Global term structure modelling using principal component analysis

Received (in revised form): 28th January, 2008

Arcady Novosyolov
PhD is a Chief Scientific Adviser at FactSet Research Systems, Inc. Dr Novosyolov's role at FactSet is to verify mathematical foundations of quantitative products and to help create solutions to critical questions in the field of risk measurement. He also teaches risk theory at Siberian Federal University. Founded in 1978, FactSet Research Systems, Inc. combines integrated financial information, analytical applications, and client service to enhance the workflow and productivity of the global investment community.

Daniel Satchkov*
CFA is Vice President at FactSet Research Systems, Inc. At FactSet he leads research and development of multi-asset class portfolio risk measurement products. Founded in 1978, FactSet Research Systems, Inc. combines integrated financial information, analytical applications, and client service to enhance the workflow and productivity of the global investment community.

*E-mail: DSatchkov@factset.com

Abstract Principal component analysis (PCA) is a technique commonly applied to the interest rate markets to describe yield curve dynamics in a parsimonious manner. Despite an increase in global investing and the growing interconnectedness of the international markets, PCA has not been widely applied to decomposing joint structure of global yield curves. Our objective is to describe the joint structure with a model that can potentially be used for scenario analysis and for estimating the risk of interest rate-sensitive portfolios. In this study, we examine three variations of the PCA technique to decompose global yield curve and interest rate implied volatility structure. We conclude that global yield curve structure can be described with 15–20 factors, whereas implied volatility structure requires at least 20 global factors. The procedure that we identify as preferable is a two-step PCA, with local curves decomposed in the first step and combined local PCs decomposed into a joint structure (PCA of PCs) in the second step. This procedure has a key advantage in that it makes any scenario analysis more meaningful by keeping local PCA factors, which have important economic interpretations as shift, twist and butterfly moves of the yield curve.
Journal of Asset Management (2008) **9,** 49–60. doi:10.1057/jam.2008.3

Keywords: *global yield curve, interest rate implied volatility, principal component analysis, PCA, LIBOR*

Principal component analysis and its use in fixed income risk management

Principal component analysis (PCA) is a common technique applied to interest rate markets to describe yield curve behaviour in a parsimonious manner. The first three PCs are frequently identified with the economically meaningful shift, twist and butterfly moves of the yield curve. Despite the common use of and the wealth of research on applying PCA to yield curve decomposition, there does not appear to be much literature on applying such decomposition to the joint global yield curve structure. As the international markets become more integrated, looking at the global yield curve environment as one interconnected structure is very important to global investors. In our study,[1] we examine three techniques of modelling such a global structure. Our objective is to accurately and parsimoniously describe the global joint structure of yield curves and implied volatility curves with a model that can, potentially, be used for scenario analysis and for estimating the risk of interest rate-sensitive portfolios.

An excellent overview of the uses of PCA in the areas of fixed income risk measurement and management can be found in Golub and Tilman (2000). The benefits of PCA may be roughly divided into three related categories: risk estimation, risk reporting and scenario analysis.

The benefit of using PCA in risk estimation involves its ability to parsimoniously describe complex structures. For the purposes of interest rate risk measurement, the yield curve can be represented as a structure that is comprised of individual key rates, viewed as random variables. PCA allows users to describe the whole distribution of key rates with a much more compact distribution of PCs, which retain most of the variability of the initial structure. If a normal distribution is assumed for key rates and portfolio returns, the volatility of portfolio returns due to interest rate risk becomes a simple function of the volatility of PCs and of a given portfolio's sensitivities to them. In addition, some risk systems employ Monte Carlo simulation to estimate the distribution of portfolio returns in the presence of nonlinearities and of other deviations from the normality assumption. Monte Carlo simulation is a computation-intensive technique whose accuracy and cost depends on the number of variables being simulated. If a system can be described with fewer variables, the cost of simulation is reduced and its accuracy is increased.

Risk reporting is simplified because practitioners can see contributions to portfolio risk from factors that are not postulated *a priori*, but are rather derived from actual market data.

There are two main benefits to using PCA for scenario analysis. First, it helps to understand the shape and dynamics of the way the yield curves move before applying shocks to them. This benefit is very similar to the benefit for risk reporting mentioned above. For example, the most common move of the yield curve is assumed to be parallel, hence the ubiquity of the duration measure and common practice of applying parallel shocks of various magnitudes to the yield curve when performing scenario analysis. In reality, the most common move is not exactly parallel, although it resembles such a move. Consider the shape of the sensitivities of original key rates to the first PC shown in Figure 1. It is clear that the most common move represented by the first PC is not parallel. Therefore, performing PCA helps the user to specify more meaningful scenarios. The second benefit for the purposes of scenario analysis arises because PCA allows users to describe joint distribution of the key rates. Therefore, the probability of any particular scenario can be determined, which is important in designing an appropriate reaction to the results of the scenario analysis.

Criteria for selecting the best model

- Model Accuracy
  The model has to explain at least 90 per cent, and preferably close to 95 per cent, of the variation in the global yield and implied volatility curves. A portion of the variation explained is defined as the ratio of total variability of the PCs to the total variability of the initial data. Mathematically, the portion of explained variation is the ratio of the trace of the diagonal matrix of the selected eigenvalues to the trace of the covariance matrix of the initial data.
- Parsimony
  The model has to describe the structure using the fewest PCs possible.
- The model allows for some economic interpretation of factors and for scenario analysis using those factors.
- The model requires the fewest possible auxiliary calculations (eg calculation of sensitivities of bonds to each of the resulting factors).

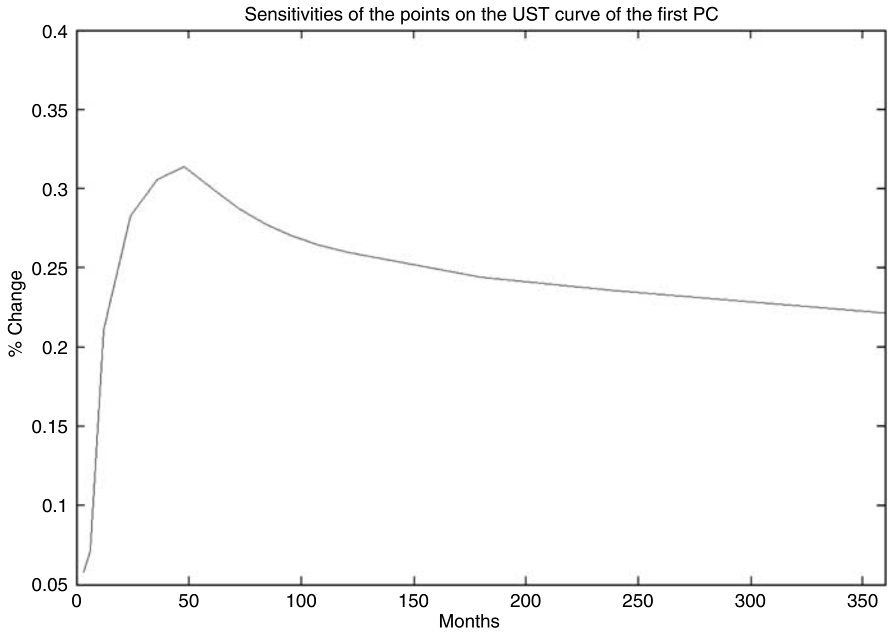

Figure 1 Sensitivities of the points on the UST curve to the first PC

Review of literature

Despite the fact that PCA for term structure is a well-established technique, we have not found many attempts to apply it to the estimation of the distribution of the joint global term structure. Malava (1999) performs direct PCA of the global term structure. The author finds that 14 are needed to explain 99 per cent of variability in the joint term structure of LIBOR USD, JPY, EUR and GBP yield curves.

Phoa (2000) briefly touches on the topic of global joint structure behaviour and uses PCA to decompose international ten-year bond yields. He concludes that global shift factor, while somewhat visible, does not explain nearly as much movement in the curves as it does in the curve-specific models. Moraux *et al.* (2002) use Common PCA (CPCA) for international curves. CPCA is, however, only designed to look for common eigenvectors and does not specify a joint distribution of the resulting factors.

Heidari and Wu (2003) look at US curves only but incorporate implied volatilities along with interest rates. They were able to describe joint term and volatility curve structures in the US using five factors (three for term structure and two for volatilities). They performed country-specific analysis only, and the potential relevance of their work is that they attempted to incorporate both interest rates and volatility into the joint framework and to reduce the number of factors necessary for describing the joint interest rate and volatility structure by working with residuals of implied volatilities after regressing out the interest rate factors.

Data

Our model is created mainly for primary assets but with a view to possibly extend it to scenario analysis for derivatives. In such a case, both Government and LIBOR curves have to be part of the model, because Treasuries and their derivatives use the UST curve while other derivatives are usually valued using the LIBOR curve. As Brooks and Yan (1999) show, there is a good deal of difference between the structure of Government and LIBOR curves in the US, which confirmed our decision to treat them as two separate curves in our analysis.

We used the following ten-spot rate curves: US Treasury, US LIBOR, EU LIBOR, Swiss LIBOR, UK Government, UK LIBOR, Japan Government, Japan LIBOR, Canada Government and Canada LIBOR. The data are daily from 1st August, 2006 to 8th August, 2007. Each curve consisted of the following 16 points: 3, 6, 12, 24, 36, 48, 60, 72, 84, 96, 108, 120, 180, 240, 300 and 360 months.

For volatilities analysis, we used the short-term (defined as three months) and long-term (ten years) implied volatilities for the following five countries: US, EU, UK, Japan and Canada. Short-term volatilities were derived from caps, while the long-term ones were derived from swaptions. The implied volatilities were derived in the following manner: the starting point is the quoted implied volatilities, for lognormal interest rates, at-the-money caps and European swaptions. In the case of caps, the individual caplet volatilities are extracted from the caps of tenors ranging from one year to 30 years. Then a natural splining technique is used to derive the full curve of three-month spot volatility. In the case of swaptions, the natural splining technique is directly applied to the implied volatilities to derive the full curve of a ten-year volatility. Like interest rates, the data are daily from 1st August, 2006 to 8th August, 2007.

Choice of data units for model estimation

Choice of units

PCA is usually done on absolute or relative (percentage) changes in the levels of interest rates. In one example however, Heidari and Wu (2003) used levels for both rates and volatilities in a single country setting. Order 1 autocorrelation of daily series of both interest rate and volatility levels is in excess of 0.9 across all curves. PCA is done on the covariance or correlation matrix of observations, and assumptions regarding the distribution of the sample used for calculation of those matrices must be satisfied. Clearly, autocorrelation of over 0.9 would violate those assumptions. Table 1 shows an example of autocorrelation of some points on the US LIBOR curve. Heidari and Wu (2003) pointed out the high levels of autocorrelation, and their work appeared to be purely descriptive, where such autocorrelation may not be a problem. On the other hand, if the objective is to estimate distributions of random variables summarising the term and volatility structures, then it presents a major problem and precludes the direct use of level data.

Table 1 Autocorrelation(1) for daily US LIBOR level data

3 months	6 months	1 year	2 years	5 years
0.96	0.95	0.94	0.96	0.97

Table 2 Autocorrelation (1) of percentage changes in daily US LIBOR

3 months	6 months	1 year	2 years	5 years
−0.21	−0.13	−0.11	−0.01	0.01

Therefore, our choice is between changes in the levels and relative (percentage) changes of spot rates. Both of those data sets appear to have low autocorrelation of order 1. Table 2 shows autocorrelation for the percentage changes in the US LIBOR curve data. Similar results are seen across other curves. So transforming the data to changes (see the section 'Percent' or 'Level' changes for more details) from levels of rates eliminates the serial correlation problem.

We must also consider whether the spot rates might be converted to returns. Our conclusion would depend on the objective of the subsequent inference. If the objective is to estimate the volatility of the return of interest rate-sensitive instruments, it would

Table 3 Kurtosis of percentage change in spot rates from 1st August, 2006 to 8th August, 2007

	Government (except Swiss)					LIBOR				
	US	Swiss LIBOR	UK	Japan	Canada	US	Euro	UK	Japan	Canada
3	6.5	**63.4**	15.7	**35.8**	9.5	**47.6**	**114.4**	**46.9**	**21.2**	**19.8**
6	5.7	**77.0**	16.2	**45.3**	9.6	9.2	**49.0**	**42.9**	**16.4**	**11.6**
12	5.7	**50.6**	10.0	**19.7**	7.6	6.7	5.1	**23.7**	**19.8**	7.0
24	5.9	**21.6**	7.5	7.2	5.6	5.2	3.6	3.7	7.5	5.7
36	5.9	**17.1**	5.5	5.1	5.1	4.6	3.3	3.6	6.4	5.0
48	5.5	**10.8**	4.9	4.6	4.3	4.3	3.0	50.0	5.3	4.4
60	5.0	7.4	4.6	4.2	4.1	4.1	3.0	48.5	4.5	4.1
72	4.6	5.8	3.9	3.9	3.8	4.0	3.0	41.2	4.1	3.8
84	4.3	5.0	3.5	3.6	3.5	3.9	2.9	45.8	3.9	3.6
96	4.1	4.8	3.4	3.4	3.4	3.8	2.9	47.1	3.7	3.5
108	4.1	4.8	3.4	3.3	3.3	3.8	2.9	44.5	3.7	3.4
120	4.0	5.0	3.5	3.3	3.3	3.8	3.0	40.6	3.6	3.4
180	3.9	5.6	3.5	3.4	3.4	4.1	3.1	36.5	3.5	3.6
240	4.1	6.0	3.4	3.4	3.3	4.0	3.0	47.4	3.5	3.7
300	4.0	5.9	3.2	3.3	3.4	4.0	3.0	51.7	3.5	3.9
360	3.9	5.5	3.4	3.4	3.6	4.1	3.0	48.4	3.5	4.2

be helpful to convert the spot rates into prices and calculate the zero coupon bond returns associated with them. If the objective is a yield curve scenario analysis, such conversion may not be necessary. The conversion should not introduce substantial differences into the quality of the estimation of the joint curve distributions, but should be a matter of final use of the results. Either choice is acceptable, but in this study we do not convert rates to prices and work with percentage changes in the spot rates themselves.[2]

Distribution of data

Looking more closely at the distribution of the data, it appears that some of the data display significant departures from normality, particularly due to the fat tails. This observation confirms the general view that high-frequency series in the financial markets display fatter tails than those suggested by the normal distribution. Some series, such as UK LIBOR, exhibit extremely fat tails. Table 3 shows particularly high values of kurtosis at the short end of the curves.

Our application of the PCA model will assume that these series are multivariate normal, which is a simplifying assumption that is made in the interest of mathematical tractability. This simplifying assumption, while necessary for analysis, must be kept in mind for any inference. For example, if PCA of these curves is used for yield curve scenario analysis, it must be noted that the probability of tail events, especially at the short end of the curve, is much higher than would be suggested by the multivariate normal distribution implied by the PCA.

'Percent' or 'Level' changes

We had to choose whether to model the dynamics using 'percent' or 'level' changes of the spot rates. The methods have similar statistical properties and show similar results in terms of explanatory power. As an example, Table 4 shows the cumulative proportion of variance explained by the first four PCs in the US Treasury and LIBOR curves, broken down by the type of the data used. Looking at Table 4, we can see that percentage changes seem to do noticeably better in the Treasury market and marginally worse in the LIBOR market. Overall, the explanatory power is similar between those data sets across curves as well. Therefore, we have to use some other criteria to choose between them.

Table 4 Cumulative proportion of variance explained in UST and US LIBOR curves by the first four PCs (per cent changes vs level changes)

Level changes		Per cent changes	
Gov't	LIBOR	Gov't	LIBOR
0.57	0.66	0.53	0.48
0.70	0.86	0.73	0.75
0.78	0.91	0.88	0.88
0.83	0.93	0.90	0.92

We decided to use percentage changes using the following reasoning. Because we are trying to build a global model, we should not allow the magnitude of interest rates to have a decisive impact. To take an extreme example, if we used level changes, Japanese curves would be practically lost, since the magnitude of changes in Japan has been much smaller. Using percentage changes is an effective way to ensure that different data sets have comparable impact. Correlation-based PCA on the level changes would be another way of making the data comparable, but we decided against that because it would disregard the difference in volatility between data sets, which is not a desirable characteristic for our purposes. Going forward, when we discuss interest rate or volatility data in this document, we are referring to data measured as daily percentage change.

Choice of model and estimation

We considered three main approaches to creating joint global structure of interest rates and implied volatilities.

Procedure 1. Create curve-specific term structure PCs for each curve in the data set. Choose three or four, and then regress implied volatility data on the chosen term structure PCs, analogous to the procedure used by Heidari and Wu (2003). Then, select some number of PCs (term structure PCs from above plus additional volatility PCs from the residual of the regression), and perform global PCA on term structure and volatility (residual) PCs derived from each curve. This procedure gives us the joint global structure that we are looking for.

Procedure 2. Perform PCA on the whole global term structure in one step, analogous to the method described by Malava (1999). Repeat for volatilities. Create a joint distribution of volatilities and term structure PCs.

Procedure 3. First, perform PCA on country-specific yield curves. Choose three or four PCs and then perform a second-step PCA on the factors chosen from the individual curves. Then repeat the whole process for joint structure of short-term and long-term volatilities.

Procedure 3 mathematically

The original data $X^{(i)}$, $i = 1, \ldots, 10$, are processed by PCA $X^{(i)} - \mu^{(i)} = Y^{(i)} U^{(i)\mathrm{T}}$, where $\mu^{(i)}$ stands for means and $U^{(i)}$ denotes the orthogonal matrix of eigenvectors of covariance matrix of $X^{(i)}$. Then n_1 columns of each matrix $Y^{(i)}$, $i = 1, \ldots, 10$, are selected and gathered in the matrix Y, which in turn is processed by PCA $Z = YV$, where V is the orthogonal matrix of eigenvectors of cov(Y). Then $\tilde{Z}$ containing n_2 columns of Z represents selected factors from the second stage. The data are restored as follows: $\tilde{Y} = \tilde{Z}\tilde{V}^T$, then $\tilde{Y}$ is split into $\tilde{Y}^{(i)}$, $i = 1, \ldots, 10$, with n_1 columns each, such that $\tilde{X}^{(i)} = \tilde{Y}^{(i)} \tilde{U}^{(i)T} + \mu^{(i)}$.

Applying to our data

$i = 1, \ldots, 10$ — is the set of 10 interest rate curves (five volatility curves with short-term and long-term volatilities combined).

$X^{(i)}$. — is the matrix of percentage change in yields (volatilities) for curve i.

Yield curve matrix $X^{(i)}$ consists of 266 observations and 16 points on each yield curve i (for volatilities, 16 points on long-term volatility curve and 16 points on short-term volatility curve are combined within each country prior to analysis, so the analogous data set consists of 266 observations and 32 points on the curve for each country).

n_1 — is the number of PCs picked at first step.

n_2 — is the number of PCs picked at the second step.

$Y^{(i)}$ — is the matrix of all PCs within the curve i. If four PCs are chosen, then this matrix is 266 observations of a four-element vector of PCs.

Y — is the matrix of chosen PCs from all countries (if four were chosen in each of ten curves in the first step, then it is 266 observations of 40 PCs).

V — is the matrix of eigenvectors that relate curve-specific PCs to the second-step global PCs (if four were chosen in each of ten curves in the first step, then it is a square matrix with dimension 40).

$\tilde{V}$ — is the matrix of the chosen eigenvectors from the second step corresponding to $\tilde{Z}$ (if $n_2 = 20$, then it contains loadings of aggregating 40 curve specific factors onto 20 global factors).

Z — is the matrix of all the global PCs (equal in size to Y).

$\tilde{Y}^{(i)}$ — is the matrix of reconstructed curve-specific PCs. (It is equal in size to $Y^{(i)}$.)

$\tilde{X}^{(i)}$ — is the matrix of the original data for curve i reconstructed using the transpose of the chosen eigenvector matrix. (It is equal in size to $X^{(i)}$.)

Statistical comparison of Procedures 2 and 3

We perform linear regression:

$$X_j^{(i)} = \alpha_j^{(i)} + B_j^{(i)} \tilde{Y}^{(i)^{\mathrm{T}}} + \varepsilon_j^{(i)} \quad \text{for each } i \text{ and } j \tag{1}$$

where i is the curve index (eg UST, US LIBOR, UK LIBOR), j the point on the curve index (eg if we have 16 points on one curve, then 16 regressions are run for it), $X_j^{(i)}$ the vector of original data in curve i at a single point j on the curve (eg for a 3-month point in US Treasury curve, it would be 266 observations of the 3-month point on the curve).

$\tilde{Y}^{(i)}$ — For Procedure 3 it is the matrix of 'observations' of curve-specific PCs re-constructed from the global PCs. This is an estimate of the matrix of original curve PCs $Y^{(i)}$ and will not be exactly equal to it with some loss of information due to discarded PCs.
For Procedure 2 it is simply the matrix of selected global PCs.

$\alpha_j^{(i)}$ — is the intercept of the regression of point j in curve i.

$B_j^{(i)}$ — is the vector of coefficients with dimension equal to the number of PCs chosen (or reconstructed in the case of Procedure 3).

$\varepsilon_j^{(i)}$ — is the vector of residual after performing regression in a single curve i at a single point on the curve j with dimension equal to number of observations.

The accuracy of the reconstruction of each point on the curve is assessed by examining R^2 statistic of this regression.[3] Table 6 contains R^2 values for Procedure 2 (one-step procedure) for 20 global PCs. Table 8 shows results of similar analysis for Procedure 3 (two-step procedure). Table 9 contains the differences in R^2 between the two procedures. In the case of Procedure 3, 20 global PCs were reconstructed back to four local curve PCs ($\tilde{Y}^{(i)}$). The main reason for such reconstruction is to tie the analysis back to the local yield curve PCs, which have economic interpretation.

Comparison of procedures

(A) — Procedure 1 was discarded after only a few tests. A surprising result was that regressing volatilities on the term structure PCs does not produce the results that Heidari and Wu (2003) saw in the US. As can be seen from Table 5, there appears to be no benefit in performing an additional step of regressing percentage changes in volatilities on the term structure PCs. It seems that there is not enough common structure between percentage changes in volatilities and interest rate data on a daily basis to make this procedure useful in the sense of reducing the number of PCs needed to explain the variation in percent volatilities. This is an interesting and somewhat unintuitive result because the expectation is that the interest rate changes and the implied volatilities of interest rates should be related. It is possible that interest rate changes and implied volatility changes are related in a nonlinear fashion, which suggests a basis for future research.

We believe that the main reason we obtained different results from Heidari and Wu (2003) is because they used (highly autocorrelated) levels, and we are using percentage changes. Using autocorrelated series can describe the data, but it does not allow for inference, which is the main objective when building a model for scenario analysis and risk estimation. This detail emphasises the importance of having a clear objective for the model.

(B) — The results that we obtained for Procedure 2 are somewhat different from Malava (1999). Fifteen PCs are necessary to explain 96 per cent of variance in the original global term structure, but 30 are required to explain 99 per cent. Volatilities have considerably less structure to them. To

Table 5 Comparison of performing PCA on volatilities separately vs performing it on the residual of regressing volatilities on the first three PCs from the term structure (US volatilities and US LIBOR term structure)

ST Vol from residual	ST Vol separately	LT Vol from residual	LT Vol separately
0.70	0.72	0.49	0.55
0.85	0.87	0.81	0.83
0.90	0.91	0.90	0.92
0.93	0.94	0.95	0.96
0.96	0.96	0.97	0.97
0.98	0.98	0.98	0.99
0.99	0.99	0.99	0.99
1.00	1.00	0.99	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00
1.00	1.00	1.00	1.00

The values shown are cumulative proportions of variation of implied volatility data explained by the PCs.

Table 6 **R^2 values of regressing of each interest rate curve, point by point, on the 20 global (Procedure 2) interest rate PCs**

	Government (except Swiss)					LIBOR				
	US	Swiss LIBOR	UK	Japan	Canada	US	Euro	UK	Japan	Canada
3	0.28	0.96	0.82	1.00	0.39	0.31	0.21	0.79	0.99	0.28
6	0.31	0.96	0.90	0.99	0.72	0.67	0.59	0.83	0.97	0.78
12	0.80	0.91	0.93	1.00	0.86	0.72	0.52	0.85	0.98	0.86
24	0.97	0.91	0.88	0.99	0.95	0.97	0.87	0.83	0.99	0.92
36	0.98	0.94	0.93	0.99	0.96	0.98	0.90	0.85	1.00	0.95
48	0.97	0.95	0.95	0.99	0.97	0.98	0.96	0.98	1.00	0.94
60	0.99	0.97	0.95	0.98	0.97	0.99	0.98	0.99	1.00	0.97
72	0.97	0.97	0.97	0.99	0.97	0.99	0.98	0.99	1.00	0.95
84	0.96	0.97	0.98	0.99	0.97	0.99	0.99	0.99	1.00	0.94
96	0.97	0.97	0.98	0.98	0.96	0.98	0.99	0.99	1.00	0.96
108	0.98	0.98	0.98	0.98	0.97	0.98	0.99	0.99	1.00	0.97
120	0.99	0.98	0.98	0.97	0.97	0.98	0.98	0.99	0.99	0.97
180	0.92	0.95	0.96	0.92	0.85	0.97	0.96	0.99	0.99	0.92
240	0.89	0.96	0.93	0.97	0.83	0.97	0.96	0.99	0.98	0.96
300	0.95	0.95	0.91	0.98	0.92	0.96	0.95	0.98	0.97	0.94
360	0.84	0.95	0.89	0.93	0.87	0.95	0.93	0.96	0.96	0.91

explain 95 per cent of variability in the global volatility curves (both short- and long-term) 33 PCs are required, and 55 PCs are required for explaining 99 per cent of the variability. The next test we performed is to regress each point on both the interest rate and volatility curve on the 15 (term structure) and 33 (volatility) PCs separately. Table 6 shows the R^2 values of regressing each point on the curve on the 15 interest rate global PCs. It is clear that the explanatory power of the model breaks down somewhat on the short end of the curve (ie the short end of the curve is somewhat disconnected from the global systematic moves in interest rates). This disconnect is likely due to the control that central banks exercise over the short end of the curve.

(C) — The first step of the two-step procedure is the most common way of doing PCA of the yield curve. Typically, the first three PCs have economic interpretation such as shift, twist and butterfly. As seen in Table 7, the first four factors explain at least 97 per cent of total variability in each individual curve.

For the second step, we tested a different number of factors and settled on 20 PCs because they explain close to 99 per cent

Table 7 **Cumulative proportion of variability explained by the first four PCs in each curve separately (Procedure 3, first step)**

	PC1	PC2	PC3	PC4
US Treasury	0.87	0.91	0.94	0.97
Swiss LIBOR	0.76	0.92	0.96	0.98
UK Government	0.77	0.92	0.97	0.98
Japan Government	0.67	0.93	0.96	0.97
Canada Government	0.86	0.93	0.96	0.98
US LIBOR	0.91	0.96	0.99	0.99
EU LIBOR	0.87	0.93	0.98	0.99
UK LIBOR	0.87	0.94	0.98	0.99
Japan LIBOR	0.83	0.95	0.98	0.99
Canada LIBOR	0.89	0.94	0.97	0.98

of variation in the combined curve-specific PCs. Eigenvectors in the second step represent loadings of curve specific PCs onto global PCs. Table 8 shows the explanatory power of the reconstructed PCs ($\tilde{Y}^{(i)}$ in Equation 1) when regressed on the original data. The main reason for reconstructing the global PCs back to curve-specific PCs was to keep the shift, butterfly and twist factors that practitioners are familiar with to facilitate more meaningful analysis using the model. Overall, the explanatory power is similar to the one-step procedure, which would be expected with the short end of the curve similarly

Table 8 R^2 of regressing 20 global PCs from the Procedure 3 (first step is four PCs)

	Government (except Swiss)					LIBOR				
	US	Swiss LIBOR	UK	Japan	Canada	US	Euro	UK	Japan	Canada
3	0.39	0.97	0.88	1.00	0.36	0.25	0.17	0.81	0.96	0.25
6	0.44	0.97	0.95	0.99	0.77	0.67	0.55	0.87	0.95	0.80
12	0.83	0.90	0.95	0.96	0.89	0.74	0.50	0.85	0.94	0.88
24	0.98	0.94	0.90	0.95	0.95	0.97	0.88	0.90	0.98	0.94
36	0.98	0.96	0.95	0.98	0.96	0.98	0.91	0.91	0.99	0.97
48	0.98	0.97	0.96	0.97	0.97	0.98	0.97	0.99	1.00	0.95
60	0.99	0.98	0.96	0.97	0.97	0.99	0.98	0.99	0.99	0.98
72	0.98	0.98	0.98	0.97	0.97	0.99	0.98	1.00	0.99	0.96
84	0.97	0.97	0.98	0.97	0.97	0.99	0.99	1.00	0.99	0.95
96	0.97	0.97	0.98	0.96	0.96	0.99	0.99	1.00	0.99	0.96
108	0.98	0.98	0.98	0.97	0.97	0.99	0.99	0.99	0.99	0.98
120	0.99	0.98	0.98	0.95	0.99	0.98	0.99	0.99	0.99	0.97
180	0.94	0.96	0.96	0.87	0.95	0.98	0.96	0.99	0.98	0.92
240	0.92	0.97	0.96	0.91	0.95	0.98	0.96	1.00	0.97	0.96
300	0.96	0.96	0.96	0.89	0.98	0.97	0.95	0.99	0.95	0.94
360	0.85	0.96	0.95	0.84	0.87	0.96	0.94	0.98	0.93	0.91

Table 9 Differences in R^2 of regressing term structure points on 20 PCs resulting from the Procedure 3 vs. 20 global PCs resulting from Procedure 2

	Government Curves (except Swiss)					LIBOR				
	US	Swiss LIBOR	UK	Japan	Canada	US	Euro	UK	Japan	Canada
3	0.11	0.01	0.06	0	−0.03	−0.06	−0.04	0.02	−0.03	−0.03
6	0.13	0.01	0.05	0	0.05	0	−0.04	0.04	−0.02	0.02
12	0.03	−0.01	0.02	−0.04	0.03	0.02	−0.02	0	−0.04	0.02
24	0.01	0.03	0.02	−0.04	0	0	0.01	0.07	−0.01	0.02
36	0	0.02	0.02	−0.01	0	0	0.01	0.06	−0.01	0.02
48	0.01	0.02	0.01	−0.02	0	0	0.01	0.01	0	0.01
60	0	0.01	0.01	−0.01	0	0	0	0	−0.01	0.01
72	0.01	0.01	0.01	−0.02	0	0	0	0.01	−0.01	0.01
84	0.01	0	0	−0.02	0	0	0	0.01	−0.01	0.01
96	0	0	0	−0.02	0	0.01	0	0.01	−0.01	0
108	0	0	0	−0.01	0	0.01	0	0	−0.01	0.01
120	0	0	0	−0.02	0.02	0	0.01	0	0	0
180	0.02	0.01	0	−0.05	0.1	0.01	0	0	−0.01	0
240	0.03	0.01	0.03	−0.06	0.12	0.01	0	0.01	−0.01	0
300	0.01	0.01	0.05	−0.09	0.06	0.01	0	0.01	−0.02	0
360	0.01	0.01	0.06	−0.09	0	0.01	0.01	0.02	−0.03	0

disconnected from the rest of the structure in the US Government, Canada Government, US LIBOR, Euro LIBOR and Canada LIBOR curves.

Comparison of Procedures 2 and 3

Term structure comparison

Overall, the picture is very similar between the two approaches. The description of the US Treasuries curve at the short end is considerably better under Procedure 3 (showing 0.11 and 0.13 improvement in the three-month and six-month rates, respectively). Procedure 3 also performs better for UK Government and LIBOR curves at both the short and long ends of the curve. At the same time, Japan Government long rates and Euro LIBOR short rates perform somewhat better under the one-step approach.

Table 10 R^2 of regressing each volatility curve point on the eight local PCs reconstructed from 20 global PCs obtained in the second step

	US ST	EU ST	UK ST	Jpn ST	Can ST	US LT	EU LT	UK LT	Jpn LT	Can LT
3 mo	1.00	1.00	0.99	0.97	0.98	0.86	0.59	0.51	0.89	0.82
6 mo	1.00	1.00	0.99	0.97	0.98	0.87	0.71	0.59	0.87	0.95
12 mo	0.93	0.94	0.94	0.91	0.96	0.89	0.76	0.66	0.84	0.79
24 mo	0.92	0.71	0.49	0.42	0.42	0.88	0.80	0.75	0.80	0.92
36 mo	0.73	0.64	0.41	0.43	0.42	0.88	0.84	0.74	0.73	0.92
48 mo	0.78	0.65	0.46	0.87	0.67	0.84	0.83	0.71	0.70	0.91
60 mo	0.78	0.69	0.52	0.90	0.79	0.87	0.82	0.68	0.74	0.92
72 mo	0.71	0.72	0.49	0.85	0.88	0.86	0.81	0.71	0.70	0.90
84 mo	0.78	0.77	0.55	0.96	0.90	0.79	0.83	0.78	0.64	0.87
96 mo	0.76	0.79	0.61	0.96	0.92	0.77	0.82	0.76	0.60	0.82
108 mo	0.75	0.81	0.65	0.96	0.95	0.85	0.86	0.78	0.70	0.84
120 mo	0.74	0.78	0.64	0.98	0.98	0.85	0.84	0.71	0.87	0.88
180 mo	0.75	0.65	0.62	1.00	0.96	0.80	0.71	0.53	0.96	0.79
240 mo	0.75	0.70	0.65	1.00	0.95	0.99	0.83	0.58	0.94	0.76
300 mo	0.80	0.76	0.71	1.00	0.98	0.96	0.78	0.54	0.98	0.84
360 mo	0.97	0.94	0.83	0.97	0.95	0.99	0.83	0.58	0.87	0.88

ST indicates short term volatility, LT indicates long term.

Overall, the two-step approach (Procedure 3) appears to have an edge (the sum of improvements in R^2 is 0.82 over all countries/all maturities), but those minor differences are unlikely to be statistically significant. The main conclusion would be that in terms of explanatory power, the procedures are very similar, and we have to consider the ease of use and implementation issues as the discriminating factors between the two procedures (Table 9).

Volatility comparison

As with term structure, Procedures 2 and 3 show virtually identical results for volatilities. Table 10 shows results for Procedure 3 applied to volatilities. Considerably lower R^2, relative to the term structure tables, means that changes in implied volatilities are less connected globally than changes in the interest rates. Improving explanatory power significantly would require something like 40 or 50 factors for volatility structure, which would make the model too unwieldy for estimation and implementation. The explanatory power is, however, adequate if we settle for modelling global volatility effects and remember that there are curve-specific volatility effects that we simply cannot include in such a framework.

Conclusions

We have compared three approaches to describing the global structure of interest rate and volatility changes. The first approach tested, described in Heidari and Wu (2003), suggests performing PCA on the interest rate term structure within the country. After term structure PCs are identified, volatilities are regressed on them and any additional PCs that explain movement in volatilities are identified by performing PCA on the residuals of this regression. This procedure was discarded because it did not reduce the number of PCs necessary to describe the term structure of volatilities in our framework and, therefore, offered no help in terms of creating a more parsimonious model.

After examining the remaining two procedures, we believe that the two-step procedure (Procedure 3), with four PCs in the first step and 20 global PCs in the second step, is the best approach for modelling the global term structure of interest rates. First, it explains slightly more of the overall variation of the yield curves over the period tested. More importantly, extracting four PCs out of each curve and then performing global analysis on them offers multiple advantages in terms of implementation:

(a) It retains curve-specific PCs, which are usually interpreted as being similar to shift, twist and butterfly moves of the yield curve. Global PCs do not have such interpretations. This feature can be very helpful for implementing the model for practical use. For example, practitioners can perform scenario analysis based on the factors they are familiar with. Risk estimation for global portfolios using such a model allows for decomposition of risk into contributions from exposure to the local shift, twist and butterfly factors. We believe that this characteristic is a crucial advantage of Procedure 3, especially because statistical modelling is too often disconnected from economic realities.

(b) The second advantage is less important, but under certain conditions it may reduce the cost of the implementation. Fewer sensitivities need to be computed for the two-step procedure. Procedure 3 allows approximation of the sensitivity of each bond to the global factor from the combination of its sensitivity to the local factor and a loading of the local factor on global factors. To calculate the sensitivity of a bond to one PC using the central difference formula, each bond has to be fully repriced twice. With the one-step procedure and 20 global PCs, this leads to 40 repricings for each interest rate-sensitive instrument. With the two-step procedure, there will be only eight. Thirty-two bond repricings may not seem like an overwhelming amount of calculation, but if the sensitivities need to be computed daily (or even intra-day), this could play a role in raising the cost of implementation. In addition, there are instruments (eg ABS/MBS) that may require simulation to be priced. Since fixed income universes tend to be large, on the order of tens of thousands of instruments, the computation cost may be far from negligible.

(c) The third advantage is a more theoretical one. Taking three or four curve-specific PCs in the first step should theoretically help eliminate the curve data noise and pass more or less systematic effects through to the second step. This method would not necessarily improve the description in sample, but could result in a more stable estimate. Testing this hypothesis requires future research.

Overall, Procedures 2 and 3 offer similar explanatory power, but Procedure 3 is preferred based on the implementation issues.

Notes

1. Authors would like to thank Bill McCoy for many helpful discussions and for providing the data. We would also like to thank Katherine McCabe and Minh Vu for helping prepare this article for publication. Any errors are ours.
2. Authors would like to thank Steven Satchell and Jason MacQueen for helpful discussion on this topic.
3. We did not use Adjusted R^2 in the evaluation because we are comparing procedures that use the same number of global independent variables.

References

Brooks, R. and Yong Yan, D. (1999) 'LIBOR vs Treasury Rate: Evidence from the Parsimonious Term Structure Model', *Journal of Fixed Income*, 9, 71–83.

Golub, B. and Tilman, L. (2000) *Risk Management: Approached for Fixed Income Markets*, John Wiley & Sons, Inc., New York, Chapter 3.

Heidari, M. and Wu, L. (2003) 'Are Interest Rate Derivatives Spanned by the Term Structure of Interest Rates?', *Journal of Fixed Income*, 13, 75–86.

Malava, A. (1999) 'Principal Component Analysis on Term Structure of Interest Rates', Helsinki University of Technology Department of Engineering Physics and Mathematics Working Paper.

Moraux, F., Perignon, C. and Villa, C. (2002) 'Common Factors in International Bond Returns Revisited: A Common Principal Component Approach', Swiss National Science Foundation Working Paper.

Phoa, W. (2000) *Yield Curve Risk Factors: Domestic and Global Contexts*, Professional's Handbook of Financial Risk Management, Butterworth-Heinemann, Oxford, Chapter 5.

