Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks[1]

Scott Joslin[2] Marcel Priebsch[3] Kenneth J. Singleton[4]

First Draft: April, 2009
This Draft: May 11, 2010

[1]We are grateful for feedback from seminar participants at Stanford University, the University of Chicago, the Federal Reserve Board, 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."

[2]MIT Sloan School of Management, sjoslin@mit.edu
[3]Department of Economics, Stanford University, priebsch@stanford.edu
[4]Graduate School of Business, Stanford University, and NBER, kenneths@stanford.edu

Abstract

This paper develops and implements an arbitrage-free *DTSM* within which we are able to econometrically identify the contributions of macroeconomic variables to variation in the market prices of level, slope, and curvature risks. A key property of our macro-*DTSM* is that macro variables may have predictive content for excess returns *over and above* the standard level, slope, and curvature factors of *DTSM*s. That is, we explicitly accommodate macro risks that are theoretically *unspanned* by bond yields. We quantify the impacts of unspanned variation in output growth and inflation in the U.S. on the market prices of yield curve risks in swap markets over the past twenty years. Our analysis reveals that both level and slope risks are priced and that the model-implied excess returns associated with exposures to these risks vary substantially over the business cycle, particularly with changes in real economic activity. We also quantify the responses of forward term premiums to macroeconomic shocks and, conversely, the response of output growth to shocks to term premiums.

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.

There is compelling descriptive evidence that compensation for bearing exposure to these risks is correlated with business cycle variables. For instance, Figure 1 displays the realized excess returns, over one-month and one-year holding periods, on a portfolio of bonds that reflects *pure* slope risk. That is, its payoff tracks movement in the slope of the U.S. swap curve, while being (locally) invariant to changes in the level or curvature of the yield curve (see Section 3.1 for details). A striking feature of these excess returns is how closely they track the negative of the growth rate of U.S. industrial production ($-GIP$) over the past twenty years.[2]

This paper develops and implements an arbitrage-free $DTSM$ within which we are able to econometrically identify the contributions of macroeconomic variables to variation in the market prices of level, slope, and curvature risks. Specifically, we quantify how variation in output growth and inflation in the U.S. influenced the market prices of these risks in swap markets over the past twenty years. Our analysis reveals that both level and slope risks are priced and that the model-implied excess returns associated with exposures to these risks vary substantially over the business cycle, particularly with changes in real economic activity. We also quantify the responses of forward term premiums to macroeconomic shocks and, conversely, the response of output growth to shocks to term premiums. Within our model, we reassess some of Chairman Bernanke's interpretations of the interplay between term

[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. We focus largely on maturities between one- and ten-years. The introduction of longer-term maturities would also likely require additional risk factors. All of our subsequent analysis is easily extended to accommodate a wider span of maturities and additional priced yield-curve risks.

[2] GIP is measured at the inception of the investments in the slope-mimicking portfolios of bonds. So Figure 1 says that output growth has substantial contemporaneous correlation with risk premiums (expected excess returns) in swap markets. This is a distinct observation from the widely documented result that the slope of yield curve *itself* has predictive content for *future* output growth. On the latter, see for example Estrella and Mishkin (1998) and Wright (2006).

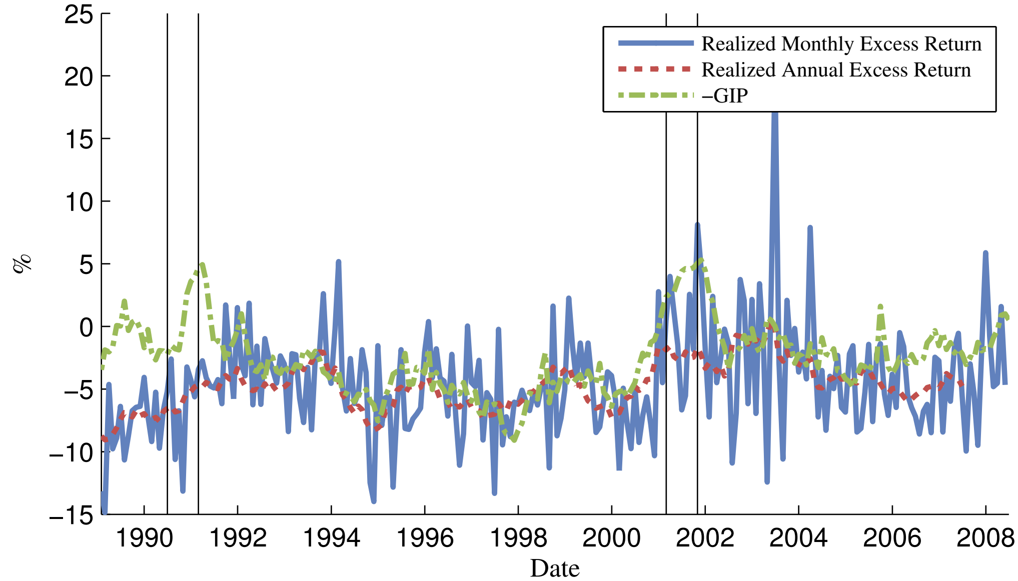

Figure 1: This figure displays the one-month and one-year realized excess returns on a portfolio of bonds with a payoff that, by construction, tracks changes in the slope of the U.S. swap curve. GIP is a smoothed version of the monthly growth rate of industrial production in the U.S.

premiums, the shape of the yield curve, and macroeconomic developments in the U.S.

A key motivating consideration in developing our macro-$DTSM$ is that macro variables have significant predictive content for excess returns *over and above* the standard level, slope, and curvature factors of $DTSM$s (Cooper and Priestley (2008) and Ludvigson and Ng (2009)).[3] This predictive content of macro information cannot be replicated in macro-$DTSM$s that have the macro variables entering directly as pricing factors determining the short-term rate. By construction, such models enforce the theoretical spanning of the macro variables by the model-implied PCs of bond yields.[4] Hence they preclude an incremental role for macro information in the determination of risk premiums beyond the yield curve factors themselves. In

[3] 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.

[4] 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 (2008), and Chernov and Mueller (2009). See Joslin and Singleton (2010) for a more in depth discussion of the properties of models with *spanned* macro risks.

contrast, our macro-$DTSM$ explicitly accommodates macro risks that are theoretically *unspanned* by bond yields and, therefore, we can quantify the impacts of these unspanned macro risks on the market prices of yield curve risks.

The state vector in our macro-$DTSM$ is comprised of a three-dimensional set of pricing factors $\mathcal{P}_t$[5] and the macro variables $M_t' \equiv (GIP_t, INF_t)$, where INF_t is CPI inflation and GIP_t is the growth rate of industrial production. Model-implied bond yields are determined by $\mathcal{P}_t$ under the pricing measure and we show that, without loss of generality, we can choose $\mathcal{P}_t$ to be the yield-curve risks of interest, the first three PCs of bond yields. The historical distribution of $(\mathcal{P}_t, M_t)$ is specified so that, while $\mathcal{P}_t$ and M_t may be correlated, large components of INF_t and GIP_t that are unspanned by bond yields are accommodated. Furthermore, forecasts of future values of $\mathcal{P}_t$ may depend on the past history of M_t, so macro factors are potentially important determinants of expected excess returns on exposures to PC risks.

Allowing the information determining expected excess returns to strictly nest the information in model-implied bond yields (equivalently, the information in $\mathcal{P}_t$) is non-standard relative to many equilibrium models of bond price determination. However, we have seen that macro information has incremental forecasting power for excess returns over and above standard pricing factors. Moreover, Cochrane and Piazzesi (2005) and Duffee (2009a) provide similarly compelling evidence of a nested information structure using yield-curve information to forecast excess returns.

One of our primary objectives is an exploration of the effects of macro uncertainty on the market prices of PC risks within a macro-$DTSM$ that enforces the nested information structure documented in these descriptive studies. Along the way, we provide theoretical conditions for a macro-$DTSM$ with a fully unconstrained information structure to have the property that the dimensionality of $\mathcal{P}_t$ is smaller than that of the state vector driving expected excess returns. This result is of interest in its own right as it provides a means for researchers studying equilibrium affine models to verify that their models are able to match the joint distribution of expected excess returns and yields on bonds documented in reduced-form analyses.

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 $DTSM$.[6] We propose a more systematic

[5]The choice of three is consistent with both the PC evidence and previous studies of yield-based $DTSM$s. See Dai and Singleton (2003) and Piazzesi (2006) for surveys of this literature.

[6]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

approach that uses likelihood-based model selection criteria to search over all 2^{15} nested parameterizations of the risk premiums on exposures to the PC risks $\mathcal{P}$.[7] In this manner we end up with the most parsimonious model that preserves the essential ingredients called for by the selection criteria to fit risk premiums in bond markets. By design, this model selection exercise requires an arbitrage-free macro-$DTSM$ to identify the relevant market prices of risk.

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 similar framework to ours 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 provide a canonical framework for incorporating unspanned information that affects expected excess returns into Gaussian $DTSM$s 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 *Observable* Gaussian Macro-$DTSM$ with *Unspanned* Macro Risks

We define a (discrete-time) $\mathcal{R}$-factor *Gaussian DTSM* as a model with an $\mathcal{R}$-dimensional vector of (possibly latent) pricing factors $\mathcal{P}_t$ that follows a Gaussian Markov processes under the pricing ($\mathbb{Q}$) distribution, and in which the short rate is affine in $\mathcal{P}_t$. More formally, letting $(\Omega, \mathbb{P}, \mathcal{F})$ be a probability space with a discrete filtration $\{\mathcal{F}_t\}$, we suppose that there is a short-rate process $\{r_t\}$ that is related to the adapted Markov process $\{\mathcal{P}_t\}$ according to:

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

high-dimensional factor models; see, for example, Ang, Dong, and Piazzesi (2007). Dai and Singleton (2000) and Bikbov and Chernov (2008) are examples of papers that follow the second strategy.

[7] What makes this computationally feasible is that our canonical form is such that we achieve convergence to the global optimum of our likelihood function nearly instantaneously. A related approach to model selection based on Bayesian posterior odds ratios has recently been proposed by Bauer (2010). We say more about the differences in our approaches in Section 3.2.

$A\mathcal{P}\mathbb{Q}:$ There is an equivalent 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 $\epsilon^{\mathbb{Q}}_{\mathcal{P}t} \sim N(0, I_{\mathcal{R}})$, $\{e^{-\sum_{s=0}^{t-1} r_{t+s}} P_t\}$ is a $\mathbb{Q}$-martingale, and the price of a τ-period zero coupon bond is $E^{\mathbb{Q}}_t[e^{-\sum_{s=0}^{\tau-1} r_{t+s}}]$.

Assumptions Ar and $A\mathcal{P}^{\mathbb{Q}}$ ensure "affine pricing" so yields on zero-coupon bonds are affine functions of the pricing factors $\mathcal{P}$ (Duffie and Kan (1996)). We let y_t denote the model-implied counterparts of the J-vector of yields on zero-coupon bonds that are being used in assessments of goodness-of-fit, where $J \geq \mathcal{R}$.

Starting from this generic Gaussian setup, we propose a new subfamily of Gaussian macro-$DTSM$s that allows for macro factors M_t to influence expected excess returns on bond portfolios, while ensuring that M_t is *unspanned* by the pricing factors $\mathcal{P}_t$. The construction of our canonical macro-$DTSM$ with these properties maintains $Ar\mathbb{Q}$. However we nest $A\mathcal{P}\mathbb{Q}$ as part of the $\mathbb{Q}$ representation of the expanded, $\mathcal{N}$-dimensional state vector $X_t' \equiv (\mathcal{P}_t', M_t')$:

$AX\mathbb{Q}:$ Under $\mathbb{Q}$, X_t follows the Gaussian process

$$\begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} = \begin{bmatrix} K^{\mathbb{Q}}_{0\mathcal{P}} \\ * \end{bmatrix} + \begin{bmatrix} K^{\mathbb{Q}}_{\mathcal{P}\mathcal{P}} & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \mathcal{P}_{t-1} \\ M_{t-1} \end{bmatrix} + \sqrt{\Sigma_X}\epsilon^{\mathbb{Q}}_{Xt}, \tag{2}$$

where $\epsilon^{\mathbb{Q}}_{Xt} \sim N(0, I_{\mathcal{N}})$ and $\Sigma_{\mathcal{P}\mathcal{P}}$ is the upper 3×3 block of Σ_X, and M_t is assumed to be non-trivial so that $\mathcal{N} > \mathcal{R}$.

The specification of our macro-$DTSM$ is completed with the assumption:

$AX\mathbb{P}:$ Under the historical measure $\mathbb{P}$, the $\mathcal{N}$-dimensional state X_t follows the unconstrained VAR

$$\begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} = \begin{bmatrix} K^{\mathbb{P}}_{0\mathcal{P}} \\ K^{\mathbb{P}}_{0M} \end{bmatrix} + \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}\mathcal{P}} & K^{\mathbb{P}}_{\mathcal{P}M} \\ K^{\mathbb{P}}_{M\mathcal{P}} & K^{\mathbb{P}}_{MM} \end{bmatrix} \begin{bmatrix} \mathcal{P}_{t-1} \\ M_{t-1} \end{bmatrix} + \sqrt{\Sigma_X}\epsilon^{\mathbb{P}}_{Xt}, \tag{3}$$

with $\epsilon^{\mathbb{P}}_{Xt} \sim N(0, I_{\mathcal{N}})$.

The implications of the $(Ar\mathbb{Q}, AX\mathbb{Q}, AX\mathbb{P})$ for pricing are identical to standard $\mathcal{R}$-factor Gaussian $DTSM$s, but their implications for how macro risks affect expected excess returns are very different. Taking these issues in order, $A\mathbb{Q}r$ together with the assumption that M_t does not feed back on $\mathcal{P}$ under $\mathbb{Q}$ (the right upper block of $K^{\mathbb{Q}}_X$ in (2), $K^{\mathbb{Q}}_{\mathcal{P}M}$, is zero) imply that bond prices are determined by an autonomous $\mathcal{R}$-factor Gaussian $DTSM$ with pricing factors $\mathcal{P}$.

Elaborating, whether $DTSM$s adopt $Ar\mathbb{Q}$ with latent pricing factors $\mathcal{P}$ as in Dai and Singleton (2000) and Duffee (2002), or a representation of r_t with a mix of latent (L_t) and observed (M_t) pricing factors as in the specification

$$r_t = \rho_0 + \rho_M \cdot M_t + \rho_L \cdot L_t \tag{4}$$

adopted by Ang, Piazzesi, and Wei (2003) and Bikbov and Chernov (2008), among others, *they all imply the same theoretical bond prices and that the same risks are priced in the underlying economies.*[8] We will subsequently exploit this fact to show that any canonical Gaussian $\mathcal{R}$-factor $DTSM$ can be normalized so that the $\mathcal{R}$ priced risks $\mathcal{P}_t$ are the first $\mathcal{R}$ PCs of bond yields.

Where our macro-$DTSM$ differs from the extant literature is in how we model the effects of macro risk on expected excess returns on bonds or, equivalently, how we model the market prices of risk (MPRs) of the pricing factors $\mathcal{P}_t$. Assumption $AX\mathbb{P}$ allows for general feedback between the pricing factors $\mathcal{P}_t$ and the macro factors M_t.[9] Moreover, M_t is unspanned by bond yields: in general, there is variation in M_t that is orthogonal to variation in $\mathcal{P}_t$ and, in particular, to variation in the first $\mathcal{R}$ PCs of bond yields.[10]

The accommodation of unspanned macro risk has potentially important implications for the properties of model-implied expected excess returns. The excess return over h periods on a bond with maturity n issued at date t is

$$rx_{t+h}^{n,h} = -y_{t+h}^{n-h}(n-h) + y_t^n\, n - y_t^h\, h. \tag{5}$$

Since (3) implies that y_{t+h}^{n-h} is forecastable by both $\mathcal{P}$ and the unspanned components of M_t, macro information has forecasting power over and above the first $\mathcal{R}$ PCs of yields. In this manner we are able to accommodate the rich variation in risk premiums documented in previous descriptive studies within an arbitrage-free, macro-$DTSM$.

[8]Though all of these models share the same priced risks, in practice the precise nature of the risks being studied will depend on the set of bond yields used in estimation. Risks may differ not only between treasury and swap rates, but also across the different splines that have been used to extract zero yields from treasury coupon-bond yields.

[9]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.

[10]More precisely, so long as Σ_X is non-singular, $\sigma(M_{it}) \subsetneq \sigma(\mathcal{B}_t \cup M_t^{(-i)})$, 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 $\{P(X_t) = E_t^{\mathbb{Q}}[g(r_{t+t_1}, r_{t+t_2}, \ldots, r_{t+t_n})] : g \in C_0\}$.

In contrast, the typical way a Gaussian $DTSM$ is extended to a macro-$DTSM$ is by assuming that, under $\mathbb{P}$, $\mathcal{P}$ follows the process[11]

$$\mathcal{P}_t = K_{0\mathcal{P}}^{\mathbb{P}} + K_{\mathcal{P}\mathcal{P}}^{\mathbb{P}} \mathcal{P}_{t-1} + \sqrt{\Sigma_{\mathcal{P}\mathcal{P}}} \epsilon_{\mathcal{P}t}^{\mathbb{P}}. \tag{6}$$

Models that assume (6) have the strong theoretical implication that every element of M_t is spanned by y_t. This is an immediate implication of the fact that bond yields in macro-$DTSM$s satisfying (4) are affine functions of $\mathcal{P}_t' = (M_t', L_t')$. Therefore, by inverting these pricing relations, M_t can be expressed as an affine function of y_t.

Indirect evidence on the empirical validity of this spanning property comes from the projection of GIP onto the first three PCs of swap yields:[12] the R^2 is only 14%. In the light of this low R^2, adopting (4) in a three-factor ($\mathcal{R} = 3$) model with output growth as an element of M_t seems tantamount to assuming that over 80% of the variation in output growth arises from measurement error. Adding $PC4$ and $PC5$ as regressors – equivalently, increasing $\mathcal{R}$ to five – only raises the R^2 for output growth to 28%, suggesting that measurement errors also have to be huge to sustain a five-factor macro-$DTSM$ with r_t given by (4).

Moreover, Gaussian macro-$DTSM$s based on (6) imply that, once one conditions on the first $\mathcal{R}$ PCs of bond yields, the macro information M_t is irrelevant for forecasting excess returns on portfolios of bonds. Compelling evidence against this restriction is presented in the descriptive analysis of Ludvigson and Ng (2009), as well as by Figure 1 and the evidence from our macro-$DTSM$s presented subsequently. Macro information has forecasting power for risk premiums over and above the information in the PCs of bond yields.

This heuristic derivation of our macro-$DTSM$ leaves many questions unaddressed, including: Are all $\mathcal{N}$-factor, Gaussian $DTSM$s with unspanned macro risks observationally equivalent to a $DTSM$ of this form? What is the minimal set of normalizations that achieve an econometrically identified model with unspanned macro risks? Are macro (e.g., inflation or output growth) risks priced in these macro-$DTSM$s? The following proposition answers the first two of these questions, with the proof given in Appendix A. We take up the last question in Section 2.2. Notationally, we let $UMA_0^{\mathcal{R}}(\mathcal{N})$denote the family of observable Gaussian $DTSM$s with $\mathcal{R}$ pricing factors and $\mathcal{N} - \mathcal{R}$ unspanned macro conditioning variables, analogous to the notation of Dai and Singleton (2000).

[11]Examples of this modeling strategy include Ang, Dong, and Piazzesi (2007), Bikbov and Chernov (2008), Chernov and Mueller (2009), and Smith and Taylor (2009). The following observations carry over to the higher-order processes adopted in some of these studies.

[12] See Section 2.1 for a detailed description of our data set.

Proposition 1 *Every canonical model*[13] *for the family* $UMA_0^{\mathcal{R}}(\mathcal{N})$ *is observationally equivalent to the following canonical Gaussian macro-DTSM: the state vector is* $X_t = (\mathcal{P}_t, M_t)$, *where* $\mathcal{P}_t$ *are the first* $\mathcal{R}$ *principal components of* y_t; $r_t = \rho_0 + \rho_{\mathcal{P}} \cdot \mathcal{P}_t$; *and the* $\mathbb{P}$ *and* $\mathbb{Q}$ *representations of* X_t *are given by (3) and (2), respectively. Moreover, there is a unique mapping between the parameters governing* $(\rho_0, \rho_{\mathcal{P}}, K_{0\mathcal{P}}^{\mathbb{Q}}, K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}})$ *and the parameter set* $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}})$, *where* $\lambda^{\mathbb{Q}}$ *denotes the* $\mathcal{R}$*-vector of ordered non-negative* $\mathbb{Q}$*-eigenvalues of* $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$,[14] *and* $r_\infty^{\mathbb{Q}}$ *denotes the long-run mean of the short rate under* $\mathbb{Q}$. *Thus, the parameter vector for the full model is* $(K_0^{\mathbb{P}}, K_1^{\mathbb{P}}, \Sigma_X, \lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}})$.

Several features of our canonical model warrant further discussion. First, we have chosen to rotate $\mathcal{P}$ so that it becomes the first $\mathcal{R}$ PCs of the yields y_t and, thus, the pricing factors and macro variables are both *observable* in the sense that they have immediately recognizable counterparts in the data. If fact, we are free to select the pricing factors to be any $\mathcal{R}$ linearly independent combinations of y_t,[15] with each such linear combination defining a different and observationally equivalent canonical model. This is a consequence of the fact that bond yields are affine functions of $\mathcal{P}$ so we can invert the pricing model to replace any given set of pricing factors, including latent factors, by portfolios of yields.

Second, a striking implication of Proposition 1 is that, regardless of a given choice of yield portfolios $\mathcal{P}$ as the pricing factors, the associated $\mathbb{Q}$ distribution is fully characterized by the parameter vector $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}})$. The matrix $\Sigma_{\mathcal{P}\mathcal{P}}$ depends, of course, on the portfolio of yields comprising $\mathcal{P}$. However, $(\lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}})$ are rotation-invariant (that is, independent of the choice of pricing factors) and, hence, are economically interpretable parameters. Appendix A gives the explicit construction of $(\rho_0, \rho_{\mathcal{P}}, K_{0\mathcal{P}}^{\mathbb{Q}}, K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}})$ from $(\Sigma_{\mathcal{P}\mathcal{P}}, \lambda^{\mathbb{Q}}, r_\infty^{Q})$ for our choice of PCs as pricing factors.

Finally, given our choice of $\mathcal{P}_t$ and the way we have normalized its $\mathbb{Q}$ distribution, our canonical model is an econometrically identified and maximally flexible model. Any other model in $UMA_0^{\mathcal{R}}(\mathcal{N})$ can be mapped uniquely to a special case of our canonical model.

[13] A canonical form for the family $UMA_0^{\mathcal{R}}(\mathcal{N})$ is a family of models (called canonical models) that are maximally flexible, in the sense that a minimal set of normalizations is imposed to achieve econometric identification.

[14] Building upon the yields-based analysis in Joslin, Singleton, and Zhu (2010), we can construct similar canonical models for the cases where the $\mathbb{Q}$ eigenvalues of $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$ are not distinct or where some of these eigenvalues are zero.

[15] One such choice would be to select $\mathcal{R}$ distinct yields on zero-coupon bonds.

2.1 Empirical Motivations and Economic Underpinnings for For the Family $UMA_0^{\mathcal{R}}(\mathcal{N})$ of Macro-$DTSM$s

From the empirical literature on affine $DTSM$s, there is growing evidence that bond yields and their associated expected excess returns are better described by a $DTSM$ in which the set of pricing factors $\mathcal{P}$ is of lower dimension than that of the state vector X_t determining risk premiums; that is, that $\mathcal{N} > \mathcal{R}$. Cochrane and Piazzesi (2005) found that there is information in the higher order PCs of bond yields that is useful for forecasting excess returns on bonds, over and above the information in the first three PCs. Duffee (2009a) finds that two of his factors (within a model with $\mathcal{N} = \mathcal{R} = 5$) play essentially no role in fitting the cross-section of bond yields, while having a quantitatively important effect on expected excess returns. These patterns are consistent with a yield-based model in which $\mathcal{N} = 5$ and $\mathcal{R} = 3$.

We set $\mathcal{R} = 3$ and normalize our macro-$DTSM$ so that $\mathcal{P}$ is the first three PCs of y_t, a collection of zero-coupon bond yields constructed from the LIBOR and swap yield curves.[16] LIBOR/swap data are used instead of U.S. Treasury data to avoid the large on- and off-the-run premiums in Treasury markets,[17] and because this is the benchmark curve underlying fixed-income operations at most large financial institutions. PCi is the i^{th} principal component of the continuously compounded 6 month, one- through five-year, seven-year, and ten-year nominal zero coupon bond yields. For our data and range of maturities, over 99% of the variation in bond yields is explained by their first three PCs, hence our choice of $\mathcal{R} = 3$.

To complete our macro-$DTSM$ we set $\mathcal{N} = 5$ and include the core CPI inflation rate from the Bureau of Labor Statistics (INF) and the growth rate of industrial production from the Federal Reserve's G.17 release (GIP) in our state vector X_t. By including $M_t' = (INF, GIP)$ we encompass the essential ingredients of standard Taylor-style policy rules and include the two macro risks that have received the most attention in prior studies of macro-$DTSM$s. The sample is monthly from January, 1989 through June, 2008.[18]

[16] On each date, these are bootstrapped from 6m LIBOR and the available subset of the 1y-10y swap rates, under the assumption of constant forward rates between maturities. LIBOR rates are provided by the British Bankers' Association (downloaded from Datastream), and monthly swap rates are mid-market rates from Bloomberg. Data were taken from the last trading day of the month. On 5 occasions (9/99, 10/99, 12/99, 4/02, and 5/04), the 1y swap rate was unavailable in Bloomberg and data from Datastream was used instead.

[17] See Duffee (1996) for a discussion of money-market effect on the short end of the Treasury curve, and Feldhutter and Lando (2007) and Krishnamurthy and Vissing-Jorgensen (2008) for evidence that U.S. Treasury yields embody a substantial, indigenous convenience premium.

[18] A disadvantage of working directly with monthly CPI inflation is that it is a highly volatile series which seems to reflect large, transitory shocks, one source of which may be measurement error.

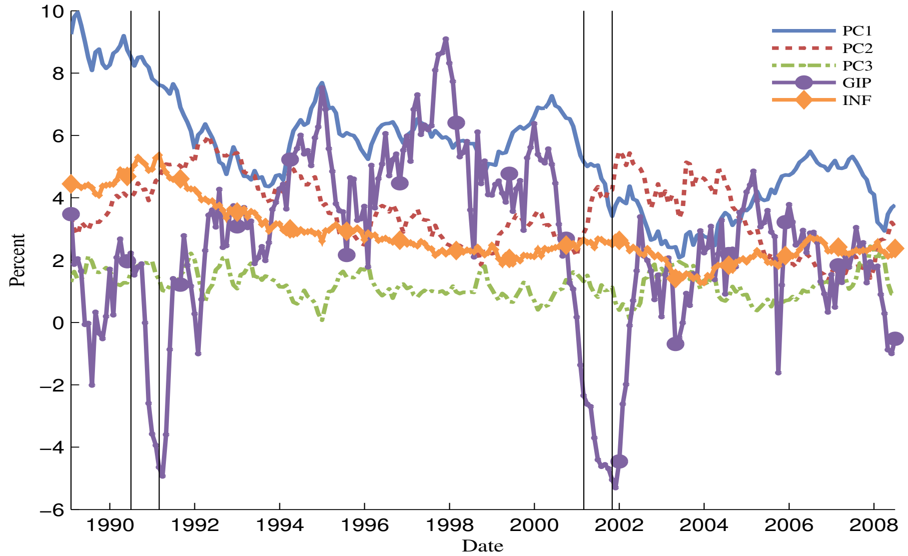

Figure 2: **Term Structure and Macro Variables** This figure plots the time series of $(PC1, PC2, PC3)$ of swap-implied zero yields and smoothed growth in industrial production and CPI inflation. The vertical bars mark $NBER$ recessions.

Normalized values of the state variables used in our analysis of macro $DTSM$s are displayed in Figure 2.[19] The LIBOR-based swap market largely came into existence in the late 1980's and this explains the start date of our sample. During this period there was a generally declining level of interest rates ($PC1$), accompanied by sizable swings in both the level and slope of the swap curve. Clearly visible in the GIP series are the recessions during the early 1990's and 2001. Inflation was relatively benign.

Though our canonical framework fully accommodates inclusion of both higher-order PCs and macro variables in X_t, for our case of swap yields, we found that adding higher-order PCs to our choice of X_t was not informative about expected excess returns. Table 1 displays the R^2s of the projections of realized excess returns

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. The growth rate of industrial production is filtered similarly. Wachter (2006) and Kim (2008) apply similar filters.

[19] 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 GIP to an annual scale. Now all variables take on values in $[-5\%, 10\%]$.

LHS \ RHS	$PC1$-$PC3$	$PC1$-$PC5$	$PC1$-$PC3$, GIP, INF
$rx^{(24)}_{t+12}$	0.20	0.32	0.39
$rx^{(60)}_{t+12}$	0.30	0.36	0.38
$rx^{(120)}_{t+12}$	0.34	0.36	0.39

Table 1: Regression R^2s from Excess Return Regressions

from holding an n-month bond from time t to time $t+12$, $rx^{(n)}_{t+12}$, onto the yield-based PCs, GIP, and INF. The R^2s from the projections onto $PC1 - PC5$ are comparable to those reported by Cochrane and Piazzesi (2005),[20] confirming that there is incremental informational content to $PC4$ and $PC5$ for forecasting excess returns. Inclusion of GIP and INF in these projections renders $PC4$ and $PC5$ statistically insignificant, and inclusion of (GIP, INF) leads to a larger increase in R^2 than inclusion of $(PC4, PC5)$. These findings suggest that it is macroeconomic risk that underlies a large part of the variation in excess returns in swap markets that is not captured by $\mathcal{P}_t$. Therefore, our subsequent analysis focuses on macro variables alone in M_t, and explores how these variables impact the MPRs in bond markets.

Fixing $\mathcal{N} = 5$, it might seem innocuous to proceed with a model in which $\mathcal{R} = 5$, as such a model fully encompasses our more restrictive formulation with $\mathcal{R} = 3$. However, Joslin, Singleton, and Zhu (2010) and Duffee (2009b) provide compelling evidence of over-fitting in their Gaussian $DTSM$s with $\mathcal{P}_t$ chosen to be the first five PCs of bond yields. With $\mathcal{R} = 5$, model-implied Sharpe ratios for various bond portfolios are implausibly large and volatile, and the fitted values of yields on bonds with maturities outside the range used in estimation take on wildly implausible, in-sample values. These issues do not arise in our more parsimonious framework with a small number of pricing factors ($\mathcal{R} = 3$) and a rich conditioning information set for risk premiums ($\mathcal{N} = 5$).

2.2 Priced Risks in the Family $UMA_0^{\mathcal{R}}(\mathcal{N})$

An inherent feature of any macro-$DTSM$ is that the only risks that are potentially priced are the shocks to the pricing factors $\mathcal{P}_t$. From (3) we obtain the drift $\mu^{\mathbb{P}}_{\mathcal{P}}(X_t)$ of $\mathcal{P}_t$ under $\mathbb{P}$. The drift of $\mathcal{P}_t$ under $\mathbb{Q}$, $\mu^{\mathbb{Q}}_{\mathcal{P}}(\mathcal{P}_t)$, comes from (2) and it depends only

[20] Our PCs are not identical to those used by Cochrane and Piazzesi, because we construct PCs using bonds with maturities out to ten years, whereas their maximum maturity is five years.

on $\mathcal{P}_t$. Combining these components gives the market price of $\mathcal{P}_t$ risk:

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

What differentiates our model from standard formulations of three-factor macro-$DTSM$s are: (i) macro information contributes incremental forecasting power for excess returns over and above our $\mathcal{P}_t$, and (ii) these macro risks are priced as distinct risks from the PCs comprising $\mathcal{P}_t$. In contrast, in macro-$DTSM$s in which M_t is spanned by the first $\mathcal{R}$ PCs, M_t is not a priced risk other than in the degenerate sense that its MPR is an affine function of the MPRs of these PCs.

Though our canonical model fully accommodates unspanned macro risks that are priced in fixed-income markets, using bond yield information alone, we cannot identify the MPRs of such unspanned risks. The $\mathbb{P}$ drift of X_t, $\mu_{Xt}^{\mathbb{P}}$, is again taken directly from (3). However, the component of $\mu_{Xt}^{\mathbb{Q}}$ associated with M_t plays no role in pricing. Hence the subvector of $\Lambda_{Xt} = \Sigma_X^{-1/2}(\mu_{Xt}^{\mathbb{P}} - \mu_{Xt}^{\mathbb{Q}})$ associated with M_t, Λ_{Mt}, is not identified from historical information on X_t and bond prices.[21]

Nevertheless, the effects of macro risks – both spanned and unspanned – on the risk premiums and expected excess returns for exposures to $\mathcal{P}$ risks are fully identified in our canonical model. Accordingly we are able to quantitatively explore the effects of shocks to unspanned macro risks on term premiums. Moreover, by accommodating the unspanned macro risks called for by the historical data, we are likely to obtain more reliable assessments of the properties of $\Lambda_{\mathcal{P}}(X_t)$ and, thereby, more reliable measurement of the MPRs of spanned macro risks as well.

2.3 Interpreting the Error From Projecting M_t onto $\mathcal{P}_t$

Given the importance of this distinction between spanned and unspanned macro risks, and the emphasis in the literature on spanned risks, it is instructive to briefly elaborate. Suppose that, as in our model, the short rate satisfies $Ar\mathbb{Q}$. However, instead of imposing (2), suppose that M_t is related to the pricing factors according to $M_t = \gamma_0 + \gamma_1 \mathcal{P}_t$. This is the theoretical spanning condition implied by (4).

To break the counterfactual empirical implication that the observed macro variables (M_t^o) are spanned by $\mathcal{P}_t$, one could introduce measurement errors: M_t^o differs from

[21] More can be said about the market prices of unspanned macro risks if securities whose payoffs depend directly on the unspanned risks are included in the analysis. For instance, the MPRs of unspanned inflation risk are potentially identified from TIPS yields, as in D'Amico, Kim, and Wei (2008) and Campbell, Sunderam, and Viceira (2009). Owing to the illiquidity and limited availability of data on TIPS, most of the extant literature on macro-$DTSM$s focuses on nominal bond yields alone, and we do so as well.

its theoretical counterpart M_t by an additive error:[22]

$$M_t = \gamma_0 + \gamma_1' \mathcal{P}_t + \nu_t, \tag{8}$$

where ν_t is say a Gaussian error satisfying $E^{\mathbb{P}}[\nu_t | M_{t-s}, \mathcal{P}_{t-s}, s = 1, 2, \ldots] = 0$. While this ensures that M_t is not literally a linear combination of the model-implied PCs, the error ν_t plays no role in pricing or in forecasting excess returns. That is, once one conditions expected excess returns on the PCs comprising $\mathcal{P}_t$, M_t has no additional forecasting power for these returns, contrary to the empirical evidence.

A more subtle way of breaking the perfect spanning condition is implicit in the framework of Kim and Wright (2005), the model cited by Chairman Bernanke in several of his assessments of the impact of the macro economy on bond market risk premiums. Kim and Wright assume that $M_t = INF_t$, and they arrive at their version of (8) for inflation by assuming that expected inflation is an affine function of 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 X_t implied by their assumptions is:

$$\begin{bmatrix} \mathcal{P}_t \\ M_t \end{bmatrix} = \begin{bmatrix} K^{\mathbb{P}}_{0\mathcal{P}} \\ \gamma_0 \end{bmatrix} + \begin{bmatrix} K^{\mathbb{P}}_{\mathcal{P}\mathcal{P}} & 0 \\ \gamma_1' K^{\mathbb{P}}_{\mathcal{P}\mathcal{P}} & 0 \end{bmatrix} \begin{bmatrix} \mathcal{P}_{t-1} \\ M_{t-1} \end{bmatrix} + \sqrt{\Sigma_X} \begin{bmatrix} \epsilon^{\mathbb{P}}_{\mathcal{P}t} \\ \eta_t \end{bmatrix}, \tag{9}$$

where $\eta_t = (\nu_t + \gamma_1' \sqrt{\Sigma_{\mathcal{P}\mathcal{P}}} \epsilon^{\mathbb{P}}_{\mathcal{P}t})$. Thus, the Kim-Wright formulation leads to a constrained special case (3) in which the history of M_t has no 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 (9) are not supported by our data.

2.4 The Likelihood Function for Our Canonical Model

We adopt the canonical form of Proposition 1. Initially, we suppose that the first three PCs of y_t ($\mathcal{P}_t$) are priced perfectly by the model and that the higher order $PC4 - PC8$ (PC^e) are priced with *i.i.d.* $N(0, \Sigma_e)$ errors. In this case, the joint (conditional) likelihood function can be written as:

$$\begin{aligned} \ell(X_t, PC^e_t | X_{t-1}; \Theta) &= \ell(PC^e_t | X_t, X_{t-1}; \Theta) \times \ell(X_t | X_{t-1}; \Theta) \\ &= \ell(PC^e_t | X_t, X_{t-1}; \lambda^{\mathbb{Q}}, r^{\mathbb{Q}}_\infty, L_X, L_e) \times \ell(X_t | X_{t-1}; K^{\mathbb{P}}_X, K^{\mathbb{P}}_0, L_X), \end{aligned} \tag{10}$$

[22]The properties of macro-$DTSM$s that assume r_t follows (4) and that M_t is measured with error are explored in Joslin and Singleton (2010).

where L_X and L_e are the Cholesky factorizations of Σ_X and Σ_e, respectively. The conditional density $\ell(X_t|X_{t-1};\Theta)$ depends on $(K_0^{\mathbb{P}}, K_X^{\mathbb{P}}, L_X)$, 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. Therefore, for any Σ_X, the $(K_0^{\mathbb{P}}, K_X^{\mathbb{P}})$ that maximize the likelihood are simply the standard OLS estimates. So, in estimation, we need only maximize the likelihood over $(L_X, r_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$ and L_e.

In formulating our likelihood function we have assumed a first-order VAR representation of our state $X_t' = (\mathcal{P}_t', GIP_t, INF_t)$. Ang, Piazzesi, and Wei (2003) and Jardet, Monfort, and Pegoraro (2009) posit higher-order VARs in studying $DTSM$s with (spanned) macro pricing factors. However, neither of these studies include inflation (a highly persistent process) nor $PC3$ in their state vector. Additionally, their much longer sample period includes the economically significant structural shifts in monetary policy in the late '70's and early '80's. In contrast, for our sample period of 1989 – 2008, choice of bond yields, and expanded state vector, a variety of formal model selection criteria all point to a first-order multivariate process (see Appendix B), consistent with our maintained assumption.

Our canonical form offers three advantages over working with a latent factor model directly as, for example, in Dai and Singleton (2000). First, the pricing factors are the observable constructs of interest, namely the PCs that have interpretations as level, slope, and curvature. Second, we end up searching over a low-dimensional parameter space and, hence, convergence to the global optimum of our likelihood function is extremely fast. Third, by having the primitive parameters be the readily interpretable long-run $\mathbb{Q}$ mean of r and the eigenvalues of $K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}$, one can often guess reasonably good starting values for $(r_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}})$. A good starting value for L_X comes directly from OLS estimation of (3).

3 Model Selection and Risk Premium Accounting

Within our canonical model there are sixty parameters governing the $\mathbb{P}$ distribution of X (those comprising $K_0^{\mathbb{P}}$, $K_X^{\mathbb{P}}$, and L_X). There are an additional four parameters governing the $\mathbb{Q}$ distribution of X ($r_\infty^{\mathbb{Q}}$ 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 $(K_0^{\mathbb{P}}, K_X^{\mathbb{P}})$ and $(K_0^{\mathbb{Q}}, K_X^{\mathbb{Q}})$ that are statistically insignificantly different from zero at a conventional significant level, and then to re-estimate the more parsimoniously parameterized constrained model.[23] The approach

[23] This was the parameter-reduction procedure followed, for example, by Dai and Singleton (2000), Ang and Piazzesi (2003), and Bikbov and Chernov (2008).

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.

For our empirical analysis we assume that all of the state variables, $X_t' = (\mathcal{P}_t', GIP, INF)$, are measured without error. So, while each individual bond yield is priced up to an additive error, the first three PCs of swap yields are assumed to be priced perfectly by our macro-$DTSM$. In Section 3.4 we verify the robustness of our findings to allowing all of the PCs to be priced up to additive errors.

3.1 Constraining the Market Prices of PC Risks

From (7) it follows that the state dependence of scaled MPRs of $\mathcal{P}$ risk, $\Sigma_{\mathcal{P}\mathcal{P}}^{1/2}\Lambda_{\mathcal{P}}(X_t)$, is governed by the matrix $\mathcal{K} \equiv (K_{\mathcal{P}X}^{\mathbb{P}} - K_{\mathcal{P}X}^{\mathbb{Q}})$, where $K_{\mathcal{P}X}^{\mathbb{P}}$ denotes the first three rows of K_X. We approach the model selection problem as one of finding the *best* set of zero restrictions on $\mathcal{K}$, where we trade off fit against the costs of over-parameterization. It has been common practice to enforce zero restrictions on the MPRs of the pricing factors in macro-$DTSM$s. We take a more systematic approach compared to say setting parameters to zero based on individual t-statistics. Furthermore, we focus on $\mathcal{K}$, rather than the parameters governing $\Lambda_{\mathcal{P}t}$, since the latter depend on the (arbitrary) choice of $\Sigma_{\mathcal{P}\mathcal{P}}^{-1/2}$. Our selection strategy is not implementable outside of a $DTSM$ as both $K_X^{\mathbb{P}}$ and $K_X^{\mathbb{Q}}$ must be econometrically identified.

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 the constraints on excess returns on the portfolio whose value changes (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 different perspective on the nature of priced risks in swap markets, and we are able to link up with the striking patterns in Figure 1.

3.2 Selecting Among 2^{15} Parameterizations of the MPRs

Since there are fifteen free parameters in $\mathcal{K}$, there are 2^{15} possible configurations of $DTSM$s with some of its entries set to zero. Though 2^{15} 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^{15} specifications examined, we compute full-information ML estimates of the parameters and then evaluate the Akaike (AIC, Akaike (1973)), Hannan and Quinn (HQIC, Hannan and Quinn (1979)), and Schwarz's Bayesian (SBIC, Schwarz (1978)) information criteria.[24] 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.[25]

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 (2009)). We also find that $PC1$, $PC2$, and INF exhibit behavior consistent with a near cointegrating relationship, whereas $PC3$ and GIP appear stationary. We do not believe that $(PC1, PC2, INF)$ literally embody unit-root components. At the same time we feel that it is essential to enforce a high degree of persistence under $\mathbb{P}$, since ML estimators of drift parameters are known to be biased in small samples. Furthermore, 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)).

With these considerations in mind, we proceed as follows. First, we search over all 2^{15} specifications of zero restrictions in $\mathcal{K}$, without consideration of the near cointegration among the PCs and INF. The resulting frontier of maximal values of the log-likelihood function achieved for each choice of the number of zero restrictions in $\mathcal{K}$ is displayed in Figure 3. The tangent points of the information criteria show that the AIC and HQIC criteria select ten zero restrictions in $\mathcal{K}$, while the SBIC criterion selects twelve restrictions (leaving five and three free parameters, respectively). All three criteria indicate that most of the parameters in $\mathcal{K}$ are not needed within our sample period to adequately describe the risk premiums on swap PC risks.

Next, to address the persistence issue, we recompute the frontier under the constraint that the largest eigenvalue of $K_X^{\mathbb{P}}$ is the same as the largest eigenvalue of $K_X^{\mathbb{Q}}$. From the extant literature we know that the largest eigenvalue of $K_X^{\mathbb{Q}}$ tends to be close to unity: one of the pricing factors exhibits near unit-root behavior. Without this constraint, the largest eigenvalue of $K_X^{\mathbb{P}}$ tends to be sufficiently below unity to imply that variation in long-dated forward term premiums is due entirely to risk premiums – expected future short-term rates out ten years or longer are

[24]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 MPRs. Nor does he address near-cointegration (see below).

[25]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.

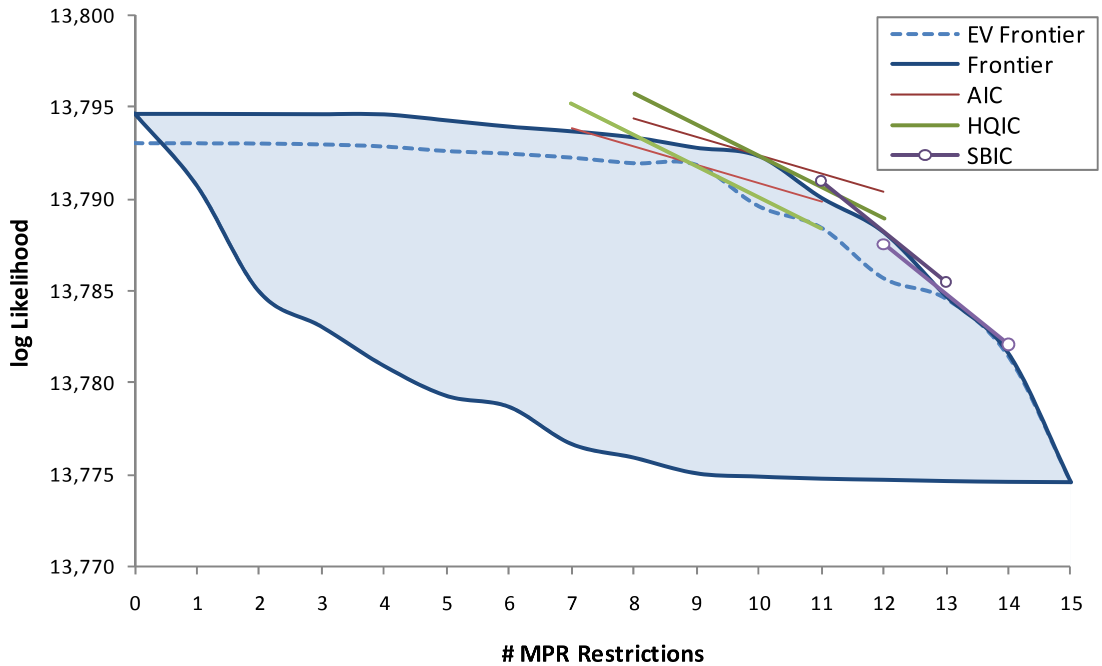

Figure 3: Frontiers of values of the log-likelihood function over 2^{15} specifications of $\mathcal{K}$, along with the tangency points for the AIC, HQIC, and SBIC information criteria.

virtually constant. This is inconsistent with surveys on interest rate forecasts (Kim and Orphanides (2005)).[26]

Moreover, the cross-section of bond yields precisely identifies the parameters of the $\mathbb{Q}$ distribution (in our case, $r^{\mathbb{Q}}$ and $\lambda^{\mathbb{Q}}$). Therefore, by enforcing this constraint, we force the $\mathbb{P}$ distribution of the state to inherit the near-cointegrating pattern for bond yields that arises under $\mathbb{Q}$, and we exploit the high degree of precision with which the cross-section of yields pins down this eigenvalue.

Both the AIC and HQIC information criteria trade our eigenvalue constraint for one zero restriction in $\mathcal{K}$. For both selection criteria, the best fitting $DTSM$ now has nine zero restrictions – six free parameters – in $\mathcal{K}$. On the other hand, because of the shape of the constrained EV frontier, the SBIC criterion now points to thirteen zero restrictions, one more than in the unconstrained model. Given that two of our three criteria point to nine zero restrictions on $\mathcal{K}$ and, as we will see, the extra parameter

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

plays a statistically significant role in subsequent analysis, we take as our preferred $DTSM$ the model selected by (AIC, HQIC).

Consistent with these selection results, the likelihood ratio (LR) statistic for testing our preferred model against the unconstrained canonical model is 6.76 with probability value 0.82. The LR statistic for testing the best fitting model against the canonical model that (only) imposes the eigenvalue constraint on $(K_X^{\mathbb{P}}, K_X^{\mathbb{Q}})$ is 3.40 with probability value 0.97.

3.3 Risk Premium Accounting: Model Comparison

To highlight the properties of the model selected from this extensive search we compare the properties of three Gaussian $DTSM$s with ($\mathcal{R} = 3, \mathcal{N} = 5$): (1) the unconstrained canonical model ($\mathcal{CM}$); (2) the canonical model with the constraint that the largest eigenvalue of $K_X^{\mathbb{P}}$ equals the corresponding eigenvalue of $K_X^{\mathbb{Q}}$ ($\mathcal{CM}_E$); and (3) the model imposing this eigenvalue constraint and the zero restrictions on $\mathcal{K}$ selected by the AIC and HQIC criteria ($\mathcal{CM}_E^R$). It turns out that the last row of $\mathcal{K}$ in model $\mathcal{CM}_E^R$ is set entirely to zero: our model selection criteria reveal that the expected excess returns over one month on exposure to curvature ($PC3$) risk are either zero or constant, but they are not state-dependent. The estimated value of the third entry of $(K_0^{\mathbb{P}} - K_0^{\mathbb{Q}})$, which governs the the constant portion of $xPC3_t$, is small with a relatively large standard error, suggesting that $PC3$ risk is not priced at all during our sample period. Therefore, in model $\mathcal{CM}_E^R$ we also set this intercept term to zero so as to gain some precision in estimating the levels of the risk premiums associated with $PC1$ and $PC2$ risks.

Maximum likelihood estimates of the parameters governing the $\mathbb{Q}$ distributions of X_t are displayed in Table 2a.[27] These estimates are very similar across the three models, regardless of the constraints imposed on our canonical model. Consistent with our earlier discussion, this says that the parameters of the $\mathbb{Q}$ distribution are determined largely by the cross-sectional restrictions on bond yields, and not by their time-series properties under the $\mathbb{P}$ distribution. Models $\mathcal{CM}_E$ and $\mathcal{CM}_E^R$ exploit this fact to restrict the degree of persistence of the state under $\mathbb{P}$.

The eigenvalues of $K_{\mathcal{PP}}^{\mathbb{P}}$ are displayed in Table 2b.[28] The largest $\mathbb{P}$-eigenvalue in the unconstrained model $\mathcal{CM}$ is smaller than in the constrained models. Although the latter difference might seem small, it is large enough to imply that expected

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

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

Parameter	$\mathcal{CM}$	$\mathcal{CM}_E$	$\mathcal{CM}_E^R$
$r_\infty^{\mathbb{Q}}$	0.1053	0.1054	0.1055
	(0.0043)	(0.0043)	(0.0043)
$\lambda_1^{\mathbb{Q}}$	0.9974	0.9974	0.9974
	(0.0002)	(0.0002)	(0.0002)
$\lambda_2^{\mathbb{Q}}$	0.9522	0.9522	0.9522
	(0.0022)	(0.0022)	(0.0022)
$\lambda_3^{\mathbb{Q}}$	0.8619	0.8619	0.8643
	(0.0116)	(0.0116)	(0.0104)

(a) Parameters of the $\mathbb{Q}$ Distribution

Parameter	$\mathcal{CM}$	$\mathcal{CM}_E$	$\mathcal{CM}_E^R$
$\lvert\lambda_1^{\mathbb{P}}\rvert$	0.9676	0.9974	0.9974
	(0.0476)	(0.0002)	(0.0002)
$\lvert\lambda_2^{\mathbb{P}}\rvert$	0.9396	0.9415	0.9443
	(0.0367)	(0.0264)	(0.0132)
$\lvert\lambda_3^{\mathbb{P}}\rvert$	0.9396	0.9415	0.9443
	(0.0367)	(0.0264)	(0.0132)
$\lvert\lambda_4^{\mathbb{P}}\rvert$	0.8613	0.8544	0.8829
	(0.0448)	(0.0363)	(0.0288)
$\lvert\lambda_5^{\mathbb{P}}\rvert$	0.8613	0.8544	0.8304
	(0.0448)	(0.0363)	(0.0164)

(b) Moduli of Eigenvalues of $K_X^{\mathbb{P}}$

Table 2: ML estimates of the $\mathbb{Q}$ parameters and of the moduli of the eigenvalues of $K_X^{\mathbb{P}}$ for models $\mathcal{CM}$, $\mathcal{CM}_E$, and $\mathcal{CM}_E^R$. Standard errors are given in parentheses.

future short-term rates out ten years or longer are virtually constant in model $\mathcal{CM}$. As discussed above, this counterfactual implication of model $\mathcal{CM}$ motivates our eigenvalue constraint in models $\mathcal{CM}_E$ and $\mathcal{CM}_E^R$.

Estimates for models $\mathcal{CM}_E$ and $\mathcal{CM}_E^R$ of the matrix $\mathcal{K}$ governing one-month expected excess returns on PC risk exposures are displayed in Table 3. Recall that the ordering of the variables in X_t, corresponding to the columns in Table 3, is $(PC1, PC2, PC3, GIP, INF)$. Focusing first on model $\mathcal{CM}_E^R$ (Table 3b), we see that the first and second rows of $\mathcal{K}$ have (statistically significant) non-zero entries, while

$\mathcal{P}$	$PC1$	$PC2$	$PC3$	GIP	INF
$PC1$	-0.0238	-0.0303	0.0072	0.0141	0.0436
	(0.0241)	(0.0226)	(0.0524)	(0.0077)	(0.0515)
$PC2$	0.0144	-0.0119	-0.0279	-0.0197	0.0112
	(0.0239)	(0.0240)	(0.0524)	(0.0079)	(0.0496)
$PC3$	-0.0233	-0.0099	-0.0096	0.0060	0.0368
	(0.0204)	(0.0180)	(0.0449)	(0.0067)	(0.0438)

(a) $\mathcal{K}$ for model $\mathcal{CM}_E$

$\mathcal{P}$	$PC1$	$PC2$	$PC3$	GIP	INF
$PC1$	-0.0394	-0.0329	0.0000	0.0189	0.0659
	(0.0150)	(0.0136)	(NA)	(0.0055)	(0.0290)
$PC2$	0.0159	0.0000	0.0000	-0.0169	0.0000
	(0.0074)	(NA)	(NA)	(0.0058)	(NA)
$PC3$	0.0000	0.0000	0.0000	0.0000	0.0000
	(NA)	(NA)	(NA)	(NA)	(NA)

(b) $\mathcal{K}$ for model $\mathcal{CM}_E^R$

Table 3: ML estimates of the matrix $\mathcal{K}$ governing expected excess returns on the PC mimicking portfolios. Standard errors are given in parentheses.

the last row is chosen to be zero by our model selection criteria. 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.[29] The finding that both level and slope risks are priced differs from the conclusion in Cochrane and Piazzesi (2008) that only level risk is priced. We attribute this difference to our having conditioned risk premiums on macro information.

The expected excess returns $xPC1_t$ and $xPC2_t$ both depend in statistically significant ways on $PC1$ and GIP. Expected excess returns on $PC1$ risk in model $\mathcal{CM}_E^R$ are also influenced by inflation and the slope of the swap curve ($PC2$). Note that, for both of these conditioning variables, the estimates for $\mathcal{CM}_E^R$ are larger (in absolute value) and more precisely estimated than their counterparts in model $\mathcal{CM}_E$.

[29] An alternative approach to analyzing risk premiums within a macro-$DTSM$ would have been to adapt the methods in Joslin, Singleton, and Zhu (2010) 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.

The positive signs on GIP and INF imply that risk premiums on $PC1$ exposures are pro-cyclical (positively correlated with GIP and INF). This can be seen graphically in Figure 4a for all three versions of our canonical macro-$DTSM$, 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 our three models about the (annualized) expected excess return on a level-mimicking portfolio, $xPC1_t$, with the biggest difference being for model $\mathcal{CM}$, particularly over the early part of our sample period. These excess returns take on their largest (absolute) values during the 1990 and 2001 recessions. More modest declines in $xPC1_t$ occur during 1994 and 2003–04. Notably, Chairman Bernanke gave speeches on deflation risk in November, 2002 and and again in July, 2003 expressing concerns about possible deflation in the U.S.[30] We see that, during this window of time, expected excess returns to bearing level risk – which is nearly cointegrated with inflation – fell sharply.

The negative sign on GIP in the second row of $\mathcal{K}$ implies that premiums on exposure to slope risk are counter-cyclical. The finding that $PC2$ risk is priced when we include macro information in X_t was anticipated by Figure 1. Nevertheless, it is striking that, after searching over 2^{15} specifications of the model-implied $xPC2_t$, our model-selection criteria place most of the weight on GIP and zero weights on $(PC2, PC3, INF)$ in characterizing expected excess returns on exposure to slope risk. The risk premium on $PC2$ risk achieves its lowest value (during our sample period) around the time of the Asian crisis in 1997–98. The slope of the swap curve was relatively flat during this period (see Figure 2). As the swap curve steepened during the recession in 2000–01 and GIP fell, $xPC2_t$ increased substantially.

ML estimates of $K_0^{\mathbb{P}}$ and $K_X^{\mathbb{P}}$ governing the $\mathbb{P}$-drift of X_t are displayed in Table 4 for model $\mathcal{CM}_E^R$.[31] Note that the non-zero coefficients on (GIP_{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 useful for forecasting future bond yields. Additionally, and not surprisingly, the coefficients on the own lags of GIP and INF are large and significantly different from zero. We conclude that the zero restrictions on the $\mathbb{P}$-drift of X_t implicit in the Kim and Wright (2005) model are strongly rejected by our data.

[30]See Chairman Bernanke's remarks "Deflation: Making Sure 'It' Doesn't Happen Here" before the National Economists Club on November 21, 2002 and "An Unwelcome Fall in Inflation?" before the Economists Roundtable on July 23, 2003.

[31]The zero entries are implied by the constraints on $\mathcal{K}$ selected by our model-selection criteria.

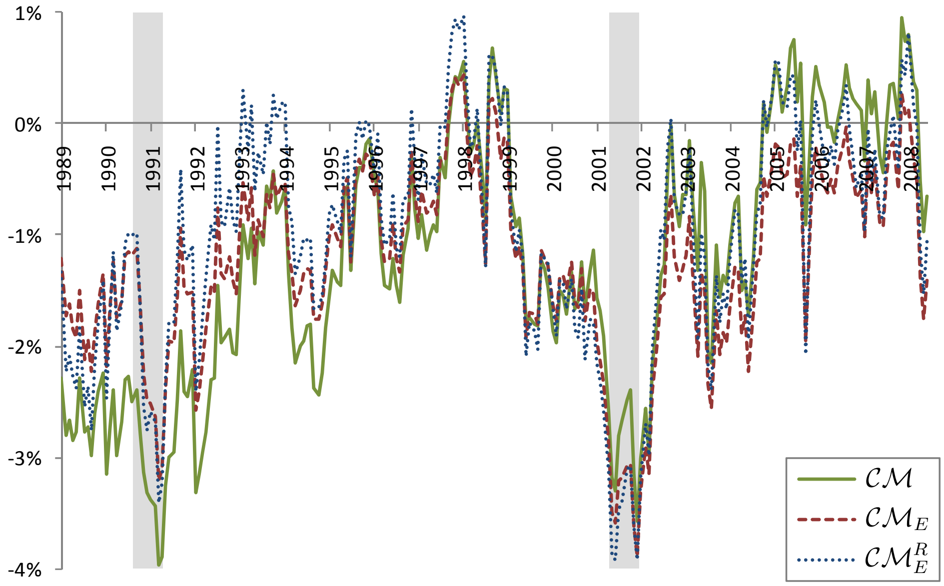

(a) Excess Return on Level-Mimicking Portfolio

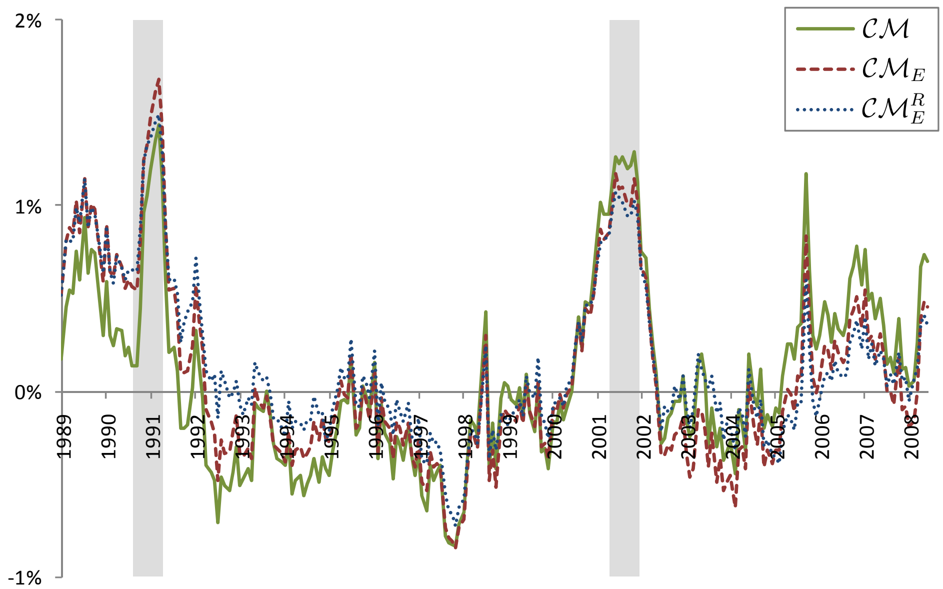

(b) Excess Return on Slope-Mimicking Portfolio

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

X	$K_0^{\mathbb{P}}$	$K_X^{\mathbb{P}}$ $PC1$	$PC2$	$PC3$	GIP	INF
$PC1$	0.0008	0.9566	−0.0031	−0.0593	0.0189	0.0659
	(0.0004)	(0.0149)	(0.0136)	(0.0026)	(0.0055)	(0.0290)
$PC2$	−0.0009	0.0059	0.9750	0.1642	−0.0169	0.0000
	(0.0005)	(0.0074)	(0.0014)	(0.0071)	(0.0058)	(NA)
$PC3$	0.0011	0.0157	−0.0066	0.8429	0.0000	0.0000
	(0.0001)	(0.0013)	(0.0020)	(0.0097)	(NA)	(NA)
GIP	0.0027	0.0373	0.1516	−0.3958	0.9109	−0.1119
	(0.0044)	(0.0881)	(0.0797)	(0.1885)	(0.0298)	(0.1832)
INF	−0.0008	0.0445	0.0160	0.0347	−0.0064	0.9105
	(0.0005)	(0.0108)	(0.0099)	(0.0235)	(0.0037)	(0.0232)

Table 4: Maximum Likelihood Estimates of $K_0^{\mathbb{P}}$ and $K_X^{\mathbb{P}}$ for Model $\mathcal{CM}_E^R$. Standard errors are reported in parentheses.

3.4 Robustness To Pricing Errors in Bond Markets

Up to this point we have assumed that the PCs comprising $\mathcal{P}_t$ are priced perfectly by our macro-$DTSM$s. The literature has often assumed instead that all bond yields (and hence their PCs) are priced up to additive measurement errors (e.g., Ang, Dong, and Piazzesi (2007), Duffee (2009a)). To assess the robustness of our analysis to relaxation of the assumption that $\mathcal{P}$ is priced perfectly by the model (observed $\mathcal{P}_t^o = \mathcal{P}_t$), we compute ML estimates of model $\mathcal{CM}$ assuming that $(y_t^o - y_t) \sim N(0, \sigma_m^2 I)$.[32] The likelihood function is constructed with the observation equation

$$y_t^o = A + B\mathcal{P}_t + \eta_t, \tag{11}$$

where $\mathcal{P}$ represents the model-implied PCs computed with loadings based on the historical bond yields, y_t^o. With the addition of measurement errors, ML estimation involves the use of the Kalman filter. To set up the Kalman filtering problem we start with our normalization with theoretical pricing factors $(\mathcal{P}_t, M_t)$. With an initial guess of $(\lambda^{\mathbb{Q}}, r_\infty^{\mathbb{Q}}, L_X)$, we construct $(K_{0\mathcal{P}}^{\mathbb{Q}}, K_{\mathcal{P}\mathcal{P}}^{\mathbb{Q}}, \rho_0, \rho_{\mathcal{P}})$. Based on the no-arbitrage pricing of bonds we then construct $A \in \mathbb{R}^J$ and $B \in \mathbb{R}^{J\times\mathcal{N}}$ with $y_t = A + B\mathcal{P}_t$. These theoretical pricing relations are linked to the data by the observation equation (11).

[32]The results presented here assume that $\Sigma_m = \sigma I_m$, consistent with the assumption in the previously cited literature.

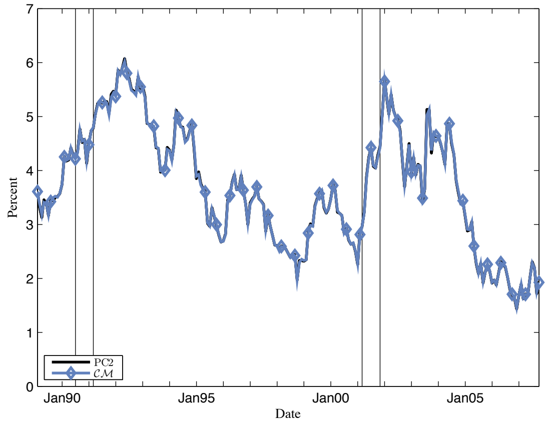

(a) $PC2$

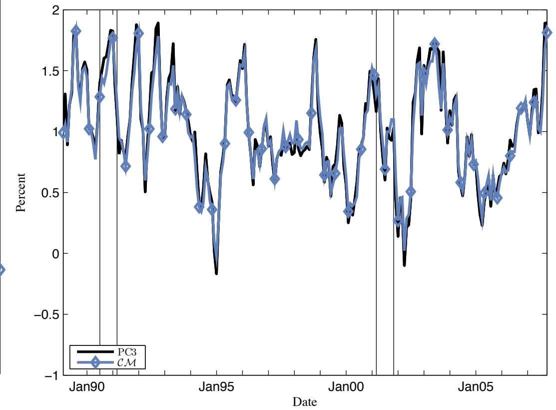

(b) $PC3$

Figure 5: Filtered $PC2$ and $PC3$ implied by model $\mathcal{CM}$.

To distinguish filtered values of the pricing factors from their theoretical and observed counterparts we use the notation $\mathcal{P}_t^F$.[33]

Figures 5a and 5b display the pairs of the observed PCs and their filtered, model-implied values $(PC2_t^o, PC2_t^F)$ and $(PC3_t^o, PC3_t^F)$. For $PC2$ these series are virtually on top of each other (this is even more so for $PC1$, not shown), and the differences for $PC3$ are small. These patterns show that the conclusions drawn from our models are robust to the introduction of pricing errors.[34]

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 6 displays the forward term premiums (FTP) based on the point estimates of model $\mathcal{CM}_E^R$ 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

[33]For our data and sample period the smoothed values ($E_T[\mathcal{P}_t]$) are similar to these filtered values.

[34]We focus on the model without pricing errors, since the introduction such errors adds computational complexity, given our eigenvalue restriction and use of model selection criteria. With pricing errors, we would need to filter for the latent states.

periods, as the level of rates fell, and are increasing in p.

Within our canonical model both forward rates and expected future one-year rates are affine functions of the state X_t: $FTP_t^{p,1} = \varsigma_0^{p,1} + \varsigma_X^{p,1} \cdot X_t$. Based on this relationship and the ML estimates of model $\mathcal{CM}_E^R$, we compute the 95% confidence bands for our estimated FTPs. The darker shaded areas in Figure 6 represent the confidence bands based on the precision of the ML estimates of $\varsigma_X^{p,1}$, and the wider, light-shaded bands reflect the sampling variability of the entire parameter set $(\varsigma_0^{p,1}, \varsigma_X^{p,1})$. 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 $\varsigma_X^{2,1}$ (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 $\mathcal{CM}_E^R$ gives quite precise measures of the dynamic properties of premiums.

The "in-2-for-1" forward term premium implied by model $\mathcal{CM}_E^R$ 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 (without a comparable eigenvalue restriction imposed) estimated using survey forecasts of future interest rates. Professional forecasters are conditioning (at least) on similar macro information as that embodied in GIP 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 $\mathcal{CM}_E^R$ comes from Figure 7 which displays standardized FTPs along with standardized versions of the PMI index constructed by the Institute for Supply Management[35] and of the Coincident Economic Index (CEI) as published by the Conference Board.[36] The CEI is constructed to be an indicator of current economic conditions. 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 CEI 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

[35]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.

[36]The Conference Board takes into account (with weight) employees on nonagricultural payrolls (0.5439), personal income less transfer payments (0.1873), industrial production (0.1497), and manufacturing and trade sales (0.1191).

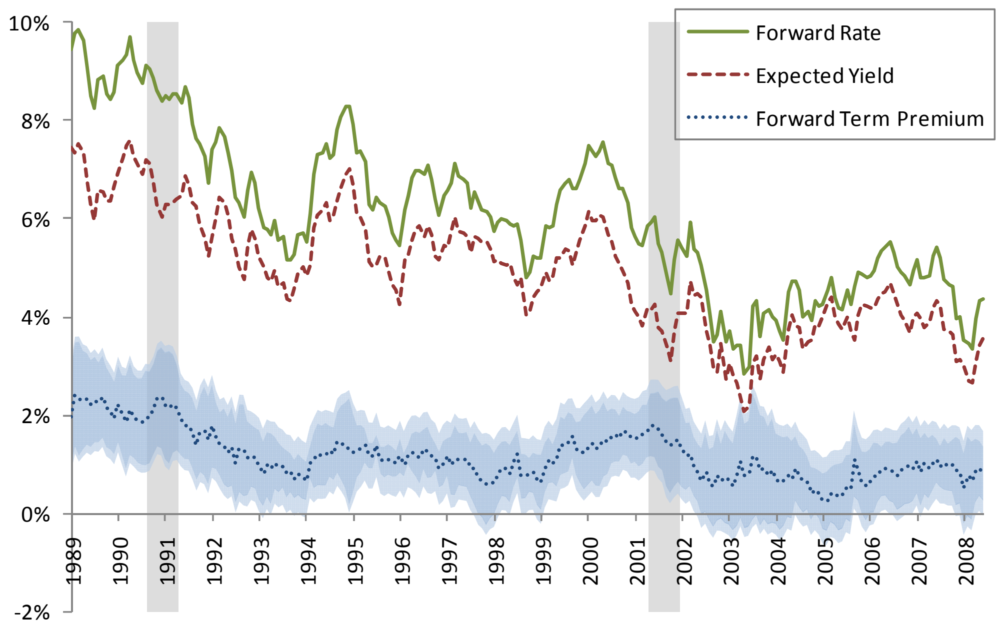

(a) In-2-For-1 Forward Term Premium

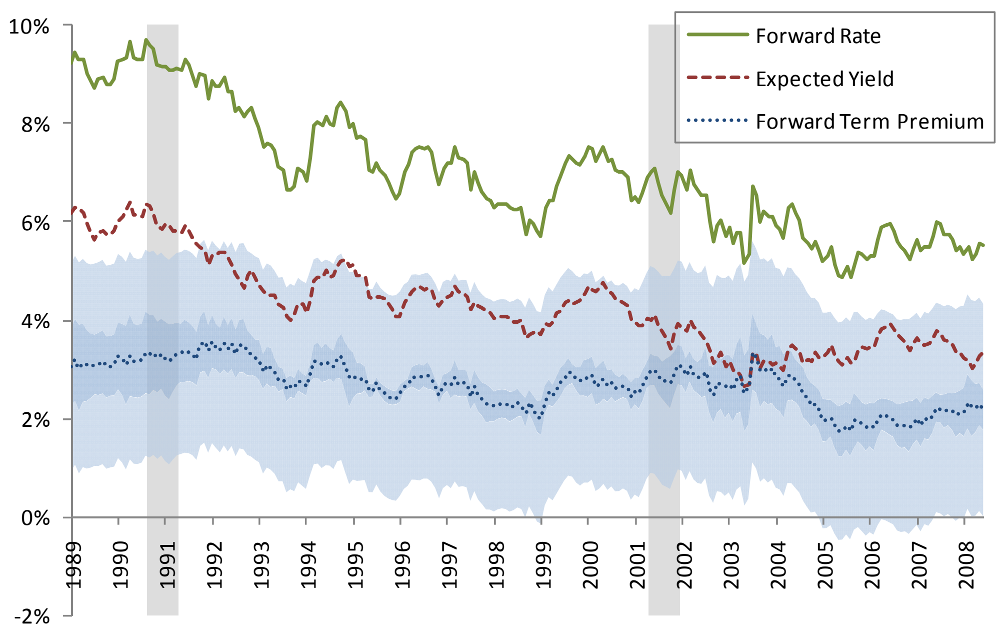

(b) In-9-For-1 Forward Term Premium

Figure 6: 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 $\mathcal{CM}_E^R$. The shaded areas are confidence bands for the forward term premiums.

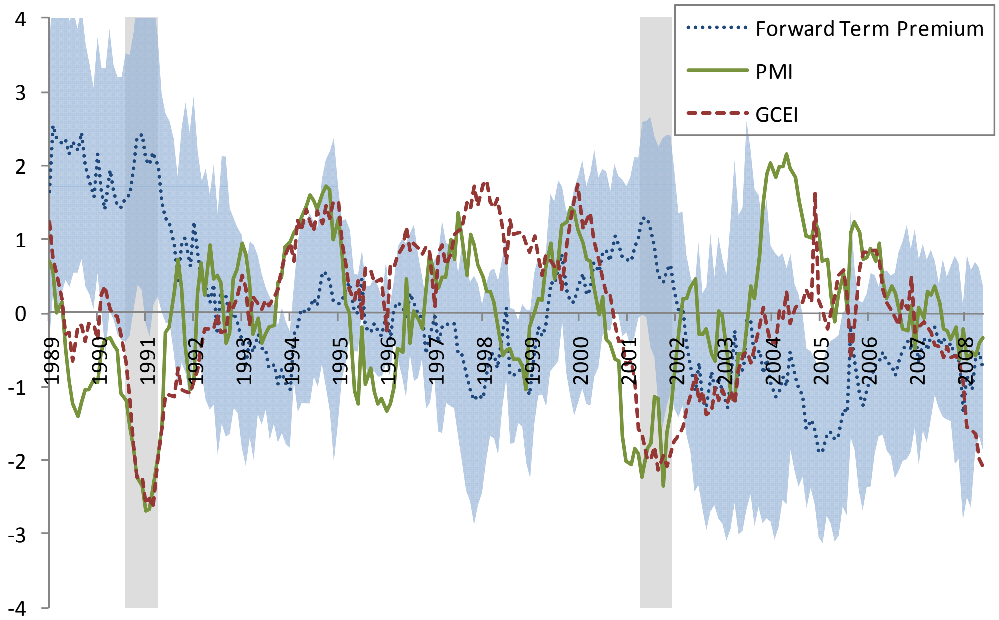

(a) In-2-For-1 Forward Term Premium

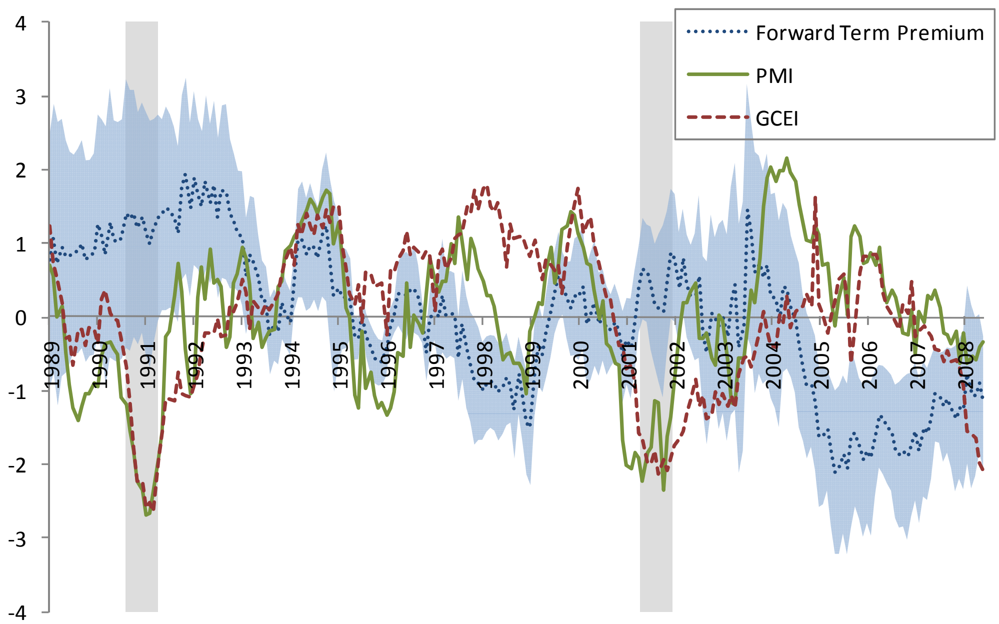

(b) In-9-For-1 Forward Term Premium

Figure 7: Standardized forward term premiums for "in-p-for-1" forward contracts, $p = 2$ and 9, implied by model $\mathcal{CM}_E^R$, 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	GIP	INF
2-for-1	−0.0069	0.3757	0.1693	0.2748	−0.1041	−0.3092
	(0.0098)	(0.1418)	(0.1945)	(0.2250)	(0.0494)	(0.2640)
5-for-1	−0.0035	0.2923	0.3836	0.2290	−0.0142	−0.2859
	(0.0123)	(0.1012)	(0.1414)	(0.1650)	(0.0292)	(0.1536)
9-for-1	0.0032	0.2235	0.4271	0.3529	−0.0189	−0.2651
	(0.0141)	(0.0898)	(0.1133)	(0.1211)	(0.0238)	(0.1346)

Table 5: Coefficients $\varsigma_0^{p,1}$ and $\varsigma_X^{p,1}$ determining the mapping between the forward term premiums $FTP_t^{p,1}$ and the state X_t in model $\mathcal{CM}_E^R$.

pessimistic sentiment, while in the latter case they were more optimistic, than what data comprising the CEI indicated about the economy's strength.

Neither $FTP^{2,1}$ nor $FTP^{9,1}$ follow an unambiguously counter-cyclical pattern. In fact, from roughly 1993 through 2000, the PMI and $FTP^{9,1}$ track each other quite closely, so forward term premiums tended to be pro-cyclical during this period. Inspection of the coefficients $\varsigma_X^{p,1}$ relating FTPs and X_t in Table 5 reveals that the negative weights on GIP and INF induce counter-cyclical movements in FTPs. 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.

Starting in late 1993 there was a substantial narrowing of the 95% confidence bands around $FTP_t^{9,1}$ (Figure 6b). This coincides with the time at which the Federal Reserve began announcing changes in its target for the federal funds rate. As documented by Swanson (2006), these announcements led to a substantial improvement in accuracy of private-sector forecasts of interest rates relative to the late 1980's. Model $\mathcal{CM}_E^R$ seems to capture well this increased precision, at least with regard to the expected future short-term rates embodied in forward term premiums.

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 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)). Given the role of swap transactions in corporate debt issuance, 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.[37] The evidence in Figure 7b is consistent with this. However it also suggests that falling forward term premiums were not driven by weak real economic activity alone.[38]

We turn next to a more in depth exploration of the macroeconomic forces underlying variation in risk premiums.

5 More on Macro Risks and the Term Structure

To delve more deeply into the effects of macroeconomic information on the shape of the swap curve it is instructive to distinguish between new information about *spanned* macro variables $(SGIP_t, SINF_t)$ – the components of (GIP_t, INF_t) explained by linear projections onto $\mathcal{P}_t$– and the unspanned residuals from these projections $(OGIP_t, OINF_t)$. As we discussed in Section 2, macro-$DTSM$s that include GIP and INF as pricing factors enforce theoretical spanning and, thereby, rule out *a priori* a role for $(OGIP, OINF)$ in forecasting risk premiums. In contrast, our family of models $UMA_0^3(5)$ allows for the possibility that $(OGIP, OINF)$ predict excess returns, and that their forecasting power is different from that of $(SGIP, SINF)$.

To gain insight into the roles of $(OGIP, OINF)$, we decompose the variation in each $FTP^{p,1}$ into two mutually orthogonal pieces: the component $S\mathcal{P}_t^p$ obtained by projecting $FTP_t^{p,1}$ onto $\mathcal{P}_t$, and the component OM_t^p obtained from projecting $FTP_t^{p,1}$ onto $(OGIP_t, OINF_t)$.[39] These two components are mutually orthogonal, since $(OGIP_t, OINF_t)$ are, by construction, orthogonal to $\mathcal{P}_t$. Figure 8 displays $(S\mathcal{P}_t^p, OM_t^p)$, for $p = 2, 9$. In both cases $S\mathcal{P}_t^p$ accounts for a majority of the variation in the FTPs, an expected result given that $S\mathcal{P}_t^P$, r_t, and y_t are all affine in $\mathcal{P}_t$. Of greater interest is the finding that OM_t^2 accounts for 25% of the variation in $FTP_t^{2,1}$, with OM_t^2 being particularly large during recessions and in the aftermath of Hurricane

[37] 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).

[38] Wright (2009) conjectures that the relatively steep decline in forward term premiums during the conundrum period may be attributable in part to declining uncertainty about future inflation rates. Such a decline would likely be reflected in the shape of the swap curve and hence in FTPs through these PCs.

[39] This decomposition is very different than the decomposition of r_t studied by Bikbov and Chernov (2008). They assume that M_t is spanned by $\mathcal{P}_t$ so $(OGIP_t, OINF_t)$ are set to zero in their model.

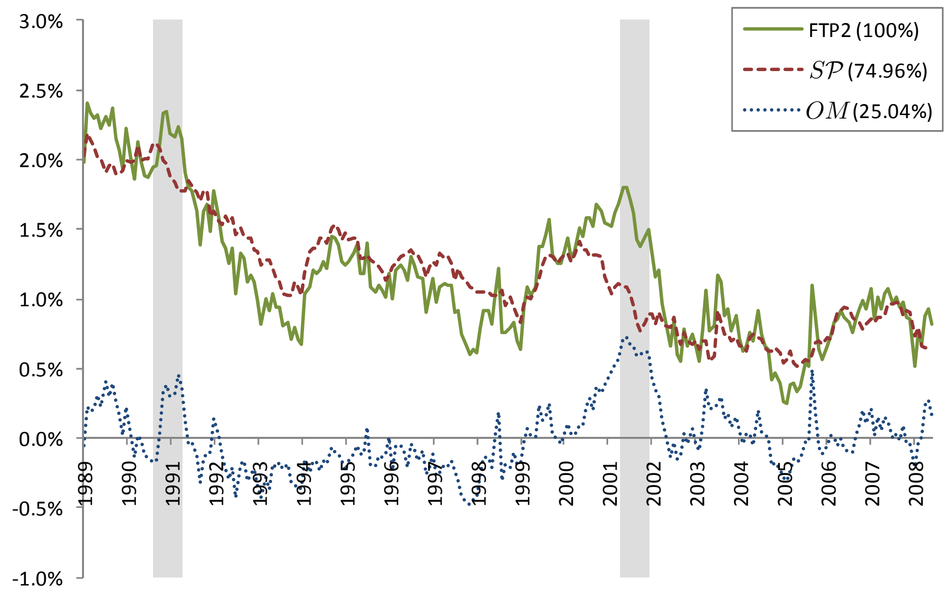

(a) Decomposition of $FTP^{2,1}$

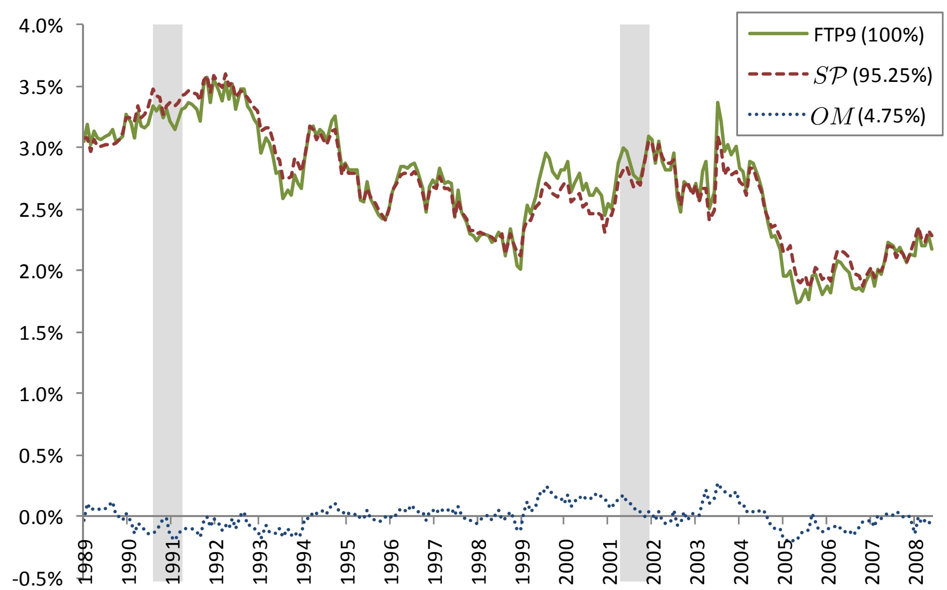

(b) Decomposition of $FTP^{9,1}$

Figure 8: Decomposition of $FTP^{p,1}$ into two orthogonal components: $S\mathcal{P}_t^p$ spanned by $\mathcal{P}_t$ and OM_t^p spanned by $(OGIP_t, OINF_t)$, based on estimates from model $\mathcal{CM}_E^R$.

Katrina in 2005 when GIP_t experienced a large, temporary decline. We conclude that models that omit a role for macro risks that are unspanned by $\mathcal{P}$ are likely to misrepresent the structure of excess returns on bond portfolio positions, at least over short- to intermediate horizons.

Turning to the decomposition of $FTP^{9,1}$ in Figure 8b, at this maturity point over 95% of the variation is explained by variation in $S\mathcal{P}_t^9$ (less than 5% is explained by OM_t^9). This does not mean, of course, that unspanned macro risks can be reliably omitted from $DTSM$s for studying long-dated term premiums. Accurate measurement of these premiums will likely depend on accurate measurement of premiums along the entire maturity spectrum, and we have just seen that unspanned macro risks are important in explaining the variation of shorter-dated premiums.

Focusing more specifically on the decline in long-dated FTPs during the period of the conundrum, the fall in $FTP^{9,1}$ is accounted for (almost) entirely by the information embodied in $\mathcal{P}$. Together, Figures 7b and 8b suggest that the conundrum is not easily explained by economic weakness as captured by PMI and $GCEI$. Similarly, when we project $FTP_t^{9,1}$ onto (GIP_t, INF_t) during our sample period the resulting R^2 is only 37%. In fact, roughly 24% of the variation in $FTP_t^{9,1}$ is due to factors that are orthogonal to the expanded information set $(SGIP_t, SINF_t, OGIP_t, OINF_t)$.

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."[40] Figure 8b shows that the decline in distant-horizon term premiums was not explained by changing unspanned macro risks. In the next section we take up the question of how much of this drop might have been explained by changes in spanned macro risks.

Yet a different perspective on the contributions of macro information to risk premiums comes from inspection of the expected excess returns on the portfolios of bonds with payoffs that are perfectly correlated with movements in the spanned macro risks $SGIP$ and $SINF$, $xSGIP_t$ and $xSINF_t$.[41] We see from Figure 9 that $xSINF$ achieved its lowest levels (largest absolute values) during the 2001 recession and again around the time of Chairman Bernanke's expressions of concerns about deflation. In this respect there is a parallel with the excess returns to the level-mimicking portfolio

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

[41] 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 $SINF$.

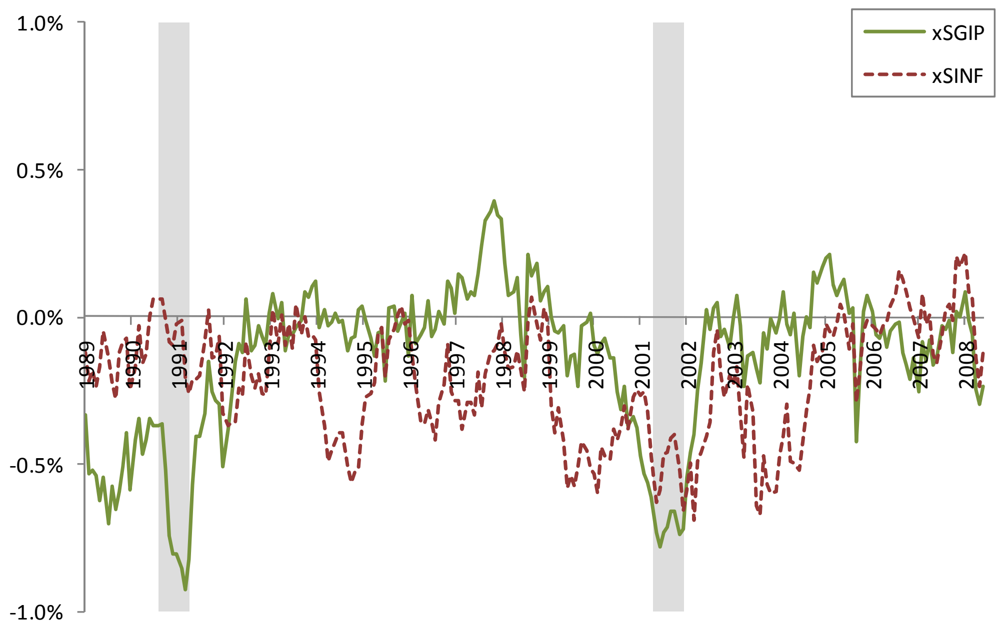

Figure 9: The monthly expected excess returns on bond portfolios with payoffs that are perfectly correlated with the components of GIP ($SGIP$) and INF ($SINF$) that are spanned by $\mathcal{P}_t$.

in Figure 4a. $xGIP$ is near zero for most of the sample period, with the exceptions of the recessions in 1990 and 2001.

Perhaps the most important take-away from Figure 9 is that both of these (annualized) premiums on spanned macro risks are small in absolute magnitude. It is (model-specific) variants of these risk premiums that most of the prior literature has been measuring as "inflation" or "output growth" risk premiums. Unless securities with macro-specific payoffs are included in pricing (e.g., TIPS for inflation risk), $DTSM$s only reveal information about the risks of the spanned components of macro risks. For our sample period and in swap markets, the premiums on spanned macro risks are small, notably smaller than the (annualized) excess returns on level and slope risks (Figure 4).

This finding leaves open the possibility that unspanned macro risks are priced and, hence, that *actual* market premiums on inflation and output growth risks are large. 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

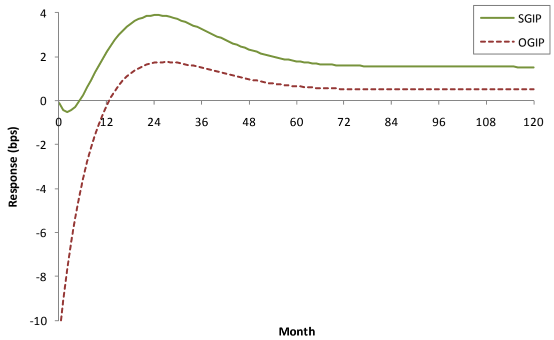

(a) $GIP \rightarrow FTP^{2,1}$

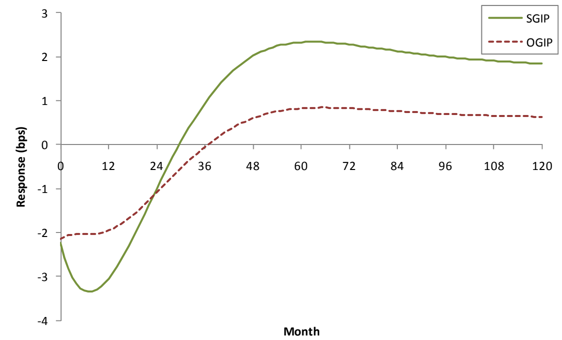

(b) $GIP \rightarrow FTP^{9,1}$

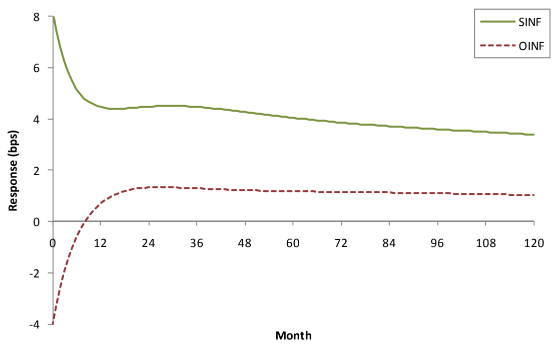

(c) $INF \rightarrow FTP^{2,1}$

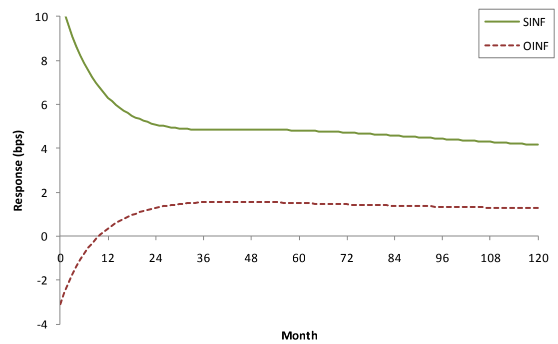

(d) $INF \rightarrow FTP^{9,1}$

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

TIPS data for the period 2000 through 2008. Consistent with our analysis, they find that this premium was negative during the first half of their sample, 2000-2004. Moreover, just as in Figure 9, their premium increased substantially around 2004 and turned positive. Without TIPS data we find that the spanned premium was just slightly positive late in our sample; using TIPS, Grishchenko and Huang (2010) find that this premium was significantly positive, particularly at the long end of the yield curve. This suggests that towards the end of our sample the premium on unspanned inflation implicit in TIPS was large relative to the premium on spanned inflation. Their analysis was outside of an arbitrage-free $DTSM$. We defer to future research an exploration of risk premiums on unspanned inflation within an extended version of our modeling framework.

5.1 Responses of Term Premiums to Innovations in M_t

In the preceding section we found that $(OGIP, OINF)$ accounted for a substantial fraction in the variation of $FTP^{2,1}$. We 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 GIP and INF. At the two-year horizon, innovations in $OGIP$ and $OINF$ have much larger effects on $FTP^{2,1}$ than innovations in $SGIP$ and $SINF$ (Figure 10a). The effects of unspanned output shocks dissipate quickly (within about one year), while the effects of spanned inflation shocks persist for several years (Figure 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 $OGIP$ and $OINF$ are small, consistent with the decomposition results in Figure 8. Figure 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. This finding is not easily reconciled, it seems, with Chairman Bernanke's explanation of the conundrum as a decline in compensation for bearing real interest rate risk. More likely, it seems, is that 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 (GIP, INF). Recall that nearly a quarter of the variation in $FTP^{9,1}$ during our sample period was attributable to such orthogonal factors.

5.2 Responses of Macro Variables to Term-Premiums Shocks

The existing theoretical and empirical literature has not reached clear-cut conclusions on the relationship between term premiums and 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 output growth and inflation to innovations in $FTP^{9,1}$ in the context of model $\mathcal{CM}_E^R$, using the model-implied VAR with ordering $(SGIP, SINF, OGIP, OINF, FTP^{9,1})$. As Figure 11 shows, a one standard deviation increase in $FTP^{9,1}$ is followed by a decline in $OGIP$ over a period of about 12 months, and has virtually no effect on spanned $SGIP$. The latter result is consistent with

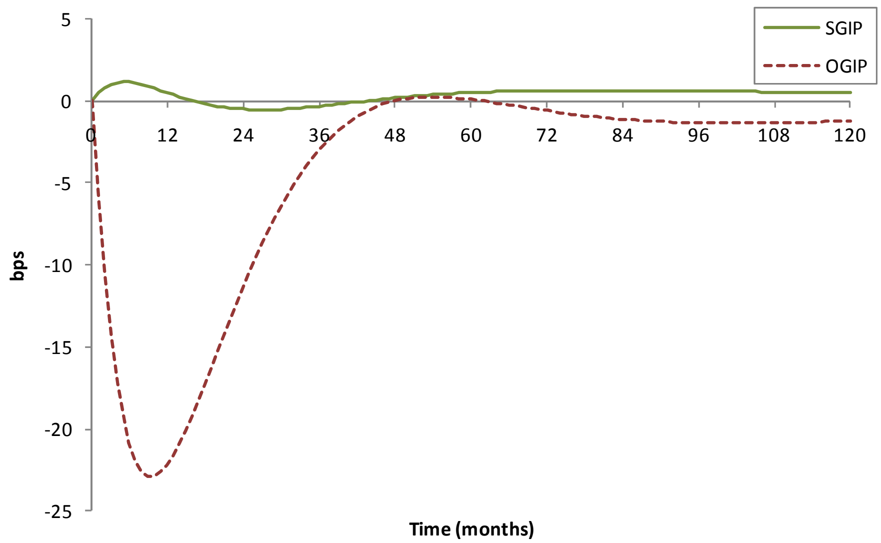

Figure 11: The impulse responses of $SGIP$ and $OGIP$ to a one standard deviation shock to $FTP^{9,1}$ implied by model $\mathcal{CM}_E^R$.

the results in Ang, Piazzesi, and Wei (2003) 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. The large impact on $OGIP$ of shocks to forward term premiums is a novel finding, one that we can identify by explicitly accommodating unspanned macro risks in model $\mathcal{CM}_E^R$.

The initial negative response of $OGIP$ is reversed after about 36 months, with the long-term impact being close to zero. This finding raises several interesting questions for future research, including: Is the mechanism underlying the response pattern of $OGIP$ the one that Bernanke articulated? Why do term-premium shocks affect the component of output growth that is orthogonal to bond yields? Do bond market risk premiums affect the compensation for bearing real output growth risk? Answering the latter question in particular will require a more elaborate pricing model and the inclusion of a richer set of financial instruments in the empirical analysis.

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.

Our modeling framework, formally developed in Appendix A, is applicable to 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, our framework is well suited to addressing a wide variety of economic questions about characteristics of risk premium in financial markets, including bond and currency markets, as well as equity markets when the latter pricing problems maps into an affine pricing model (e.g., Bansal, Kiku, and Yaron (2009)). The robust means by which we are able to restrict the dimensionality of expected excess returns might be particularly advantageous in multi-market settings, since such restrictions implicitly lead to reductions on the dimensionality of the parameter space. 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 Derivation of Results in Section 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 Y_t such that $r_t = r_\infty^{\mathbb{Q}} + 1 \cdot Y_t$ and

$$\Delta Y_t = \mathrm{diag}(\lambda^{\mathbb{Q}}) Y_{t-1} + \sqrt{\Sigma_Y^0}\epsilon_t^{\mathbb{Q}}$$

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

To derive a canonical version of (2), let $B^0(\tau)$ be the loadings given by $\dot{B}^0 = \mathrm{diag}(\lambda^{\mathbb{Q}})B^0 - 1$, $B^0(0) = 0$. 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 $B\Sigma_Y^0 B^\top$. In order that (2) is satisfied, it must be that $\Sigma_Y^0 = (B^\top)^{-1}\Sigma_{\mathcal{P}} B^{-1}$.

Now let $A^0(t)$ solve $\dot{A}^0 = \frac{1}{2}(B^0)^\top \Sigma_Y^0 B^0 - r_\infty^{\mathbb{Q}}$, $A^0(0) = 0$. 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: $K_X^{\mathbb{Q}} = B(\mathrm{diag}(\lambda^{\mathbb{Q}})B^{-1}$, $K_0^{\mathbb{Q}} = -(K_1^{\mathbb{Q}})^{-1}a$, $\rho_0 = r_\infty^{\mathbb{Q}} - 1^\top B^{-1}a$, and $\rho_1 = (B^\top)^{-1}1$.

Since (2) is an invariant transformation of an identified, canonical model, we know that (2) is also identified and canonical. The underlying parameters are $(r_\infty^{\mathbb{Q}}, \lambda^{\mathbb{Q}}, L_X)$, where L is the Cholesky factorization of Σ_X.

B Order Selection of Autoregressive Models

We compute three well-known information criteria based on unrestricted VARs with one through twelve lags (the maximum lag length was chosen so that both seasonal and annual effects would be captured): Akaike's information criterion (AIC), Hannan and Quinn's information criterion (HQIC), and Schwarz's Bayesian information criterion (SBIC). When additional lags are included as explanatory variables, the in-sample fit improves; the information criteria trade off this gain in likelihood against the additional number of parameters introduced. The recommended lag length is that at which the information criteria attain their minimum values. As Table 6 shows, all three criteria are minimized at a lag length of 1, suggesting that a first-order Markov structure fits our data well.

Lags	AIC	HQIC	SBIC
0	-41.6785	-41.6475	-41.6019
1	-54.1559^{*}	-53.9703^{*}	-53.6961^{*}
2	-54.1031	-53.7628	-53.2601
3	-54.0842	-53.5892	-52.858
6	-53.8453	-52.8861	-51.4695
9	-53.6811	-52.2579	-50.1558
12	-53.3595	-51.4721	-48.6846

Table 6: Assessments of Lag Length in VAR Models. The criteria-impled optimal lag lengths are indicated by '$*$'.

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 the portfolio $P_t^w = \sum_{i=1}^N w_i P_t^{n_i}$ 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^w_t:

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

This is a weighted average of the returns on the individual zero coupon bonds. Now, it follows from Le, Singleton, and Dai (2009) that $P^{n_i}_t = \exp(-A_{n_i} - B_{n_i}X_t)$, and further that

$$E^{\mathbb{P}}[P^{n_i-1}_{t+1}/P^{n_i}_t] = \exp\{B_{n_i-1}[(K^{\mathbb{Q}}_0 - K^{\mathbb{P}}_0) + (K^{\mathbb{Q}}_1 - K^{\mathbb{P}}_1)X_t] + r_t\}.$$

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

$$\frac{\sum_i a_i/n_i B_{n_i-1}[(K^{\mathbb{P}}_0 - K^{\mathbb{Q}}_0) + (K^{\mathbb{P}}_1 - K^{\mathbb{Q}}_1)X_t]}{|\sum_i a_i/n_i|}.$$

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

$$PCj_t = \sum_{i=1}^{N} \ell^j_i y^{n_i}_t = \sum_{i=1}^{N} \ell^j_i (A_{n_i}/n_i + B_{n_i}/n_i X_t)$$

it follows that $\sum_i \ell^j_i B_{n_i}/n_i$ 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

$$(K^{\mathbb{P}}_0 - K^{\mathbb{Q}}_0) + (K^{\mathbb{P}}_1 - K^{\mathbb{Q}}_1)X_t$$

scaled by $|\sum_i \ell^j_i/n_i|$. 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 (2003). What Does the Yield Curve Tell us About GDP Growth? Working Paper, Columbia University.

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.

Bansal, R. and I. Shaliastovich (2007). Risk and Return in Bond, Currency, and Equity Markets. Working Paper, Duke University.

Bansal, R. and A. Yaron (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance 59*, 1481–1509.

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 (2008). No-Arbitrage Macroeconomic Determinants of the Yield Curve. Working Paper, London Business School.

Breitung, J. (2002). Nonparametric tests for unit roots and cointegration. *Journal of Econometrics 108*, 343–363.

Campbell, J. and J. Cochrane (1999). By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy 107*(2), 205–251.

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.

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.

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. (1996). Idiosyncratic Variation in Treasury Bill Yields. *Journal of Finance 51*, 527–552.

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.

Estrella, A. and F. Mishkin (1998). Predicting U.S. Recessions: Financial Variables as Leading Indicators. *Review of Economic Studies 80*, 45–61.

Feldhutter, P. and D. Lando (2007). Decomposing Swap Spreads. Working Paper, Copenhagen Business School.

Giese, J. (2008). Level, Slope, Curvature: Characterizing the Yield Curve in a Cointegrated VAR Model. *Economics 2*, 1–21.

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.

Henry, O. T. and K. Shields (2004). Is there a unit root in inflation? *Journal of Macroeconomics 26*, 481–500.

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 (2009). No-Arbitrage Near-Cointegrated VAR(p) Term Structure Models, Term Premiums and GDP Growth. Working Paper, Banque de France.

Joslin, S. and K. Singleton (2010). What Gaussian Macro-*DTSM*s *Cannot* Tell us About the Macro Economy. Working Paper, Sloan School of Management.

Joslin, S., K. Singleton, and H. Zhu (2010). A New Perspective on Gaussian DTSMs. Working Paper, Stanford University.

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.

Krishnamurthy, A. and A. Vissing-Jorgensen (2008). The Aggregate Demand for Treasury Debt. Working Paper, Kellogg School, Northwestern University.

Le, A., K. Singleton, and Q. Dai (2009). Discrete-time Affine$^{\mathbb{Q}}$ Term Structure Models with Generalized Market Prices of Risk. Working Paper, Stanford University.

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

Lopez, C. and J. Reyes (2007). Real Interest Rate Stationarity and Per Capita Consumption Growth Rate. Working Paper, University of Cincinnati and University of Arkansas.

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. (2006). Affine Term Structure Models. Working Paper, Stanford University.

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. H. and M. W. Watson (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking 39*(s1), 3–33.

Swanson, E. (2006). Have Increases in Federal Reserve Transparency Improved Private Sector Interest Rate Forecasts? *Journal of Money, Credit, and Banking 38*, 791–819.

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

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

Wright, J. (2006). The Yield Curve and Predicting Recessions. Working Paper, Federal Reserve Board Working Paper 2006-07.

Wright, J. (2009). Term Premiums and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset. Working Paper, Johns Hopkins University.