



SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY



CREATES Research Paper 2010-14

An Asset Pricing Approach to Testing General Term Structure Models including Heath-Jarrow-Morton Specifications and Affine Subclasses

Bent Jesper Christensen and Michel van der Wel

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

An Asset Pricing Approach to Testing General Term Structure
Models including Heath-Jarrow-Morton Specifications and Affine
Subclasses*

Bent Jesper Christensen[†]
University of Aarhus and CREATES

Michel van der Wel
Erasmus University Rotterdam and CREATES

March 25, 2010

*We are grateful to Bob Jarrow for useful comments, and to the Center for Research in Econometric Analysis of Time Series (CREATES), funded by the Danish National Research Foundation, the Danish Center for Accounting and Finance (D-CAF), and the Danish Social Science Research Council for research support.

[†]Corresponding author: Bent Jesper Christensen, School of Economics and Management, University of Aarhus, 322 University Park, 8000 Aarhus C, Denmark. Phone: +45 8942 1547, Email: bjchristensen@creates.au.dk

An Asset Pricing Approach to Testing General Term Structure Models including Heath-Jarrow-Morton Specifications and Affine Subclasses

Abstract

We develop a new empirical approach to term structure analysis that allows testing for time-varying risk premia and for the absence of arbitrage opportunities based on the drift restriction within the Heath, Jarrow and Morton (1992) framework. As in the equity case, a zero intercept condition is tested, but in addition to the standard bilinear term in factor loadings and market prices of risk, the relevant mean restriction in the term structure case involves an additional nonlinear (quadratic) term in factor loadings. We estimate our general model using likelihood-based dynamic factor model techniques for a variety of volatility factors, and implement the relevant likelihood ratio tests. Our factor model estimates are similar across a general state space implementation and an alternative robust two-step principal components approach. The evidence favors time-varying market prices of risk. Most of the risk premium is associated with the slope factor, and individual risk prices depend on own past values, factor realizations, and past values of other risk prices, and are significantly related to the output gap, consumption, and the equity risk price. The absence of arbitrage opportunities is strongly rejected with one or two factors in the model, but not with three or more factors.

Keywords: arbitrage, bond aging effect, dynamic factor model, macroeconomic conditioning variables, nonlinear drift restriction, state space model, time-varying risk premia, yield curve model

JEL classification: G12, C32.

1 Introduction

The Treasury market is among the largest financial markets in the world, with \$7,888 billion worth of domestic debt securities outstanding at the end of 2008, compared to the total stock market capitalization of \$11,738 billion.¹ Treasuries are held by many different financial institutions, including commercial and investment banks, pension funds, insurance companies, and hedge funds, as well as private investors, for purposes of diversification, immunization, asset allocation, market timing, and risk management. The weight of government bonds in many portfolios has increased dramatically after the onslaught of the financial crisis. In the relatively liquid Treasury market, where individual issues of comparable maturity and contractual terms are fairly close substitutes, market conditions at any given point in time are effectively summarized by the yield curve. Indeed, bonds, notes, and bills are always quoted in terms of their corresponding yields.

In this paper, we use an asset pricing approach to test the relevant efficient market pricing conditions directly on the sequence of consecutive observed yield curves. As in the equity case, a zero intercept condition is tested, but in addition to the standard bilinear term in factor loadings and market prices of risk, the relevant mean restriction in the term structure case involves a nonlinear (quadratic) term in the loadings, and the test is applied to yield changes appropriately adjusted for both average slope (or yield spread) and local slope (or bond aging effect, cf. Litterman and Scheinkman (1991)) of the yield curve. With these extensions to the asset pricing approach, testing can proceed in a fashion parallel to multivariate testing of equity pricing models. Indeed, a tension between apparent non-zero intercepts (indicative of mispricing) and potentially omitted factors, similar to that common in the equity case, arises empirically in our term structure setting.

The idea that the yield curve captures current market conditions provides the foundation of the Ho and Lee (1986) and Heath, Jarrow and Morton (1992) (henceforth HJM) approach to general interest rate modeling, where the entire yield curve acts as the state variable in the dynamic term structure model. An assumption is made about the shape of the volatility functions governing the stochastic evolution through time of the curve. To preclude arbitrage opportunities, a restriction linking drifts and volatilities of yields through market prices of risk is required. At this point, the distribution of all future interest rates is specified, both under the physical and the risk neutral measure, hence facilitating derivative pricing. Calibration to the entire current term structure is feasible, since the shape of the current yield curve is not restricted by parameters and state variables, and this makes the approach overwhelmingly popular among practitioners, for pricing, trading, and hedging bonds and interest rate sensitive claims. The idea is to condition current prices on all the information

¹Source: Bank for International Settlements (BIS) Quarterly Review, June 2009, for the bond market figure, and the World Federation of Exchanges, the sum of the NYSE Euronext (U.S.), NASDAQ OMX, and American Stock Exchange total capitalizations for the stock market.

in the current yield curve, but reduce the dependence on potentially obsolete parameter estimates based on past observations. In spite of this, the framework has rarely been analyzed econometrically. In particular, it has not been tested whether risk premia are time-varying in the HJM framework, and whether the no-arbitrage drift condition is satisfied in practice in the market place.

We build an econometrically tractable dynamic term structure model consistent with the HJM framework, and we also consider the potential reduction to well-known affine subclasses. The general model allows for an arbitrary number of latent factors, and a potentially time-varying risk premium is associated with each of these. The model is specifically formulated at the level of yields to maturity, since this is how market prices are quoted, and because factor loadings are more reasonably taken to be similar for consecutive (changes in) yields of equal maturity than for a time series of returns to a bond that by definition becomes shorter as maturity is approached (unlike a stock). We develop an asset pricing approach to the empirical analysis of yields as opposed to returns, using likelihood-based methods and a state space implementation. We specify latent state variables such that market prices of risk are given as conditional expectations of these. In the special case of serially uncorrelated state variables, the state space model reduces to the classical factor analysis applied by Roll and Ross (1980) to stock market data. We use the Kalman filter to handle the generalization to a hidden Markov process for the state variables. To verify robustness to departures from distributional assumptions, we compare with results from an alternative two-step principal component-based approach. The analysis allows studying the appropriate number of latent factors to be included in order to explain term structure movements in this framework, as well as testing whether risk premia are in fact time-varying, and whether the no-arbitrage drift restriction is satisfied in the data. Here, we consider the unsmoothed Fama-Bliss fixed maturity panel of monthly U.S. zero-coupon Treasury yields of maturities ranging from three months through ten years, over the period 1985 through 2000.

Allowing for time-varying risk premia is important for a number of reasons. In particular, this is so even if the main interest is in derivative pricing and hedging, and hence in the risk neutral distribution, although this is obtained by setting all risk premia equal to zero. On the one hand, derivatives may be priced based only on the current yield curve and the volatility functions governing future changes in yields. Under the no-arbitrage restriction, the drifts of future yields of all maturities are known functions of the volatilities. In this sense, the interest in whether risk premia are time-varying under the physical measure is perhaps less obvious, and this may explain why the hypothesis has not been tested before in the literature. On the other hand, the volatility functions are required in any case for the pricing of bonds as well as derivatives. A natural way to obtain the volatility functions is to estimate them econometrically. This could be done using data on yields, derivatives, or both. Estimation based on yields over time necessarily uses the model under the physical measure,

and so risk premia enter. Thus, to bring yield data to bear on volatility estimation, risk premia must be accommodated in the empirical approach. Simply setting premia equal to zero would lead to inconsistent estimates of the volatility functions in yield data. Similarly, incorrectly restricting risk premia to be constant through time would lead to inconsistent volatility estimates, and hence the importance of testing for whether risk premia are in fact time-varying. Of course, if all volatility function parameters are backed out as implicit parameters from current data on derivative prices, only, then potentially time-varying risk premia may be ignored entirely, see, e.g., Flesaker (1993), and Amin and Morton (1994). However, the precision of volatility estimates is generally enhanced greatly by including data under the physical measure, in particular, yields to maturity, and thus the econometric approach of the present paper.

Allowing for time-varying risk premia is important when testing for the absence of arbitrage opportunities, as well. For one, as already argued, volatility estimates would be inconsistent if risk premia were inappropriately constrained to be constant, and as the volatilities enter into the drift specification under the no-arbitrage condition, the test of this would be misspecified. Secondly, even if consistent volatility estimates were used, the test on the drift restriction would be misspecified if the wrong (constant in time) form of market prices of risk were entered into the condition tested. An additional issue relates to the interpretation of results of tests. Thus, rejection of the arbitrage restriction may arise spuriously if the maintained model is misspecified. In the present application, this could occur, e.g., if the number of latent factors included were insufficient, or if market prices of risk were incorrectly taken to be constant. This reinforces the importance of our general approach allowing for an arbitrary number of factors, and testing for time-varying risk premia and no arbitrage for each given number of factors. Indeed, what we are dealing with is nothing but the usual joint hypothesis problem of market efficiency tests, well-known from the asset pricing literature using stock market data. In particular, any rejection of the arbitrage restriction could either be interpreted as an indication of the presence of arbitrage opportunities, or, if no arbitrage is a maintained hypothesis throughout, as suggesting that more factors be added, or that risk premia be allowed to vary across time, if they are not already. Either alternative would render the model under the null inappropriate for pricing purposes, and hence the critical importance of having a procedure for testing the no-arbitrage condition.

The HJM framework for term structure modeling is general and so includes, e.g., the popular affine models and all standard short rate models. In the short rate approach, a stochastic process assumption is adopted for the short rate of interest. When combined with a rule for the measure change from the physical to the risk neutral, e.g., based on a risk premium specification from equilibrium theory, then this delivers a model for the shape and stochastic evolution through time of the entire term structure of interest rates of all maturities. The literature on short rate models is huge, and a complete survey well

beyond the scope of this paper, but some of the seminal contributions in the area are the Gaussian model of Vasicek (1977), the general equilibrium model of Cox, Ingersoll and Ross (henceforth CIR, 1985a), the resulting square root model of CIR (1985b), the general affine model framework of Duffie and Kan (1996), along with the empirical analysis of this in Dai and Singleton (2000), and the essentially affine models of Duffee (2002). In all these models, the issue of whether or not risk premia are time-varying is a relevant one, hence providing a link to our research. In most cases, risk premia are set as given parametrized (e.g., affine) functions of state variables, see, e.g., CIR (1985b), Dai and Singleton (2002), Duffee (2002), and Duarte (2004). Stanton (1997) estimates this functional relation nonparametrically. Another empirical literature explores the expectations hypothesis and extensions, see, e.g., Campbell and Shiller (1991). An early contribution on time-varying risk premia in this area is Fama (1984), and Bams and Wolff (2003) is a recent study using a state space framework. In the present paper, we allow for time-varying risk premia that are adapted to the driving Wiener processes in the HJM framework and are dependent through time according to a Markov transition scheme.

In addition to time-varying risk premia, many short rate models allow for time-varying volatility, starting with CIR (1985b). In the present paper, we focus instead on unrestricted shape in the maturity direction of the volatility function, while keeping it constant through time, and dynamics enter through the risk premium transition equation. More general specifications could in principle be considered, since each short rate model can be extended within the HJM framework by letting the volatilities of yields of all maturities coincide with those in the short rate model in question and leaving the initial term structure shape unrestricted. This prescription highlights the generality of the HJM approach, although restrictions on yield curve shapes are implied also in this framework, once a process specification for risk premia is adopted.

Papers closely related to ours are few. Bliss and Ritchken (1996) and de Jong and Santa Clara (1999) consider a specific HJM-style model in a dynamic factor model setting, but they impose the no-arbitrage drift restriction throughout, rather than testing it, which is part of our main focus. Their model involves only a single market price of risk, and de Jong and Santa Clara (1999) restrict this to be proportional to the volatility of the short rate, as in CIR (1985b), whereas Bliss and Ritchken (1996) do not estimate it at all. In contrast, we work with a vector of potentially time-varying risk premia, and as a second part of our focus test for the special case that premia are constant. Recently, Jeffrey et al. (2004) consider nonparametric estimation of the volatility function, but also do not test the no-arbitrage drift condition, or estimate any premia.

Our empirical results show that time-varying risk premia are preferred over constant premia. We find that at least three or four factors are necessary to explain term structure movements, consistent with the bulk of the empirical literature, dating back to Litterman

and Scheinkman (1991). Our results regarding the drift condition are broadly consistent with the absence of arbitrage opportunities. When only one or two latent factors are included in our model, the arbitrage restriction is strongly rejected at all conventional levels. When three or four latent factors are included, there is much weaker evidence of arbitrage opportunities. The difference in results is of interest in its own right. In particular, researchers working with one or two factors, only, may find apparent signs of arbitrage opportunities in the market place, whereas the proper understanding of their results is actually that the model is misspecified, with too few factors. This is reminiscent of the tension between non-zero Jensen's alphas and omitted factors in the asset pricing literature using stock market data. On the other hand, even if no arbitrage is maintained throughout, our results do not yield any evidence that more than four factors (or other generalizations) are required. While our implementation uses a state space formulation and (Gaussian) distributional assumptions, our dynamic factor model estimates are confirmed using an alternative robust two-step principal components-based approach. We document that even with three or four factors our tests have power, and so in particular the failure to get strong rejection of no arbitrage is meaningful. Our estimates show that most of the risk premium is associated with the slope factor, and individual time-varying market prices of risk depend on own past values, factor realizations, and past values of risk prices associated with other factors. In addition, market prices of risk are significantly related to a number of relevant macroeconomic conditioning variables, in particular the output gap suggested by Cooper and Priestley (2009), consumption, and the equity risk price or Sharpe ratio.

Our work establishes the empirical importance of the nonlinear (quadratic) component of the no arbitrage condition and of suitable slope-adjustment of the yield change data. In particular, we show that if the average and local slopes (yield spread and bond aging, respectively) are not correctly accounted for, or if the nonlinear component is left out of the arbitrage condition, so that the test coincides with that for Ross' (1976) arbitrage pricing theory (APT) applicable to returns even in the present term structure case, then this changes the qualitative conclusions regarding arbitrage opportunities compared to those from our test appropriately accounting for the nonlinear (in loadings) drift term and applied to slope-adjusted yield changes.

The rest of the paper is laid out as follows. In Section 2, we present the model and the main hypotheses. In Section 3, we present in turn the data, empirical methodology, estimation results, and hypothesis tests. In a separate analysis in Section 4, we consider the sensitivity of our general approach to a number of relevant issues arising in practice, such as possible errors in model specification, data, distributional assumptions, etc. We investigate the empirical validity of standard parametrized volatility structures from the literature, including affine subclasses, and results using the alternative robust two-step principal components-based approach are presented and compared to those from the previous

section. Section 5 concludes.

2 The Model

The analysis is set in the Heath, Jarrow and Morton (1992) (henceforth HJM) framework for yields as opposed to forward rates and using the Brace and Musiela (1994) parametrization where term to maturity τ rather than maturity date enters as a separate argument in the continuously compounded zero coupon yield to maturity $y(t, \tau)$ at time t . The yield curve dynamics are given by the infinite dimensional stochastic differential equation

$$dy(t, \tau) = \alpha(t, \tau) dt + \sigma(t, \tau)' dW_t, \quad (1)$$

with drift α and yield volatility function σ . Writing d for the dimension of the driving Wiener process W_t , the dimension of $\sigma(t, \tau)$ is $d \times 1$. Thus, the mapping $\tau \mapsto y(t, \cdot)$ gives the yield curve at date t , and (1) shows how this evolves through calendar time. The no-arbitrage drift condition in this setting is

$$\alpha(t, \tau) = \frac{1}{\tau} (y(t, \tau) - y(t, 0)) + \frac{\partial y}{\partial \tau}(t, \tau) + \sigma(t, \tau)' \lambda_t + \frac{\tau}{2} \sigma(t, \tau)' \sigma(t, \tau), \quad (2)$$

where λ_t is the d -vector of market prices of risk. As HJM derived the relevant drift condition in an alternative parametrization with maturity date $t + \tau$ instead of term to maturity τ as a separate argument, and for instantaneous forward rates instead of the yield curve, we present a brief derivation of (2) in Appendix A. The first two terms in (2) are not present in the HJM version of the drift condition. The first term, $(y(t, \tau) - y(t, 0)) / \tau$, is an average slope or yield spread, and appears because we consider yields rather than forward rates. The second term, $\partial y(t, \tau) / \partial \tau$, is a local slope, i.e., the yield curve differentiated in the maturity direction, and appears because we consider constant terms to maturity rather than constant maturity dates. Litterman and Scheinkman (1991) decomposed bond returns into a locally deterministic bond aging effect and the return on constant maturity zeroes, and the local slope reflects the former. The third term in (2) is the risk premium, given by the volatility functions multiplied by the relevant market prices of risk. For derivative pricing purposes, changing measure from the physical to the risk neutral is equivalent to setting λ_t equal to zero. The final term, linear in maturity and quadratic in volatility, replaces a term involving an integral in the HJM representation.

Consider panel data on yields $y_t = (y_{t, \tau_0}, \dots, y_{t, \tau_m})$, $t = 1, \dots, T$. Thus, there are $m + 1$ observed yields, corresponding to terms to maturity $\tau_0 < \dots < \tau_m$, and T is the number of time periods. By a discrete time (Euler) approximation to (1), allowing for idiosyncratic

(e.g., measurement) error $\varepsilon_{t,i}$ in the i 'th yield, and imposing the drift condition (2), we have

$$y_{t,\tau_i} - y_{t-1,\tau_i} = \frac{1}{\tau_i - \tau_0} (y_{t-1,\tau_i} - y_{t-1,\tau_0}) + \frac{y_{t-1,\tau_i} - y_{t-1,\tau_{i-1}}}{\tau_i - \tau_{i-1}} + b'_i \lambda_{t-1} + \frac{\tau_i}{2} b'_i b_i + b'_i w_t + \varepsilon_{t,i}, \quad (3)$$

$i = 1, \dots, m$, where $\sigma(t, \tau_i) = b_i$, i.e., we henceforth take the volatility function to be time-invariant. This is more relevant in the current parametrization with fixed term to maturity than in the HJM parametrization. The d -dimensional vector w_t of driving factors corresponds to increments in the Brownian motions W_t . Note that the shortest observed term to maturity τ_0 in practice is positive, so $y(t, 0)$ from (2) is replaced by y_{t,τ_0} , and one cross-section dimension is lost. Thus, the shortest maturity τ_0 is not considered on the left hand side, but on the right hand side it is used in calculation of the first term (the average slope, or yield spread) for all longer maturities, and also in the calculation of the second term (the local slope) for the second-shortest maturity τ_1 . Because y_{t,τ_0} enters these calculations, we measure it by the 3-month T-bill yield, which should be a more market based rate than, say, the 1-month yield, and a good short rate proxy, see Chapman, Long, and Pearson (1999).

It is natural to collect yield data on the left hand side by defining the slope-adjusted yield changes

$$\tilde{y}_{t,\tau_i} = y_{t,\tau_i} - y_{t-1,\tau_i} - \frac{1}{\tau_i - \tau_0} (y_{t-1,\tau_i} - y_{t-1,\tau_0}) - \frac{y_{t-1,\tau_i} - y_{t-1,\tau_{i-1}}}{\tau_i - \tau_{i-1}}, \quad (4)$$

i.e., the raw yield changes adjusted for both average and local slope, and form the m -vector $\tilde{y}_t = (\tilde{y}_{t,\tau_1}, \dots, \tilde{y}_{t,\tau_m})$. The resulting model has factor structure,

$$\tilde{y}_t = \mu_t + B w_t + \varepsilon_t, \quad (5)$$

where the i 'th row of the $m \times d$ matrix B is given by b'_i , and $\varepsilon_t = (\varepsilon_{t,1}, \dots, \varepsilon_{t,m})'$. Thus, B gives the loadings on the common covariance-generating factors w_t , and $\mu_t = (\mu_{1t}, \dots, \mu_{mt})'$ contains the conditional means of the slope-adjusted yield changes \tilde{y}_t , given information through $t - 1$. Without loss of generality, we specify $\text{var}(w_t) = I_d$, since any covariance terms may be absorbed in B . In addition, we assume that the idiosyncratic errors are contemporaneously uncorrelated, i.e., $\text{var}(\varepsilon_t) = \Psi$ is a diagonal matrix, and that w_t and ε_t are independent of each other and across time (note that $w_t = \int_{t-1}^t dW_s$). In the system (5), the no-arbitrage drift restriction (2) is recast as

$$\mu_{it} = b'_i \lambda_{t-1} + \frac{\tau_i}{2} b'_i b_i. \quad (6)$$

Note the similarity to the arbitrage pricing theory (APT) of Ross (1976). Following Roll and Ross (1980), a standard approach to testing the APT is to apply the classical factor analysis to excess stock returns, thus estimating the loadings, say B , then test the APT as a cross-

sectional restriction on the mean excess returns μ , in particular, $\mu = B\lambda$, for suitable market prices of risk λ . There are two important differences in our term structure case. First, the factor analysis structure (5) applies to the appropriately slope-adjusted yield changes defined in (4) above, not to raw yields, spreads, or yield changes. Secondly, the theory restriction tested is not that means be linear in loadings, as in the APT, viz. $\mu_{it} = b'_i\lambda_{t-1}$. Instead, the no-arbitrage term structure restrictions include the additional terms $\frac{\tau_i}{2}b'_ib_i$ in (6) above, which are nonlinear (indeed, quadratic) in loadings. In this context, the original APT test simply drops the second (quadratic) term.

The no-arbitrage condition should be tested in term structure analysis. Unfortunately, the empirical literature has invariably left out this step. If the condition is rejected, then the model is simply wrong, since it admits arbitrage. In particular, the volatility function (loadings) or the price of risk specification is too simplistic. Once a specification is determined where the no-arbitrage condition is not rejected, it is possible to use the framework to explore the dynamic properties of the vector of risk prices λ_{t-1} . To test the no-arbitrage restriction, a number of alternative approaches suggest themselves. Two-step procedures may be used, in the first step estimating μ by the sample means and B either by principal components or the classical factor analysis. In the second step, risk prices are estimated by cross-sectional regression based on (6). In a sense, the first step uses constant risk premia, but time-varying premia may be picked up in the second step by running the cross-sectional regressions period by period. We develop a full information maximum likelihood one-step procedure that instead allows introducing time-varying risk prices from the outset. In the special case of constant risk prices λ , the procedure is equivalent to expanding the classical factor analysis likelihood function by the mean specification (6) above. With time-varying risk prices the state space form is used, and the prediction error decomposition of the likelihood function is calculated based on the innovations from the Kalman filter. In both one- and two-step approaches, this description corresponds to the restricted model, and the unrestricted alternative specified for testing purposes adds maturity-specific intercepts α_i to the drift specification (6), producing

$$\mu_{it} = \alpha_i + b'_i\lambda_{t-1} + \frac{\tau_i}{2}b'_ib_i. \quad (7)$$

Here, the no-arbitrage null hypothesis is $H_0 : \alpha = 0$, where $\alpha = (\alpha_1, \dots, \alpha_m)'$. Our main focus is on testing this hypothesis, which can now proceed much as in multivariate testing of equity pricing models, as well as on determining the appropriate number of factors, d , and testing for time-varying market prices of risk λ_t .

3 Empirical Results

3.1 Data and Summary Statistics

The data set we use is the unsmoothed Fama-Bliss fixed maturity panel over the period 1985 through 2000. It is constructed by Diebold and Li (2006), who use end of month prices for U.S. Treasuries from the CRSP government files to obtain unsmoothed Fama and Bliss (1987) forward rates. With this procedure the authors obtain a fixed maturity dataset for a cross-section of 17 maturities: 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, and 120 months. We have for the cross-sectional dimension of the yields $m + 1 = 17$, and thus for the dimension of the vector of slope-adjusted yield changes we have $m = 16$.

Summary statistics appear in Table 1. The first column shows term to maturity. In the remainder of the table, the left portion is summary statistics for yields, and the right portion is for slope-adjusted yield changes. The yield portion of the table is similar to Diebold and Li (2006). Thus, on average, the term structure of interest rates is upward sloping, from 5.6% to 7.3% (column labelled *Mean*), while the term structure of volatilities tends to be downward sloping (column *Sd*). For each maturity, there is a considerable spread between the minimum and maximum observation, indicating a great deal of variation over the period. Finally, there is strong positive autocorrelation at the one and 12 month lags, and even at lag 30 for maturities of two years and above. Turning to the slope-adjusted yield changes in the right hand side of the table, the means are much smaller, and very close to zero. The term structure of volatilities is now essentially flat for maturities of 12 months and above, and serial correlation is much less than for yields and not systematically positive for lags 12 and 30. This suggests that the factor model approach is more appropriate for slope-adjusted yield changes.

Figure 1 provides a three-dimensional view of our data on raw yields and slope-adjusted yield changes. In both cases, the curves appear somewhat flat in the term to maturity direction, consistent with level shifts being of greater order of magnitude than slope and curvature shifts. The stronger autocorrelation in yields shows up as more systematic swings in the calendar date dimension for raw than for slope-adjusted yield changes. The vertical axes are different because of the smaller size of the slope-adjusted yield changes.

Table 1: Summary Statistics of Yields and Slope-Adjusted Yield Changes

The table reports summary statistics for the yields and slope-adjusted yield changes from the unsmoothed Fama-Bliss dataset. We look at the yield curve for 17 maturities over the period January 1985 up to December 2000. The slope-adjusted yield changes are obtained according to the equation

$$\tilde{y}_{t,\tau_i} = y_{t,\tau_i} - y_{t-1,\tau_i} - \frac{1}{\tau_i - \tau_0}(y_{t-1,\tau_i} - y_{t-1,\tau_0}) - \frac{y_{t-1,\tau_i} - y_{t,\tau_{i-1}}}{\tau_i - \tau_{i-1}},$$

for $i = 1, \dots, m$, and where y_{t,τ_i} is the i -th observed yield at time t , with corresponding time to maturity τ_i , with $\tau_0 = 3$, $\tau_1 = 6$ and $\tau_m = 120$. In the table we show mean, standard deviation (*Sd*), minimum (*Min*), maximum (*Max*) and three autocorrelation coefficients, 1 month ($\hat{\rho}(1)$), 1 year ($\hat{\rho}(12)$) and 30 months ($\hat{\rho}(30)$).

Summary Statistics														
τ_i	Yields							Slope-Adjusted Yield Changes						
	Mean	Sd	Min	Max	$\hat{\rho}(1)$	$\hat{\rho}(12)$	$\hat{\rho}(30)$	Mean	Sd	Min	Max	$\hat{\rho}(1)$	$\hat{\rho}(12)$	$\hat{\rho}(30)$
3	5.630	1.484	2.732	9.131	0.978	0.569	-0.079							
6	5.785	1.479	2.891	9.324	0.976	0.555	-0.042	-0.119	0.273	-1.209	0.561	0.132	0.047	0.050
9	5.907	1.488	2.984	9.343	0.973	0.545	-0.005	-0.105	0.283	-1.239	0.609	0.175	0.037	-0.026
12	6.067	1.497	3.107	9.683	0.969	0.539	0.021	-0.120	0.319	-1.452	0.723	0.109	0.058	-0.091
15	6.225	1.500	3.288	9.988	0.968	0.527	0.060	-0.123	0.314	-1.156	0.716	0.207	0.052	-0.072
18	6.308	1.492	3.482	10.188	0.965	0.513	0.089	-0.096	0.312	-1.123	0.870	0.234	0.053	-0.084
21	6.375	1.480	3.638	10.274	0.963	0.502	0.115	-0.088	0.315	-1.029	0.780	0.193	0.080	-0.079
24	6.401	1.460	3.777	10.413	0.960	0.481	0.133	-0.070	0.327	-1.141	0.948	0.200	0.042	-0.102
30	6.550	1.458	4.043	10.748	0.957	0.479	0.190	-0.086	0.329	-1.168	0.831	0.210	0.021	-0.090
36	6.644	1.435	4.204	10.787	0.956	0.471	0.226	-0.073	0.329	-1.086	0.824	0.206	0.031	-0.103
48	6.838	1.435	4.308	11.269	0.951	0.457	0.294	-0.072	0.337	-1.109	0.869	0.149	0.043	-0.086
60	6.928	1.426	4.347	11.313	0.951	0.464	0.336	-0.060	0.330	-1.098	0.741	0.147	0.009	-0.094
72	7.082	1.453	4.384	11.653	0.953	0.454	0.372	-0.064	0.322	-1.066	0.768	0.137	-0.008	-0.082
84	7.142	1.422	4.352	11.841	0.948	0.448	0.391	-0.055	0.323	-1.365	0.822	0.110	-0.020	-0.080
96	7.226	1.410	4.433	11.512	0.954	0.468	0.417	-0.054	0.310	-1.158	0.772	0.098	-0.030	-0.064
108	7.270	1.425	4.429	11.664	0.953	0.475	0.426	-0.050	0.308	-1.176	0.663	0.097	-0.036	-0.070
120	7.254	1.428	4.443	11.663	0.953	0.467	0.428	-0.043	0.313	-1.176	0.776	0.071	-0.013	-0.072

3.2 Estimation Methodology

The factor structure in (5) combined with the potential dynamics of the risk prices in (6) suggests a dynamic factor model approach for estimation and inference. To proceed further, the process for λ_t must be specified. Following HJM, we assume that the risk price process is adapted to the filtration generated by the driving Wiener processes. In particular, in the discrete time formulation of the econometric model, we specify that risk prices be affine in a suitable vector of latent state variables, say x_t , i.e.,

$$\lambda_t = a + Ax_t, \quad (8)$$

where a is $d \times 1$ and A is $d \times d$, and x_t is measurable with respect to the w_t -process. Thus, constant risk premia are the special case $A = 0$. Smooth time-variation in risk premia is generated by ensuring this property for the state variable sequence, and analysis under the risk neutral measure (for derivative pricing) is facilitated by setting both a and A equal to zero. In the absence of arbitrage opportunities, i.e., in the system (5)-(6), and aside from idiosyncratic noise ε_t , time variation in the slope-adjusted yield change \tilde{y}_{t,τ_i} is generated by $b'_i(\lambda_{t-1} + w_t)$, so it is natural to specify the latent state vector as

$$x_t = \lambda_{t-1} + w_t. \quad (9)$$

Thus, state variables are driven by the w_t process, and also reflect past risk prices. Inserting (8) into (9) produces

$$x_t = a + Ax_{t-1} + w_t, \quad (10)$$

showing that latent state variables are governed by a vector autoregressive (VAR) system forming a hidden Markov process underlying the observed yield data. In particular, the state vector reflects the full shock $w_t = \int_{t-1}^t dW_s$, and, in addition, depending on the structure of the transition matrix A , a portion of the past state. In the special case $A = I_d$, the state variables x_t are exactly tracking the driving Wiener processes $W_t = \sum_s w_s$. Presumably, more empirically relevant specifications involve less weight on past shocks, e.g., through eigenvalues below unity in A . Note that from (8) and (10), the market prices of risk are given by

$$\lambda_t = E_t(x_{t+1}), \quad (11)$$

i.e., the expected next period state variables, conditional on the sequence of states (or, equivalently, w_s) through t , thus allowing a natural economic interpretation of our latent x_t .

With these specifications, our model is

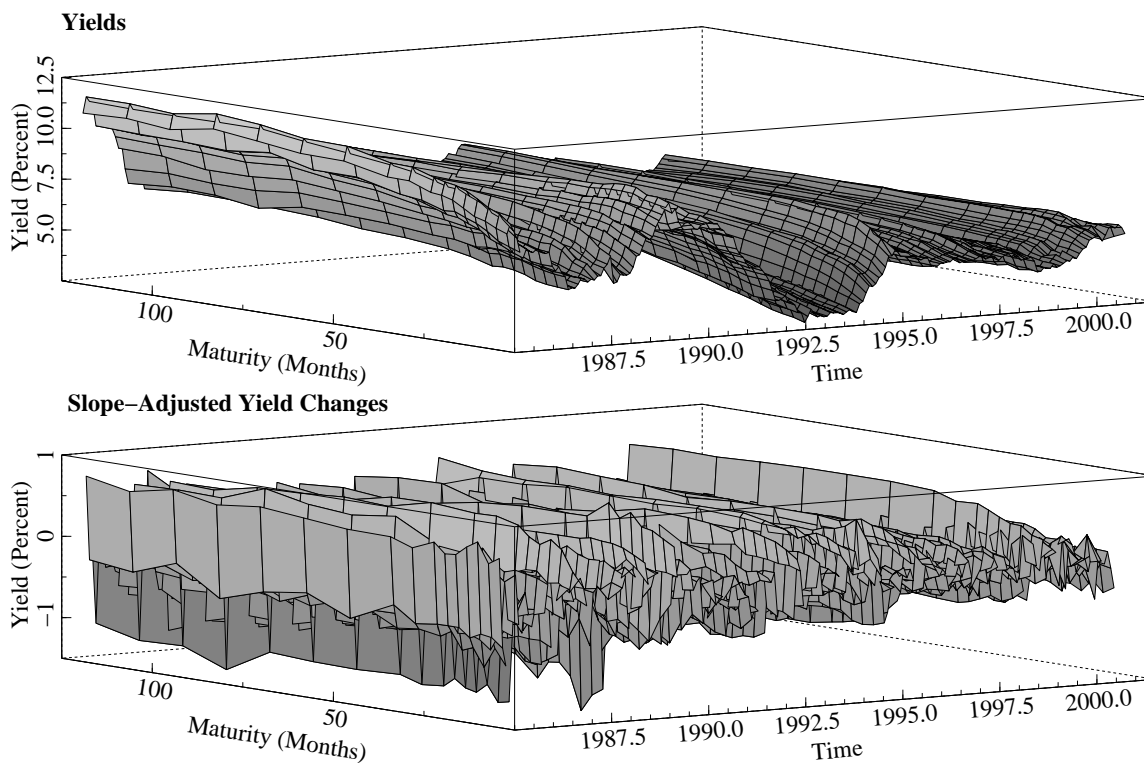
$$\begin{aligned} \tilde{y}_t &= \alpha + \text{vec}_{1:m} \{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t, \end{aligned} \quad (12)$$

Figure 1: Plot of Yields and Slope-Adjusted Yield Changes

In this figure we show a 3-dimensional plot of both the yields and slope-adjusted yield changes from the unsmoothed Fama-Bliss dataset. We look at the yield curve for 17 maturities over the period January 1985 up to December 2000. The slope-adjusted yield changes are obtained according to the equation

$$\tilde{y}_{t,\tau_i} = y_{t,\tau_i} - y_{t-1,\tau_i} - \frac{1}{\tau_i - \tau_0}(y_{t-1,\tau_i} - y_{t-1,\tau_0}) - \frac{y_{t-1,\tau_i} - y_{t,\tau_{i-1}}}{\tau_i - \tau_{i-1}},$$

where y_{t,τ_i} is the i -th observed yield at time t , with corresponding time to maturity τ_i .



with \tilde{y}_t the m -vector collecting all slope-adjusted yield changes \tilde{y}_{t,τ_i} at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, B the $m \times d$ loading matrix whose columns are the volatility functions corresponding to the d factors, b'_i the i 'th row of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ , and the elements of w_t standard normal. The upper equation in (12) corresponds to equations (5), (7), and (9). As discussed in Section 2, intercepts $\alpha = (\alpha_1, \dots, \alpha_m)'$ are included to measure the deviation from the arbitrage restriction (6). In the restricted model imposing no arbitrage, α is set equal to 0 in (12). By considering the state variable x_t rather than w_t and λ_{t-1} separately in the model, we are able to introduce time-variation in the risk prices λ_t without introducing additional noise terms. From (8), risk prices are adapted to the driving Wiener processes, as in HJM, with cross-effects (non-diagonal A) allowed. If λ_t had involved separate sources of uncertainty not present in w_t , then the physical and risk neutral ($\lambda_t = 0$) measures would not be equivalent.

The model is in state space form, and (12) gives the m -dimensional measurement equation and d -dimensional state transition equation. Assuming normal distributions for the state shocks w_t and noise term ε_t , the Kalman filter can be employed (see Harvey (1989) and Durbin and Koopman (2001); Appendix B provides a brief summary of the Kalman filter recursions) to filter the latent states x_t . The filter provides estimates of the unobserved x_t conditionally on data through $t - 1$, i.e., the ‘predicted’ state $E_{t-1}[x_t]$ (this and the following expectations using square brackets are conditional on data through $t - 1$, as opposed to the true driving w_s), and conditionally on the expanded information set including current (time t) data \tilde{y}_t , the ‘filtered’ state $E_t[x_t]$. This allows calculating the innovations (prediction errors) in the data sequence and the corresponding (prediction error decomposition of) the conditional likelihood function given the initial observation \tilde{y}_0 . We maximize this likelihood to obtain our parameter estimates. In addition, we use the filtered and predicted states to obtain estimates of our risk prices λ_{t-1} and covariance generating factors w_t . In particular, as $E_{t-1}[w_t] = 0$, an estimate of λ_{t-1} is provided by the predicted state, $\hat{\lambda}_{t-1} = E_{t-1}[x_t] = a + AE_{t-1}[x_{t-1}]$. The covariance generating factors are estimated off the expectations revisions from predicted to updated or filtered states, $\hat{w}_t = E_t[x_t] - E_{t-1}[x_t]$.

For a given dimension of covariance generating factors $d = 1, 2, 3, 4$ we estimate four variations of our model. The first distinction is between constant and time-varying risk prices. Constant risk prices correspond to the special case of no state transition, $A = 0$. In this case, the vector of market prices of risk is given by $\lambda = a$. Time-varying risk prices is the case where A is left free. A second distinction is between the presence or absence of arbitrage opportunities. Under the no-arbitrage restriction, $\alpha = 0$ is imposed on the top equation in (12). In the unrestricted case, with α free, a is set equal to 0 for identification, as it enters the mean of the slope-adjusted yield changes through x_t .

Finally, we discuss the restrictions that we need to impose to ensure identifiability for

all these models and the initialization of the Kalman filter. As Geweke and Zhou (1996) point out, a general identification issue in factor models is that the factors can be rotated (say, pre-multiplied by an orthogonal matrix P) and the factor loadings (the B -matrix in our case) adjusted for this (post-multiplied by P') to yield an observationally equivalent model. A way to overcome this is to impose $m \times (m - 1)/2$ restrictions on the matrix B such that this transformation is no longer possible. In practice we do this by restricting the portion above the diagonal in the upper square $m \times m$ part of B to consist of zeros,

$$B = \begin{pmatrix} \sigma_{1,1} & 0 & \dots & 0 \\ \sigma_{2,1} & \sigma_{2,2} & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \sigma_{m,1} & \sigma_{m,2} & \dots & \sigma_{m,d} \end{pmatrix}, \quad (13)$$

with $\sigma_{i,j}$ the (i, j) 'th element of B . Next, to initialize the Kalman filter, we specify the distribution of the first unobserved state. In general, one can write

$$x_1 \sim N(\mu_{x_1}, \Sigma_{x_1}). \quad (14)$$

Different types of initializations can be implemented by changing μ_{x_1} and Σ_{x_1} . When not much is known about the initial state, one can set its variance Σ_{x_1} very large, and take $\mu_{x_1} = 0$. As in our case the dynamics of x_t are in fact a VAR system, we can use the properties of its unconditional distribution to initialize the system (see Appendix B). In this case, μ_{x_1} and Σ_{x_1} depend on a and A , and reversing the sign of a factor (state variable) is not simply counteracted by a reversal in sign of the associated column of B , but also impacts both a and A and hence the initialization of the system.

3.3 Estimates

In Table 2 we report the estimated market prices of risk and likelihood values for the restricted and unrestricted models with constant and time-varying risk premia. This is done separately for $d = 1, 2, 3$, and 4 factors. In the first model, the single-factor case, the restricted version has a negative market price of risk, estimated at about -2.0 in the constant λ case. In the time-varying case, the mean value of λ_t given by $\mu = (I - A)^{-1}a$ is reported. This takes about the same value, and both are strongly significant (t -values in parentheses). The difference in degrees of freedom between the restricted and unrestricted models in the constant λ case is $48 - 33 = 15$, corresponding to $m = 16$ parameters saved in α by imposing the no-arbitrage restriction, whereas λ is a new parameter introduced in the restricted model. The difference is 15 in the time-varying λ_t case, too, since again 16 parameters α are saved and in this case the parameter a is introduced when passing to the restricted model,

whereas A is included both with and without the no-arbitrage condition.

In the two-factor model, $d = 2$, the first risk price is smaller in magnitude, about -1.3 , but the second risk price is an order of magnitude larger, and both are strongly significant. Of course, the total impact on the risk premium is the risk price times the relevant volatility or loading. A graphical depiction of the estimated loadings b_i for each factor as function of maturity is shown in Figure 2. From the second exhibit in Panel (A) ($d = 2$), the first factor is a level factor, and the second factor with numerically larger risk price (at -27 , see Table 2) is seen to be a slope factor. From the third exhibit ($d = 3$) of Figure 2, the third factor is a curvature factor, reaching a maximum or hump at 12 (maturity 6 years). Upon introducing this, the ordering of the second and third factor is reversed, as seen in the figure. The reversal is also reflected in Table 2, the case $d = 3$, i.e., it is still the slope factor that carries the largest risk price, now at -30 , whereas also the price on the curvature factor (at -4) exceeds that on the level factor in magnitude. Similarly, when a fourth factor is introduced, this is seen from Figure 2 to be an additional curvature factor, with hump at 5 (maturity 18 months). From Table 2, $d = 4$, this factor, too, gets a higher risk price than the level factor (15 in the constant and 6 in the time-varying risk price case). Still, the risk price on the slope factor remains the largest (for $d = 4$, the slope factor is the fourth factor, see Figure 2). Throughout, average risk prices in the restricted models are quite similar across constant and time-varying λ_t specifications.

In Table 3 we show the estimated risk price dynamics, as given by the matrix A . In the single-factor model, the autocorrelation coefficient of the risk price process is .20 in the unrestricted model and .23 with the no-arbitrage restriction imposed. The numbers are very similar for the first factor (the upper left corner of A) in the higher-dimensional models. When introducing the slope factor, it seems to get a lower coefficient, below .04 for $d = 2$ and 3 (recall from Figure 2 that for $d = 3$ the slope factor is the third factor), and also low (.02) in the restricted four factor model (where the slope factor is the fourth factor). There is an indication of negative dependence on own lag for the third factor (second curvature factor) in the four-factor model. Some of the off-diagonal elements in A are of the same order of magnitude as the diagonal elements, showing strong interdependence between the risk price dynamics processes. In the three and four factor models, the risk price on the last (slope) factor depends strongly (coefficient about .25 in magnitude and significant) on the lagged risk price on the first (level) factor. To interpret the results, note that a slope factor with loading increasing in maturity tends to get a negative risk price. An observationally equivalent model is obtained by reversing the sign of both the corresponding column in B and the relevant coordinate of λ . Accounting for the change of sign in both loadings and risk price for the slope factor between the three and four factor models, the results show that an increase in yields in one period reduces the slope of the yield curve in the next.

Table 2: Estimates of Time-Varying Risk Premia

In this table we report the estimated risk prices. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (denoted with *Constant* λ in the table, to this end we set $A = 0$) and time-varying λ_t (*TV* λ , $A \neq 0$). For both of these variations we estimate an unrestricted version (*Unrestr*, where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction (*Restr*, $\alpha = 0$). In the case with constant risk premia we report the estimate, in the case with time-varying risk premia we report the estimated mean (in both cases this is $\mu = (I - A)^{-1}a$ from the above equation). The t -statistics are given below the estimates, the asterisks (*/**) denote significance at the 5%/1% level. In addition we report the log likelihood, number of parameters and Akaike information criterion (*AIC*).

		Estimates of Time-Varying Risk Premia – $d = 1, 2$							
		Number of Risk Factors							
		$d = 1$				$d = 2$			
		Constant λ		TV λ		Constant λ		TV λ	
		Unrestr	Restr	Unrestr	Restr	Unrestr	Restr	Unrestr	Restr
Factor 1 (λ_{1t})	Est/ μ		-2.04**		-2.03**		-1.27**		-1.27**
			-27.4		-21.4		-15.2		-12.2
Factor 2 (λ_{2t})	Est/ μ						-27.1**		-27**
							-10		-9.99
Factor 3 (λ_{3t})	Est/ μ								
Factor 4 (λ_{4t})	Est/ μ								
Log likelihood		2684	1394	2688	1428	3575	3548	3582	3553
# Parameters		48	33	49	34	63	49	67	53
AIC		-5272	-2722	-5277	-2789	-7025	-6999	-7030	-7000

Table 2 Cont'd: Estimates of Time-Varying Risk Premia – $d = 3, 4$									
Number of Risk Factors									
$d = 3$									
$d = 4$									
		Constant λ		TV λ		Constant λ		TV λ	
		Unrestr	Restr	Unrestr	Restr	Unrestr	Restr	Unrestr	Restr
Factor 1 (λ_{1t})	Est/ μ		-1.23** -15.3		-1.23** -12.8		-1.23** -15.3		-1.23** -12.8
Factor 2 (λ_{2t})	Est/ μ		-4.11* -2.3		-4.04* -2.29		4.21* 2.24		-4.17* -2.5
Factor 3 (λ_{3t})	Est/ μ		-30** -9.59		-30.7** -9.48		15.3 1.06		6.21 1.25
Factor 4 (λ_{4t})	Est/ μ						-25.9** -2.83		29.4** 8.76
Log likelihood		3691	3680	3706	3694	3722	3705	3749	3737
# Parameters		77	64	86	73	90	78	106	94
AIC		-7227	-7231	-7240	-7242	-7263	-7253	-7286	-7285

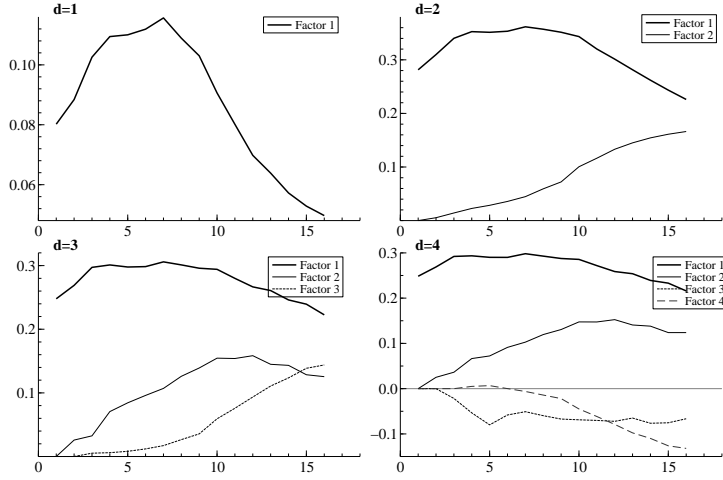
Figure 2: Volatility Functions and Covariance-Generating Factors

In this figure we show estimates of the volatility functions $b_i = \sigma(t, \tau_i)$ and a time series plot of the estimated covariance-generating factors. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t, τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). We show the volatility functions and covariance-generating factors for the model with time-varying risk premia and the no-arbitrage restriction imposed (thus $A \neq 0$ and $\alpha = 0$). The estimates in Panel (A) document the estimated B matrix for $d = 1, 2, 3, 4$. The time series plots in Panel (B) document the estimated w_t time series for each of the factors in the four factor case $d = 4$ imposing the no-arbitrage restriction and allowing for time-varying risk prices.

(A) Volatility Functions (B)



(B) Covariance-Generating Factors (w_t)




Table 3: Estimated Risk Price Dynamics

In this table we report the estimated risk price dynamics. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities τ_i , $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions with row element b'_i , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). We report the estimated A matrix for the time-varying risk price case (thus where $A \neq 0$). The t -statistics are given below the estimates, asterisks (*/**) denote significance at the 5%/1% level.

Panel A: VAR Estimates $d = 1$			Panel B: VAR Estimates $d = 2$				
	Unrestr	Restr	Unrestricted		Restricted		
	Factor 1	Factor 1	Factor 1	Factor 2	Factor 1	Factor 2	
Factor 1 ($\lambda_{1,t}$)	0.196** 2.72	0.226** 8.38	Factor 1 ($\lambda_{1,t}$)	0.242** 3.38	-0.0173 -0.233	0.239** 2.76	-0.0167 -0.254
			Factor 2 ($\lambda_{2,t}$)	-0.104 -1.39	0.0389 0.518	-0.0937 -1.02	0.0323 0.474

Panel C: VAR Estimates $d = 3$						
	Unrestricted			Restricted		
	Factor 1	Factor 2	Factor 3	Factor 1	Factor 2	Factor 3
Factor 1 ($\lambda_{1,t}$)	0.195** 2.62	-0.0251 -0.312	0.0184 0.24	0.205** 2.77	-0.0393 -0.476	0.0318 0.434
Factor 2 ($\lambda_{2,t}$)	0.0887 0.727	0.206* 2.32	-0.15 -1.4	0.0865 1.04	0.201* 2.25	-0.139 -1.72
Factor 3 ($\lambda_{3,t}$)	-0.265** -3.07	0.109 0.98	0.00698 0.0874	-0.253** -3.04	0.109 1.2	-0.0012 -0.0153

Panel D: VAR Estimates $d = 4$								
	Unrestricted				Restricted			
	Factor 1	Factor 2	Factor 3	Factor 4	Factor 1	Factor 2	Factor 3	Factor 4
Factor 1 ($\lambda_{1,t}$)	0.209** 2.75	-0.0211 -0.191	-0.107 -0.372	0.125 0.495	0.186* 2.49	-0.0649 -0.701	-0.142 -1.6	-0.0277 -0.341
Factor 2 ($\lambda_{2,t}$)	0.0932 0.805	0.19 1.77	-0.0775 -0.204	-0.143 -0.449	0.00557 0.0515	0.181 1.13	-0.318 -1.45	0.169 1.8
Factor 3 ($\lambda_{3,t}$)	-0.134 -0.284	0.184 0.615	-0.214 -0.199	0.463* 2.48	-0.143 -1.42	-0.0299 -0.145	-0.301 -1.93	-0.0124 -0.0917
Factor 4 ($\lambda_{4,t}$)	-0.21 -0.689	0.0227 0.0516	0.031 0.274	-0.209 -0.193	0.255** 2.81	-0.0647 -0.778	-0.0102 -0.0608	0.0202 0.252

Panel (B) of Figure 2 shows the fitted time series of covariance-generating factors from the restricted models imposing the no arbitrage condition and allowing for time-varying risk prices. Consistent with the model assumptions, the factors are all of the same order of magnitude, they appear roughly serially uncorrelated, and move around a zero level. The corresponding estimated time-varying market prices of risk appear in Figure 3. Evidently, these are more strongly serially dependent, especially the first (level) factor, consistent with the results from Table 3. Each risk price moves around the stable non-zero level μ reported in the table, and in particular is not trending. The levels and volatilities are different across the four risk prices, in contrast to the fitted factors from the previous figure.

Figure 4 shows in panel (A) the combined effect on slope-adjusted yield changes of both loadings and risk prices, i.e., this is the contribution to the total risk premium $B\lambda_t$ from each risk factor, by maturity. From the figure, by far the majority of the risk premium stems from the pricing of the slope factor, especially for longer maturities, and this is so for all models with two or more factors. Panel (B) of the figure shows the evolution over time of the total risk premium $b'_i\lambda_t$ for four selected maturities τ_i (6, 12, 60, and 120 months). Risk premia move in similar fashions across different maturities, and in each case the movement is around a stable level that differs by maturity due to the slope.

In Figure 5, Panel (A), we show the estimated intercepts α from the unrestricted models. They are strongly decreasing in maturity for all four dimensions of the vector of factors and both for the constant and time-varying risk price specifications. It is this pattern that must be captured by $B\lambda$ in the restricted models. This is exactly achieved by the estimated risk premia in the previous Figure 4, where the slope factor is particularly useful in picking up the required pattern to preclude arbitrage opportunities. Finally, Panel (B) of Figure 5 shows the estimated idiosyncratic variances in Ψ , by maturity. Errors are largest in the long end of the term structure in case of the single-factor model. With two or more factors, errors are smaller, and actually decreasing in maturity, except for the very longest. Note that the vertical axes are different, and that adding the third and fourth factors continues to reduce the magnitude of errors. This is consistent with the likelihood value and Akaike information criteria (AIC) of both unrestricted and restricted models in the previous table, that also pointed to the models with at least three or even four factors. The four factor model might be slightly over-parametrized, in that the idiosyncratic error variance is brought to zero for $i = 5$ (maturity 18 months) in Panel (B) of Figure 5, exactly corresponding to the hump in the factor loading on the second curvature factor in Panel (A) of Figure 2, i.e., one factor is fitted to this key yield.

3.4 Tests of the No-Arbitrage Restriction

We are now ready to look at the results from a test of the no-arbitrage condition $H_0 : \alpha = 0$ in (12).

Figure 3: Time Series Plot of Estimated Market Prices of Risk

In this figure we show a time series plot of the estimated risk prices. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). We show the estimated risk premia for the model with time-varying risk premia and the no-arbitrage restriction imposed (thus $A \neq 0$ and $\alpha = 0$). We show the estimated λ time series that is obtained for each of the four factors in the case $d = 4$, and provide 95% confidence boundaries.

Risk Price Time Series for $d = 4$ with Confidence Boundaries





Figure 4: Time Series Plot of Risk Prices Interacted with Loading Matrix

In this figure we show how the estimated risk prices and loadings interact to form risk premia for different maturities. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). We show the estimated risk prices for the model with time-varying risk premia and the no-arbitrage restriction imposed (thus $A \neq 0$ and $\alpha = 0$), interacted with the loading matrix B . In Panel (A) we document the estimated mean risk prices $\mu = (I - A)^{-1}a$ premultiplied by the factor loading matrix B that is obtained for the four different numbers of factors $d = 1, 2, 3, 4$. In Panel (B) we focus on the case of four factors, and show the estimated risk prices λ_t pre-multiplied with loading matrix B for four maturities (6 months, and 1, 5 and 10 years).

(A) Loading Matrix B times Mean Risk Price ($\mu = (I - A)^{-1}a$) for $d = 1, 2, 3, 4$



(B) Loading Matrix B times Risk Price λ_t for $d = 4$, Selected Maturities





Figure 5: Unrestricted Intercept and Measurement Variance

In this figure we show the estimates of the mean and measurement error variance of the HJM model with time-varying risk premia. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). In Panel (A) we show the estimates of α for both the constant and time-varying risk premia case when the no-arbitrage restriction is not imposed (thus when $\alpha \neq 0$). In Panel (B) we show the idiosyncratic variance Ψ for the model with time-varying λ_t and the no-arbitrage restriction imposed (thus $A \neq 0$ and $\alpha = 0$).

(A) Unrestricted Intercept (α)



(B) Idiosyncratic Variance (Ψ)




Table 4: LR-Tests of No-Arbitrage Restriction and Time-Varying Risk Premia

In this table we report results of tests of the no-arbitrage drift restriction and time-varying risk prices. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model with the no-arbitrage drift restriction ($\alpha = 0$). In the top part of the table (*Restr vs Unrestr*) we provide LR-tests of the no-arbitrage drift restriction (i.e. a test whether $\alpha = 0$) for both the case that risk premia are constant ($A = 0$) and time-varying ($A \neq 0$). In the bottom part (*Const vs TV λ*) we provide LR-tests of the restriction that λ is constant rather than time-varying (i.e. a test whether $A = 0$) for both the case without and with the no-arbitrage drift restriction imposed ($\alpha \neq 0$ and $\alpha = 0$ respectively). The p -value of rejecting the null hypothesis of the test is reported below the test statistic. Asterisks (*/**) denote significant rejection at the 5%/1% level.

LR-Tests of No-Arbitrage Restriction				
	Number of Risk Factors			
	$d = 1$	$d = 2$	$d = 3$	$d = 4$
Restr vs. Unrestr				
Constant λ	2580** 1.00	53.7** 1.00	22.2 0.947	34.1** 0.999
TV λ	2518** 1.00	58.4** 1.00	23.6* 0.965	24.8* 0.984
Const vs TV λ				
Unrestricted	7.25** 0.993	13.5** 0.991	30.2** 1.00	54.7** 1.00
Restricted	68.7** 1.00	8.78 0.933	28.7** 0.999	64.1** 1.00

In Table 4 we show the results for the one through four factor models. The top half of the table shows tests of the no-arbitrage condition separately for the constant and time-varying risk price specifications. Numbers under the test statistics are p -values for rejecting the no-arbitrage condition. For $d = 1$ and 2, the condition is strongly rejected. The test fails to reject the no-arbitrage condition at level 5% in the three-factor model with constant risk price. In the three-factor model with time-varying risk price, the test fails to reject no-arbitrage at 1%, although not at 5%, and this is so in the corresponding four-factor model with time-varying risk price, too (except here the test rejects in the constant λ case).

The bottom half of the table shows tests of the hypothesis of constant risk price ($A = 0$) against the time-varying case, separately for the unrestricted and the restricted (no-arbitrage) models. All tests except that in the restricted two-factor model strongly reject. Thus, the evidence favors the specification with time-varying risk prices.

The results are interesting, from a number of perspectives. First, on the methodological side, the outcome of the no-arbitrage test depends strongly on the number of factors in the estimated model. Thus, it is important to determine the number of factors correctly before proceeding to the no-arbitrage test. If the market is in fact driven by three common covariance-generating factors, the wrong conclusion about arbitrage is drawn in our data if too few factors are used in estimation. Secondly, on the substantive side, given the strong prior that three (or maybe even four) factors matter for interest rates, based on both the literature (e.g., Litterman and Scheinkman (1991)) and our results, and given that the evidence supports the time-varying over the constant risk-price specification (lower portion of table), it is interesting to see that the test fails to reject the absence of arbitrage opportunities at 1% once at least three factors are included in the time-varying risk price model. This softens the evidence against market efficiency compared to that apparently present when using too few factors. Note also that the failure to produce a strong rejection is not a general lack of power of our procedure. In particular, the tests do reject at 1% in several of the other cases ($d = 1$ and 2, and $d = 4$ with constant λ).

Overall, we conclude that market prices of risk are time-varying, and that there is no strong evidence of arbitrage opportunities. We now turn to an investigation of the underlying source of this time-variation in market prices of risk.


3.5 Macroeconomic Sources of Time-Varying Risk Prices

We relate our time-varying market prices of risk to a number of macroeconomic variables. There is recent evidence that macroeconomic variables explain bond market premia, see, e.g., Wachter (2006), Joslin, Priebsch and Singleton (2009), Ludvigson and Ng (2009), and the factors underlying the bond market, see, e.g., Ang and Piazzesi (2003), Diebold, Rudebusch and Aruoba (2006) and Duffee (2008). Stimulated by this evidence, we explore the potential macroeconomic underpinning of premia in our model by regressing our estimated market


Figure 6: Time Series Plot of Macroeconomic Variables

In this figure we show time series plots of the macroeconomic variables that we consider. We show Output Gap, Industrial Production, Nonfarm Payroll Employment, Consumption Expenditures and a Stock Market Risk Price. Industrial Production and Personal Consumption Expenditure data are collected from the Federal Reserve Bank of St. Louis Economic Database (FRED) and the Nonfarm Payroll Employment data is obtained from the Bureau of Labor Statistics. The output gap is obtained by detrending industrial production using a nonlinear trend. For Industrial Production, Employment and Consumption, we show monthly relative changes. We use the difference between the (annualized) monthly return of the S&P500 index and the 3m T-bill yield scaled by an estimate of the monthly S&P500 index volatility to get a stock market risk price.


(A) Output Gap




(B) Industrial Production




(C) Nonfarm Payroll Employment



(D) Consumption Expenditures on Nondurable Goods



(E) Market Price of Risk based on Stock Returns



prices of risk one by one on variables suggested in the literature. The most explanatory power is achieved using the variables exhibited in Figure 6 as regressors. Indeed, the three most important variables turn out to be the monthly relative change in industrial production, the monthly change in nonfarm payroll employment, and the stock market risk price or Sharpe ratio, defined as the monthly S&P 500 excess return above the 3-month Treasury yield, relative to the sum of daily absolute returns over the month. Also significant in some of our regressions are the two other variables shown, the output gap, defined as deviations of industrial production from a nonlinear trend, as in Cooper and Priestley (2009), and the monthly relative change in nondurable consumption expenditures. In other experiments, we also tried the CPI inflation rate, a set of macro factors defined as the eight leading principal components extracted from the 132 macroeconomic variables in the Stock and Watson (2002) data (appropriately transformed and detrended to stationarity, following the article), the credit spread (Baa less Aaa rates from Moody's) and term spread (10 years less 3-months Treasury yields), and, finally, the short rate (3-month yield) itself. None of these alternative variables yielded significance at the 1% level in any univariate or multivariate regressions of the type reported below, and two of the macro principal component factors (the second and fifth of the eight) were the only significant regressors at the 5% level in univariate regressions explaining our first, second, and fourth risk prices. Thus, we focus on the five most powerful explanatory variables in the following.

Results from regressing our market prices of risk on the five macroeconomic variables appear in Table 5. Each panel shows results for one of the risk prices from the estimated four factor model, starting with the first risk price in Panel A. The five univariate regressions are reported in the first five columns of each panel. The last column reports the multivariate regression of the relevant market price of risk on all five macroeconomic variables. From the results in Panel A, the market price of the level factor is significantly related to industrial production, employment, and the equity risk price at the 1% level, to consumption at the 5% level, and to the output gap at the 10% level (p -value of 6.5% in two-sided test). In the multivariate regression (last column in Panel A of Table 5), consumption and the equity risk price remain significant at the 5% and 1% levels, respectively, and the regression explains 17% of the variation in the market price of risk on the level factor, up from 9.0% in the highest of the univariate regressions, that on the equity risk price. From Panel (D) of Figure 6, there are apparent outliers in the first few months of 2000 in the nondurable consumption data series, and when excluding these periods from consideration, the t -value for consumption changes to 3.1 in the univariate regression, compared to 2.0 in the table, indicating that the relation to the risk premium is not spurious and purely driven by the outliers. We conclude that the pricing of level factor risk in the bond market is significantly related to macroeconomic conditions, in particular the equity risk price and consumption.

Panel B of Table 5 shows the similar results from regressing the risk price of the second

Table 5: Regression of Risk Prices on Macroeconomic Variables

In this table we regress the time-varying risk prices from our model on the macroeconomic variables that we consider. The risk price estimates used for constructing the correlation matrix correspond to the restricted four factor model with time-varying risk premia. The macroeconomic variables we consider are Output Gap, Industrial Production, Nonfarm Payroll Employment, Consumption Expenditures on Nondurable Goods and a Stock Market Risk Price. In each of the four panels we show these regressions for each of the four factors, omitting the first and the last observation. We consider six regressions for all panels, where we first regress the relevant risk price on the individual macro variables, and then jointly. Robust Newey and West (1987) t -statistics are given below the estimates, and the asterisks (*/**) denote significance at the 5%/1% level. In addition we report the adjusted R^2 and the number of observations.

Panel A: Risk Price 1 Regression						
Dependent Variable: Risk Price 1						
	(1)	(2)	(3)	(4)	(5)	(6)
Output Gap	0.0189 1.85					0.015 1.55
Ind Prod		0.103** 3.06				0.0397 1
Employment			0.0003** 2.61			0.000148 1.17
Nond Cons Exp				0.00427* 1.98		0.00457* 2.27
Risk Price					-0.075** -4.66	-0.0741** -4.57
Intercept	-1.23** -69.2	-1.26** -62.1	-1.28** -48.2	-1.24** -59.8	-1.21** -67.7	-1.26** -61.2
R^2	0.020	0.046	0.030	0.029	0.090	0.168
#Obs	190	190	190	190	190	190

Panel B: Risk Price 2 Regression						
Dependent Variable: Risk Price 2						
	(1)	(2)	(3)	(4)	(5)	(6)
Output Gap	0.0316* 2.43					0.026* 2.2
Ind Prod		0.128 1.96				0.0752 1.08
Employment			0.000474* 2.12			0.000285 1.25
Nond Cons Exp				0.00281 0.827		0.00198 0.583
Risk Price					-0.0248 -0.966	-0.0157 -0.631
Intercept	-4.17** -159	-4.2** -137	-4.26** -91.5	-4.18** -137	-4.16** -135	-4.25** -91.4
R^2	0.024	0.031	0.033	0.005	0.004	0.065
#Obs	190	190	190	190	190	190

Panel C: Risk Price 3 Regression						
Dependent Variable: Risk Price 3						
	(1)	(2)	(3)	(4)	(5)	(6)
Output Gap	0.00814 0.815					0.00889 0.984
Ind Prod		0.0364 0.705				0.0415 0.738
Employment			-0.0000370 -0.222			-0.000135 -0.807
Nond Cons Exp				0.00221 1.09		0.00212 0.925
Risk Price					-0.00691 -0.284	-0.00583 -0.236
Intercept	6.21** 376	6.2** 324	6.21** 203	6.2** 333	6.21** 339	6.22** 216
R^2	0.002	0.003	0.000	0.004	0.000	0.011
#Obs	190	190	190	190	190	190

Panel D: Risk Price 4 Regression						
Dependent Variable: Risk Price 4						
	(1)	(2)	(3)	(4)	(5)	(6)
Output Gap	0.0181 1.35					0.0132 1.04
Ind Prod		0.0926* 2.05				0.0246 0.539
Employment			0.000379** 2.66			0.000267 1.9
Nond Cons Exp				0.00346 1.68		0.00355 1.89
Risk Price					-0.06** -3.48	-0.0585** -3.64
Intercept	29.4** 1351	29.4** 1174	29.3** 1021	29.4** 1194	29.4** 1310	29.3** 1179
R^2	0.015	0.031	0.041	0.016	0.048	0.111
#Obs	190	190	190	190	190	190

factor, a curvature factor (see Figure 2, Panel (A), $d = 4$), on the macroeconomic variables. The output gap and employment are significant at 5% in univariate regressions, and the output gap remains significant in the multivariate regression. This explains 6.5% of the variation in the risk price, i.e., the level factor risk price is more strongly related to the macroeconomic factors than the curvature risk price, and different macroeconomic forces determine the two, with output gap in the price of curvature, and equity risk price and consumption in the price of level risk.

Panel C shows that the second curvature factor is unrelated to all of the macroeconomic variables. Finally, from Panel D, the important slope factor has an associated market price of risk that is significantly related especially to the equity risk price and employment, and also to industrial production. In the multivariate regression, the equity risk price remains significant at the 1% level, and the regression explains 11% of the variation in the price of slope risk.

Overall, we conclude that our estimated market prices of risk are related to relevant macroeconomic variables, including the output gap suggested by Cooper and Priestley (2009), along with the equity risk price or Sharpe ratio, and changes in consumption. In particular, the output gap explains the price of curvature risk, whereas the price of slope risk is explained by the equity risk price, and the latter along with consumption changes explains the price of level factor risk. The results suggest that our estimates represent true economic pricing mechanisms in the bond market.

4 Sensitivity Analysis

In this section, we investigate the sensitivity of our general approach to a number of relevant issues arising in practice, such as errors in model specification, data, distributional assumptions, etc.

4.1 Sensitivity to Model Misspecification: Missing Nonlinear Term

Our approach is to test for no arbitrage at the level of yields, which is how traded instruments are quoted. In this case the relevant condition is (6), and the test is on the conditional means of slope-adjusted yields. An alternative analysis of the returns of the heterogeneous bonds, bills, and notes that comprise the market would instead focus on the standard APT condition omitting the quadratic form in loadings in (6). As a first diagnostic we examine the behavior of the test based on the same data as in Section 3 when the quadratic form in factor loadings is incorrectly left out. In particular, in equation (12) the term $\text{vec}_{1:m}\{\hat{b}'_i\hat{b}_i\tau_i/2\}$ is omitted. The rest of the analysis is identical to that of Section 3.

Panel A of Table 6 is directly comparable to Table 4. In both cases, the tests reject the

Table 6: Sensitivity Analysis of LR-Tests for No-Arbitrage Restrictions

In this table we provide three sensitivity analyses of the LR-tests of the no-arbitrage restriction. In all cases we estimate variations of the following model:

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t.\end{aligned}$$

In Panel A we report our LR-tests for the case where we omit the nonlinear $\text{vec}_{1:m}\{b'_i b_i \tau_i / 2\}$ term, such that the test is the same as according to the Arbitrage Pricing Theory (APT). In Panel B we use raw yield changes Δy_t rather than slope-adjusted yield changes \tilde{y}_t . In Panel C we restrict B to be of the Nelson-Siegel functional form. In all cases, in the top part of the table (*Restr vs Unrestr*) we provide LR-tests of the no-arbitrage drift restriction (i.e. a test whether $\alpha = 0$) for both the case that risk premia are constant ($A = 0$) and time-varying ($A \neq 0$). In the bottom part (*Const vs TV λ*) we provide LR-tests of the restriction that λ is constant rather than time-varying (i.e. a test whether $A = 0$) for both the case without and with the no-arbitrage drift restriction imposed ($\alpha \neq 0$ and $\alpha = 0$ respectively). The p -value of rejecting the null hypothesis of the test is reported below the test statistic. Asterisks (*/**) denote significant rejection at the 5%/1% level.

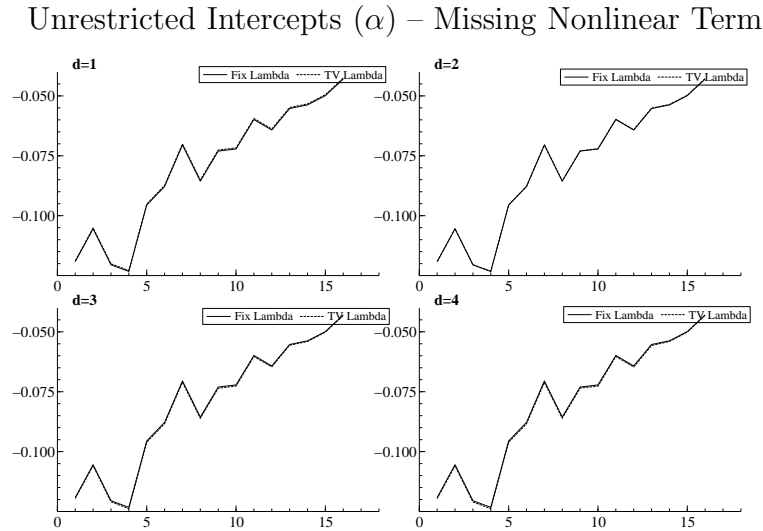
Panel A: LR-Tests – Missing Nonlinear Term				
	Number of Risk Factors			
	$d = 1$	$d = 2$	$d = 3$	$d = 4$
Restr vs. Unrestr				
Constant λ	204** 1.00	131** 1.00	64.5** 1.00	47.8** 1.00
TV λ	204** 1.00	131** 1.00	64.0** 1.00	46.1** 1.00
Const vs TV λ				
Unrestricted	7.25** 0.993	13.5** 0.991	30.2** 1.00	54.7** 1.00
Restricted	7.2** 0.993	13.6** 0.991	30.7** 1.00	56.4** 1.00
Panel B: LR-Tests – Missing Slope-Adjustment				
	Number of Risk Factors			
	$d = 1$	$d = 2$	$d = 3$	$d = 4$
Restr vs. Unrestr				
Constant λ	2930** 1.00	77.5** 1.00	21.0 0.927	50.7** 1.00
TV λ	2859** 1.00	85.4** 1.00	21.4 0.935	22.4* 0.966
Const vs TV λ				
Unrestricted	8.13** 0.996	19.6** 0.999	32.5** 1.00	72.2** 1.00
Restricted	79.4** 1.00	11.7* 0.981	32.1** 1.00	101** 1.00
Panel C: LR-Tests – Restricted B				
	Number of Risk Factors			
	$d = 1$	$d = 2$	$d = 3$	$d = 4$
Restr vs. Unrestr				
Constant λ	3444** 1.00	932** 1.00	217** 1.00	217** 1.00
TV λ	3333** 1.00	924** 1.00	205** 1.00	104** 1.00
Const vs TV λ				
Unrestricted	7.21** 0.993	7.5 0.888	36.4** 1.00	131** 1.00
Restricted	118** 1.00	15.5** 0.996	48.1** 1.00	244** 1.00

Figure 7: Estimated Intercept Terms in Sensitivity Analysis Model

In this figure we provide the estimated intercept term in the time-varying risk premia HJM model for a model from our sensitivity analysis. We estimate the following model for the slope-adjusted yield changes \tilde{y}_{t,τ_i} :

$$\begin{aligned}\tilde{y}_t &= \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t,\end{aligned}$$

with \tilde{y}_t the vector that collects all slope-adjusted yield changes at time t for different maturities $\tau_i = \tau_1, \dots, \tau_m$, $x_t = \lambda_{t-1} + w_t$, B the loading matrix whose columns are the volatility functions, b'_i the i -th row element of B , $\text{vec}_{1:m}\{\cdot\}$ denoting the $m \times 1$ vector with the typical element given in brackets, ε_t idiosyncratic noise with diagonal variance matrix Ψ and w_t standard normal. For the number of variance generating factors $d = 1, 2, 3, 4$, we estimate 4 models. We estimate the model with a constant λ_t (to this end we set $A = 0$) and time-varying λ_t ($A \neq 0$). For both of these variations we estimate an unrestricted version (where $\alpha \neq 0$ but we set $a = 0$ to ensure identifiability) and a model imposing the no-arbitrage drift restriction ($\alpha = 0$). In the above model we omit the $\text{vec}_{1:m}\{\hat{b}'_i \hat{b}_i \tau_i / 2\}$ term from the specification. In the figure below we report the estimates of α for both the constant and time-varying risk price case when the no-arbitrage restriction is not imposed (thus when $\alpha \neq 0$).



restriction $\alpha = 0$ at 1% if $d = 1$ or 2 factors are used. The difference is that in Panel A of Table 6, the same happens with $d = 3$ or 4 factors. This shows the importance of including the nonlinear $\text{vec}_{1:m}\{\hat{b}'_i \hat{b}_i \tau_i / 2\}$ term in the empirical analysis. Without this, strong rejection (at 1%) is erroneously obtained in these data. The intuition is that while precluding arbitrage opportunities in equities requires that the means be spanned by the volatility matrix through coefficients (risk prices) that enter linearly, in the term structure case it is the means of the appropriately slope-adjusted yield changes less the critical nonlinear function of volatilities given by $\text{vec}_{1:m}\{\hat{b}'_i \hat{b}_i \tau_i / 2\}$ that must be spanned in this manner. The simplified APT-style test ignoring this term may point to arbitrage opportunities that do not exist.

Figure 7 shows the unrestricted intercepts when ignoring the additional nonlinear volatility term. Regardless of d , these correspond exactly to the means from the summary statistics for slope-adjusted yield changes. The means are generally increasing in maturity, in sharp

contrast with the declining pattern of the unrestricted intercepts in the correctly specified model in Panel (A) of Figure 5. It is the latter means that may be spanned by the columns of B , at least to the extent of avoiding rejection at 1%, and from the difference in patterns it is not surprising that the means in Figure 7 may not be similarly spanned.

These comparisons demonstrate the empirical importance of the detailed structure of the no-arbitrage condition applied to slope-adjusted yield changes. Next, we consider the possibility that this adjustment is missed altogether.

4.2 Sensitivity to Data Error: Missing Slope-Adjustment

The slope-adjusted yield changes of (4) look very similar to raw yield changes, defined as $\Delta y_{t,\tau_i} = y_{t,\tau_i} - y_{t-1,\tau_i}$. The only differences between the two are given by the average slope or yield spread $(y_{t-1,\tau_i} - y_{t-1,\tau_0})/(\tau_i - \tau_0)$ and local slope or bond aging $(y_{t-1,\tau_i} - y_{t-1,\tau_{i-1}})/(\tau_i - \tau_{i-1})$ terms, which empirically tend to be decreasing in i as the step length between the maturities increases. As a second experiment we examine what happens if our approach is applied directly to raw yield changes $\Delta y_{t,\tau_i}$, without slope-adjustment.

Like Panel A, also Panel B of Table 6 may be directly compared and contrasted with Table 4. Qualitatively, the conclusions of tests again agree when using one or two factors, only. This time the main difference is that when using raw yield changes, we fail to reject even at 5% in the model with time-varying risk prices (the failure occurs for levels 6.5% and better). Again, due to the empirical relevance of the specification with three factors and time-varying risk prices, the difference is important. What it means is that the empirical analysis should be conducted carefully on properly slope-adjusted yield changes. Otherwise, the wrong conclusions may be drawn regarding the presence or absence of arbitrage opportunities.

Combining with the results of the previous subsection, we conclude that both the non-linear term in the drift restriction and the proper slope-adjustment of the yield changes are required in the empirical analysis.

4.3 Sensitivity to Framework: Restricted Volatility Factors and Affine Subclasses

Well-known affine term structure models arise by appropriately restricting the general volatility factors. As an additional illustration of the sensitivity to model misspecification, we consider testing this type of restrictions. In particular, we consider restrictions generating the affine models of Ho and Lee (1986) and Hull and White (1990), as well as more general volatility restrictions.

The volatility functions in Figure 2 behave like factor loadings in the set-up of the general model. The estimated functions seem to give rise to a level, slope and curvature interpretation. If this is indeed the case, alternative specifications may greatly reduce the number

of parameters. In the case where we have one covariance generating factor, $d = 1$, we could for example obtain the forward rate volatility $\sigma_f(\tau_i)$ from the Ho and Lee (1986) setting: $\sigma_f^1(\tau_i) = \sigma_1$, i.e., flat forward rate volatility. Integrating and dividing by the time to maturity provides a yield volatility of exactly σ_1 in this case for all maturities: $\sigma^1(\tau_i) = \sigma_1$. Our volatility matrix B would be a vector: $\sigma_1 \times 1_m$, i.e., the first factor is a level factor. For $d = 2$ we could take the first volatility factor as in the Ho-Lee case, and the second from the Hull and White (1990) model: $\sigma_f^2(\tau_i) = \sigma_2 \exp(-\gamma_1 \tau_i)$, exponentially declining forward rate volatility, with γ_1 a parameter to be estimated. Similar derivations as before produce the yield volatility $\sigma^2(\tau_i) = \sigma_2(1 - \exp(-\gamma_1 \tau_i))/\gamma_1 \tau_i$ in this case, i.e., the second factor is a slope factor, and the level and slope loadings or volatility functions only use up three degrees of freedom through $(\sigma_1, \sigma_2, \gamma_1)$. Given the similarity of these first two volatility functions with the Nelson and Siegel (1987) model we take a similar hump-shaped yield volatility as in that model for the next dimension. For the full four-factor case $d = 4$ we set

$$\sigma(t, \tau_i) = \begin{pmatrix} \sigma_1 \\ \sigma_2 \frac{1 - e^{-\gamma_1 \tau_i}}{\gamma_1 \tau_i} \\ \sigma_3 \frac{1 - (1 + \gamma_2 \tau_i) e^{-\gamma_2 \tau_i}}{\gamma_2^2 \tau_i} \\ \sigma_4 \frac{1 - e^{-2\gamma_3 \tau_i}}{2\gamma_3 \tau_i} \end{pmatrix}, \quad (15)$$

and in case $d < 4$ we use the first d elements of this vector.

Panel C of Table 6, again directly comparable with Table 4, provides the LR-test results when we estimate the time-varying HJM model with the restricted volatility factors. In this case the no-arbitrage restrictions are rejected for all tests. On the methodological side, the analysis shows the importance of using our approach to testing the no-arbitrage restrictions in the general HJM framework, without parametric restrictions imposed on the factor loadings. On the substantive side, the results provide evidence against the affine subclasses considered.

4.4 Sensitivity to Distributional Assumptions: Principal Components Analysis


To see how robust our results are to alternative distributional assumptions we also estimate two model variations where we use principal components analysis (PCA). In a first step we run a principal component analysis on the demeaned slope-adjusted yield changes. In particular, writing the PCA as $\tilde{y}_t - \hat{\mu} = Bw_t + \varepsilon_t$, where $\hat{\mu}$ denotes the sample mean, we also obtain estimates for the volatility matrix B and the covariance generating factors w_t in this framework. For comparison with the state space model, we use a version of the PCA with $\text{var}(w_t) = I_d$, absorbing the factor variances and covariances into B .

In Panel (A) of Figure 8 we show the estimated loadings on the covariance generating factors. The pattern is very similar to that in Panel (A) of Figure 2. With two factors,


Figure 8: Volatility Functions and Covariance-Generating Factors – PCA

In this figure we show a time series plot of estimates of the volatility functions $b_i = \sigma(t, \tau_i)$ and the estimated covariance-generating factors. We obtain these by running a principal component analysis on the slope-adjusted yield changes \tilde{y}_{t, τ_i} . The estimates in Panel (A) document the estimated B matrix (from the equation $\tilde{y}_t - \hat{\mu} = Bw_t + \varepsilon_t$, with $\hat{\mu}$ the sample mean of the slope-adjusted yield changes) for the dimension of the driving Wiener process W_t equal to $d = 1, 2, 3, 4$. The time series plots in Panel (B) document the estimated w_t .

(A) Volatility Functions (B) – PCA



(B) Covariance-Generating Factors (w_t) – PCA



the first captures the level, and the second is a slope factor. The third factor captures the curvature or hump. With only one factor, the loading pattern in the $d = 1$ portion of the figure may appear to be explaining part of the hump, but the vertical axes are different, and this is in fact the same loading pattern as that for the level factor in the specifications with higher d . In general, the PCA loadings on the first k factors are unchanged as more factors are added, a property not shared with the previous state space analysis. Note also that the loadings on the level factor are of the order .3 in the PCA, compared to about .08 in the state space case with $d = 1$ in Figure 2, although the shape is similar. The reason for the difference is that the loadings exhibited in the state space case are those from the model with the no-arbitrage condition imposed, and they have to capture both variance-covariance properties and means of the data, by the nature of the restriction where B enters into both first and second moments. For higher d , the impact of the restriction on the level factor is less since with more factors it is easier to span the means by the columns of B , and so the loadings on the level factor for $d = 2$ and higher in Panel (A) of Figure 2 are more similar to those from the PCA in Panel (A) of Figure 8.

In Panel (B) of Figure 8 we provide a time series plot of the fitted covariance-generating factors \hat{w}_t . Again, the picture is similar to that in the state space case, Panel (B) of Figure 2. For example, the large positive spike in the fourth (slope) factor in 1986 in the state space case is matched by a positive spike in the second principal component similarly associated with downward sloping loadings.

The estimated volatility functions B and covariance generating factors w_t allow us to take another look at the obtained risk prices. To this end, the regression of the sample estimate of the unconditional mean of the slope-adjusted yield changes $\hat{\mu}$ on the estimated volatility matrix from the principal component analysis \hat{B} gives these risk prices. Following equation (6), we need to subtract $\hat{b}'_i \hat{b}_i \tau_i / 2$ from each cross-sectional observation. The regression in this case is

$$\hat{\mu} - \text{vec}_{1:m} \{ \hat{b}'_i \hat{b}_i \tau_i / 2 \} = \hat{B} \lambda + \eta. \quad (16)$$

An important thing to note is that this procedure is only an approximate way to get the risk prices. The estimates of the volatility functions are obtained in the first step through principal components. Thus, the information from the no-arbitrage relation estimated in the second step is not utilized when the volatility functions and covariance generating factors are obtained in the first. The state space approach allows imposing the drift restriction from the outset, so that estimation is subject to this.

In Table 7 we show the estimated risk prices when we estimate equation (16) by generalized least squares (GLS), to account for heterogeneity. The covariance used to weigh the observations is estimated using the PCA results: $\widehat{\text{var}}[\eta] = \hat{B} \hat{B}' + \hat{\Psi}$, with $\hat{\Psi}$ diagonal with elements $\hat{\Psi}_{ii} = (1/T) \sum_t \hat{\varepsilon}_{it}^2$ for $i = 1, \dots, m$. The risk prices we obtain follow a similar pattern to those in the state space model in Table 2. In both cases, the first (level) risk

Table 7: Risk Price Estimates – Two-Step PCA and GLS Approach

In this table we report the estimated risk prices and a test of the no-arbitrage drift restriction using a two-step approach. In the first step we run a principal component analysis (PCA) on the demeaned slope-adjusted yield changes \tilde{y}_{t,τ_i} . Specifically, writing the PCA as $\tilde{y}_t - \hat{\mu} = Bw_t + \varepsilon_t$, with $\hat{\mu}$ the sample mean of the slope-adjusted yield changes, we obtain \hat{B} . Then, in the second step, we run a cross-sectional GLS regression to obtain estimates of the risk premia. Specifically, we estimate

$$\hat{\mu} - \text{vec}_{1:m}\{\hat{b}'_i \hat{b}_i \tau_i / 2\} = \hat{B}\lambda + \eta,$$

with \hat{b}'_i the i -th row element of \hat{B} and $\text{vec}_{1:m}\{\cdot\}$ defined as the vector with the typical element in brackets for $i = 1, \dots, m$. The table documents the estimated λ for the dimension of the driving Wiener process W_t equal to $d = 1, 2, 3, 4$. To account for heteroskedasticity and cross-correlation we use GLS. The covariance is estimated using the results from the principal components analysis: $\widehat{\text{var}}[\eta] = \hat{B}\hat{B}' + \hat{\Psi}$, with $\hat{\Psi}$ diagonal with elements $\hat{\Psi}_{ii} = (1/T) \sum_t \hat{\varepsilon}_{it}^2$ for $i = 1, \dots, m$. The t -statistics are given below the estimates, the asterisks (*/**) denote significance at the 5%/1% level. In addition we report the adjusted R^2 .

Risk Prices – Two-Step PCA & GLS				
	Number of Risk Factors			
	$d = 1$	$d = 2$	$d = 3$	$d = 4$
Factor 1 (λ_1)	-7.09 -0.426	-7.66 -0.976	-8.08 -1.86	-8.06 -1.93
Factor 2 (λ_2)		24.3** 3.03	22.8** 5.15	23.1** 5.46
Factor 3 (λ_3)			14.1* 2.99	13.3* 3.05
Factor 4 (λ_4)				4.00 0.925
R^2	0.012	0.421	0.752	0.783
#Obs	192	192	192	192

price is negative regardless the number of factors included in the model. With two factors, $d = 2$, the first risk price is similar to the $d = 1$ case, whereas the second risk price (on the slope factor) is an order of magnitude larger (in an absolute sense). Keeping in mind that the ordering of the factors in the table switches in the state space case but not in the PCA case, it is seen that the patterns of risk prices is similar across the two cases, with the slope factor carrying the largest risk price.

To obtain time-varying estimates of the risk prices in this framework we estimate a slightly altered version of the above. Based on (12) and (9), the relevant cross-sectional regression in period t is

$$\tilde{y}_t - \hat{B}\hat{w}_t - \text{vec}_{1:m}\{\hat{b}'_i\hat{b}_i\tau_i/2\} = \hat{B}\lambda_{t-1} + \eta_t, \quad (17)$$

again estimated by GLS using $\widehat{\text{var}[\eta]}$ from above. As in this case the regressand is time-varying, our risk prices will be so, too. In Figure 9 we show the time series of estimated risk prices. The evolution of risk prices resembles that from the state space model, Figure 3. The time series behavior of each risk price is quite homogeneous through the period, except perhaps for the first risk price, where variability is greater in the early part of the period, and this is true for both the state space and PCA versions.

5 Conclusion

Our general dynamic factor model approach facilitates analysis of several relevant issues in term structure analysis, such as the number of factors, the shapes of the volatility functions, the no-arbitrage drift restriction, the dynamics of risk premia, and their relation to macroeconomic conditioning variables. On the methodological side, the implementation clearly demonstrates the tractability of our likelihood based state space approach. Formally, Gaussian distributional assumptions are adopted, but we demonstrate consistency of our results with those obtained using robust alternative procedures. Furthermore, even without Gaussianity, the Kalman filter generates minimum mean squared error linear predictions of the latent state variables. On the substantive side, our results document the importance of time-varying risk premia in the Heath, Jarrow and Morton (1992) framework, and we fail to reject the absence of arbitrage opportunities at level 1%, although not at 5%, provided three or more factors are included in the model specification. Most of the risk premium is associated with the important slope factor. Movements in the price of slope risk depend strongly on past movements in both the slope price itself and in the price of level risk, with an increase in yield levels leading to a subsequent decrease in slope. The price of slope risk is strongly related to the equity risk price. The price of level risk is strongly related to both the equity risk price and consumption growth, whereas the price of curvature risk is


Figure 9: Time Series Plot of Risk Prices – Two-Step Approach

In this figure we show a time series plot of the estimated time-varying risk prices using a two-step approach. In the first step we run a principal component analysis (PCA) on the demeaned slope-adjusted yield changes \tilde{y}_{t,τ_i} . Specifically, writing the PCA as $\tilde{y}_t - \hat{\mu} = Bw_t + \varepsilon_t$, with $\hat{\mu}$ the sample mean of the slope-adjusted yield changes, we obtain \hat{B} and \hat{w}_t . Then, in the second step, we run a cross-sectional GLS regression at each point in time to obtain estimates of the risk prices. Specifically, we estimate

$$\tilde{y}_t - \hat{B}\hat{w}_t - \text{vec}_{1:m}\{\hat{b}'_i\hat{b}_i\tau_i/2\} = \hat{B}\lambda_t + \eta_t,$$

with \hat{b}'_i the i -th row element of \hat{B} and $\text{vec}_{1:m}\{\cdot\}$ defined as the vector with the typical element in brackets for $i = 1, \dots, m$. The time series plots below document the estimated λ_t for the dimension of the driving Wiener process W_t equal to $d = 1, 2, 3, 4$. To account for heteroskedasticity and cross-correlation we use GLS. The covariance is estimated using the results from the principal components analysis: $\widehat{\text{var}}[\eta_t] = \hat{B}\hat{B}' + \hat{\Psi}$, with $\hat{\Psi}$ diagonal with elements $\hat{\Psi}_{ii} = (1/T) \sum_t \varepsilon_{it}^2$ for $i = 1, \dots, m$. We show the estimated λ time series that is obtained for each of the four factors in the case $d = 4$.

Risk Prices Time Series for $d = 4$ – Two-Step PCA and GLS Approach



strongly related to the output gap variable suggested by Cooper and Priestley (2009). Our work establishes the empirical importance of the nonlinear (quadratic in loadings) term in the no arbitrage condition and of proper treatment of bond aging and yield spread through adjustment of yield changes for local and average slope of the yield curve.

Judging from our results, derivative pricing in the HJM framework should not be based on model specifications with only one or two factors, and should not restrict the shape of the volatility function to the standard parametrized forms associated with the affine subclasses considered, even though both simplifications are extremely popular in practical applications. More generally, the indication is that the framework should not be adopted for derivative pricing based on a volatility function and an initial yield curve, only, without a preceding analysis of the appropriate structure involving dynamic risk premium specification and testing for arbitrage opportunities under the physical measure.

Appendix A

Here, we derive the appropriate no-arbitrage drift condition for the yield dynamics with fixed term to maturity. There are two main differences between our setting and that in HJM: (i) we study yields to maturity, while HJM consider instantaneous forward rates; and (ii) we use the parametrization from Brace and Musiela (1994) with fixed term τ to maturity, whereas HJM use a fixed maturity date T in their notation.

The HJM specification of the dynamics of instantaneous forward rates takes the form

$$df(t, T) = \alpha_f(t, T) dt + \sigma_f(t, T)' dW_t,$$

i.e., an infinite dimensional SDE for the forward curve $f(t, \cdot)$, where α_f and σ_f are the forward rate drift and volatility functions, t is calendar time, and T the fixed maturity date. Equivalently, we assume that under the physical measure the yield curve dynamics can be written as

$$dy(t, \tau) = \alpha(t, \tau) dt + \sigma(t, \tau)' dW_t,$$

which is (1). Here, $\tau \geq 0$ indicates term to maturity. For the traded bond with maturity date T , we may write $T = t + \tau$, i.e., term to maturity τ shrinks as calendar date t increases, which is the bond aging effect. Write $P(t, T) = \exp(-(T - t)y(t, T - t))$ for the zero coupon bond price. By Itô's lemma, using $T = t + \tau$, bond price dynamics are

$$\begin{aligned} \frac{dP(t, T)}{P(t, T)} &= -\tau dy(t, \tau) + \frac{1}{2}\tau^2 \sigma(t, \tau)' \sigma(t, \tau) dt + y(t, \tau) dt + \tau \frac{\partial y}{\partial \tau}(t, \tau) dt \\ &= \left(-\tau \alpha(t, \tau) + \frac{1}{2}\tau^2 \sigma(t, \tau)' \sigma(t, \tau) + y(t, \tau) + \tau \frac{\partial y}{\partial \tau}(t, \tau) \right) dt - \tau \sigma(t, \tau)' dW_t \\ &= \alpha_p(t, \tau) dt + \sigma_p(t, \tau)' dW_t, \end{aligned}$$

where α_p and σ_p are the expected return and return volatility of the bond. The absence of arbitrage opportunities requires the existence of a market price of risk process λ_t such that

$$\alpha_p(t, \tau) = r_t + \sigma_p(t, \tau)' \lambda_t,$$

where $r_t = y(t, 0)$ is the short rate at time t . Inserting for the bond drift and volatility from above, we get the yield drift condition under the physical measure,

$$\alpha(t, \tau) = \frac{1}{\tau} (y(t, \tau) - y(t, 0)) + \frac{\partial y}{\partial \tau}(t, \tau) + \frac{1}{2} \tau \sigma(t, \tau)' \sigma(t, \tau) + \sigma(t, \tau)' \lambda_t,$$

which is (2).

Appendix B

Here we briefly describe the state space approach and the Kalman filter, following Harvey (1989) and Durbin and Koopman (2001), that we apply to our model (12), restated as

$$\begin{aligned} \tilde{y}_t &= \alpha + \text{vec}_{1:m} \{b'_i b_i \tau_i / 2\} + Bx_t + \varepsilon_t, \\ x_t &= a + Ax_{t-1} + w_t. \end{aligned}$$

The model is in state space form, with measurement equation given by the first equation above, and state transition equation given by the second. In the special case of a static model, $A = 0$, the model collapses to

$$\tilde{y}_t = \alpha + \text{vec}_{1:m} \{b'_i b_i \tau_i / 2\} + Ba + Bw_t + \varepsilon_t,$$

which is the classical factor analysis model with special structure on the means. If no arbitrage is not imposed, α is estimated by the sample average of $\tilde{y}_t - \text{vec}_{1:m} \{b'_i b_i \tau_i / 2\}$, and the average risk price $\lambda = a$ is unidentified. With no arbitrage imposed, $\alpha = 0$, estimation of risk prices proceeds as a GLS regression of $\tilde{y}_t - \text{vec}_{1:m} \{b'_i b_i \tau_i / 2\}$ on the columns of the loadings matrix B . In both cases, the relevant variance-covariance matrix is that from the classical factor analysis, $\Sigma = BB' + \Psi$, and since the estimates of α respectively $\lambda = a$ depend on B (through the quadratic term in b_i in the unrestricted case, and in the restricted case in addition through regression on B and GLS using Σ), the estimates are inserted in the likelihood function and the resulting concentrated likelihood is maximized iteratively with respect to B and Ψ .

In the dynamic case with unrestricted state transition A , the Kalman filter recursions provide estimates of the unobserved state x_t given the data, with a distinction between conditioning on two different information sets. In each recursive step the estimate of the unobserved state at time t is first made based on the information up to the previous time

point, that is $\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_{t-1}$. We denote this estimate as $x_{t|t-1} = E[x_t | \tilde{y}_1, \dots, \tilde{y}_{t-1}] = E_{t-1}[x_t]$ and refer to it as the ‘predicted’ state. The uncertainty around this estimate is measured by the covariance matrix $P_{t|t-1} := \text{var}_{t-1}[x_t]$. When the new observation \tilde{y}_t becomes available we can update our information set to $\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_{t-1}, \tilde{y}_t$, and obtain an updated estimate of the unobserved state. We denote this updated estimate as $x_{t|t} = E[x_t | \tilde{y}_1, \dots, \tilde{y}_t] = E_t[x_t]$ and refer to it as the ‘filtered’ state; the uncertainty around this estimate is measured by $P_{t|t} := \text{var}_t[x_t]$.

Due to the properties of the Gaussian ε_t and w_t the predicted and filtered estimates of the unobserved state, and the uncertainty around these, can be calculated in a simple recursive manner. Given that for time $t - 1$ a filtered estimate $x_{t-1|t-1}$ (and error variance $P_{t-1|t-1}$) of the state has been obtained, we get a predicted estimate and error variance for the state at the next time t through

$$\begin{aligned} x_{t|t-1} &= a + Ax_{t-1|t-1}, \\ P_{t|t-1} &= AP_{t-1|t-1}A' + \text{var}(w_t) = AP_{t-1|t-1}A' + I_d, \end{aligned}$$

with I_d the d -dimensional identity matrix. After the new observation \tilde{y}_t comes in we can update these estimates and obtain the filtered estimate and error variance for the state using

$$\begin{aligned} x_{t|t} &= x_{t|t-1} + P_{t|t-1}B'F_t^{-1}v_t, \\ P_{t|t} &= P_{t|t-1} - P_{t|t-1}B'F_t^{-1}BP_{t|t-1}, \end{aligned}$$

where we define

$$\begin{aligned} v_t &= \tilde{y}_t - \tilde{y}_{t|t-1} = \tilde{y}_t - (\alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_{t|t-1}), \\ F_t &= \text{var}_{t-1}(v_t) = BP_{t|t-1}B' + \text{var}(\varepsilon_t) = BP_{t|t-1}B' + \Psi. \end{aligned}$$

To initialize the Kalman filter we use the properties of the process of the latent state x_t to get the unconditional distribution (see also the discussion in Section 3.2). We can write the dynamics of the state as

$$x_t = (I - A)(I - A)^{-1}a + Ax_{t-1} + w_t \Leftrightarrow x_t = (I - A)\mu + Ax_{t-1} + w_t \Leftrightarrow x_t - \mu = A(x_{t-1} - \mu) + w_t,$$

with $\mu = (I - A)^{-1}a$. For the unconditional variance of x_t , denoted Σ_x , we get

$$\Sigma_x = A\Sigma_x A' + I_d,$$

which we can solve using the properties of the vectorization operator vec to give:

$$\text{vec}(\Sigma_x) = [I_{d^2} - (A \otimes A)]^{-1} \text{vec}(I_{d^2}),$$

with \otimes the Kronecker product. Thus, the filter is initialized by setting $x_{0|0} = \mu$ and $P_{0|0} = \Sigma_x$.

The vectors a and α and the matrices B , Ψ , A contain the parameters that need to be estimated. Given normality of the ε_t and w_t we can use the output from the above Kalman recursions. In general, the logarithm of the likelihood for a model with conditional probability density function $p(\tilde{y}_t|\tilde{y}_1, \dots, \tilde{y}_{t-1})$ is given by

$$\log L(\tilde{y}_1, \dots, \tilde{y}_T) = \sum_{t=1}^T \log p(\tilde{y}_t|\tilde{y}_1, \dots, \tilde{y}_{t-1}),$$

where $p(\tilde{y}_1|\tilde{y}_0)$ is according to the initialization of the system (see above), and T is the number of observations. Due to normality of all distributions and the additivity in the system, the latent state x_t conditional on $\tilde{y}_1, \dots, \tilde{y}_{t-1}$ is normal with mean $x_{t|t-1}$ and variance $P_{t|t-1}$. For \tilde{y}_t we can write

$$\tilde{y}_t = \alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + B(x_t - x_{t|t-1}) + Bx_{t|t-1} + \varepsilon_t,$$

from which we get that the distribution of \tilde{y}_t conditional on $\tilde{y}_1, \dots, \tilde{y}_{t-1}$ is normal with mean $\alpha + \text{vec}_{1:m}\{b'_i b_i \tau_i / 2\} + Bx_{t|t-1}$ and variance matrix F_t . Thus, using the definition of v_t above, we can write the log-likelihood for the state space model as

$$\log L(\tilde{y}_1, \dots, \tilde{y}_T) = -Tm/2 - \sum_{t=1}^T \log |F_t|/2 - \sum_{t=1}^T v'_t F_t^{-1} v_t / 2.$$

This is maximized numerically, re-running the filter at each trial value of the parameter vector. Standard errors are estimated off the squareroots of the diagonal elements of the negative inverse Hessian at the optimum.

References

- Amin, K. and A. Morton (1994). Implied Volatility Functions in Arbitrage-Free Term Structure Models. *Journal of Financial Economics* 35, p141–180.
- Ang, A. and M. Piazzesi (2003). A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables. *J. Monetary Economics* 50, p745–787.
- Bams, D. and C. Wolff (2003). Risk Premia in the Term Structure of Interest Rates: A Panel Data Approach. *International Financial Markets, Institutions and Money* 13, p211–236.
- Bliss, R. and P. Ritchken (1996). Empirical Tests of Two State-variable Heath-Jarrow-Morton Models. *Journal of Money, Credit and Banking* 28, p452–476.

- Brace, A. and M. Musiela (1994). A Multifactor Gauss Markov Implementation of Heath, Jarrow, and Morton. *Mathematical Finance* 4, p563–576.
- Campbell, J. and R. Shiller (1991). Yield Spreads and Interest Rate Movements: A Bird’s Eye View. *Review of Economic Studies* 58, p495–514.
- Chapman, D., J. Long, and N. Pearson (1999). Using Proxies for the Short Rate: When Are Three Months Like an Instant? *Review of Financial Studies* 12, p763–806.
- Cooper, I. and R. Priestley (2009). Time-Varying Risk Premiums and the Output Gap. *Review of Financial Studies* 22, p2801–2833.
- Cox, R., J. Ingersoll, and S. Ross (1985a). An Intertemporal General Equilibrium Model of Asset Prices. *Econometrica* 53, p363–384.
- Cox, R., J. Ingersoll, and S. Ross (1985b). A Theory of the Term Structure of Interest Rates. *Econometrica* 53, p385–407.
- Dai, Q. and K. Singleton (2000). Specification Analysis of Affine Term Structure Models. *Journal of Finance* 55(5), p1943–1978.
- Dai, Q. and K. Singleton (2002). Expectation Puzzles, Time-Varying Risk Premia, and Affine Models of The Term Structure. *Journal of Financial Economics* 63, p415–441.
- de Jong, F. and P. Santa-Clara (1999). The Dynamics of the Forward Interest Rate Curve: a Formulation With State Variables. *Journal of Financial and Quantitative Analysis* 34, p131–157.
- Diebold, F. and C. Li (2006). Forecasting the Term Structure of Government Bond Yields. *Journal of Econometrics* 130, p337–364.
- Diebold, F., S. Rudebusch, and S. Aruoba (2006). The Macroeconomy and the Yield Curve. *Journal of Econometrics* 131, p309–338.
- Duarte, J. (2004). Evaluating an Alternative Risk Preference in Affine Term Structure Models. *Review of Financial Studies* 17, p379–404.
- Duffee, G. (2002). Term Premia and Interest Rate Forecasts in Affine Models. *Journal of Finance* 57, p405–443.
- Duffee, G. (2008). Information in (and not in) the Term Structure. Working Paper.
- Duffie, D. and R. Kan (1996). A Yield-Factor Model of Interest Rates. *Mathematical Finance* 6, p379–406.
- Durbin, J. and S. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford Statistical Science Series.
- Fama, E. (1984). The Information in the Term Structure. *Journal of Financial Economics* 13, p509–528.

- Fama, E. and R. Bliss (1987). The Information in Long-Maturity Forward Rates. *American Economic Review* 77, p680–692.
- Flesaker, B. (1993). Testing the Heath-Jarrow-Morton/Ho-Lee Model of Interest Rate Contingent Claims Pricing. *Journal of Financial and Quantitative Analysis* 28, p483–495.
- Geweke, J. and G. Zhou (1996). Measuring the Pricing Error of the Arbitrage Price Theory. *Review of Financial Studies* 9, p557–587.
- Harvey, A. (1989). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Heath, D., R. Jarrow, and A. Morton (1992). Bond Pricing and the Term Structure of Interest Rates: a New Methodology for Contingent Claims Valuation. *Econometrica* 60, p77–105.
- Ho, T. and S. Lee (1986). Term Structure Movements and Pricing Interest Rate Contingent Claims. *Journal of Finance* 41, p1011–1029.
- Hull, J. and A. White (1990). Pricing Interest-Rate-Derivative Securities. *Review of Financial Studies* 3, p573–592.
- Jeffrey, A., D. Kristensen, O. Linton, T. Nguyen, and P. Phillips (2004). Nonparametric Estimation of a Multifactor Heath-Jarrow-Morton Model: An Integrated Approach. *Journal of Financial Econometrics* 2, p251–289.
- Joslin, S., M. Priebsch, and K. Singleton (2009). Risk Premium Accounting in Macro-Dynamic Term Structure Models. Working Paper.
- Litterman, R. and J. Scheinkman (1991). Common Factors Affecting Bond Returns. *Journal of Fixed Income* 1, p54–61.
- Ludvigson, S. and S. Ng (2009). Macro Factors in Bond Risk Premia. *Review of Financial Studies* (forthcoming).
- Nelson, C. and A. Siegel (1987). Parsimonious Modelling of Yield Curves. *Journal of Business* 60-4, p473–489.
- Newey, W. K. and K. D. West (1987). A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55, 703–8.
- Roll, R. and S. Ross (1980). An Empirical Investigation of the Arbitrage Pricing Theory. *Journal of Finance* 35, p1073–1103.
- Ross, S. (1976). The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* 13, p341–360.
- Stanton, R. (1997). A Nonparametric Model of Term Structure Dynamics and the Market Price of Interest Rate Risk. *Journal of Finance* 52, p1973–2002.

- Stock, J. H. and M. Watson (2002). Macroeconomic Forecasting using Diffusion Indexes. *Journal of Business and Economic Statistics* 20, 147–62.
- Vasicek, O. (1977). An Equilibrium Characterization of the Term Structure. *Journal of Financial Economics* 5, p177–188.
- Wachter, J. (2006). A Consumption-Based Model of the Term Structure of Interest Rates. *Journal of Financial Economics* 79, p365–399.

Research Papers 2010



- 2010-01: Anders Bredahl Kock and Timo Teräsvirta: Forecasting with nonlinear time series models
- 2010-02: Gunnar Bårdsen, Stan Hurn and Zoë McHugh: Asymmetric unemployment rate dynamics in Australia
- 2010-03: Jesper Rangvid, Maik Schmeling and Andreas Schrimpf: Cash Flow-Predictability: Still Going Strong
- 2010-04: Helle Bunzel and Walter Enders: The Taylor Rule and “Opportunistic” Monetary Policy
- 2010-05: Martin M. Andreasen: Non-linear DSGE Models and The Optimized Particle Filter
- 2010-06: Søren Johansen and Bent Nielsen: Discussion of The Forward Search: Theory and Data Analysis by Anthony C. Atkinson, Marco Riani, and Andrea Ceroli
- 2010-07: Giuseppe Cavaliere, Anders Rahbek and A.M. Robert Taylor: Bootstrap Sequential Determination of the Co-integration Rank in VAR Models
- 2010-08: Peter R. Hansen and Asger Lunde: Estimating the Persistence and the Autocorrelation Function of a Time Series that is Measured with Error
- 2010-09: Tom Engsted, Thomas Q. Pedersen and Carsten Tanggaard: Pitfalls in VAR based return decompositions: A clarification
- 2010-10: Torben G. Andersen and Luca Benzoni: Stochastic Volatility
- 2010-11: Torben B. Rasmussen: Affine Bond Pricing with a Mixture Distribution for Interest Rate Time-Series Dynamics
- 2010-12: Martin M. Andreasen and Bent Jesper Christensen: The SR Approach: a new Estimation Method for Non-Linear and Non-Gaussian Dynamic Term Structure Models
- 2010-13: Peter Reinhard Hansen, Zhuo (Albert) Huang and Howard Howan Shek: Realized GARCH: Complete Model of Returns and Realized Measures of Volatility
- 2010-14: Bent Jesper Christensen and Michel van der Wel: An Asset Pricing Approach to Testing General Term Structure Models including Heath-Jarrow-Morton Specifications and Affine Subclasses