Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Journal of Banking & Finance 31 (2007) 3125–3144
Journal of BANKING & FINANCE
www.elsevier.com/locate/jbf

Yield-factor volatility models

Christophe Pérignon *, Daniel R. Smith

Faculty of Business Administration, Simon Fraser University, 8888 University Drive, Burnaby, BC, Canada V5A 1S6

Received 6 December 2005; accepted 15 November 2006
Available online 24 January 2007

Abstract

The term structure of interest rates is often summarized using a handful of yield factors that capture shifts in the shape of the yield curve. In this paper, we develop a comprehensive model for volatility dynamics in the level, slope, and curvature of the yield curve that simultaneously includes level and GARCH effects along with regime shifts. We show that the level of the short rate is useful in modeling the volatility of the three yield factors and that there are significant GARCH effects present even after including a level effect. Further, we find that allowing for regime shifts in the factor volatilities dramatically improves the model's fit and strengthens the level effect. We also show that a regime-switching model with level and GARCH effects provides the best out-of-sample forecasting performance of yield volatility. We argue that the auxiliary models often used to estimate term structure models with simulation-based estimation techniques should be consistent with the main features of the yield curve that are identified by our model.

JEL classification: E43; C32; C51

Keywords: Yields; GARCH; Level effect; Regime shifts

1. Introduction

The term structure of interest rates is often summarized using a handful of yield factors that capture shifts in the shape of the yield curve, i.e., changes in the overall level, slope,

* Corresponding author. Tel.: +1 604 291 3471; fax: +1 604 291 4920.
E-mail addresses: cperigno@sfu.ca (C. Pérignon), drsmith@sfu.ca (D.R. Smith).

0378-4266/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2006.11.016

and curvature of the yield curve (Litterman and Scheinkman, 1991). This factor decomposition provides a parsimonious representation of the term structure and is extensively used in fixed-income derivative pricing (Driessen et al., 2003), to model the linkages between interest rates and macroeconomic variables (Ang and Piazzesi, 2003), or to estimate term structure models (Brandt and Chapman, 2003). Despite the wide application of yield factors in financial economics, very little is known about their volatility dynamics. In this paper, we study the dynamics of yield-factor conditional volatility.

The key features of our econometric model are inspired by early models of the short-term interest rate. The main conclusion from this literature is that a level effect, in which the volatility is a positive function of the level of interest rates, GARCH effects, and regime shifts are required to adequately model the short-rate volatility. The dependence of interest rate volatility on the level of the short rate is systematically studied in Chan et al. (1992, hereafter CKLS). The first model that combines both level and GARCH effects for the short-rate volatility is proposed in Longstaff and Schwartz (1992). Further, Brenner et al. (1996) show that models that include both GARCH and level effects are better able to predict volatility than models that only include one of these effects. Gray (1996) extends the GARCH-level model to allow for multiple regimes in short-rate volatility and concludes that all three effects are needed to adequately model interest rate volatility.

Less research has been devoted to understanding the joint-dynamics of the level, slope, and curvature of the yield curve, which we call the yield factors. Pérignon and Villa (2006) show that the volatility of the yield factors is the primary source of time variation in the covariance matrix of interest rates. The role of conditional heteroscedasticity in the dynamics of the volatility of the yield factors has been highlighted by Christiansen and Lund (2002) and Christiansen (2004). The empirical challenge of estimating a level effect in a multi-factor framework has been tackled by Boudoukh et al. (1998), Brandt and Chapman (2003) and Christiansen (2005a). Another strand of research examines the role of regime shifts in the dynamics of yield-factor volatilities. Using international data, Kugler (1996) and Ang and Bekaert (2002a,b) estimate a two-state regime-switching VAR model of the level and the slope factors with a constant covariance matrix in each regime (i.e., without any level or GARCH effects). Christiansen (2004) extends the latter approach by fitting a two-state regime-switching ARCH model to the US level and the slope factors. A broad conclusion of this research is that regime shifts are a central feature of yield-factor volatilities.

Our primary contribution is the development of a comprehensive model for yield factors that simultaneously includes level and GARCH effects along with regime shifts in the factor volatilities. This model contrasts with comprehensive univariate short-rate models (Gray, 1996; Smith, 2002) and with bi- or tri-variate yield-factor models accounting for one or two of the above volatility features (Christiansen and Lund, 2002; Christiansen, 2004; Christiansen, 2005a). To the best of our knowledge, our multivariate model of the level, slope, and curvature of the yield curve is the first one to jointly include all these volatility features.

We model the variance of the three yield factors as a function of the short term interest rate for two reasons. First, empirically, the volatility of all three yield factors tends to be higher when short-term interest rates are higher. Second, we demonstrate theoretically that within the class of affine term structure models in which the volatility of all state variables is determined by a single state variable, i.e., the $\mathbb{A}_1(3)$ class in the terminology of Dai and Singleton (2000), both the level of the short-term interest rate and the variance of all three yield factors are determined by this very same state variable. As a result, we model the variance of all three yield factors using a level effect in which conditional volatility is a nonlinear function of the level of short-term interest rates. In our level effect, the short rate serves as a proxy for the unobserved latent volatility state variable. Further, we endeavor to combine both GARCH and level effects, while maintaining the traditional interpretation that a GARCH $(1,1)$ model implies an ARMA $(1,1)$ representation for the squared residual. In addition, our model explicitly includes regime shifts, a feature which has been demonstrated to be important in fitting short-term interest rates. Each model we consider is nested within this encompassing model, so we are able to directly measure the marginal contribution of each component of the model.

Using monthly US bond yields over the 1970–2002 period, we show that all three yield factors display a significant level effect. In particular, we find that the level effect for the slope factor is better captured by the overall level of interest rates rather than by the level of the slope factor. A similar conclusion is reached for the curvature factor. Further, the three yield factors exhibit strong GARCH effects even after including a level effect. We also examine regime-switching models that recognize different regimes in the volatility of the yield factors. We find that allowing for regime shifts dramatically improves the model's fit and strengthens the level effect. Finally, we show that a regime-switching model with level and GARCH effects provides the best out-of-sample forecasting performance of yield volatility.

Our empirical results have some important implications for the estimation of term structure models. We argue that the auxiliary models often used to estimate term structure models with simulation-based estimation techniques should be consistent with the main features of the yield curve that are identified by our model. We also show how the yield-factor models presented in this paper can provide a common platform to compare term structure models. The remainder of the paper proceeds as follows. Section 2 details the model, Section 3 describes the data, Section 4 presents the empirical results, Section 5 discusses the main implications of our results for estimating and comparing term structure models, and Section 6 offers some concluding comments.

2. Model development

2.1. Single-regime models

We are interested in modeling the dynamics of the following three yield factors: the level (L) of interest rates, the slope (S) of the yield curve, and the curvature (C) of the yield curve. We associate the level factor with the short rate, the slope factor with the difference between the long rate and the short rate, and the curvature factor with the sum of the short and long rates minus twice the intermediate rate. The latter transformation is a discrete measure of the second derivative of the yield curve, which captures its curvature. We model linear combinations of the yields, rather than yields directly, for two reasons. First, we find that the yield factors are much less correlated than the underlying yields, which allows us to use a more parsimonious model that captures all the important features of the yield curve. Second, we get a better in-sample fit using yield factors than by using yields.

Denote by F_{it} the ith yield factor, with $i = L, S, C$, whose dynamics is modeled as

$$\mathrm{d}F_{it} = (a_{0i} + a_{1i} \cdot F_{it})\mathrm{d}t + \sigma_{it} F_{jt}^{\gamma_i} \mathrm{d}W_{it}, \tag{1}$$

where W_{it} is a standard Brownian motion and W_{Lt}, W_{St}, and W_{Ct} may be correlated.[1] This model is inspired by early models of the short rate which include both mean reversion and allow the conditional volatility to be a function of the level of the short rate (Cox et al., 1985, among others). When allowing for a level effect in a multi-factor framework, one can either model the residual volatility of a given factor as a function of the value of this very factor ($j = i$) or of the level factor ($j = L$). We argue that the approach based on the level factor should be preferred. First, the own-level approach appears rather inconsistent since the residual standard deviation becomes negative when the value of the factor (slope or curvature) is negative, which is clearly impossible. Second, given that the level effect is so important in modeling volatility of short-term interest rates, it is likely to also affect the volatilities of the slope and curvature of the yield curve. As will be seen below, this conjecture is born out by the data. Third, the approach based on the level factor is consistent with the popular affine model with N state variables, among which one drives the state variable volatilities. Indeed, if the yield factors are defined as linear combinations of bond yields, as they are in the literature, the conditional volatility of all the yield factors depends on the same single source of risk (see the Appendix available on request for a proof).[2]

In our empirical work, we discretize the process in Eq. (1) as

$$\Delta F_{it} = \alpha_{0i} + \alpha_{1i} \cdot F_{it-1} + e_{it} \tag{2}$$

for $i = L, S, C$ and $t = 1, \ldots, T$. We approximate $\mathrm{d}W_{it}$, which is normally distributed with variance $\mathrm{d}t$, by a normally distributed innovation e_{it}. We decompose the conditional variance of e_{it} into the product of two terms:[3]

$$E(e_{it}^2 | \psi_{t-1}) = \delta_{it}^2 = \sigma_{it}^2 F_{jt-1}^{2\gamma_i}, \tag{3}$$

with ψ_{t-1} denoting the information set at time $t-1$. This specification allows heteroscedasticity to enter through a time-varying coefficient σ_{it}^2, which depends on past shocks on the residuals factors, and through the level effect. It is useful to rewrite the residual in Eq. (2) as $e_{it} = \sigma_{it} F_{jt-1}^{\gamma_i} z_{it}$ with z_{it} being i.i.d. $N(0,1)$. In this context, σ_{it}^2 is the conditional variance of the scaled residual $v_{it} = e_{it}/F_{jt-1}^{\gamma_i} = \sigma_{it} z_{it}$. We model the conditional variance of v_{it} as a GARCH process:

$$\sigma_{it}^2 = \beta_{0i} + \beta_{1i} \cdot v_{it-1}^2 + \beta_{2i} \cdot \sigma_{it-1}^2. \tag{4}$$

As is common in the literature, this GARCH specification assumes that v_{it}^2 evolves as an ARMA(1,1) process yielding a GARCH model for the scaled residual v_{it} (Hamilton, 1994, p. 666). This model is related to the short-rate stochastic volatility literature where the

[1] We also considered a model with a full VAR in which the current value of each yield factor is driven by all lagged yield factors. This alternative specification introduces six extra parameters but only marginally improves the value of the log-likelihood function.

[2] In our model, we do not explicitly enforce the no-arbitrage restriction since we do not specify the price of risk. It is possible that Eq. (1) is consistent with no-arbitrage in some economy but a formal proof is beyond the scope of this paper. We favor modelling flexibility over formally imposing the no-arbitrage constraint. Further, to the extent that no-arbitrage is satisfied in the data – as is likely for US Treasury bonds – it will also be satisfied in our estimated model (Diebold and Li, 2006).

[3] As pointed out by a referee, although we model volatility as the product of the level of the short rate and fitted volatility from a GARCH model, volatility is a deterministic function of the history of bond yields since GARCH models are deterministic. A possible extension would be to allow the volatility to have its own independent source of risk, like in the "unspanned stochastic volatility" literature (Collin-Dufresne and Goldstein, 2002).

conditional volatility of the short rate is modeled as $\sigma_t^2 F_{t-1}^{2\gamma}$ and the conditional volatility of the scaled residual σ_t^2 follows an autoregressive process. Here, we also model the conditional volatility of the scaled residual but use a GARCH model. Our specification for σ_{it}^2 is related to the univariate model of short-rate volatility developed by Brenner et al. (1996):

$$\sigma_t^2 = \beta_0 + \beta_1 \cdot e_{t-1}^2 + \beta_2 \cdot \sigma_{t-1}^2. \tag{5}$$

This specification models the variance of v_t as a function of lagged values of e_{t-1}^2 which is inconsistent with the traditional motivation of the GARCH model. Our model is also slightly different from Longstaff and Schwartz (1992), Brenner et al. (1996) and Gray (1996) who model the variance of e_{t-1} as a GARCH model augmented with the level effect:

$$\delta_t^2 = \beta_0 + \beta_1 \cdot e_{t-1}^2 + \beta_2 \cdot \delta_{t-1}^2 + \beta_3 \cdot F_{Lt-1}^{2\gamma_L}. \tag{6}$$

We estimate the yield-factor volatility models using the constant conditional correlation model of Bollerslev (1990). We assume that ΔF_t is a tri-dimensional vector of the changes in yield factors with conditional mean vector $\mu_t = E(\Delta F_t|\psi_{t-1})$ and conditional covariance matrix $\Sigma_t = H_t^{1/2}\rho H_t^{1/2}$, where ρ is a (3×3) conditional correlation matrix and H_t is a (3×3) diagonal matrix with conditional volatility of the ith factor on the ith element of the principal diagonal. We estimate the parameter vector using quasi-maximum likelihood, where $\ln L = \sum_{t=1}^{T} \ln f(e_t|\psi_{t-1})$ is the quasi-log-likelihood function, and $f(e_t|\ \psi_{t-1})$ is the probability density function of the multivariate normal density with mean 0 and covariance matrix $\Sigma_{t|t-1}$. The initial observation is assumed to be drawn from the unconditional distribution of ΔF_t.

Many classical models for interest rates are nested in our model. Firstly, a multivariate homoscedastic-AR(1) model, labelled as the NO GARCH-NO LEVEL model, can be derived by assuming that the residual volatility of each factor is constant through time ($\beta_{ki} = 0$, $i = L, S, C$ and $k = 1,2$, and $\gamma_i = 0$, $i = L, S, C$). Secondly, a multivariate version of the CKLS model, which is called the LEVEL model, is obtained by assuming that σ_{it}^2 is constant. In the latter model, the volatility remains time-varying but depends solely on the level of the factor ($\beta_{ki} = 0, i = L, S, C$ and $k = 1,2$). Thirdly, a multi-factor model, labelled as the GARCH model, allows σ_{it}^2 to follow a GARCH process but does not permit volatility to be a function of the level of the factors ($\gamma_i = 0, i = L, S, C$). Finally, the unrestricted version of the model, which is referred to as the GARCH-LEVEL model, permits σ_{it}^2 to vary through time as new information arrives and the residual volatility to depend on the level of the factors.

2.2. Regime-switching models

We extend our basic model to allow for different regimes in the volatility of the yield factors. There are economic reasons to believe that the entire yield curve is subject to regime shifts caused by transitions between periods of economic expansion and contraction. This approach is also motivated by the extensive empirical literature suggesting that regime-switching models describe historical interest rates better than single-regime models (Hamilton, 1988; Gray, 1996; Smith, 2002). Further, Ang and Bekaert (2002a,b), Bansal and Zhou (2002), and Dai et al. (2006) show that regime shifts are also important in capturing the dynamics of interest rates using multi-factor term structure models.

We denote by S_t the random state of the world at time t which can take two values, $s_t = \{1,2\}$, where 1 denotes the "high-volatility regime" and 2 the "low-volatility regime". We assume that these regimes are common to the level, slope, and curvature of the yield curve. This assumption is primarily to keep the state space parsimonious.[4] Moreover, a simple perusal of the absolute values of the factor residuals indicates that the interesting high-volatility episodes of one series also appear in the other two series. To keep the model simple we allow only the unconditional mean and volatility of each series to be state-dependent. The conditional mean is given by

$$\Delta F_{it} = \alpha_{0is_t} + \alpha_{1i} \cdot F_{it-1} + e_{it} \tag{7}$$

and the conditional volatility of the scaled residual is

$$\hat{\sigma}^2_{it|s_t,s_{t-1}} = \beta_{0is_t} + \beta_{1i} \cdot \hat{v}^2_{it-1} + \beta_{2i} \cdot \hat{\sigma}^2_{it-1|s_{t-1}}. \tag{8}$$

This specification of the GARCH model follows Dueker (1997) by defining $\hat{\sigma}^2_{it|s_t,s_{t-1}} = E(v_t^2|S_t = s_t, S_{t-1} = s_{t-1}, \psi_{t-1})$ and

$$\hat{v}_{it-1} = \sum_{s_{t-1},s_{t-2}=1}^{2} P(S_{t-1} = s_{t-1}, S_{t-2} = s_{t-2}|\psi_{t-1}) v_{it|s_{t-1}}. \tag{9}$$

Note that for any time point t, the conditional volatility depends only on the regimes in the current period and in the previous period. The dependence of lagged volatility on states in previous periods is integrated out by substituting the entire path dependent $\sigma^2_{it-1|s_{t-1},s_{t-2},\ldots}$ with $\hat{\sigma}^2_{it-1|s_{t-1}} = E(v^2_{t-1}|S_{t-1} = s_{t-1}, \psi_{t-1})$:

$$\hat{\sigma}^2_{it|s_t} = \sum_{s_{t-1}=1}^{2} P(S_{t-1} = s_{t-1}|S_t = s_t, \psi_{t-1}) \sigma^2_{it|s_t,s_{t-1}} \tag{10}$$

$$P(S_{t-1} = s_{t-1}|S_t = s_t, \psi_{t-1}) = \frac{P(S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1})}{\sum_{s_{t-1}=1}^{2} P(S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1})}. \tag{11}$$

The transition between the two latent states is modeled as a first-order Markov process with constant transition probabilities.[5] We define $\xi_{t|\tau}$ as

$$\xi_{t|\tau} = \begin{bmatrix} P(S_t = 1, S_{t-1} = 1|\psi_\tau) \\ P(S_t = 1, S_{t-1} = 2|\psi_\tau) \\ P(S_t = 2, S_{t-1} = 1|\psi_\tau) \\ P(S_t = 2, S_{t-1} = 2|\psi_\tau) \end{bmatrix} \tag{12}$$

with ψ_τ denoting three possible information sets. For $\tau = t-1$, we get the forecast probabilities, which are used to construct the log-likelihood function; for $\tau = t$, we get the filtered probabilities, which are a product of the updating algorithm; for $\tau = T$, we get the full-sample smoothed probabilities, which use all information and are helpful when making inference regarding states. The transition matrix from one point to another is

[4] Without this assumption, the state space would enlarge to $2^3 = 8$ regimes. A similar assumption is made by Kugler (1996), Ang and Bekaert (2002a,b), and Christiansen (2004).

[5] Alternatively, the transition probabilities may depend on the level of interest rates (Gray, 1996). However, as acknowledged by Ang and Bekaert (2002a, p. 172), multi-factor models with time-varying transition probabilities are likely to be overparameterized, which leads to many insignificant coefficients in the probability terms.

given by $\xi_{t|t-1} = P\xi_{t-1|t-1}$. Imbedded in this formula is the property that the previous regime is integrated out at each point in time. The transition matrix P is given by

$$P = \begin{bmatrix} p & p & 0 & 0 \\ 0 & 0 & 1-q & 1-q \\ 1-p & 1-p & 0 & 0 \\ 0 & 0 & q & q \end{bmatrix}, \tag{13}$$

where $p = P(S_t = 1|S_{t-1} = 1)$ and $q = P(S_t = 2|S_{t-1} = 2)$. We follow Hamilton (1994) and set the initial probability vector $\xi_{1|0}$ to the ergodic steady state probabilities. The conditional density $f(\Delta F_t|S_t = s_t, S_{t-1} = s_{t-1}, \psi_{t-1})$ is a multivariate normal density with conditional mean $\mu_{t|s_t,s_{t-1},t-1} = \{\Delta\hat{F}_{it}\}_{i=L,S,C}$ and conditional covariance matrix $\hat{\Sigma}_{t|s_t,s_{t-1},t-1} = \hat{H}^{1/2}_{t|s_t,s_{t-1}}\rho\hat{H}^{1/2}_{t|s_t,s_{t-1}}$. Although the states are latent, the forecast probabilities can be used to calculate the joint density of ΔF_t and the states as

$$\begin{aligned} f(\Delta F_t, S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1}) &= f(\Delta F_t|S_t = s_t, S_{t-1} = s_{t-1}, \psi_{t-1}) \\ &\quad \times P(S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1}). \end{aligned} \tag{14}$$

The marginal density of ΔF_t is found by integrating the joint density of ΔF_t over all possible states and is given by $f(\Delta F_t|\psi_{t-1}) = \sum_{s_t,s_{t-1}=1}^{2} f(\Delta F_t, S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1})$. The log-likelihood function is calculated as $\ln L = \sum_{t=1}^{T} \log f(\Delta F_t|\psi_{t-1})$ and is maximized to estimate the parameters. Finally, the updated filter probabilities of the latent states (the appropriate elements of $\xi_{t|t}$) can be obtained using the definition of the conditional probability:

$$P(S_t = s_t, S_{t-1} = s_{t-1}|\psi_t) = \frac{f(\Delta F_t, S_t = s_t, S_{t-1} = s_{t-1}|\psi_{t-1})}{f(\Delta F_t|\psi_{t-1})}. \tag{15}$$

The various models fitted in this section recognize diverse sources of conditional heteroscedasticity in the yield factors: (1) In the Regime-Switching (RS)-NO GARCH-NO LEVEL model, conditional heteroscedasticity can only be driven by switches between regimes ($\beta_{ki} = 0$, $i = L, S, C$ and $k = 1, 2$, and $\gamma_i = 0$, $i = L, S, C$); (2) In the RS-LEVEL model, conditional heteroscedasticity comes from either time-variation in the level of interest rates or from switches between regimes ($\beta_{ki} = 0$, $i = L, S, C$ and $k = 1, 2$); (3) In the RS-GARCH model, conditional heteroscedasticity is driven by serial correlation in volatility or by switches between regimes ($\gamma_i = 0$, $i = L, S, C$); (4) In the RS-GARCH-LEVEL model, conditional heteroscedasticity comes from the three different sources of time-variation.

3. Data and specification tests

We use the Fama and Bliss (1987) monthly data on US Treasury zero-coupon bond yields over the 1970:01–2002:12 period. We consider monthly data to get a long enough sample period, in particular for the 10-year yield. Note that the longest series of daily and weekly data is from the FRED database and only starts in 1982 for the 3-month yield, and the CRSP Daily US Treasury Database does not include the 10-year yield. We denote by $y_t^{(\tau)}$ the zero-coupon bond yield with a τ-month maturity observed at time t. We associate the level factor with the 3-month yield ($F_{Lt} = y_t^{(3)}$), the slope factor with the difference between the 120-month yield and the 3-month yield ($F_{St} = y_t^{(120)} - y_t^{(3)}$), and the curvature

Table 1
Descriptive statistics

	Yield factors (F_i)			Yield factor changes (ΔF_i)		
	Level	Slope	Curvature	Level	Slope	Curvature
Mean	6.495	1.359	−0.043	−0.017	0.008	0.006
Variance	7.756	2.040	0.571	0.373	0.273	0.213
Skewness	1.047	−0.597	0.044	−1.727	1.181	−0.345
Kurtosis	4.436	3.193	4.030	16.459	10.656	8.496
BJ	106.39	24.11	17.62	2981.14	964.75	497.07
Corr(F_L,F_i)	–	−0.618	−0.407			
Corr(F_S,F_C)	–	–	0.140			
Corr(ΔF_L,Δ F_i)				–	−0.807	−0.021
Corr(ΔF_S,ΔF_C)				–	–	−0.185
CACorr(Level$_{t-1}$,i)	0.976	−0.576	−0.395	0.129	−0.091	−0.019
CACorr(Slope$_{t-1}$,i)	−0.594	0.933	0.134	−0.050	0.036	−0.067
CACorr(Curv$_{t-1}$,i)	−0.408	0.182	0.812	−0.223	0.170	−0.169
BP$_1$	376.24	345.01	258.20	6.56	0.50	11.34
BP$_{12}$	3325.19	2063.31	1051.63	47.60	24.91	51.51

Notes: This table presents the mean, variance, skewness, and kurtosis of the three yield factors (level F_L, slope F_S, curvature F_C) and yield-factor changes (level ΔF_L, slope ΔF_S, curvature ΔF_C). BJ stands for the Bera–Jarque normality test, Corr for correlation, CACorr for first-order cross-autocorrelation, and BP$_1$ and BP$_{12}$ for the Box–Pierce test with one and twelve lags respectively. The latter two statistics are distributed as chi-squared with 1 and 12 degrees of freedom and the 5% critical values are 3.84 and 21.03, respectively.

factor with a linear transformation of the short, medium, and long-term yields ($F_{Ct} = y_t^{(3)} - 2y_t^{(24)} + y_t^{(120)}$).[6]

Table 1 presents some descriptive statistics for the yield factors and the yield-factor changes. For each series, we provide the first four moments, the Bera–Jarque normality test, the correlation with other factors (or factor changes), the first-order (cross-)autocorrelation, and the Box–Pierce statistics to test for the *k*th-order autocorrelation of the series. We observe that the factor series are strongly autocorrelated and depart from normality. The level factor turns out to be negatively correlated with the other two factors, while the slope and curvature factors exhibit a positive correlation. The factor changes are negatively correlated among each other and are far from being normal since their distributions are clearly leptokurtic. Both the autocorrelation estimates and the Box–Pierce test suggests that the fact or changes are much less persistent than the factor levels.

The time series of each factor is plotted in the left three panels of Fig. 1. While the level factor is always positive, both the slope and curvature factors are occasionally negative. There are several episodes when the yield curve is downward-sloping: the 1973 OPEC oil crisis, during a significant portion of the 1979–1982 monetary experiment (i.e., period during which the Federal Reserve focused primarily on reducing the rate of growth of monetary aggregates, rather than targeting interest rates, in an effort to reduce inflation), 1989, and towards the end of 2000. Note also that the monetary experiment had a great impact on the curvature of the term structure. The panels on the right of Fig. 1 display the absolute

[6] We use the 3-month yield to proxy the level factor because the evolution of the 1-month yield is known to be idiosyncratic (Duffee, 1996). The 2-year maturity is an important intermediate maturity given the fact that the term structure of volatility peaks at this maturity (Dai and Singleton, 2003). We proxy the long-term rate by the 10-year maturity since this yield is less subject to liquidity problems than longer-maturity yields.

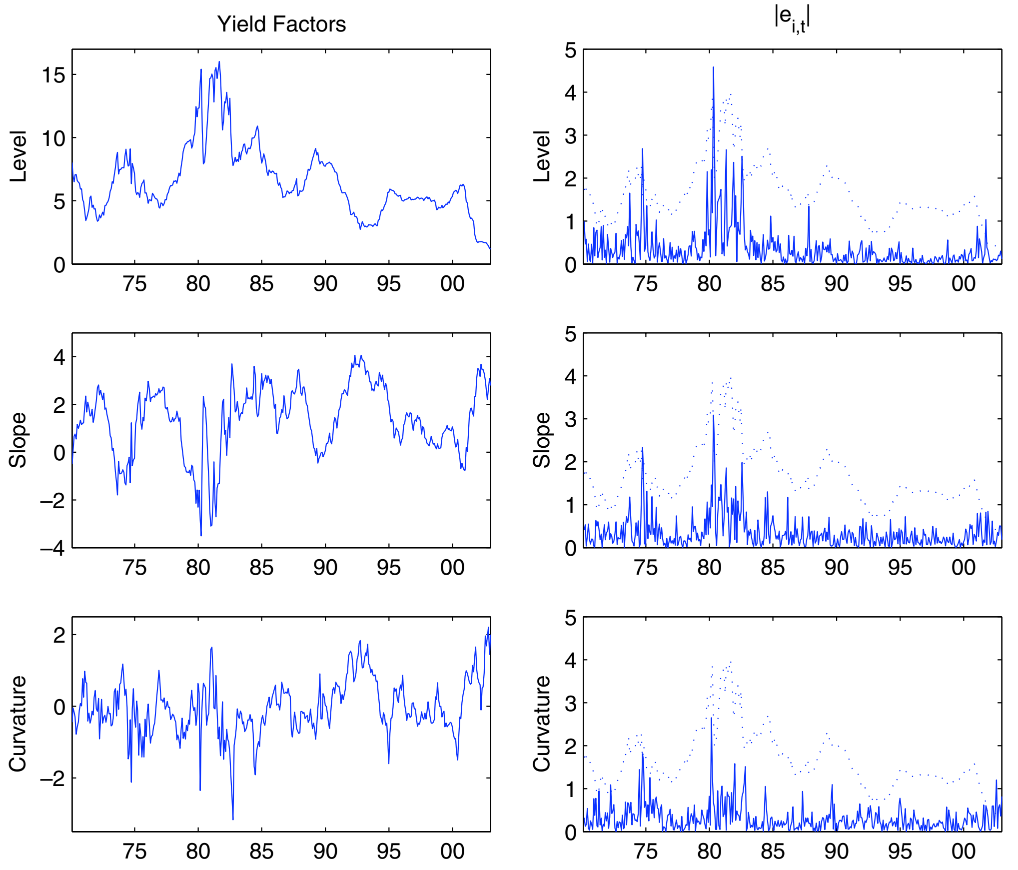

Fig. 1. Value and volatility of the level, slope, and curvature factors. *Notes*: the three panels on the left display the values of the three yield factors. The level factor is associated with the 3-month yield, the slope factor with the difference between the 120-month yield and the 3-month yield, and the curvature factor with a linear transformation of the short (1-month), medium (24-month), and long-term (120-month) yields. The three panels on the right display the absolute value of the factor residuals (solid line) with the level factor (dashed line).

value of the factor residuals and we see that all three series exhibit volatility clustering. This suggests that an appropriate model for the yield-factor volatility should include ARCH effects. Moreover, factor volatilities appear to depend on the level of interest rates as attested by the superimposed level series. Indeed, the volatility of both the slope and curvature factors tends to be high when interest rates are high. This suggests that a diffusion model, in which volatility is a positive function of the level of interest rates, may be able to account for this effect.

4. Empirical results

4.1. Single-regime models

In this section, we report the results of fitting the competing models, i.e., the NO LEVEL-NO GARCH, LEVEL, GARCH, and GARCH-LEVEL models, to the US term structure of interest rates over the 1970–2002 period. We begin by estimating the univariate version of the four models. Table 2 reports for each yield factor the parameter estimates and the Bollerslev and Wooldridge (1992) robust standard-errors. The first column of Table 2 reports the estimates of the homoscedastic model. There appears to be mean reversion in all three yield factors ($\hat{\alpha}_{1i} < 0$, $i = L, S, C$), though not statistically significant for the level. When a level effect is introduced, the log-likelihood function increases dramatically ($\Delta \ln L_L = 136.89$, $\Delta \ln L_S = 53.68$, and $\Delta \ln L_C = 11.40$). The elasticity parameter on the level factor $\hat{\gamma}_L$ is equal to 1.202, which is slightly less than the point estimates in CKLS (1992) and Brenner et al. (1996). Interestingly, $\hat{\gamma}_S$ is about half of $\hat{\gamma}_L$, and, in turn, $\hat{\gamma}_C$ is half of $\hat{\gamma}_S$. The elasticity parameter is significant for all three yield factors. If the level effect is modeled using the level of each factor, instead of the overall level of interest rates, the fit of the LEVEL model is significantly reduced ($\ln L_S = -297.58$ and $\ln L_C = -234.91$, not reported in the tables). Results of the GARCH model suggest that explicitly modeling the serial correlation in volatility leads to a superior fit. Further, the variance processes exhibit high persistence, though the persistence is lower for the slope and curvature factors ($\hat{\beta}_1 + \hat{\beta}_2 = 0.990$ for the level, 0.932 for the slope, and 0.907 for the curvature). Finally, the GARCH-LEVEL model gives rise to lower estimates of the elasticity parameters than in the LEVEL model. The marginal contribution of the GARCH effect turns out to be stronger than the marginal contribution of the LEVEL effect for all three yield factors.

Table 3 reports the parameter estimates and robust standard-errors for the multivariate models. The main difference between the specifications presented in Table 2 and the present specifications is that the residual factors can now be correlated. The point estimates for the correlation coefficients ($\rho_{L,S}$, $\rho_{L,C}$, and $\rho_{S,C}$) are negative in all models, while their magnitude varies across models. Allowing the residuals to be correlated improves the fit of each model. For example, the sum of the log-likelihoods for the three univariate GARCH models is equal to −602.24 and the log-likelihood for the trivariate GARCH model is as high as −477.18. However, correlation matters not only for fitting purposes but it also strongly impacts the point estimate for the elasticity parameters. Indeed, for the LEVEL and GARCH-LEVEL models, the point estimate of γ_i drops considerably for the level and slope factors after accounting for correlation. On the other hand, the point estimate of γ_i increases for the curvature factor.

Comparing the various multivariate volatility models yields some interesting conclusions. Consistent with the univariate results, the GARCH effect seems to dominate the level effect. When only a GARCH effect is introduced, the value of the log-likelihood function increases by 193, which exceeds the increase observed when a level effect is introduced (i.e., 126). However, the level effect only requires three extra parameters while the GARCH model requires six extra parameters. Another interesting observation is that the value of the elasticity parameter is weakened when a GARCH effect is introduced, though it remains significant for the level and slope factors. In conclusion, it seems that one needs both level and GARCH effects to adequately model yield-factor volatilities.

4.2. Regime-switching models

Parameter estimates and robust-standard errors for the regime-switching models are reported in Table 4. Regime-switching models outperform the models estimated in Table 3, which is not overly surprising given that the latter models have been estimated under the

Table 2
Parameter estimates for univariate models

	NO LEVEL-NO GARCH			LEVEL			GARCH			GARCH-LEVEL		
	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.
α_0	0.135	0.089	−0.005	−0.000	0.074	0.013	0.114*	0.042	−0.015	0.033*	0.033	−0.016
	(0.109)	(0.056)	(0.021)	(0.012)	(0.039)	(0.020)	(0.055)	(0.043)	(0.018)	(0.016)	(0.039)	(0.018)
α_1	−0.022	−0.066*	−0.182**	− 0.002*	−0.045**	−0.149**	− 0.023	−0.042*	−0.112**	−0.005*	−0.036	−0.115**
	(0.020)	(0.028)	(0.042)	(0.001)	(0.017)	(0.038)	(0.013)	(0.021)	(0.025)	(0.002)	(0.020)	(0.026)
β_0	0.368**	0.263**	0.194**	2.521*	20.40**	63.64**	5.894	15.95	23.61**	0.144	1.720	14.07
	(0.070)	(0.038)	(0.025)	(1.008)	(6.458)	(23.87)	(3.435)	(9.862)	(7.473)	(0.192)	(1.891)	(15.77)
β_1	–	–	–	–	–	–	0.294**	0.232*	0.352**	0.226*	0.160*	0.335**
							(0.079)	(0.102)	(0.100)	(0.109)	(0.073)	(0.091)
β_2	–	–	–	–	–	–	0.696**	0.700**	0.555**	0.752**	0.774**	0.562**
							(0.080)	(0.118)	(0.078)	(0.115)	(0.102)	(0.080)
γ	–	–	–	1.202**	0.641**	0.296**	–	–	–	1.106**	0.575*	0.156
				(0.107)	(0.087)	(0.105)				(0.334)	(0.230)	(0.340)
$\ln L_i$	−365.59	−298.68	−237.60	−228.70	−245.00	−226.20	−199.81	−216.15	−186.28	−185.27	−210.32	−185.94

Notes: The general model used is the GARCH-LEVEL model:

$$\Delta F_{it} = \alpha_{0i} + \alpha_{1i} \cdot F_{it-1} + e_{it}$$

for $i = L,S,C$ and $t = 1, \ldots, T$. The conditional volatility of e_{it} is $E(e_{it}^2|\psi_{t-1}) = \sigma_{it}^2 F_{Lt-1}^{2\gamma_i}$. We model σ_{it}^2, the volatility of the scaled residual $v_{it} = e_{it}/Y_{Lt-1}^{\gamma_i}$, as

$$\sigma_{it}^2 = \beta_{0i} + \beta_{1i} \cdot v_{it-1}^2 + \beta_{2i} \cdot \sigma_{it-1}^2.$$

The restricted versions of the model impose the following constraints: NO LEVEL-NO GARCH model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$, and $\gamma_i = 0$, $i = L,S,C$; LEVEL model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$; GARCH model: $\gamma_i = 0$, $i = L,S,C$. For each univariate model, we report the maximum likelihood parameter estimates, the value of the log-likelihood function, and Bollerslev–Wooldridge robust standard errors in parentheses. β_0 parameters are multiplied by 1000 here and in Tables 3 and 4.

** and * denote statistical significance at the 1% and 5% levels, respectively.

Table 3
Parameter estimates for multivariate models

	NO GARCH-NO LEVEL			LEVEL			GARCH			GARCH-LEVEL		
	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.
α_0	0.123*	0.074*	−0.002	0.027	0.055*	0.015	0.120*	0.043	−0.011	0.046	0.025	−0.013
	(0.058)	(0.035)	(0.021)	(0.027)	(0.028)	(0.020)	(0.056)	(0.039)	(0.0174)	(0.046)	(0.021)	(0.017)
α_1	−0.021	−0.054**	−0.159**	−0.007	−0.036**	−0.144**	−0.022	−0.039*	−0.117**	−0.007	−0.031*	−0.119**
	(0.011)	(0.015)	(0.042)	(0.005)	(0.011)	(0.037)	(0.013)	(0.016)	(0.026)	(0.008)	(0.013)	(0.027)
β_0	0.368**	0.263**	0.194	6.002**	38.62**	54.80**	7.902	9.154	21.15**	0.562	4.797	9.327
	(0.070)	(0.038)	(80.03)	(1.959)	(10.61)	(20.11)	(7.082)	(11.159)	(6.865)	(0.494)	(5.292)	(10.77)
β_1	–	–	–	–	–	–	0.141	0.096	0.355**	0.091	0.090	0.338**
							(0.114)	(0.066)	(0.098)	(0.055)	(0.067)	(0.092)
β_2	–	–	–	–	–	–	0.812**	0.852 **	0.570**	0.846**	0.826**	0.572**
							(0.135)	(0.121)	(0.078)	(0.090)	(0.152)	(0.086)
γ	–	–	–	0.969**	0.468**	0.338**	–	–	–	0.800**	0.314	0.254
				(0.080)	(0.073)	(0.103)				(0.187)	(0.167)	(0.342)
$\rho_{L,S}$	−0.808 **	–	–	−0.712**	–	–	−0.661**	–	–	−0.651**	–	–
	(0.032)			(0.035)			(0.038)			(0.036)		
$\rho_{i,C}$	−0.047	−0.155	–	−0.107	−0.161*	–	−0.162*	−0.132	–	−0.175**	−0.133	–
	(0.100)	(0.096)		(0.083)	(0.079)		(0.065)	(0.073)		(0.062)	(0.070)	
$\ln L$		−670.637			−544.633			−477.183			−458.124	

Notes: The general model used is the GARCH-LEVEL model:

$$\Delta F_{it} = \alpha_{0i} + \alpha_{1i} \cdot F_{it-1} + e_{it}$$

for $i = L,S,C$ and $t = 1,\ldots,T$. The conditional volatility of e_{it} is $E(e_{it}^2|\psi_{t-1}) = \sigma_{it}^2 F_{Lt-1}^{2\gamma_i}$. We model σ_{it}^2, the volatility of the scaled residual $v_{it} = e_{it}/F_{Lt-1}^{\gamma_i}$, as

$$\sigma_{it}^2 = \beta_{0i} + \beta_{1i} \cdot v_{it-1}^2 + \beta_{2i} \cdot \sigma_{it-1}^2.$$

The restricted versions of the model impose the following constraints: NO LEVEL-NO GARCH model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$, and $\gamma_i = 0$, $i = L,S,C$; LEVEL model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$; GARCH model: $\gamma_i = 0$, $i = L,S,C$. For each multivariate model, we report the maximum likelihood parameter estimates, the value of the log-likelihood function, and Bollerslev–Wooldridge robust standard errors in parentheses.

** and * denote statistical significance at the 1% and 5% levels, respectively.

Table 4
Parameter estimates for regime-switching models

	RS-NO GARCH-NO LEVEL			RS-LEVEL			RS-GARCH			RS-GARCH-LEVEL		
	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.	Level	Slope	Curv.
$\alpha_{0,S1}$	−0.004	0.087	−0.097	−0.013	0.108*	0.052	−0.017	0.125	−0.065	−0.019	0.142	0.014
	(0.158)	(0.120)	(0.091)	(0.031)	(0.047)	(0.060)	(0.097)	(0.079)	(0.053)	(0.071)	(0.079)	(0.064)
$\alpha_{0,S2}$	0.059	0.024	0.019	0.099	0.006	−0.002	0.071**	0.011	0.001	0.118	0.014	−0.008
	(0.040)	(0.022)	(0.016)	(0.067)	(0.020)	(0.024)	(0.007)	(0.011)	(0.015)	(0.234)	(0.015)	(0.021)
α_1	−0.009	−0.026*	−0.132**	− 0.012	−0.030*	−0.144**	− 0.010**	−0.024**	−0.117**	−0.017	−0.033	−0.128**
	(0.007)	(0.011)	(0.029)	(0.009)	(0.012)	(0.040)	(0.001)	(0.009)	(0.024)	(0.032)	(0.020)	(0.039)
$\beta_{0,S1}$	1171.6**	690.2**	665.8**	5.775**	45.36**	87.75*	413.0**	313.8**	268.0**	5.131*	44.46*	62.06
	(139.3)	(103.8)	(122.1)	(2.212)	(14.29)	(37.18)	(81.36)	(61.28)	(97.53)	(2.283)	(18.50)	(34.38)
$\beta_{0,S2}$	100.9**	104.9**	93.32**	1.137*	10.88**	12.41*	5.256	18.80	11.27	0.190	5.737	2.926
	(15.94)	(9.720)	(11.13)	(0.500)	(3.701)	(6.139)	(3.485)	(14.05)	(5.832)	(0.255)	(5.610)	(2.436)
β_1	–	–	–	–	–	–	0.240*	0.094*	0.315**	0.137	0.046	0.251**
							(0.098)	(0.047)	(0.095)	(0.077)	(0.046)	(0.082)
β_2	–	–	–	–	–	–	0.116**	0.186**	0.138**	0.194**	0.144	0.128**
							(0.022)	(0.048)	(0.046)	(0.053)	(0.109)	(0.044)
γ	–	–	–	1.125**	0.560**	0.414**	–	–	–	1.038**	0.503**	0.384**
				(0.103)	(0.088)	(0.118)				(0.115)	(0.097)	(0.139)
$\rho_{L,S}$	−0.636**	–	–	−0.607**	–	–	−0.665**	–	–	−0.610**	–	–
	(0.036)			(0.038)			(0.040)			(0.039)		
$\rho_{i,C}$	−0.216**	−0.096	–	− 0.163**	−0.179**	–	−0.205**	−0.125	–	−0.198**	−0.149*	–
	(0.056)	(0.059)		(0.062)	(0.063)		(0.060)	(0.065)		(0.060)	(0.059)	
p_{S1}		2.147**			2.002**			2.441**			2.253**	
		(0.519)			(0.396)			(0.650)			(0.522)	
p_{S2}		3.735**			2.418**			3.329**			2.739**	
		(0.615)			(0.352)			(0.514)			(0.510)	
$\ln L$		−463.894			−417.453			−445.539			−408.805	

Notes: The general model used is the Regime-Switching GARCH-LEVEL model:

$$\Delta F_{it} = \alpha_{0is_t} + \alpha_{1i} \cdot F_{it-1} + e_{it}$$

for $i = L,S,C$ and $t = 1,\ldots,T$, where $s_t = 1$ denotes the "high-volatility regime" and $s_t = 2$ the "low-volatility regime". The conditional volatility of e_{it} is $E(e_{it}^2|\psi_{t-1}) = \hat{\sigma}^2_{it|s_t,s_{t-1}} F_{Lt-1}^{2\gamma_i}$ where $\hat{\sigma}^2_{it|s_t,s_{t-1}} = E(v_t^2|S_t = s_t, S_{t-1} = s_{t-1}, \psi_{t-1})$ is modeled as

$$\hat{\sigma}^2_{it|s_t,s_{t-1}} = \beta_{0is_t} + \beta_{1i} \cdot v_{it-1}^2 + \beta_{2i} \cdot \hat{\sigma}^2_{it-1|s_{t-1}}.$$

The probability of staying in state j is modeled as $P(S_t = j|S_{t-1} = j) = \exp(p_{Sj})/(1 + \exp(p_{Sj}))$. The restricted versions of the model impose the following constraints: RS-NO LEVEL-NO GARCH model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$, and $\gamma_i = 0$, $i = L,S,C$; RS-LEVEL model: $\beta_{ki} = 0$, $i = L,S,C$ and $k = 1,2$; RS-GARCH model: $\gamma_i = 0$, $i = L,S,C$. For each model, we report the maximum likelihood parameter estimates, the value of the log-likelihood function, and Bollerslev–Wooldridge robust standard errors in parentheses.

** and * denote statistical significance at the 1% and 5% levels, respectively.

assumption of only one regime. Allowing for multiple regimes dramatically improves the fit of all four models. Interestingly, when volatility is allowed to switch between low and high-volatility regimes, the level effect is strengthened and the volatility persistence drops significantly. Further, the performance of the RS-LEVEL model is higher than the performance of the RS-GARCH model, whereas the single-regime GARCH model outperforms the single-regime LEVEL model.[7] Because of its lack of parsimony (it requires the estimation of 26 parameters) and its log-likelihood value, the RS-GARCH level is dominated by the RS-LEVEL model according to the Bayesian information criterion ($\text{BIC}_{\text{RS-LEVEL}} = 972.48$ vs. $\text{BIC}_{\text{RS-GARCH}} = 1046.59$). In the same way, the RS-LEVEL model is also preferred to the general RS-GARCH-LEVEL model ($\text{BIC}_{\text{RS-GARCH-LEVEL}} = 991.07$).

Given the highly parameterized nature of the RS-GARCH-LEVEL model, it is natural to question the parameter stability because of concerns of possible in-sample overfitting. To test the stability of the parameter estimates, we estimate the model in two non-overlapping equal-sized subsamples of 16.5 years. We find that the parameter estimates are remarkably similar in each subperiod.[8] Some differences between subperiods are worth mentioning though. All level effects are slightly larger in the first subperiod, which is not surprising since this period includes the monetary experiment. Further, there is essentially no level effect for the curvature factor in the second subperiod.

The four panels of Fig. 2 contain plots of the smoothed probabilities of the high-volatility state for the four regime-switching models we consider. The probabilities have been computed using the smoothing algorithm of Kim (1994). Because of their multi-factor nature, our models exploit complementary information on the slope and curvature of the term structure. Our models identify all the major well-known episodes of extreme volatility: the 1973 OPEC oil crisis and its aftermath, the 1979–1982 monetary experiment, the October 1987 stock market crash, and the Russian Ruble devaluation in August 1998. Interestingly, the two models that include a level effect identify two additional high-volatility episodes. The first high volatility episode is intermittent and occurs during the 1991–1994 period in which the Federal Reserve held short rates down to combat weakness in the economy and then raised its federal funds rate by 125 basis points during the first four months of 1994 (Campbell, 1995). The second high volatility regime starts in 2001 and coincides with the massive mortgage refinancing wave that occurred in the US (Duarte, forthcoming).

4.3. Out-of-sample volatility forecasting

Forecasting the covariance matrix of the bond yields is quite challenging in practice given the large dimension of a typical term structure. Indeed, standard multivariate econometric models (e.g. GARCH) require the estimation of a large number of parameters when the number of yields N is large. We show in this section how the yield-factor models presented in this paper can be used to forecast yield volatilities and covariances in a parsimonious way. Recall that our yield-factor models are based on $K = 3$ yield factors, $F_t^\top = [F_{Lt}\ F_{St}\ F_{Ct}]$, that are linearly related to three yields, the short ($y_t^{(3)}$), intermediate ($y_t^{(24)}$), and long ($y_t^{(120)}$) yields. Consequently, these three yields are perfectly linearly related

[7] The finding that the RS-GARCH model is dominated by the RS-LEVEL model is consistent with Christiansen (2005b).

[8] To conserve space, we omit the subperiod parameter estimates, but these are available from the authors upon request.

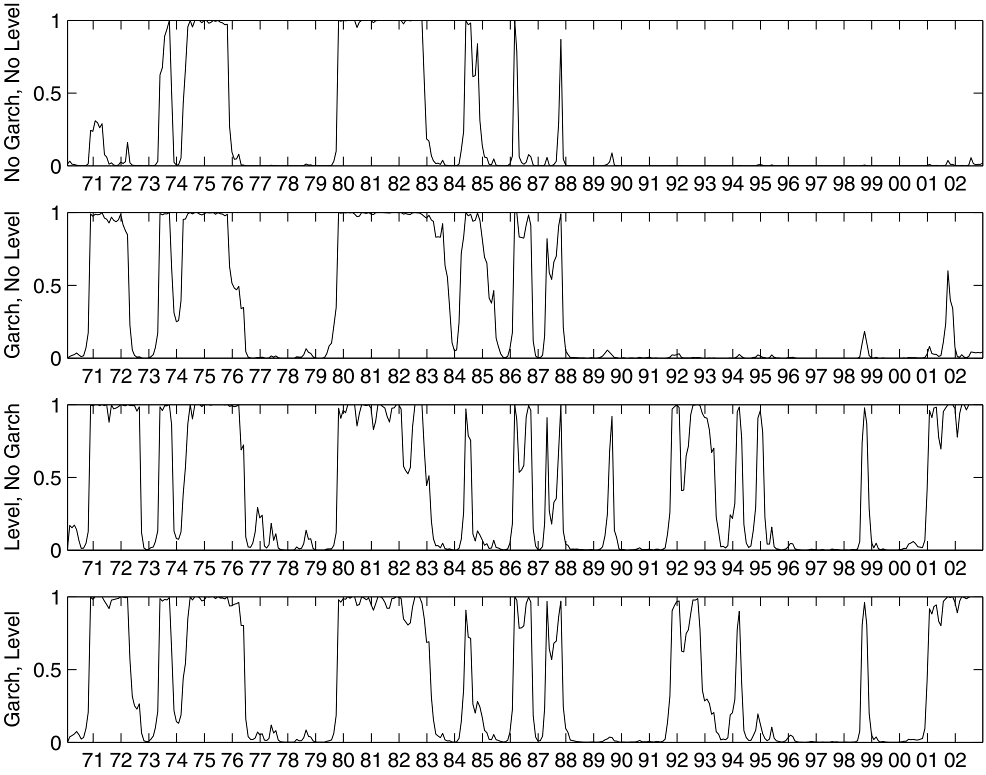

Fig. 2. Smoothed probability of high-volatility state $P(S_t = 1|\psi_T)$. *Notes*: The four panels contain the time series of the smoothed probabilities that the level factor is in the high-volatility regime at time *t* according to the Regime Switching (RS) NO GARCH-NO LEVEL model, the RS-LEVEL model, the RS-GARCH model, and the RS-GARCH-LEVEL model. The smoothed probability is based on the entire sample: $P(S_t = 1|\psi_T)$.

to the yield factors whereas yields with different maturities are imperfectly linearly related to the yield factors.

We assess the ability of the yield-factor model to forecast yield volatility using six US Treasury zero coupon bond yields: the three yields used to construct the yield factors, as well as the 6-month $y_t^{(6)}$, the 12-month $y_t^{(12)}$, and the 5-year $y_t^{(60)}$ yields. Thus, each yield can be expressed as a function of the yield factors:[9]

$$
\begin{aligned}
y_t^{(3)} &= [1 \quad 0 \quad 0] \cdot F_t \\
y_t^{(6)} &= \left[c_L^{(6)} \quad c_S^{(6)} \quad c_C^{(6)} \right] \cdot F_t + \eta_t^{(6)} \\
y_t^{(12)} &= \left[c_L^{(12)} \quad c_S^{(12)} \quad c_C^{(12)} \right] \cdot F_t + \eta_t^{(12)} \\
y_t^{(24)} &= [1 \quad 0.5 \quad -0.5] \cdot F_t \\
y_t^{(60)} &= \left[c_L^{(60)} \quad c_S^{(60)} \quad c_C^{(60)} \right] \cdot F_t + \eta_t^{(60)} \\
y_t^{(120)} &= [1 \quad 1 \quad 0] \cdot F_t.
\end{aligned}
$$

[9] The derivation of the fourth equation is as follows: $[1 \quad 0.5 \quad -0.5] \cdot F_t = F_{Lt} + 0.5 \cdot F_{St} - 0.5 \cdot F_{Ct} = y_t^{(3)} + 0.5 \cdot (y_t^{(120)} - y_t^{(3)}) - 0.5 \cdot (y_t^{(3)} - 2y_t^{(24)} + y_t^{(120)}) = y_t^{(24)}$.

Table 5
Out-of-sample volatility forecasting

	NO LEVEL-NO GARCH	LEVEL	GARCH	GARCH-LEVEL
Single regime	220.11	221.44	360.93*	345.75
Regime switching	352.22	321.65	387.57	412.29*

Note: This table presents the value of the out-of-sample log-likelihood functions of each model. We estimate all parameters using the first half of our sample period (16.5 years) and compute the log-likelihood function over the second half of the sample holding all parameters constant. *denotes the highest log-likelihood function value across models.

Rewriting this system of equations using matrix notation, we obtain $Y_t = C \cdot F_t + \eta_t$. We now have a link between the conditional covariance matrix of the yields and the conditional covariance matrix of the yield factors:

$$\mathrm{Cov}_{t-1}(Y_t) = C \cdot \mathrm{Cov}_{t-1}(F_t) \cdot C^\top + \Omega \tag{16}$$

where $\Omega = \mathrm{Cov}(\eta_t)$ is the sparse residual covariance matrix, i.e., most of its elements are zeros. In our case, since $N = 6$ and $K = 3$, there are only six unique non-zero terms in Ω.

In a first step, we estimate all the parameters of each single-regime yield-factor volatility model, as well as their regime-switching counterparts, the factor loading covariance matrix C, and the residual covariance matrix Ω using the first 16.5 years of our sample period.[10] To compare the model fit to the yields, we use the out-of-sample log-likelihood computed over the subsequent 16.5 years using the parameters estimated in the first 16.5 years. In particular, the yields are multivariate normally distributed with mean vector $E_{t-1}(Y_t) = C \cdot E_{t-1}(F_t)$ and conditional covariance matrix given in Eq. (16).

We report the values of the out-of-sample log-likelihoods in Table 5. This metric is useful to compare the ability of models to fit the entire covariance matrix. We see that the regime-switching model with level and GARCH effects provides the best out-of-sample forecasting performance of yield volatility. This result is due to the stability of the coefficient estimates between the first and second part of the sample period (see Section 4.2). Our results indicate that modeling volatility with a GARCH process and regime shifts leads to superior out-of-sample performance. Further, we note that the LEVEL and RS-LEVEL models work worst out-of-sample. As a comparison, we also estimate an orthogonal GARCH model following Alexander (2001). The orthogonal GARCH model assumes that the observed data can be linearly transformed into a set of uncorrelated components by means of an orthogonal matrix. The value of the out-of-sample log-likelihood function of this model is 221.25, which is smaller than the values obtained from most of the yield-factor models.

5. Implications for estimating and comparing models

A standard approach to estimate, test, and compare non-nested term structure models is to rely on the Efficient Method of Moments (hereafter EMM) of Bansal et al. (1995) and Gallant and Tauchen (1996). This estimation method can be applied to any class of term

[10] The free coefficients in C are estimated by regressing the actual yields on the yield factors. The R-squares of these regressions are all in excess of 0.995. We estimate the free parameters in Ω as the sample covariance matrix using the residuals from these regressions.

Table 6
Which features of volatility are most important?

	Single regime models	Regime-switching models
NO GARCH-NO LEVEL	−670.637	−463.894
	[12]	[20]
LEVEL	−544.633	−417.453
	[15]	[23]
GARCH	−477.183	−445.539
	[18]	[26]
GARCH-LEVEL	−458.124	−408.805
	[21]	[29]

Note: This table presents the value of the log-likelihood function and the number of estimated parameters (in square brackets) for all considered models. We can measure the improvement in statistical fit that arises from adding either a level effect, some GARCH effects, or a second volatility regime.

structure models: affine models (Duffie and Kan, 1996 and Dai and Singleton, 2000), quadratic models (Ahn et al., 2002 and Leippold and Wu, 2002), and regime-switching models (Bansal and Zhou, 2002 and Dai et al., 2006). [11]

EMM requires that an auxiliary model be used to estimate the conditional density of observed interest rates. The standard approach is to use a VAR-based semi-nonparametric ARCH model as the auxiliary model. This auxiliary model offers a high level of flexibility since it is based on a high-order Hermite series expansion. However, it does not capture the main features of the Treasury yields, which are summarized in Table 6. For each multi-factor model considered, we report the value of the log-likelihood function and the number of estimated parameters. The largest improvement in statistical fit from adding a single feature to the homoscedastic model comes from adding a second regime. The second most important feature of the data is the addition of a level effect. The third key feature of the data is the GARCH effect. However, the auxiliary model typically used in EMM estimation only captures low-order ARCH effects and by neglecting the level, GARCH and regime shifts EMM ignores these central features of the data.[12] Instead, EMM relies on a large set of scores that do not always have a clear economic interpretation (e.g., high-order Hermite polynomial parameters). Consequently, we argue that more research must be carried out to come up with an estimation technique for term structure models that captures the salient features of the yield data in a parsimonious way.

The yield-factor models presented in this paper can also provide a common platform to compare term structure models. Indeed, our yield factors are simply linear transformations of the bond yields themselves, specifically $F_t = D \cdot Y_t$ with $D = [1 \quad 0 \quad 0; \quad 1 \quad 0 \quad -1; \quad 1 \quad -2 \quad 1]$.[13] Further, in the important class of affine term structure models the yields are

[11] Recent applications of EMM to term structure models can be found in Dai and Singleton (2000), Ahn et al. (2002), Bansal and Zhou (2002), Ahn et al. (2003), Brandt and Chapman (2003), Bansal et al. (2004) and Duffee and Stanton (2004).

[12] Dai and Singleton (2000, p. 1960), argue that one may only need ARCH effects when modeling LIBOR and swap rates as opposed to Treasury yields. In their analysis of U.S. Treasury yields, Ahn et al. (2002), Bansal and Zhou (2002) and Bansal et al. (2004) only allow for ARCH effects and consider only up to five lags. Differently, one of the auxiliary models used by Ahn et al. (2003) allows for GARCH effects. Moreover, Brandt and Chapman (2003) suggest to use an auxiliary model with GARCH and level effects but they only consider one volatility regime.

[13] Notice that Y_t now only includes three maturities, $Y_t = [y_t^{(3)} y_t^{(24)} y_t^{(120)}]^\top$.

affine functions of a small number of state variables: $Y_t = A + B \cdot X_t$. Within the three-factor affine class, our yield factors are simply a linear transformation of the true latent factors: $F_t = D \cdot A + D \cdot B \cdot X_t$ or $X_t = (D \cdot B)^{-1}(F_t - D \cdot A)$ and D is independent of the model and its parameters. Understanding the dynamics of F_t tells us much about admissible term structure models. One particularly useful application is to assess the adequacy of a given structural term structure model. When we estimate a term structure model by maximum likelihood applied to the yields directly, the magnitude of the log-likelihood allows us to compare two models indicating which model fits the data better. However, the log-likelihood itself gives no guidance as to whether the model adequately captures the important features of the data. On the other hand, when estimating models by EMM the overidentifying condition indicates if a model adequately captures the dynamics of yields implied by the auxiliary model. Fortunately, the log-likelihood of our preferred yield-factor model provides a benchmark since the log-likelihood of the yield factors is related to the log-likelihood of the yields themselves by $\ln L(F_t) = \ln L(Y_t) - T \cdot \ln|D|$ where $|\cdot|$ denotes the determinant, and in our case $|D| = 2$. If the log-likelihood of a term structure model is much less than $\ln L(F_t) + T \cdot \ln|D|$ this suggests that the model fails short in capturing the main features of the data. Formal statistical comparisons can be facilitated using the non-nested likelihood ratio test of Vuong (1989).

6. Conclusion

In this paper, we develop a comprehensive model for volatility dynamics in the level, slope, and curvature of the yield curve that simultaneously includes level and GARCH effects along with regime shifts. First, we show that the level of the short rate is useful in modeling the volatility of the level, slope, and curvature factors of the yield curve. Second, there are significant GARCH effects present even after including a level effect. Third, when the volatility is allowed to switch from low to high-volatility regimes, the model's fit improves dramatically, the level effect is strengthened, and the volatility persistence drops significantly. Fourth, we show that a regime-switching model with level and GARCH effects provides the best out-of-sample forecasting performance of yield volatility.

Our empirical results have some important implications for the estimation of term structure models. In particular, the auxiliary models often used to estimate term structure models with simulation-based estimation techniques should capture the main features of the yield curve that are identified by our model: level and GARCH effects along with regime shifts. Moreover, the encouraging results obtained with our regime-switching models strengthen the need for including regime-shifts in theoretical term structure models.

Future research may develop a pricing model for interest rate derivatives in which the underlying term structure is modeled using our yield-factor model. Another potential application of our framework would be to compare the out-of-sample volatility estimates from yield-factor models against estimates from models that use interest rate derivatives to forecast volatilities.

Acknowledgement

We thank Avi Bick, Peter Christoffersen, Bob Dittmar, Jefferson Duarte, René Garcia, Eric Ghysels, Phil Gray, Peter Klein, Andrey Pavlov, Eric Renault, Elvezio Ronchetti, seminar participants at Simon Fraser University, University of Geneva, University of New South Wales, and University of Technology Sydney, and participants at the 2004 Australasian Meeting of the Econometric Society, Northern Finance Association 2004 Conference, and Canadian Econometrics Study Group 2005 Conference for their comments. We are grateful to two anonymous referees and an associate editor for helpful suggestions. We also thank Robert Bliss for providing us with data and his programs for constructing yield curves. We gratefully acknowledge financial support from the Social Sciences and Research Council of Canada. We are responsible for all errors and omissions.

References

Ahn, D.H., Dittmar, R.F., Gallant, A.R., 2002. Quadratic term structure models: Theory and evidence. Review of Financial Studies 16, 459–485.

Ahn, D.H., Dittmar, R.F., Gallant, A.R., Gao, B., 2003. Purebred or hybrid: Reproducing the volatility in term structure dynamics. Journal of Econometrics 116, 147–180.

Alexander, C., 2001. Orthogonal GARCH. In: Alexander, C. (Ed.), Mastering Risk. Prentice Hall, London, pp. 21–38.

Ang, A., Bekaert, G., 2002a. Regime switches in interest rates. Journal of Business and Economic Statistics 20, 163–182.

Ang, A., Bekaert, G., 2002b. Short rate nonlinearities and regime switches. Journal of Economic Dynamics and Control 26, 1243–1274.

Ang, A., Piazzesi, M., 2003. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. Journal of Monetary Economics 50, 745–787.

Bansal, R., Zhou, H., 2002. Term structure of interest rates with regime shifts. Journal of Finance 57, 1997–2043.

Bansal, R., Gallant, A.R., Hussey, R., Tauchen, G., 1995. Nonparametric estimation of structural models for high-frequency currency market data. Journal of Econometrics 66, 147–180.

Bansal, R., Tauchen, G., Zhou, H., 2004. Regime shifts, risk premiums in the term structure, and the business cycle. Journal of Business and Economic Statistics 22, 396–409.

Bollerslev, T., 1990. Modeling the coherence in short run nominal exchange rates: A multivariate generalized ARCH model. Review of Economics and Statistics 72, 498–505.

Bollerslev, T., Wooldridge, J.M., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. Econometric Reviews 11, 143–172.

Boudoukh, J., Richardson, M., Stanton, R., Whitelaw, R.F., 1998. The stochastic behavior of interest rates: Implications from a multifactor, nonlinear continuous-time model. Working Paper, New York University and Berkeley.

Brandt, M.W., Chapman, D., 2003. Comparing multifactor models of the term structure. Working Paper, Boston College and Duke University.

Brenner, R.J., Harjes, R., Kroner, R., 1996. Another look at models of short term interest rates. Journal of Financial and Quantitative Analysis 31, 85–107.

Campbell, J.Y., 1995. Some lessons from the yield curve. Journal of Economic Perspectives 9, 129–152.

Chan, K.C., Karolyi, G.A., Longstaff, F.A., Sanders, A.B., 1992. An empirical comparison of alternative models of the short-term interest rate. Journal of Finance 52, 1209–1227.

Christiansen, C., 2004. Regime switching in the yield curve. Journal of Futures Markets 24, 315–336.

Christiansen, C., 2005a. Multivariate term structure models with level and heteroscedasticity effects. Journal of Banking and Finance 29, 1037–1057.

Christiansen, C., 2005b. Level-ARCH short rate models with regime switching: Bivariate modeling of US and European short rates. Working Paper, Aarhus School of Business.

Christiansen, C., Lund, J., 2002. Revisiting the shape of the yield curve: The effect of interest rate volatility. Working Paper, Aarhus School of Business.

Collin-Dufresne, P., Goldstein, R.S., 2002. Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility. Journal of Finance 57, 1685–1730.

Cox, J.C., Ingersoll, J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–408.

Dai, Q., Singleton, K.J., 2000. Specification analysis of affine term structure models. Journal of Finance 55, 1943–1978.
Dai, Q., Singleton, K.J., 2003. Term structure modeling in theory and reality. Review of Financial Studies 16, 631–678.
Dai, Q., Singleton, K.J., Yang, W., 2006. Regime shifts in a dynamic term structure model of the US treasury yields. Working Paper, Stanford University.
Diebold, F.X., Li, C., 2006. Forecasting the term structure of government bond yields. Journal of Econometrics 130, 337–364.
Driessen, J., Klaassen, P., Melenberg, B., 2003. The performance of multi-factor term structure models for pricing and hedging caps and swaptions. Journal of Financial and Quantitative Analysis 38, 635–672.
Duarte, J., forthcoming. The causal effect of mortgage refinancing on interest-rate volatility: Empirical evidence and theoretical implications. Review of Financial Studies.
Dueker, M.J., 1997. Markov switching in GARCH processes and mean-reverting stock market volatility. Journal of Business and Economic Statistics 15, 26–35.
Duffee, G., 1996. Idiosyncratic variation of treasury bill yields. Journal of Finance 51, 527–552.
Duffee, G., Stanton, R., 2004. Estimation of dynamic term structure models. Working Paper, Berkeley.
Duffie, D., Kan, R., 1996. A yield-factor model of interest rates. Mathematical Finance 6, 379–406.
Fama, E.F., Bliss, R.R., 1987. The information in long-maturity forward rates. American Economic Review 77, 680–692.
Gallant, A.R., Tauchen, G., 1996. Which moments to match? Econometric Theory 12, 657–681.
Gray, S., 1996. Modeling the conditional distribution of interest rates as a regime switching process. Journal of Financial Economics 42, 27–62.
Hamilton, J.D., 1988. Rational expectations econometric analysis of changes in regime: An investigation of the term structure of interest rates. Journal of Economic Dynamics and Control 12, 385–423.
Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press, Princeton NJ.
Kim, C.J., 1994. Dynamic linear models with Markow-switching. Journal of Econometrics 60, 1–22.
Kugler, P., 1996. The term structure of interest rates and regime shifts: Some empirical results. Economics Letters 50, 121–126.
Leippold, M., Wu, L., 2002. Asset pricing under the quadratic class. Journal of Financial and Quantitative Analysis 37, 271–295.
Litterman, R., Scheinkman, J., 1991. Common factors affecting bond returns. Journal of Fixed Income 1, 54–61.
Longstaff, F.A., Schwartz, E.S., 1992. Interest rate volatility and the term structure: A two-factor general equilibrium model. Journal of Finance 47, 1259–1282.
Pérignon, C., Villa, C., 2006. Sources of time variation in the covariance matrix of interest rates. Journal of Business 79, 1535–1549.
Smith, D.R., 2002. Markov-switching and stochastic volatility diffusion models of short-term interest rates. Journal of Business and Economic Statistics 20, 183–197.
Vuong, Q.H., 1989. Likelihood ratio tests for model selection and non-nested hypotheses. Econometrica 57, 307–333.