striking that, after searching over 2^{19} specifications of the model-implied $xPC2_t$, our model-selection criteria place most of the weight on GRO and zero weights on $(PC2, INF)$ in characterizing expected excess returns on exposure to slope risk. These risk premiums also depend on the curvature ($PC3$) of the Treasury curve.

The premium on $PC2$ risk achieves its lowest value, and concurrently the premium on $PC1$ risk achieves its highest value, during 2004/05. Between June, 2004 and June, 2006 the Federal Reserve increased its target Federal Funds rate by 4% (from 1.25% to 5.25%). Yields on ten-year Treasuries actually fell during this time, leading to a pronounced flattening of the yield curve, what Chairman Greenspan referred to as a conundrum. We revisit these patterns subsequently.

ML estimates of $K_0^{\mathbb{P}}$ and $K_Z^{\mathbb{P}}$ governing the $\mathbb{P}$-drift of Z_t are displayed in Table 3 for model $\mathcal{CM}_E^0$. The zero entries in rows $PC2$ and $PC3$ are implied by the constraints

Z	$K_0^{\mathbb{P}}$	$K_Z^{\mathbb{P}}$ $PC1$	$PC2$	$PC3$	GRO	INF
$PC1$	0.0002	0.9552	0.0054	−0.0479	0.1560	0.0856
	(0.0000)	(0.0098)	(0.0106)	(0.0030)	(0.0364)	(0.0195)
$PC2$	−0.0004	0.0040	0.9691	0.0648	−0.1641	0
	(0.0001)	(0.0055)	(0.0017)	(0.0277)	(0.0336)	
$PC3$	0.0007	0.0150	0.0012	0.8771	0	0
	(0.0001)	(0.0015)	(0.0023)	(0.0114)		
GRO	0.0005	−0.0014	0.0217	−0.0526	0.9191	−0.0125
	(0.0004)	(0.0069)	(0.0075)	(0.0292)	(0.0230)	(0.0152)
INF	−0.0005	0.0232	0.0060	0.0831	0.0705	0.9347
	(0.0004)	(0.0115)	(0.0117)	(0.0499)	(0.0405)	(0.0261)

Table 3: Maximum Likelihood Estimates of $K_0^{\mathbb{P}}$ and $K_Z^{\mathbb{P}}$ for Model $\mathcal{CM}_E^0$. Standard errors are reported in parentheses.

on $\mathcal{K}$ selected by our model-selection criteria. Note that a zero constraint on the effect a variable has on an excess returns in Table 2b results in the variable having the same effect on the $\mathbb{P}$-forecasts as $\mathbb{Q}$-forecasts (i.e., $K^{\mathbb{Q}}_{1\mathcal{P},ij} = K^{\mathbb{P}}_{1\mathcal{P},ij}$). Since, by construction, the macro-factors do not incrementally affect the $\mathbb{Q}$-expectations of the PCs, this means a zero constraint on the excess returns in Table 2b causes the macro-variable to have no effect on the $\mathbb{P}$-forecasts of the PC.

The non-zero coefficients on (GRO_{t-1}, INF_{t-1}) in the rows for $(PC1, PC2)$ are all statistically different from zero at conventional significance levels, confirming our earlier findings outside of a $DTSM$ that macro information is *incrementally* useful for forecasting future bond yields after conditioning on $\{PC1, PC2, PC3\}$. Additionally, the coefficients on the own lags of GIP and INF are large and significantly different from zero, as expected given the high degree of persistence in these series. The Kim and Wright (2005) model implicitly sets all of the coefficients in the last two columns of Table 3 (under GRO and INF) to zero, and these restrictions are clearly rejected by our data.

4 Forward Term Premiums

Excess holding period returns on portfolios of individual bonds reflect the risk premiums for every segment of the yield curve up to the maturity of the underlying bond. A different perspective on market risk premiums comes from inspection of the forward term premiums, the differences between forward rates for a q-period loan to be initiated in p periods and the expected yield on a q-period bond purchased p periods from now. Figure 5 displays the forward term premiums (FTP) based on the point estimates of model $\mathcal{CM}_E^0$ for "in-p-for-1" loans (one-year loans initiated in p years) for $p = 2$ and 9. These premiums tend to drift downward during our sample periods, as the level of rates fell, and are increasing in p.

Within the family $UMA_0^R(N)$ both forward rates and expected future one-year rates are affine functions of the state Z_t: $FTP_t^{p,1} = \varsigma_0^{p,1} + \varsigma_Z^{p,1} \cdot Z_t$. Based on this relationship and the ML estimates of model $\mathcal{CM}_E^0$, we compute the 95% confidence bands for our estimated FTPs. The darker shaded areas in Figure 5 represent the confidence bands based on the precision of the ML estimates of $\varsigma_Z^{p,1}$, and the wider, light-shaded bands reflect the sampling variability of the entire parameter set $(\varsigma_0^{p,1}, \varsigma_Z^{p,1})$. For the case of $FTP_t^{2,1}$, the two confidence bands roughly coincide, implying that most of the imprecision in estimating $FTP_t^{2,1}$ derives from sampling variability in $\varsigma_Z^{2,1}$ (forward premium dynamics). In contrast, for the longer-horizon premium $FTP_t^{9,1}$, most of the imprecision derives from sampling variability in $\varsigma_0^{9,1}$, the level of this premium. In fact, the state-dependent component of $FTP_t^{9,1}$ is measured more precisely than its counterpart for $FTP_t^{2,1}$ over much of our sample period. Though pinning down the level of these term premiums is challenging, model $\mathcal{CM}_E^0$ gives quite precise measures of the dynamic properties of premiums.

The "in-2-for-1" forward term premium implied by model $\mathcal{CM}_E^0$ exhibits comparable high-frequency (i.e., shorter than business cycle frequency) variation as the "in-2-for-0.25" forward term premium computed by Kim and Orphanides (2005). Their premium was inferred from a three-factor Gaussian $DTSM$ model estimated using survey forecasts of future interest rates. Professional forecasters are conditioning (at least) on similar macro information as that embodied in GRO and INF, and so we find it reassuring that our implied forward term premiums show similar patterns.

Additional insight into the properties of the term premiums in model $\mathcal{CM}_E^0$ comes from Figure 6 which displays standardized FTPs along with standardized versions of GRO and the PMI index constructed by the Institute for Supply Management.[28]

[28] The PMI index is a composite index for the five business cycle indicators new orders, production, employment, supplier deliveries, and inventories, each with a weight of 20%. It reflects the sentiment of its membership about future activity in the manufacturing sector of the U.S. economy.

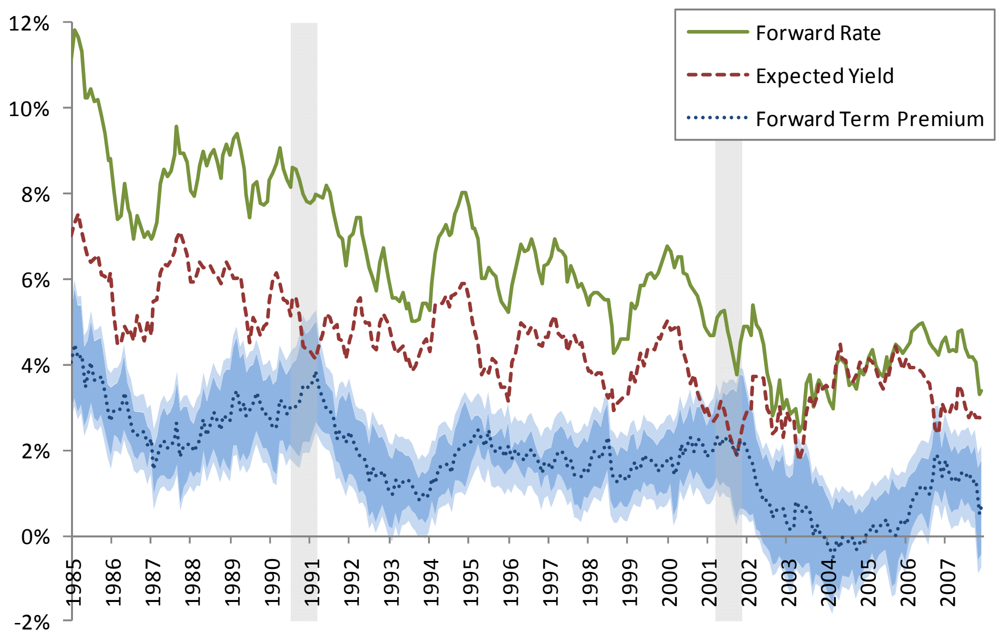

(a) In-2-For-1 Forward Term Premium

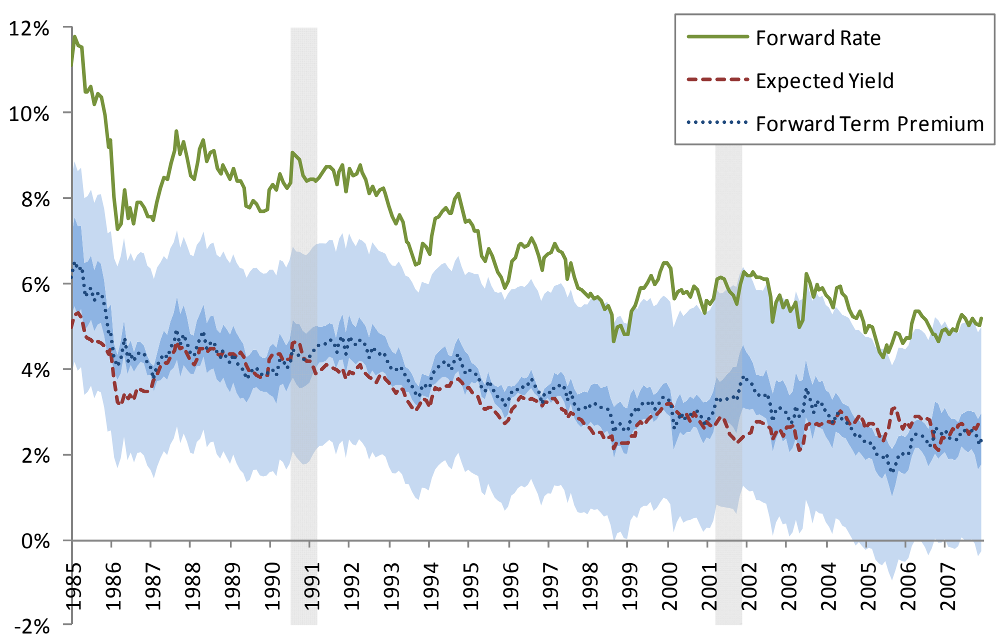

(b) In-9-For-1 Forward Term Premium

Figure 5: Decomposition of forward rates into expected future spot rates and forward term premiums for "in-p-for-1" forward contracts, $p = 2$ and 9, implied by model $\mathcal{CM}_E^0$. The shaded areas are confidence bands for the forward term premiums.

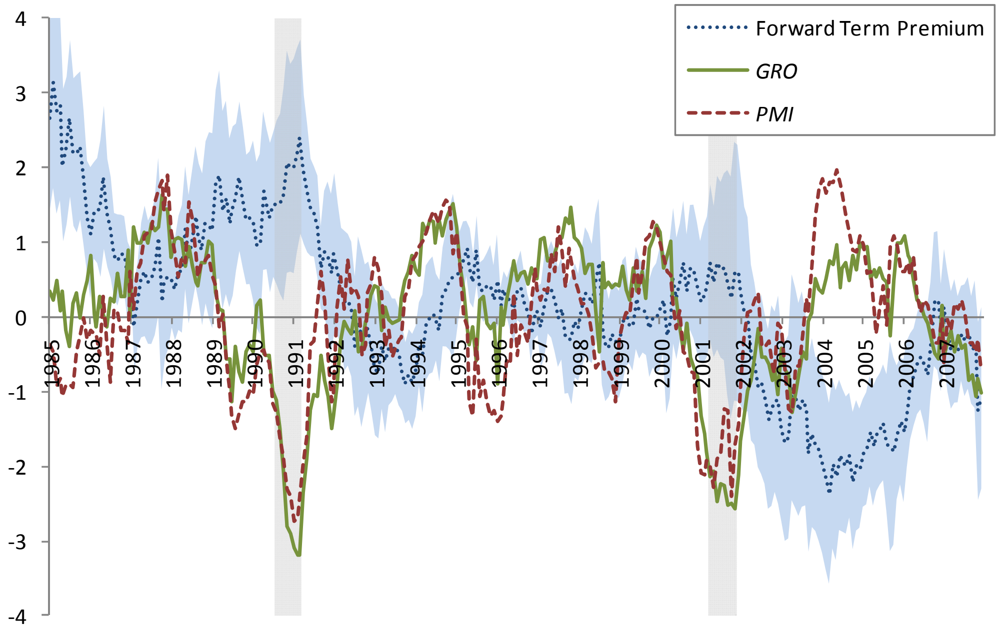

(a) In-2-For-1 Forward Term Premium

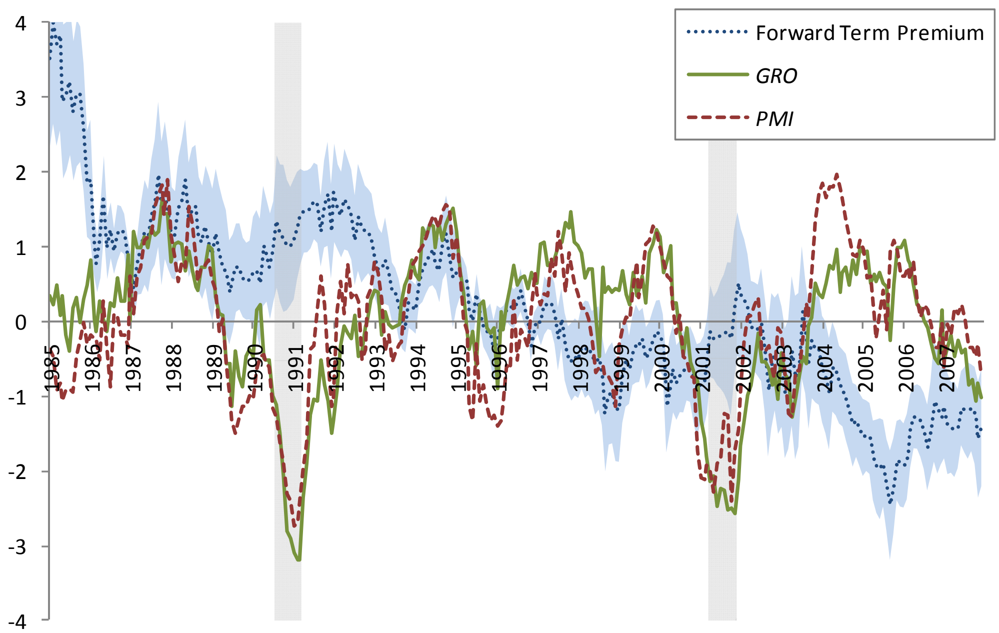

(b) In-9-For-1 Forward Term Premium

Figure 6: Standardized forward term premiums for "in-p-for-1" forward contracts, $p = 2$ and 9, implied by model $\mathcal{CM}_E^0$, plotted against the standardized Purchase Managers' Index (PMI) and smoothed growth rate in the Confidence Board's Coincident Economic Indicators Index. The shaded area is the 95% confidence band on FTP.

	const	$PC1$	$PC2$	$PC3$	GRO	INF
2-for-1	−0.0078	0.6285	0.0119	0.1636	−1.1025	−0.4183
	(0.0092)	(0.1435)	(0.1868)	(0.2992)	(0.4388)	(0.2361)
5-for-1	−0.0049	0.5713	0.3874	0.1988	−0.0748	−0.3186
	(0.0128)	(0.0979)	(0.1197)	(0.1931)	(0.2124)	(0.1288)
9-for-1	0.0026	0.4546	0.4782	0.4449	−0.2023	−0.2918
	(0.0155)	(0.0838)	(0.1040)	(0.1520)	(0.1821)	(0.1160)

Table 4: Coefficients $\varsigma_0^{p,1}$ and $\varsigma_Z^{p,1}$ determining the mapping between the forward term premiums $FTP_t^{p,1}$ and the state Z_t in model $\mathcal{CM}_E^0$.

Though the PMI is sometimes viewed as a leading indicator, and is followed by the Federal Reserve in setting monetary policy (Koenig (2002)), during our sample period the PMI and GRO track each other closely. Two exceptions are the period of the Asian crisis in the late 1990's and the 2003–04 period. In the former case managers expressed a more pessimistic sentiment, while in the latter case they were more optimistic than what GRO indicated about the economy's strength.

Consistent with prior studies and many economic theories, forward term premiums are high during the depths of recessions (see the shaded NBER recession periods in Figure 6). However, neither $FTP^{2,1}$ nor $FTP^{9,1}$ follow an unambiguously counter-cyclical pattern. In fact, during 1993 through 2000, there are subperiods when the PMI and $FTP^{9,1}$ track each other quite closely, suggesting that forward term premiums were pro-cyclical during these subperiods. Inspection of the coefficients $\varsigma_Z^{p,1}$ relating FTPs and Z_t in Table 4 reveals that the negative weights on GIP and INF induce counter-cyclical movements in FTPs.[29] However, all three PCs have statistically significant, positive effects on $FTP^{9,1}$. $PC1$ in particular followed a pro-cyclical path during this period (Figure 2), and the FTPs reflect a blending of the influences of the priced level and slope risks. Together these findings suggest that there were important economic forces driving term premiums that were orthogonal to output growth and inflation.

Turning to the post-2000 sample, two periods stand out when there were particularly large differences between $FTP_t^{9,1}$ and the business cycle indicators: around the peak of the dot-com equity market bubble and the period of the bond market "conundrum" during 2005 – 2006. At the time of the bursting of the dot-com bubble, the economic indicators showed substantial weakness in the economy, while FTPs

[29]Complementary evidence that real economic activity affects expected excess returns on short-dated federal funds futures positions is presented in Piazzesi and Swanson (2008).

remained high. This counter-cyclical pattern is plausible given the sharp drop-off in output growth and the associated increased risk related to the debt financing of corporations at this time. Speculative grade default rates in the U.S. reached their highest level during 2001/02 since the 1990/91 recession (Moody's (2009)). So it is not surprising that forward term premiums were high during both of these recessions.

Several authors have attributed the behavior of long-term rates during the conundrum period to sharp declines in forward term premiums.[30] The evidence in Figure 6b is consistent with this view. Note that $FTP^{2,1}$ was rising throughout 2004 and 2005, while long-dated forward term premiums ($FTP^{9,1}$) were falling. This disparate pattern raises the question of whether movements in term premiums were driven by real economic activity or concerns about inflation, or whether other economic forces were in play. We turn next to a more in depth exploration of the contributions of spanned and unspanned macro risks to variation in risk premiums.

5 Spanned and Unspanned Macro Risks

To delve more deeply into the effects of macroeconomic information on the shape of the Treasury curve we explore how new information about the macroeconomy impacts the conditional variances of the $FTP^{p,1}$, $Var_{t-1}[FTP_t^{p,1}] = \varsigma_Z^{(p,1)\prime}\Sigma_Z\varsigma_Z^{(p,1)}$. We focus on the conditional variances of the FTPs because of the near-cointegration of $PC1$ and inflation. Conditional variances remove the conditional means (and hence the trend-like components) and, thereby, isolate how surprise changes in the components of Z_t contribute to surprise moves in the forward term premium.

Our interest is in the proportion of $Var_{t-1}[FTP_t^{p,1}]$ attributable to innovations in the unspanned growth ($OGRO_t$) and inflation ($OINF_t$). Accordingly, we adopt the ordering $Z_t' = (\mathcal{P}_t', M_t')$ and use the Cholesky factorization of Σ_Z to orthogonalize the shocks to Z_t. We also compute the proportion of $Var_{t-1}[FTP_t^{p,1}]$ attributable to the component of INF_t that is orthogonal to both the yield PCs ($\mathcal{P}_t$) and GRO_t.

The decompositions of $Var_{t-1}[FTP_t^{p,1}]$ for p ranging from one to ten years are displayed in Figure 7. Line OM shows the contribution of the unspanned macro components ($OGRO_t, OINF_t$) to the $Var_{t-1}[FTP_t^{p,1}]$. Notably, nearly 90% of the variance of the one-year forward term premium $FTP^{1,1}$ is attributable to news about these unspanned macro risks. Equally striking is how the proportions attributable to unspanned macro risks decline as the forward initiation date of the loan increases, with

[30] Recent papers on this issue, using both reduced-form and structural pricing models, include Rudebusch, Swanson, and Wu (2006), Cochrane and Piazzesi (2008), Bandholz, Clostermann, and Seitz (2007), and Backus and Wright (2007).

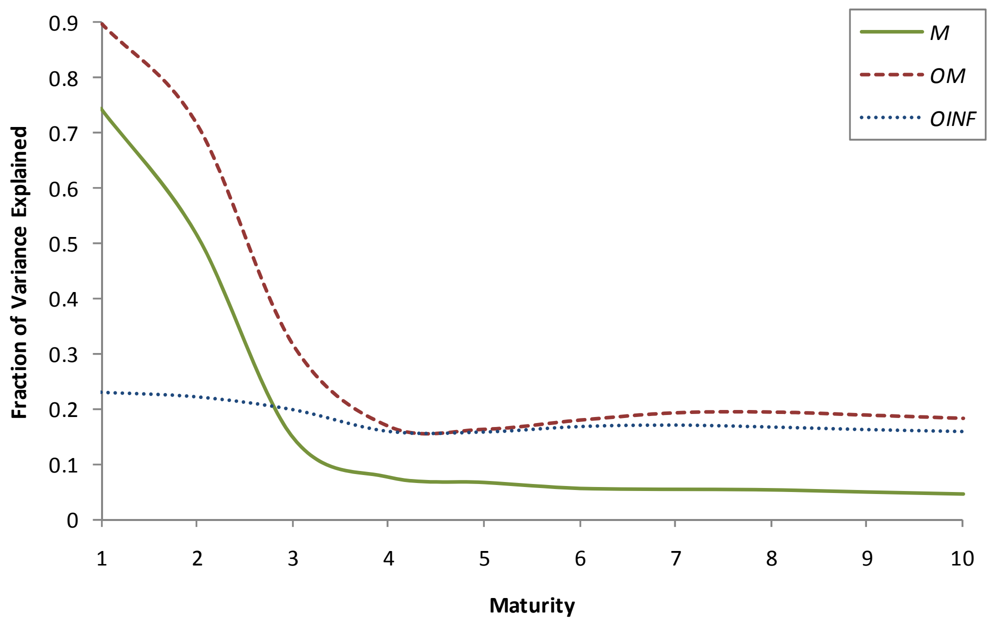

Figure 7: The proportions of $Var_{t-1}[FTP_t^{p,1}]$ attributable to innovations in $(OGRO_t, OINF_t)$ (OM) and to $OINF_t^{\perp}$ $(OINF)$, respectively, within a VAR with ordering $(\mathcal{P}_t, M_t)$. The solid line (M) is the proportion of $Var_{t-1}[FTP_t^{p,1}]$ attributable to innovations in M_t within a VAR with ordering $(M_t, \mathcal{P}_t)$. The horizontal axis is the maturity of the underlying bond, p. All estimates are based on model $\mathcal{CM}_E^0$.

OM plateauing at about 20% beyond the four-year horizon. These patterns suggest that, out to about three years, a majority of the variation in forward term premiums is attributable to surprise changes in the unspanned macro factors $(OGRO_t, OINF_t)$. On the other hand, for forward loans far in the future, close to 80% of the innovation variances of the $FTP^{p,1}$s are attributable to innovations in the yield factors $\mathcal{P}$.

Using our orthogonalization scheme, line $OINF$ measures the proportion of $Var_{t-1}[FTP_t^{p,1}]$ attributable to innovations in INF that are orthogonal to both the pricing factors $\mathcal{P}$ and real economic activity GRO. This measure of inflation risk has a roughly uniform impact on forward term premiums for loan initiation dates ranging from one to ten years. From year four onwards the lines OM and $OINF$ in Figure 7 are virtually indistinguishable. This says that, to the extent that unspanned macro risk is affecting the longer dated forward term premiums, the effects are attributable entirely to inflation risk.

So far in this discussion we have focused on unspanned macro risks. By reversing the ordering in our state vector to $(M_t, \mathcal{P}_t)$ we can measure the proportions of the innovation variances $Var_{t-1}[FTP_t^{p,1}]$ that are attributable to the macro factors M_t (the spanned and unspanned components combined). From the line M in Figure 7 we see that the contributions of the macro factors follow a pattern similar to the contributions of their unspanned components, but that the former lie uniformly below the latter. That is, combining the spanned and unspanned components of M_t *lowers* the explanatory power of the total macro series. This suggests that it is the unspanned components of output and inflation risks that are most relevant for understanding surprise changes in forward term premiums, and not total macro risks M_t.

We can also decompose variation in the level of $FTP_t^{p,1}$ into the component $S\mathcal{P}$ associated with variation in $\mathcal{P}_t$ (computed from the projection of $FTP_t^{p,1}$ onto $\mathcal{P}_t$) and the orthogonal component associated with variation in the unspanned M_t. In this decomposition, the component $S\mathcal{P}$ largely recovers the near unit-root component, analogous to a cointegrating regression.[31] This can be seen from the decomposition of $FTP^{9,1}$ in Figure 8, which shows that over 95% of the variation in $FTP_t^{9,1}$ is explained by variation in $S\mathcal{P}_t^9$ (less than 5% is explained by $(OGRO_t, OINF_t)$).

Nevertheless, it is decompositions like the one in Figure 8 that are perhaps most relevant for interpreting specific episodes in history, such as Greenspan's conundrum during 2004 - 2006. The fall in $FTP^{9,1}$ during this period is accounted for (almost) entirely by the information embodied in $\mathcal{P}$. Taken together, Figures 6b and 8 suggest that the conundrum is not easily explained by economic weakness as captured by PMI and GRO. When we project $FTP_t^{9,1}$ onto (GRO_t, INF_t) during our sample period the resulting R^2 is only 37%.

This number understates the effect of M_t on $FTP_t^{9,1}$ if, for instance, $PC1$ reflects in part variation in the Central Bank's long-term target inflation rate.[32] To conservatively control for this possibility, we project $FTP_t^{9,1}$ onto $(SGRO_t, SINF_t, OGRO_t, OINF_t)$, where $(SGRO_t, SINF_t)$ denotes spanned M_t. While the R^2 does in fact increase, it remains the case that roughly 24% of the variation in $FTP^{9,1}$ is due to factors that are orthogonal to this expanded information set. These calculations reinforce a message from Figures 7: a substantial percentage of the variation in $FTP_t^{9,1}$ was due to economic forces that were orthogonal to the business cycle variables (GRO_t, INF_t).

Yet a different perspective on the contributions of macro information to risk premiums comes from inspection of the expected excess returns on the portfolios

[31] By focusing on conditional variances we are removing the conditional means (and hence the trend-like components) and, thereby, isolating how surprise changes in the components of Z_t contribute to surprise moves in the forward term premium.

[32] This was suggested to us by some Federal Reserve economists, and this view is reflected in the interpetation Rudebusch and Wu (2008) give to their latent risk factors.

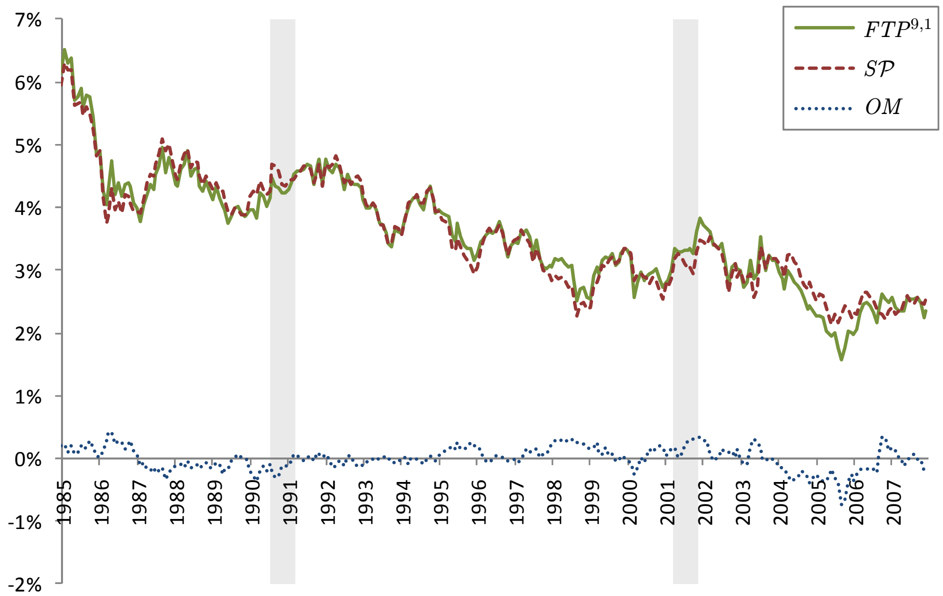

Figure 8: Decomposition of $FTP^{9,1}$ into the component spanned by $\mathcal{P}_t$ ($S\mathcal{P}$) and the component spanned by orthogonal macro factors $(OGRO_t, OINF_t)$ (OM). Estimates are based on model $\mathcal{CM}_E^0$.

of bonds with payoffs that are perfectly correlated with movements in the spanned macro risks $SGIP$ and $SINF$, $xSGIP_t$ and $xSINF_t$.[33] $xSINF$ achieved its lowest levels (largest absolute values) during the 1990 and 2001 recessions (Figures 9). In this respect there is a parallel with the excess returns to the level-mimicking portfolio in Figures 4a. $xGRO$ is near zero for most of the sample period, with the exceptions of the periods of recession.

Inflation risk premiums, as measured by $xSINF_t$, were small, typically ranging between zero and twenty basis points (in absolute value), outside of recessions. This finding leaves open the possibility that unspanned macro risks were priced and, hence, that *actual* market premiums on inflation and output growth risks were larger. Using data in inflation-indexed bonds, Hordahl and Tristani (2007) found that inflation risk premiums were insignificantly different from zero for the Eurozone. Grishchenko and Huang (2010) study the inflation risk premium in the US using TIPS data for the period 2000 through 2008. Consistent with our analysis, they find that this premium

[33]Both $SGIP$ and $SINF$ are affine functions of $\mathcal{P}$. Using this fact and our construction of excess returns on portfolios representing pure exposures to level, slope, and curvature risks, we computed model-implied expected excess returns on pure exposures to $SGIP$ and $SINF$.

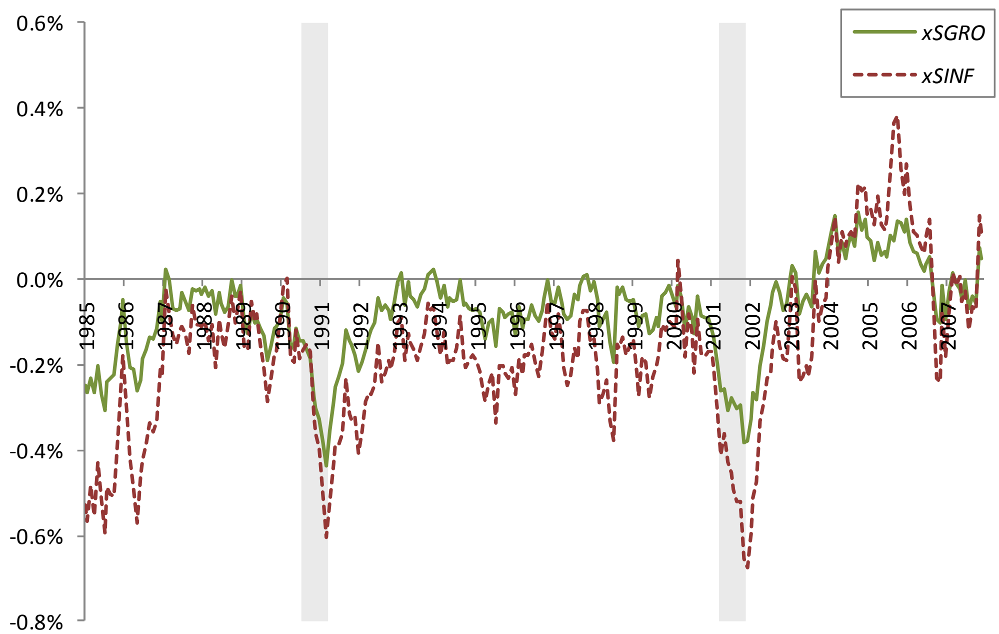

Figure 9: The expected excess returns on bond portfolios whose payoffs are perfectly correlated with the components of GRO and INF that are spanned by $\mathcal{P}_t$ ($SGRO$ and $SINF$). Returns are for monthly holding periods and are expressed on an annualized scale.

was negative during the first half of their sample, 2000-2004. Moreover, just as in Figures 9, their premium increased substantially around 2004 and ranged between fifteen and thirty basis points during the subsequent two years. Without TIPS data, and using a somewhat different measure of inflation risk premium, we find very similar results. This raises the interesting question of whether, during this period, most of the inflation risk premium was associated with spanned inflation risk.[34]

As Figures 7 demonstrates, most of the variation in shocks to forward term premiums, particularly for longer horizons, is associated with economic risk that is orthogonal to our measures of real economic activity and inflation. We now dig deeper into the impacts of unspanned macro risks on the forward term premiums by examining the impulse response functions of $FTP_t^{p,1}$, $p = 2, 9$, to shocks to the spanned and unspanned components of GRO and INF. At the two-year horizon,

[34]Their analysis was outside of an arbitrage-free $DTSM$. We defer to future research an exploration of risk premiums on unspanned inflation using TIPS data, as many institutional and measurement issues complicate the introduction of TIPS into a $DTSM$.

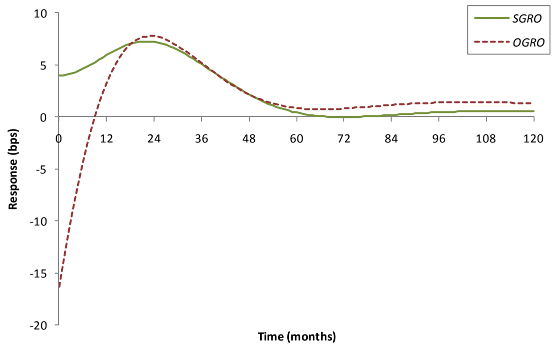

(a) $GIP \rightarrow FTP^{2,1}$

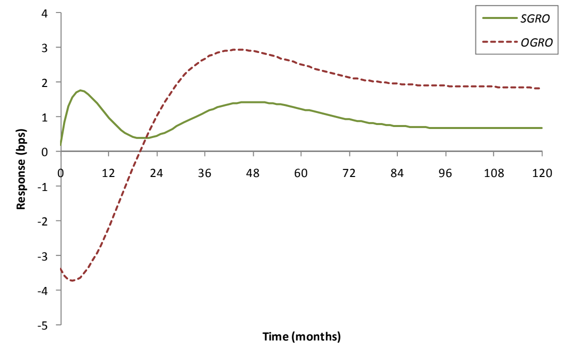

(b) $GIP \rightarrow FTP^{9,1}$

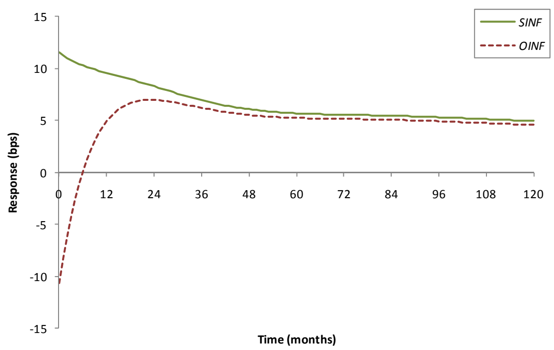

(c) $INF \rightarrow FTP^{2,1}$

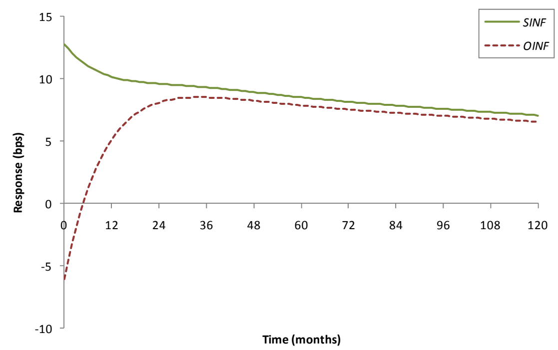

(d) $INF \rightarrow FTP^{9,1}$

Figure 10: Each panel plots the impulse responses of FTPs to shocks to either $(SGIP, OGRO)$ or $(SINF, OINF)$.

an innovation in $OGRO$ has a much larger effect on $FTP^{2,1}$ than an innovation in $SGRO$ (Figures 10a). The effects of unspanned output shocks dissipate quickly (within about one year), while the effects of unspanned inflation shocks persist for several years (Figures 10c). This difference is no doubt attributable to the near cointegration of INF with the priced risk factors $(PC1, PC2)$.

For $FTP^{9,1}$ the responses to both $OGRO$ and $OINF$ are small, consistent with the decomposition results in Figures 8. Figures 10d shows a large impact of $SINF$ on $FTP^{9,1}$ suggesting that it is largely spanned inflation risk, and not spanned output risk, that explains the variation in $FTP^{9,1}$ attributable to the macro factors M_t. In speaking about the conundrum, Chairman Bernanke asserted that "a substantial portion of the decline in distant-horizon forward rates of recent quarters can be attributed to a drop in term premiums. ... the decline in the premium since last

June 2004 appears to have been associated mainly with a drop in the compensation for bearing real interest rate risk."[35] The patterns in Figures 8 and 10 are not easily reconciled, it seems, with Chairman Bernanke's explanation of the conundrum. More likely, the decline in long-dated forward term premiums was a consequence of changes in spanned inflation risks or changes in economic factors that were orthogonal to (GRO, INF). Recall that nearly a quarter of the variation in $FTP^{9,1}$ during our sample period was attributable to such orthogonal factors.

Symmetric to this discussion is the interesting question of how changes in term premiums affect real economic activity. Bernanke, in his 2006 speech, argues that a higher term premium will depress the portion of spending that depends on long-term interest rates and thereby will have a dampening economic impact. In linearized New Keynesian models in which output is determined by a forward-looking IS equation (such as the model of Bekaert, Cho, and Moreno (2010)), current output depends only on the expectation of future short rates, leaving no role for a term premium effect. Time-varying term premiums do arise in models that are linearized at least to the third order (e.g., Ravenna and Seppala (2007b)). We examine the response of real economic activity and inflation to innovations in $FTP^{p,1}$ ($p = 2, 9$) in the context of model $\mathcal{CM}_E^0$, using the model-implied VAR with ordering $(FTP^{p,1}, SGRO, SINF, OGRO, OINF)$.

Initially, a one standard deviation increase in $FTP^{2,1}$ has a large negative impact on $OGRO$ over a period of about 12 months, and has virtually no effect on $SGRO$ (Figures 11a). The effects of innovations in $FTP^{9,1}$ on both $SGRO$ and $OGRO$ are small at all horizons. These responses, which are a manifestation of large and significant coefficients on GRO in the near-horizon forward term premium parameters in Tables 4, are consistent with the economic linkages set forth by Chairman Bernanke. However, the effects on economic activity arise from short- to medium-term risk premiums, not long-dated premiums, and the effects are virtually entirely through unspanned real economic activity.

The absence of effects on $SGRO$ is consistent with the results in Ang, Piazzesi, and Wei (2006) that term premiums are insignificant in predicting future GDP growth within a Gaussian $DTSM$ that enforces the theoretical spanning of GDP growth by bond yields. What their model, and similar models that enforce spanning, do not accommodate is our finding that term premium shocks do affect growth through their effects on unspanned real activity.

Innovations in $FTP^{2,1}$ have small effects on both spanned and unspanned inflation. In contrast, shocks to $FTP^{9,1}$ have a large negative effect on $OINF$ and a positive

[35] See his speech before the Economic Club of New York on March 20, 2006 titled "Reflections on the Yield Curve and Monetary Policy."

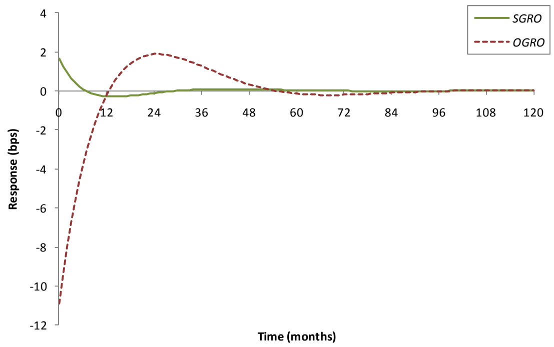

(a) $FTP^{2,1} \rightarrow GRO$

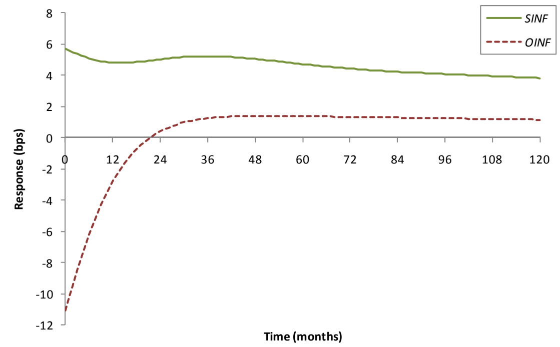

(b) $FTP^{9,1} \rightarrow INF$

Figure 11: The impulse responses of $SGRO$ and $OGRO$ to a one standard deviation shock to $FTP^{2,1}$, and of $SINF$ and $OINF$ to a one standard deviation shock to $FTP^{9,1}$, based on model $\mathcal{CM}_E^0$.

impact effect on $SINF$ (Figures 11b). The effects on $OINF$ die out after about a year, while the effects on $SINF$ persist for several years.

6 Concluding Observations

This paper develops and estimates an arbitrage-free, Gaussian $DTSM$ in which the state vector includes macroeconomic variables that are not perfectly spanned by contemporaneous bond yields, and in which these macro variables have significant predictive content for excess returns on bonds over and above the information in bond yields. We show that there is a canonical representation of this model that lends itself to easy interpretation and for which the global maximum of the likelihood function can be attained essentially instantaneously.

Perhaps the most striking aspects of our empirical results are that shocks to unspanned real economic activity and inflation have large effects on term premiums in US Treasury markets and, symmetrically, shocks to forward term premiums have substantial effects on real economic activity primarily through their effects on unspanned real output growth. To our knowledge, ours is the first paper to identify the important contributions to portfolio risk of unspanned macro risks, largely because the extant literature on macro-$DTSM$s rules out, by construction, any role for these unspanned risks.

Our findings raise several intriguing questions for future research. Unspanned macro risks, particularly real economic risks, had large effects on forward term premiums over short- to intermediate-term horizons. What were the economic sources of these unspanned risks? We also find that a substantial proportion of the variation in long-dated forward term premiums is attributable to economic factors that are orthogonal to both spanned and unspanned output and inflation. It was evidently these orthogonal risks that largely explained the decline in term premiums during the period of the conundrum. What is the nature of these macro risks that are so important in Treasury markets and yet are orthogonal to output growth and inflation? Since the onset of the financial crisis there has been considerable discussion about the roles of global imbalances and disruptions in the financial intermediation sectors. Our modeling framework provides a means of systematically examining these possibilities within an arbitrage-free $DTSM$.

More generally, properties of the fitted historical distributions of bonds and pricing factors in our macro-$DTSM$ are very different than what is implied by both a factor-VAR model or an unconstrained version of our canonical model. In particular, our model-selection search led to eleven constraints on the parameters governing the market prices of risk. A practical consequence of these constraints is that the persistence properties of bond yields, and hence the relative importance of expected future spot rates versus forward term premiums, are very different in our preferred model than in its unconstrained counterpart. This suggests that, when estimating macro-$DTSM$s, one should undertake similar model-selection exercises to systematically reduce the dimension of the parameter space, as this might similarly mitigate small-sample bias problems.

Our framework (formally developed in Appendix B) can be applied in any Gaussian pricing setting in which security prices or yields are affine functions of a set of pricing factors $\mathcal{P}_t$ and the relevant state vector embodies information (over and above the past history of $\mathcal{P}$) that is useful for forecasting $\mathcal{P}_t$ under the physical measure $\mathbb{P}$. Accordingly, it is well suited to addressing a variety of economic questions about risk premiums in bond and currency markets, as well as in equity markets when the latter pricing problems maps into an affine pricing model (e.g., Bansal, Kiku, and Yaron (2009)). Though neither the state variables nor the pricing factors exhibit time-varying volatility in the settings examined in this paper, our basic framework and its computational advantages are likely to extend to affine models with time-varying volatility. Exploration of this extension is deferred to future research.

Appendices

A Bond Pricing in $GDTM$s

Under $Ar\mathbb{Q}$ and $A\mathcal{P}\mathbb{Q}$, the price of an m-month zero-coupon bond is given by

$$D_{t,m} = E_t^{\mathbb{Q}}[e^{-\sum_{i=0}^{m-1} r_{t+i}}] = e^{\mathcal{A}_m + \mathcal{B}_m \cdot \mathcal{P}_t}, \tag{9}$$

where $(\mathcal{A}_m, \mathcal{B}_m)$ solve the first-order difference equations

$$\mathcal{A}_{m+1} - \mathcal{A}_m = \left(K_0^{\mathbb{Q}}\right)' \mathcal{B}_m + \frac{1}{2}\mathcal{B}_m' \Sigma_{\mathcal{P}\mathcal{P}} \mathcal{B}_m - \rho_0 \tag{10}$$

$$\mathcal{B}_{m+1} - \mathcal{B}_m = \left(K_1^{\mathbb{Q}}\right)' \mathcal{B}_m - \rho_1 \tag{11}$$

subject to the initial conditions $\mathcal{A}_0 = 0, \mathcal{B}_0 = 0$. See, for example, Dai and Singleton (2003). The loadings for the corresponding bond yield are $A_m = -\mathcal{A}_m/m$ and $B_m = -\mathcal{B}_m/m$.

B Derivation of Results in Sections 2

Proof of Proposition 1: From Joslin, Singleton, and Zhu (2010) we know that, for any $A_0^{\mathbb{Q}}(3)$ pricing model with distinct, real eigenvalues of the feedback matrix of the risk factors, there exists a three-dimensional, latent state vector X_t such that $r_t = r_\infty^{\mathbb{Q}} + 1 \cdot X_t$ and

$$\Delta X_t = \text{diag}(\lambda^{\mathbb{Q}}) X_{t-1} + \sqrt{\Sigma_X^0} \epsilon_t^{\mathbb{Q}}$$

for some 3×3 matrix Σ_X^0 and vector $\lambda^{\mathbb{Q}}$ of eigenvalues of the feedback matrix governing X, with $\epsilon_t^{\mathbb{Q}} \sim N(0, I)$.

To derive a canonical version of (1), let $B^0(\tau)$ be the loadings given by $\Delta B^0 = \text{diag}(\lambda^{\mathbb{Q}}) B^0 - 1$, $B^0(0) = 0$. Let $B_{PC}^0(i) = \sum -\ell_i B^0(\tau_i)/\tau_i$, where PCi has loading ℓ_i on yield maturity τ_i. Let B be the 3×3 matrix with i^{th} row given by $B_{PC}^0(i)$. It follows that the covariance matrix of the innovations to the PCs is $B\Sigma_X^0 B'$. In order that (1) is satisfied, it must be that $\Sigma_X^0 = (B')^{-1} \Sigma_{\mathcal{P}} B^{-1}$.

Now let $A^0(t)$ solve $\Delta A^0 = \frac{1}{2}(B^0)^\top \Sigma_X^0 B^0 - r_\infty^{\mathbb{Q}}$, $A^0(0) = 0$. Define $A_{PC}^0(i) = \sum -\ell_i A_y^0(\tau_i)/\tau_i$. Let a be the 3×1 vector with i-th entry $A_{PC}^0(i)$. Then $\mathcal{P}_t = a + BX_t^0$. From an invariant affine transformation it follows that: $K_X^{\mathbb{Q}} = B(\text{diag}(\lambda^{\mathbb{Q}})B^{-1}$, $K_0^{\mathbb{Q}} = -(K_1^{\mathbb{Q}})^{-1} a$, $\rho_0 = r_\infty^{\mathbb{Q}} - 1' B^{-1} a$, and $\rho_1 = (B')^{-1} 1$.

Since (1) is an invariant transformation of an identified, canonical model, we know that (1) is also identified and canonical. The underlying parameters are $(r^{\mathbb{Q}}_\infty, \lambda^{\mathbb{Q}}, L_X)$, where L is the Cholesky factorization of Σ_X.

C Returns on Generalized Mimicking Portfolios

Consider a collection of N yields, $\{y_t^{n_1}, \ldots, y_t^{n_N}\}$, and a given linear combination $y_t^a = \sum_{i=1}^N a_i y_t^{n_i}$ of these yields (y_t^a could be a principal component, or the projection of a macro variable onto the yields). Our first goal is to find weights $\{w_i\}_{i=1}^N$ such that value $P_t^w = \sum_{i=1}^N w_i P_t^{n_i}$ of a portfolio of zero coupon bonds locally tracks changes in y_t^a; that is,

$$\frac{dP_t^w}{dy_t^a} = \sum_{i=1}^N \frac{dP_t^w}{dy_t^{n_i}} \frac{dy_t^{n_i}}{dy_t^a} = 1 \tag{12}$$

Since, by definition, $P_t^{n_i} = \exp(-n_i y_t^{n_i})$, we have $dP_t^{n_i}/dy_t^{n_i} = -n_i P_t^{n_i}$. Therefore, (12) can be rewritten as

$$-\sum_{i=1}^N w_i n_i P_t^{n_i} \frac{1}{a_i} = 1$$

which will hold for weights

$$w_i = -\frac{a_i}{N n_i P_t^{n_i}}$$

Next, consider the one-period excess return on portfolio P_t^w:

$$\frac{\sum_i w_i (P_{t+1}^{n_i - 1} - e^{r_t} P_t^{n_i})}{|\sum_i w_i P_t^{n_i}|} = \frac{-\sum_i a_i/n_i (P_{t+1}^{n_i-1}/P_t^{n_i} - e^{r_t})}{|\sum_i a_i/n_i|}.$$

This is a weighted average of the returns on the individual zero coupon bonds. Now, it follows from Le, Singleton, and Dai (2010) that $P_t^{n_i} = \exp(-A_{n_i} - B_{n_i} Z_t)$, and further that

$$E^{\mathbb{P}}[P_{t+1}^{n_i-1}/P_t^{n_i}] = \exp\{B_{n_i-1}[(K_0^{\mathbb{Q}} - K_0^{\mathbb{P}}) + (K_1^{\mathbb{Q}} - K_1^{\mathbb{P}})Z_t] + r_t\}.$$

Therefore, to a first-order approximation, the expected excess return on portfolio P_t^w is given by

$$\frac{\sum_i a_i/n_i B_{n_i-1}[(K_0^{\mathbb{P}} - K_0^{\mathbb{Q}}) + (K_1^{\mathbb{P}} - K_1^{\mathbb{Q}})Z_t]}{|\sum_i a_i/n_i|}.$$

Since we rotate our model such that the first $\mathcal{R}$ elements of Z_t correspond to the first $\mathcal{R}$ principal components of yields, and since by definition,

$$PCj_t = \sum_{i=1}^{N} \ell_i^j y_t^{n_i} = \sum_{i=1}^{N} \ell_i^j (A_{n_i}/n_i + B_{n_i}/n_i Z_t)$$

it follows that $\sum_i \ell_i^j B_{n_i}/n_i$ is the selection vector for the j^{th} element, $j \in \{1, \ldots, \mathcal{R}\}$. Thus, under the further approximation that $B_{n_i-1} \approx B_{n_i}$, the expected excess return on the portfolio mimicking PCj, $xPCj$, is given by the j^{th} row of

$$(K_0^{\mathbb{P}} - K_0^{\mathbb{Q}}) + (K_1^{\mathbb{P}} - K_1^{\mathbb{Q}})Z_t$$

scaled by $|\sum_i \ell_i^j/n_i|$. While an approximation for the one-period expected excess return in discrete time, this relationship is exact for the instantaneous expected excess return in the continuous-time limit.

References

Akaike, H. (1973). Information Theory and an Extension of the Likelihood Principle. In B. Petrov and F. Csaki (Eds.), *Proceedings of the Second International Symposium of Information Theory*. Budapest: Akademiai Kiado.

Ang, A., S. Dong, and M. Piazzesi (2007). No-Arbitrage Taylor Rules. Working Paper, Columbia University.

Ang, A. and M. Piazzesi (2003). A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables. *Journal of Monetary Economics 50*, 745–787.

Ang, A., M. Piazzesi, and M. Wei (2006). What Does the Yield Curve Tell us About GDP Growth? *Journal of Econometrics 131*, 359–403.

Backus, D. and J. Wright (2007). Cracking the Conundrum. *Brookings Papers on Economic Activity 38*, 293–329.

Bandholz, H., J. Clostermann, and F. Seitz (2007). Explaining the US Bond Yield Concundrum. Working Paper, MPRA Paper No. 2386, Munich Personal RePEc.

Bansal, R., D. Kiku, and A. Yaron (2009). An Empirical Evaluation of the Long-Run Risks Model for Asset Price. Working Paper, Duke University.

Barillas, F. (2010). Can we Exploit Predictability in Bond Markets? Working Paper, New York University.

Bauer, M. (2010). Term Premia and the News. Working Paper, University of California, San Diego.

Bekaert, G., S. Cho, and A. Moreno (2010). New-Keynesian Economics and the Term Structure. *Journal of Money, Credit, and Banking 42*, 33–62.

Bikbov, R. and M. Chernov (2010). No-Arbitrage Macroeconomic Determinants of the Yield Curve. *Journal of Econometrics forthcoming*.

Campbell, J., A. Sunderam, and L. Viceira (2009). Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds. Working Paper, Harvard University.

Chernov, M. and P. Mueller (2009). The Term Structure of Inflation Expectations. Working Paper, London Business School.

Clarida, R., J. Gali, and M. Gertler (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *Quarterly Journal of Economics 115*, 147–180.

Cochrane, J. and M. Piazzesi (2005). Bond Risk Premia. *American Economic Review 95*, 138–160.

Cochrane, J. and M. Piazzesi (2008). Decomposing the Yield Curve. Working Paper, Stanford University.

Collin-Dufresne, P., R. Goldstein, and C. Jones (2008). Identification of Maximal Affine Term Structure Models. *Journal of Finance LXIII*, 743–795.

Cooper, I. and R. Priestley (2008). Time-Varying Risk Premiums and the Output Gap. *forthcoming, Review of Financial Studies*.

Dai, Q. and K. Singleton (2000). Specification Analysis of Affine Term Structure Models. *Journal of Finance 55*, 1943–1978.

Dai, Q. and K. Singleton (2002). Expectations Puzzles, Time-Varying Risk Premia, and Affine Models of the Term Structure. *Journal of Financial Economics 63*, 415–441.

Dai, Q. and K. Singleton (2003). Term Structure Dynamics in Theory and Reality. *Review of Financial Studies 16*, 631–678.

D'Amico, S., D. Kim, and M. Wei (2008). Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices. Working Paper, Working Paper No 248, Bank for International Settlements.

Diebold, F., G. Rudebusch, and S. B. Aruoba (2006). The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach. *Journal of Econometrics 31*, 309–338.

Duffee, G. (2002). Term Premia and Interest Rate Forecasts in Affine Models. *Journal of Finance 57*, 405–443.

Duffee, G. (2009a). Information In (and Not In) The Term Structure. Working Paper, Working Paper, Johns Hopkins University.

Duffee, G. (2009b). Sharpe Ratios in Term Structure Models. Working Paper, Johns Hopkins University.

Duffie, D. and R. Kan (1996). A Yield-Factor Model of Interest Rates. *Mathematical Finance 6*, 379–406.

Giese, J. (2008). Level, Slope, Curvature: Characterizing the Yield Curve in a Cointegrated VAR Model. *Economics: The Open-Access, Open-Assessment E-Journal 2*, 2008–28.

Grishchenko, O. and J. Huang (2010). Inflation Risk Premium: Evidence from the TIPS Market. Working Paper, Penn State University.

Hannan, E. and B. Quinn (1979). The Determination of the Order of an Autoregression. *Journal of Royal Statistical Society, B 41*, 190–195.

Hansen, L. and S. Richard (1987). The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models. *Econometrica 55*, 587–613.

Hordahl, P. and O. Tristani (2007). Inflation Risk Premia in the Term Structure of Interest Rates. Working Paper, Working Paper 734, European Central Bank.

Jardet, C., A. Monfort, and F. Pegoraro (2010). No-Arbitrage Near-Cointegrated VAR(p) Term Structure Models, Term Premia and GDP Growth. Working Paper, Banque de France.

Joslin, S. (2006). Can Unspanned Stochastic Volatility Models Explain the Cross Section of Bond Volatilities? Working Paper, MIT.

Joslin, S., A. Le, and K. Singleton (2010). The Conditional Distribution of Bond Yields Implied by Gaussian Macro-Finance Term Structure Models. Working Paper, MIT.

Joslin, S., K. Singleton, and H. Zhu (2010). A New Perspective on Gaussian DTSMs. Working Paper, forthcoming, *Review of Financial Studies*.

Kim, D. and A. Orphanides (2005). Term Structure Estimation with Survey Data on Interest Rate Forecasts. Working Paper, Federal Reserve Board.

Kim, D. and J. Wright (2005). An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-term Yields and Distant-Horizon Forward Rates. Working Paper, Discussion Series 2005-33, Federal Reserve Board.

Kim, D. H. (2008). Challenges in Macro-Finance Modeling. Working Paper, Federal Reserve Board, Washington D.C.

Koenig, E. (2002). Using the Purchasing Managers' Index to Assess the Economy's Strength and the Likely Direction of Monetary Policy. *Federal Reserve Bank of Dallas Economic and Financial Policy Review 1*, 1–14.

Le, A., K. Singleton, and Q. Dai (2010). Discrete-time Affine^{Q} Term Structure Models with Generalized Market Prices of Risk. *Review of Financial Studies 23*, 2184–2227.

Litterman, R. and J. Scheinkman (1991). Common Factors Affecting Bond Returns. *Journal of Fixed Income 1*, 54–61.

Ludvigson, S. and S. Ng (2009). Macro Factors in Bond Risk Premia. *Review of Financial Studies*.

Lutkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Springer, Berlin and Heidelberg.

Moody's (2009). Monthly Default Report. Working Paper, Moody's Investors Service.

Phillips, P. and J. Yu (2005). Jacknifing Bond Option Prices. *Review of Financial Studies 18*, 707–742.

Piazzesi, M. (2005). Bond Yields and the Federal Reserve. *Journal of Political Economy 113*, 311–344.

Piazzesi, M. and E. Swanson (2008). Futures Prices as Risk Adjusted Forecasts of Monetary Policy. *Journal of Monetary Economics 55*, 677–691.

Ravenna, F. and J. Seppala (2007a). Monetary Policy, Expected Inflation, and Inflation Risk Premia. Working Paper, Bank of Finland.

Ravenna, F. and J. Seppala (2007b). Monetary Policy, Expected Inflation, and Inflation Risk Premia. Working Paper, Bank of Finland Research Paper 18·2007.

Rudebusch, G., E. Swanson, and T. Wu (2006). The Bond Market "Conundrum" from a Macro-Finance Perspective. *Monetary and Economic Studies 24*, 83 – 128.

Rudebusch, G. and T. Wu (2008). A Macro-Finance Model of the Term Structure, Monetary Policy, and the Economy. *Economic Journal 118*, 906–926.

Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics 6*, 461–464.

Smith, J. and J. Taylor (2009). The Term Structure of Policy Rules. *Journal of Monetary Economics 56*, 907–919.

Stock, J. and M. Watson (1999). Forecasting Inflation. *Journal of Monetary Economics 44*, 293–335.

Tang, C. and S. Chen (2009). Parameter Estimation and Bias Correction for Diffusion Processes. *Journal of Econometrics 149*, 65–81.

Taylor, J. (1999). A Historical Analysis of Monetary Policy Rules. In J. Taylor (Ed.), *Monetary Policy Rules*. University of Chicago Press.

Wachter, J. A. (2006). A Consumption-based Model of the Term Structure of Interest Rates. *Journal of Financial Economics 79*(2), 365–399.

Woodford, M. (2003). *Interest Rates and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Wright, J. (2009). Term Premiums and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset. Working Paper, forthcoming, American Economic Review.