CONTRIBUTION N° 20

A STOCHASTIC YIELD CURVE MODEL FOR ASSET / LIABILITY SIMULATIONS

PAR / BY

James. A. TILLEY

Etats Unis / United States

UN MODELE STOCHASTIQUE DE COURBE DES TAUX POUR DES SIMULATIONS ACTIFS - ENGAGEMENTS

66 UN MODELE STOCHASTIQUE DE COURBE DES TAUX POUR DES SIMULATIONS ACTIFS - ENGAGEMENTS

JAMES A. TILLEY, F. S. A, PH. D.

RESUME

Cet article présente un modèle stochastique de courbes des taux du Trésor des Etats - Unis d'Amérique, destiné à être utilisé pour des simulations actifs - engagements applicables aux compagnies d'assurance, aux banques et autres institutions de dépôt, et aux caisses de retraite. Le prix de nombre d'actifs et de passifs de valeur nominale exprimée en dollars américains, et payables dans cette monnaie, est fixé par rapport à une courbe des taux du Trésor. Le modèle stochastique est empirique - il est dérivé dune analyse statistique de courbes historiques des taux, et non des principes fondamentaux des sciences économiques et financières. Moyenant les ajustements appropiés, pour assurer une fixation des prix sans arbitrage, le modèle peut- être utilisé pour évaluer des options débitrices et des *cash* flows dépendant du taux d'intérêt général.

Ce modèle est élaboré en exprimant des courbes des taux comme des superpositions linéaires de polynômes orthonormés, puis en effectuant une analyse en série chronologique des coefficients d'expansion et une analyse statistique des résidus de l'ajustement, pour spécifier les propriétés stochastiques de l'évolution des courbes des taux. Le résultat de cette analyse suggère qu'un haut degré d'immunisation serait réalisable en ajustant les indices qui caractérisent les réponses de prix des actifs et des engagements aux évolutions des quatre caractéristiques principales de la forme d'une courbe des taux : niveau, pente, courbure et ondulation.

Les données de la courbe des taux analysée présentent une réversion moyenne et des distributions à fort traînage, des résidus. Une combinaison de deux distributions normales s'adapte très bien aux résidus. Une analyse théorique des propriétés de stabilité du modèle est présentée, et un test de simulation de ce modèle est utilisé pour examiner son comportement sur une période de 100 ans.

A STOCHASTIC YIELD CURVE MODEL FOR ASSET / LIABILITY SIMULATIONS

BY JAMES A. TILLEY, F. S. A, PH. D.

ABSTRACT

In this paper a stochastic model of the U. S. Treasury yield curve is developed for use in. asset/liability simulations applicable to insurance companies, banking and other' depository institutions, and pension funds. Many assets and liabilities denominated and payable in U. S. dollars are priced relative to the Treasury yield curve. The stochastic model is empirical - it is derived from a statistical analysis of historical yield curves, not from basic principles of economics and finance. With appropriate adjustments to ensure arbitrage-free pricing, the model could be used to value debt options and general interest rate contingent cash flows.

The model is developed by expressing yield curves as linear superpositions of orthonormal polynomials and then performing time series analysis of the coefficients of expansion and statistical analysis of the residuals of the fit to specify the stochastic properties of the evolution of yield curves. The results of these analyses suggest that a high degree of immunization should be achievable through matching indexes that characterize asset and liability price responses to changes in four primary shape characteristics of a yield curve : its level, tilt, warp, and undulation.

The yield curve data analyzed exhibit mean reversion and fat-tailed distributions of the residuals. A mixture of two normal distribulions fits the residuals very well. A theoretical analysis of the stability properties of the model is presented and a simulation test of the model was used to examine its behavior over a 100 - year period.

1 - BACKGROUND AND INTRODUCTION

Pension and life actuaries in the United States and Canada began using interest rate simulations extensively during the 1980s as a tool for pricing and valuing various products and financial security programs. A review of stochastic interest rate models derived from basic principles of economics and finance was presented by Sharp [1] at the 23rd International Congress of Actuaries in Helsinki. Such one - factor and two-factor models have had some, but limited, success in describing the richness of yield curve movements observed in the financial market place. The wak presented in this paper stemmed from a need to model the observed dynamics of yield curves better than has been possible with purely theoretical models. An empirical approach based on an analysis of historical yield curves was deemed appropriate for the purpose of long - term asset/liability simulations, but could be of questionable value to traders interested in achieving short-run arbitrage profits.

A valuable byproduct of a study of the dynamics of yield curves is insight into the number of risk indexes that must be controlled in order to achieve acceptable asset/liability immunization. A considerable amount of research has been done in this area by both academicians and practitioners. Of particular importance is the work by Harris [2] that utilizes principal component and factor analysis of changes in yield curves to derive two duration measures - "twist" duration and "concave" duration - in addition to the classical "parallel - shift" duration. Harris's studies have demonstrated that immunizing against twist and concave shifts in the yield curve in addition to parallel shifts often achieves significantly better risk control than immunizing against parallel shifts alone.

Parallel duration quantifies the price response of an asset or liabiity to a first-order shift in the level of interest rates. It has become fashionable in the practice of portfolio management in the United States to talk about the "convexity" of assets, particularly assets with option features such as mortgage - backed securities and callable corporate bonds. The convexity index is a risk measure that quantifies the price response to a second - order shift in the level of interest rates.

If none of the asset or liabiity cash flows are contingent on interest rates, it can be shown that immunizing parallel duration and convexity is equivalent to immunizing against first - ader shocks to both the level and the slope of the term structure of interest rates. However, it is a reasonable conjecture that better risk control in the case of interest-contingent cash flows is attainable by matching asset and liability responses to first-order shocks in both the level and the slope of the term structure than by matching asset and liability parallel durations and convexities. The method of orthonormal polynomials used in this paper to break apart a yield curve into mutually orthogonal components suggests that a high degree of immunization should be achievable through matching risk indexes based on price responses to first - order changes in four primary shape characteristics of yield curves : their level, tilt, warp, and undulation. The immunixation studies needed to support this conjecture are beyond the scope of this paper and will be the subject of further investigations.

An important decision in the study of interest rate dynamics is the choice of what to model. In a recent monograph, Coleman, Fisher, and Ibbotson [3] discuss considerations in whether to analyze the yield curve that expresses the yields of bonds trading at par, the term structure that expresses the yields of zero - coupon bonds (namely, spot rates derivable from the yield curve), or forward rates that are derivable from spot rates. They selected forward rates for their analyses, whereas many other researchers have studied spot rates. From a theoretical viewpoinl, curves of spot rates and forward rates are both preferred to the yield curve because, unlike the yields defining the yield curve, spot rates and forward rates relate to discounting a single cash flow occurring at one point in time to some earlier point in time. However, yield curves are studied in this paper in keeping with its emphasis on practice. The yield curve, not the term structure, directly describes the market made by traders of U. S. government securities, and is thus what is directly observable in the financial market place.

Section 2 describes how the method of orthonormal polynomials can be used to decompose a yield curve unambiguously into its constituent parts and then presents the results of "fitting" 101 yield curves utilizing this method. Section 3 presents the results of a time series regression analysis of the fitting parameters developed from the analysis in Section 2, and leads to the specification and estimation of a stochastic yield curve model. In Section 4, stability properties of the model are analyzed and the results of a simulation test are discussed. Section 5 summarizes the paper and indicates directions for future research.

2 - FITTING HISTORICAL YIELD CURVES

The yield curves analyzed in this study came from the U.S. government security data base maintained by traders at Morgan Stanley & Co. in New York. The yields are based on the traders' bid prices at 3:00 p.m. for "on - the - run" U.S. Treasury bills, notes, and bonds. The data covers the period from December 16, 1981 to August 16, 1989 at intervals of four weeks, giving a total of 101 yield curves. Only one of the interval-: ending Wednesdays occurred when the markets were closed - November 11,1987 - and this was replaced by data for the Wednesday of the following week There was one missing yield in each of three yield curves - the three missing data points were replaced by averages of the yields on both sides of the missing points. All yields are expressed as percentages on a bond-equivalent basis (that is, convertible semi-annually).

The yield curves in the data base are specified at maturities of three months, six months, and one, two, three, four, five, seven, ten, twenty, and thirty years. This maturity range is mapped onto the unit interval [0,1] by first taking natural logarithms of the maturities expressed in years and then linearly rescaling the log - maturity range to be one unit in length. The transformation places the three-month maturity at zero and the thirty - year maturity at unity. Each yield curve so mapped onto the unit interval is assumed to be linear between the points at which it is defined, and thus is continuous and piecewise linear on [0,1].

Appendix 1 describes how a complete set of mutually orthogonal polynomials can be defined on [0,1]. By normalizing such polynomials, a complete set of orthonormal polynomials is obtained. They are complete in the sense that any continuous function f(x) defined on [0,1] can be expanded as an infinite series

$$f(x) = \sum_{n=0}^{\infty} a_n q_n(x)$$

where $q_n(x)$ represents the orthonormal polynomial of degree n and the various a_n are the coefficients of the expansion. The a_n are computed from the following equation :

$$a_n = \int_0^1 f(x)\, q_n(x)\, dx$$

The orthonormal polynomial method of "fitting" yield curve functions differs from other polynomial fitting methods in a few important respects. First, each coefficient in the expansion is independent of the other caefficients - this follows from the orthogonality property of the polynomials. Thus, the coefficient of the nth polynomial in the expansion does not depend on what other polynomials have been included or excluded. If one wants to improve upon an approximation to the yield curve based on the first n polynomials by adding higher order polynomials, none of the coefficients in the original approximation change in value when more terms are included. The decomposition of a given yield curve into these orthonormal polynomials is unique.

Second, the completeness property assures that any yield curve of the type described above can be fit exactly if an infinite number of terms is used Any yield curve, no matter how strangely shaped, can be accomodated within this framework. From a practical standpoint, the question is one of how many terms are needed to achieve an acceptable approximation to substantially all of the yield curves encountered in a historical universe. The approach is useful only if a small number of terms is required.

The 101 yield curves comprising the historical sample were decomposed in terms of orthonormal polynomials up to and including degree 10, and the errors of fit were computed for successive degrees of approximation. If the approximation to the yield curve $f(x)$ defined on $[0,1]$ that is obtained by including all orthonormal polynomials with degrees up to and including n is denoted by $f_n(x)$, the mean square error of fit, E_n^2, is given by:

$$E_n^2 = \int_0^1 [f_n(x) - f(x)]^2 \, dx$$

With the yields $f_n(x)$ and $f(x)$ both expressed as percentages, the means and standard deviations across all 101 yield curves of the root mean square errors E_n, expressed in basis points (hundredths of a percent) for $n = 0,1,\ldots,10$, are given in Table 1.

Table 1

Means and Standard Deviations of the Root Mean Square Errors of Fit

Order of Approximation	Mean of E_n	Standard Deviation of E_n
0	74.5	32.6
1	17.2	10.5
2	9.6	4.8
3	4.7	2.4
4	4.0	2.3
5	3.5	1.9
6	2.8	1.4
7	2.6	1.4
8	2.5	1.2
9	2.1	1.0
10	2.0	1.0

Table 1 clearly shows that substantial improvements in fit are obtained with each successive degree of approximation through polynomials of degree three. It then takes another five or six terms to halve the already small error. It is likely that an approximation to yield curve dynamics satisfactory for asset/liability simulation purposes and for immunization can be obtained by using only the $q_0(x)$, $q_1(x)$, $q_2(x)$ and $q_3(x)$ components of the fit. A more detailed examination of the root mean square errors of fit is given in Table 2 where the number of yield curves for which the error of fit falls within specified bands is given. These distributions of E_n for various n confirm that several terms after $n = 3$ are needed to improve the fit markedly.

Table 2

Distributions of the Root Mean Square Errors of Fit to Historical Yield Curves

Error Band*	Order of Approximation 0	1	2	3	4	5	6	7	8	9	10
0-1	0	0	0	0	1	5	6	7	9	12	14
1-2	0	0	0	5	9	9	16	17	18	36	43
2-3	0	0	4	11	22	27	47	51	52	39	34
3-4	0	1	6	25	28	30	23	18	14	7	6
4-5	1	3	8	25	20	15	3	3	4	5	2
5-6	0	3	5	14	11	8	2	3	3	1	2
6-7	0	1	12	12	6	5	3	1	0	1	0
7-8	1	2	2	5	2	1	0	0	0	0	0
8-9	1	1	16	2	0	0	0	0	1	0	0
9-10	0	5	7	1	1	0	0	0	0	0	0
10-	98	85	41	1	1	1	1	1	0	0	0
Total	101	101	101	101	101	101	101	101	101	101	101

* The bands are measured in basis points. Each band is inclusive of its lower end point and exclusive of its upper end point.

Before proceeding to the time series regression analysis of the four coefficients in the expansion of the yield curves to order three, it is instructive to examine the decomposition of an actual yield curve into its orthonormal components. An inspection of graphs of the first four orthonormal polynomials plotted on the interval [0,1] suggests that the shape characteristics associated with the zeroth, first, second, and third order polynomials are appropriately labelled level, tilt, warp, and undulation respectively. The "undulation" term is somewhat presumptive of the lack of importance of terms of order four and higher, because the orthonormal polynomials of degree four and higher also look like "waves". Table 3 gives the decomposition of the March 7, 1984 yield curve into its level, tilt, warp, and undulation parts. The first four coefficients in the polynomial expansion for that yield curve are 11.19, - 0.9339, - 0.09308, and 0.1390, respectively, to four significant figures.

Table 3

Four - Component Decomposition of the March 7.1984 Yield Curve

Maturity in Years	Bond-Equivalent Yields in Percent						
	Level	Tilt	Warp	Undulation	Total	Actual	Difference
0.25	11.19	-1.62	-0.21	0.37	9.73	9.63	0.10
0.5	11.19	-1.15	-0.05	-0.06	9.93	10.00	-0.07
1	11.19	-0.68	0.05	-0.16	10.40	10.28	0.12
2	11.19	-0.21	0.10	-0.07	11.01	11.05	-0.04
3	11.19	0.06	0.10	0.02	11.38	11.35	0.03
4	11.19	0.26	0.10	0.08	11.63	11.65	-0.02
5	11.19	0.41	0.08	0.12	11.81	11.89	-0.08
7	11.19	0.63	0.06	0.16	12.04	12.09	-0.05
10	11.19	0.88	0.01	0.15	12.23	12.20	0.03
20	11.19	1.34	-0.11	-0.07	12.36	12.36	0.00
30	11.19	1.62	-0.21	-0.37	12.23	12.31	-0.08

The contours of the values in the level, tilt, warp, and undulation columns in Table 3 indicates that the *shape* labels are appropriate. The column identified as "Total" is the third order approximation (four polynomials) of the "Actual" yield curve, and the "Difference" column displays for each maturity the error of the fit

3 • TIME SERIES ANALYSIS OF ORTHONORMAL POLYNOMIAL COEFFICIENTS

In order to specify and estimate the stochastic dynamics of the yield curve, it is necessary to perform a regression analysis of the coefficients in the expansion of the historical time series of yield curves in terms of orthonormal polynomials. Each of the 101 yield curves in the historical universe was fitted up to and including the third order polynomial. This resulted in a vector a of four fitting coefficients for each yield curve and thus a time series of 101 fitting vectors at $\mathbf{a}_t$, $t = 0,1,\ldots,100$. Of the several forms of regression attempted, the one with the greatest explanatory power was an autoregressive process of order two, with the vector $\mathbf{a}_t$ depending on the vector $\mathbf{a}_{t-1}$ at lag one and on the vector $\mathbf{a}_{t-2}$ at lag two. A greater degree of stationarity in the variance of the process was achieved by replacing the level coefficient (the zeroth element of the vector a) by its natural logarithm. The transformed vectors are denoted by a'. Including a constant term k in the regressions leads to the following model of the stochastic dynamics of the yield curve:

$$\mathbf{a}'_t = \mathbf{k} + \mathbf{R}^{(1)}\mathbf{a}'_{t-1} + \mathbf{R}^{(2)}\mathbf{a}'_{t-2} + \mathbf{e}_t$$

The matrices $\mathbf{R}^{(1)}$ and $\mathbf{R}^{(2)}$ are autoregression parameters for lags one and two, respectively, and the errors or residuals of the regression are denoted by the random vectors e_t.

Standard confidence tests based on Student's t statistics and F statistics were used to determine which elements of the constant vector k and which elements of the matrices $\mathbf{R}^{(1)}$ and $\mathbf{R}^{(2)}$ are significant at the 5% level. Table 4 lists the parameters retained at a nonzero level, the standard errors of the estimates, and the associated t statistics The hypothesis that a given parameter is equal to zero can be accepted at a 5% significance level if the absolute value of its associated t statistic is less than 1.985.

Table 4

Time Series Analysis of the Yield Curve Fitting Coefficients

Regression Parameter	Parameter Estimate	Standard Error	t Statistic
k_0	0.1000	0.0490	2.041
k_1	-0.1044	0.0365	-2.859
k_2	0.3046	0.1015	3.001
k_3	-0.0082	0.0088	-0.938
$r_{00}^{(1)}$	1.0836	0.0957	11.105
$r_{11}^{(1)}$	0.9907	0.1033	9.593
$r_{13}^{(1)}$	-0.9182	0.3388	-2.710
$r_{20}^{(1)}$	-0.1536	0.0482	-3.189
$r_{22}^{(1)}$	0.7788	0.1034	7.529
$r_{31}^{(1)}$	-0.0449	0.0170	-2.639
$r_{33}^{(1)}$	0.4667	0.0965	4.837
$r_{00}^{(2)}$	-0.1309	0.0966	-1.355
$r_{11}^{(2)}$	-0.2260	0.0988	-2.289
$r_{22}^{(2)}$	-0.1577	0.0995	-1.584
$r_{33}^{(2)}$	0.1844	0.0959	1.922

The t tests in Table 4 suggest that the parameters for the autoregressive term at lag two are very weak, since in three of the four cases the statistical test cannot be accepted at the 5% level. In each of those situations, however, the residuals of the fit exhibited insignificant serial correlation at several lags only if the lag-two parameters were retained at a nonzero value in the regression. The m e element in the constant vector that failed the t test at the 5% level was retained at a nonzero value for reasons of "symmetry" - namely, because all other constants are present in the model at nonzero values. The values of R^2 for the autoregressions of the log - level, tilt, warp, and undulation coefficients are 0.952, 0.846, 0.715, and 0.721, respectively.

The sample estimates of the standard deviations of the residuals e_t are (0.0467, 0.1464, 0.0726, 0.0358). The means of the residuals are equal to zero due to the inclusion of nonzero constant terms in the regression analysis. The sample estimate of the contemporaneous correlation matrix of the residuals is given in Table 5. Chi - square tests of goodness of fit were performed on the distribution of the residuals to determine whether they could be considered to be normally distributed. For the residuals associated with the regrsssions of the time series of the log -1 evel and tilt coefficients, the chi - square tests of the hypotheses of normality could not be rejected at the 5% level, but the tests of normality for the warp and undulation residuals were easily rejected at the 5% level. The failure of the tests was directly traceable to the fat tails of the distributions.

Table 5

Sample Estimate of the Contemporaneous Correlation Matrix of the Residuals

	e_0	e_1	e_2	e_3
e_0	1	0.156	- 0.282	- 0.022
e_1	0.156	1	0.386	- 0.227
e_2	- 0.282	0.386	1	0.426
e_3	- 0.022	- 0.227	0.426	1

Kon [4] has found that a discrete mixture of normal distributions can explain the observed kurtosis in the distribution of daily rates of return f a certain common stocks and stock indexes. Accordingly, a separate mixture of two normal distributions, both with zero mean, was fit to each series of residuals by constraining the standard deviation of the mixture to equal the sample estimate of the residual standard deviation, and then choosing the mixing probability and the ratio of the standard deviations of the two normal densities to minimize the chi-square statistic in the goodness - of - fit test of the assumed probability density function of the residuals.

After accounting for the two fewer degrees of freedom in the tests of the mixture of two normals than in the tests of a single normal density, it was found that the tilt, warp, and undulation residuals are fit to a very high degree of confidence by a mixture of two normals, but that the log - level residuals are still better fit by a single normal density. The fitted values of the mixing probabilities, expressed as percentages, are 74%, 82%, and 90% for the tilt, warp, and undulation residuals, respectively. These percentages are the weights for the normal density having the smaller standard deviation, and their complements - 26%, 18%, and 10%, respectively - are the weights for the normal density having the larger standard deviation. The fitted values of the ratios of the larger standard deviation to the smaller standard deviation are 2.50, 3.30, and 3.75 for the tilt, warp, and undulation residuals, respectively.

In order to check for stationarity or lack of it in the model estimation, the historical sample of 101 yield curves was separated into two data sets : the first 50 yield curves and the last 51 yield curves. Including only the regression parameters that emerged with nonzero values from the analysis of the entire sample of 101 yield curves, regressions were performed for each of the subsets of the full data sample. The regression. parameters and residual standard deviation for the time series of log - level coefficients showed a high degree of stability, but the estimates of most of the other regression parameters and residual standard deviations, and of the contemporaneous correlation matrix of the residuals, differed noticeably between the two data subsets. In all cases, however, the parameter estimates from both data subsets lay within the 95% confidence intervals of the corresponding parameter estimates based on the full data sample.

4 - ANALYSIS AND TEST OF THE STOCHASTIC MODEL

Instability is a difficulty often encountered when stochastic simulations are performed using a yield curve model developed from a statistical analysis of historical interest rates. For example, a lognormal model of interest rate movements will produce "runaway" interest rates in a relatively short period of time, say 10 to 20 years, unless an arbitrary ceiling is *imposed,* together with a rule that states how the ceiling absorbs or reflects interest rates that strike it. Difficulties with respect to the shape of the yield curve - for example, yield curves becoming too positively or negatively sloped - also seem to arise with disturbingly high frequency in empirical models unless arbitrary constraints are imposed.

It has been recognized by many researchers that a property known as "mean reversion" can cure the types of problems described above. In qualitative terms, mean reversion can be thought of as a restoring force that causes a variable which wanders away from its long - run mean value to return to that mean value. The farther that the variable strays from the mean value, the stronger becomes the restoring face. The theoretical one - factor model of the term structure of interest rates proposed by Cox, Ingersoll, and Ross [5] includes two parameters that characterize the mean reversion property : one specifying the location of the mean and the other establishing the strength of the reversion. Two questions naturally arise with respect to the specification and estimation of the yield curve model described in Section 3. First, is there statistically significant evidence of mean reversion in the data ? Second, if there is, does it eliminate the difficulties described above ? The first question is answered best by examining the theoretical properties of Equation [1], while the second is addressed best by performing a simulation test of the model.

Mean reversion will occur if the expected value of Equation (1) has a stable fixed point A fixed point (namely, $\mathbf{a}'_t = \mathbf{a}'_{t-1} = \mathbf{a}'_{t-2}$) will occur if the matrix $I - \mathbf{R}^{(1)} - \mathbf{R}^{(2)}$ is invertible. In that case, the fixed point, in the expected value sense, will be $\mathbf{a}'^{*} = (I - \mathbf{R}^{(1)} - \mathbf{R}^{(2)})^{-1}\, k$. Necessary and sufficient conditions for the fixed point to be stable are derived in Appendix 2. The fixed point coefficients f a the regression model estimated in Section 3 are $a'_0{}^* = 2.114$, $a'_1{}^* = -0.7070$, $a'_2{}^* = -0.05314$, and $a'_3{}^* = 0.06749$, corresponding to a stable fixed - point yield curve at a level of 8.28% with positive slope and a spread of 209 basis points between the 30 - year yield and the three - month yield.

In the case of the natural logarithm of the level coefficient, it is easy to derive the necessary and sufficient conditions for stability from the general result stated in Appendix 2 because the regression equation for the natural logarithm of the level coefficient does not involve any of the other three coefficients. Three conditions must be met simultaneously for stability to occur : $r_{00}^{(1)} + r_{00}^{(2)} < 1$, $r_{00}^{(2)} - r_{00}^{(1)} < 1$ and $r_{00}^{(2)} > -1$. From Table 4, it can be seen that the three conditions are satisfied by the regression estimates of the parameters. The condition that is closest to being violated is $r_{00}^{(1)} + r_{00}^{(2)} < 1$. In order to study further the issue of whether the historical yield curve data supports the existence of weak reversion to the mean level of interest rates, several additional regressions were performed in which the value of $r_{00}^{(1)}$ was fixed and the estimate of $r_{00}^{(2)}$ was determined. The value of $r_{00}^{(1)}$ was varied in increments of 0.01 across the interval [0.80, 1.25], and in all cases, the value of $r_{00}^{(1)} + r_{00}^{(2)}$ varied between 0.94 and 0.96, strongly confirming that the data supports the existence of weak reversion to the mean level of interest rates.

The importance of mean reversion had a strong bearing on the manner in which the model specification and estimation were carried out. A casual inspection of the autocorrelalion function for the time series of the natural logarithms of the level coefficients suggested that differencing the time series would be necessary to achieve stationarity. First - order differences are often taken in an autoregression analysis of time series of interest rate data, but such an approach forgoes the possibility of mean reversion. Over the short run, the dynamics of interest rates are rather insensitive to the presence or absence of weak mean reversion, but over the long run, interest rate movements will be stable only if mean reversion is present, even if very weakly. Accordingly, great care was exercised in order not to difference any of the time series "prematurely" - in actuality, none of the series needed to be differenced at all !

A simulation test using the last two yield curves in the historical sample of 101 yield curves as the initial conditions was conducted in order to examine the stability of the model over a very long period. Yield curves at four - week intervals for 100 years (1,300 yield curves) were generated using the regression model parameters, standard deviations, and contemporaneous correlation matrix of residuals that were reported in Section 3 . Each yield curve was transformed into its corresponding term structure - namely, spot rates of interest - and then those spot rates were transformed into forward rates of interest in order to verify that all spot and forward rates associated with the 1300 simulated yield curves were positive.

The level of interest rates in the 100 - year simulation ranged from 4.24% to 13.41%. The minimum and maximum spreads between the 30 - year yield and the three - month yield were - 2.52% (inverted yield curve) and 4.80% (normal yield curve). The average spread was 1.86%. Only 41 of the 1,300 simulated yield curves were inverted. These statistics for the simulated sample of yield curves show that the reversion to the mean level of interest rates is somewhat too strong because the simulated range of variation is smaller than has actually been observed in the U.S. Treasury markets. Similarly, the frequency of inverted yield curves is much too small, and the range of spreads between 30 - year yields and three - month yields is modestly too large, relative to actual interest rate history in the United States. Such deficiencies in the estimated stochastic model should not be surprising for two reasons : first, the model is based on a short (400 weeks) and partly uncharacteristic (only eight inverted yield curves) history of yield curves, and second, the standard errors of the regression parameter estimates are not small.

The model parameters, standard deviations, and correlation matrix of the residuals can be adjusted appropriately to produce simulations with characteristics suited to the purpose of a particular application. For example, the sum, $r_{00}^{(1)} + r_{00}^{(2)}$ can be set closer to unity in order to increase the range of interest rate variation. Also, the standard deviation of the e_1 residuals can be increased, and other parameters adjusted to increase the a'_1 * fixed point, in order to increase the probability of inverted yield curves.

5 - SUMMARY AND CONCLUSIONS

The method of orthonormal polynomials was used to decompose U.S. Treasury yield curves into their constituent parts. The method has the advantages of avoiding arbitrariness in identifying yield curve shape characteristics and in being able to fit very closely any yield curve, no matter how perversely shaped, provided a sufficient number of fitting coefficients is used. In practice, very few components and associated cœfficients are needed to fit yield curves accurately enough for purposes of asset/liability management. Utilizing only four components, an average root mean square error of fit of less than five basis points was achieved over the data sample studied. The four components, labelled by reference to their shape characteristics, are : level, tilt, warp, and undulation. The results of the yield curve analysis suggest that using four risk indexes or "duration" measures corresponding to the four primary shape attributes should be sufficient to obtain nearly full asset/liability immunization against changes in the yield curve.

A time series analysis of the first four fitting coefficients was conducted, and it led to the specification and estimation of a second - order vector autoregressive model of yield curve dynamics. The two most si cant findings of the regression analysis are : (1) the existence of mean reversion in the fitting coefficients, with the result that yield curves tend to revert over the long run to a normal (positively - sloped) form centered on a level estimated to be near 8.3%, and (2) the existence of fat-tailed distributions of the residuals of fit to the tilt, warp, and undulation cœfficients that are explainable by a mixture of two normal distributions. Unfortunately, the values of many of the regression parameters and residual covariances do not appear to be stable over time.

The main purpose of this paper has been to expose a methodology for constructing a stochastic yield curve model appropriate for asset/liability simulations of financial institutions and financial security schemes. It is hoped that others will apply the methodology to develop yield curve models for other countries. A promising area of further research is immunization studies of both fixed and interest-contingent assets and liabilities in order to test the hypothesis that indexes associated with the level, tilt, warp, and undulation shape characteristics of yield curves indeed capture most of the risk of changes in interest rates.

APPENDIX 1

A good reference for the material in this appendix is Morse and Feshbach [6].

The orthonormal polynomials $q_n(x)$ introduced in Section 2 are given by $q_n(x) = (2n+1)^{1/2} Q_n(x)$ for n = 0,1,..., where the $Q_n(x)$ are related to the well - known Legendre polynomials $P_n(x)$ through the equation $Q_n(x) = P_n(1-2x)$. Assuming a unit weight function, the Legendre polynomials are mutually orthogonal on the interval [-1, 1], and the polynomials $Q_n(x)$ are mutually orthogonal on the interval [0,1], as expressed by

$$\int_0^1 Q_m(x)\,Q_n(x)\,dx = 0, \quad \text{for } m \neq n.$$

The normalization integrals for the $Q_n(x)$ are

$$\int_0^1 [\,Q_n(x)\,]^2\,dx = (2n+1)^{-1}$$

The $Q_n(x)$ take on the following values at the boundaries : $Q_n(0) = 1$ and $Q_n(1) = (-1)^n$. The polynomials $Q_n(x)$ satisfy the following recursion relation :

$$(n+1)\,Q_{n+1}(x) + n\,Q_{n-1}(x) = (2n+1)(1-2x)\,Q_n(x).$$

The first eight orthogonal polynominals $Q_n(x)$ are :

$$Q_0(x) = 1$$
$$Q_1(x) = 1-2x$$
$$Q_2(x) = 1-6x+6x^2$$
$$Q_3(x) = 1-12x+30x^2-20x^3$$
$$Q_4(x) = 1-20x+90x^2-140x^3+70x^4$$
$$Q_5(x) = 1-30x+210x^2-560x^3+630x^4-252x^5$$
$$Q_6(x) = 1-42x+420x^2-1680x^3+3150x^4-2772x^5+924x^6$$
$$Q_7(x) = 1-56x+756x^2-4200x^3+11550x^4-16632x^5+12012x^6-3432x^7.$$

APPENDIX 2

The condition for the fixed point of the second - order vector difference equation

$$a'_t = k + R^{(1)}\,a'_{t-1} + R^{(2)}\,a'_{t-2} \qquad \text{(A-1)}$$

to be asymptotically stable is derived by first expressing the difference equation in state form as a first-order vector difference equation. This is accomplished by defining b_t as a four elements by a'_{t-1}. Similarly, define k' as a column vector of length eight, with the top four elements given by k and the bottom four elements equal to zero. Finally, define an eight-by-eight matrix R, partitioned into four - by - four matrices as follows: $R^{(1)}$ is in the upper left. $R^{(2)}$ is in the upper right, I (the identity matrix) is in the lower left, and 0 (the zero matrix) is in the lower right. Equation (A- 1) can now be expressed in state, form as:

$$b_t = k' + R\, b_{t-1}. \qquad (A\text{-}2)$$

Starting with an arbitrary initial vector b_0, Equation (A-2) is iterated n times to give

$$b_n = \left(\sum_{m=0}^{n-1} R^m\right) k' + R^n b_0. \qquad (A\text{-}3)$$

The limit of Equation (A-3) as n approaches infinity will be finite, and will be independent of the initial vectot b_0, if and only if the limit of R^n as n approaches infinity is the zero matrix. The derivation is continued under the assumption that all the eigenvalues of R are distinct. Then, if T is the modal matrix having columns equal to the eigenvectors of R, it follows that T is nonsingular and that $T^{-1} R T = D$ where D is the diagonal matrix with the eigenvalues of R along its principal diagonal. Because $R^n = T D^n T^{-1}$, a necessary and sufficient condition for the limit of R^n as n approaches infinity to equal the zero matrix is that all the eigenvalues of R lie inside the unit circle in the complex plane ; namely, that all the eigenvalues of R have modulus less than unity. In that situation, the limit of Equation (A-3) is the asymptotically stable fixed point $b^* = (I - R)^{-1} k'$, or, in the equivalent form displayed in Section 4, $a'^* = (I-R^{(1)}-R^{(2)})^{-1}k$.

If not all the eigenvalues of R are distinct, one proceeds by determining the nonsingular matrix S such that $S^{-1} R S = J$, where J is in Jordan Canonical Form. The proof continues in a manner similar to the case when R is diagonalizable. The result is the same : the fixed point is asymptotically stable if and only if all the eigenvalues of R lie inside the unit circle in the complex plane.

REFERENCES

[1]) K. P. Sharp. "Stochastic Models of Interest Rates". Proceedings of the 23rd International Congress of Actuaries - Subject 5, Helsinki, 1988.

[2] G. J. Harris. "Multi-factor Duration and Immunization". Applied Research 6, J. P. Morgan Investment Management Inc., 1986.

[3] T. J. Coleman. L. Fisher, and R. G. Ibbotson. U.S. Treasury Yield Curves 1926-1988. Moody's Investors Service Inc., 1989.

[4] S. J. Kon. "Models of Stock Returns - A Comparison". The Journal of Finance 39, 1984.

[5] J. Cox, J. Ingersoll, and S. Ross. "A Theory of the Term Structure of Interest Rates". Econometrica 53, 1985.

[6] P. M. Morse and H. Feshbach. Methods of Theoretical Physics Part 1. McGraw-Hill Book Company, 1953.