# A New Approach to the Decomposition of Yield Curve Movements for Fixed Income Attribution

*This paper presents a new decomposition algorithm that is robust and straightforward to understand, and that does not have the drawbacks of existing techniques.*

**Andrew Colin, Ph.D.**
*is currently Head of Fixed Income Research for the StatPro Group and Adjunct Professor in the Faculty of Business at Queensland University of Technology, Brisbane. He was educated at Glasgow Academy and Sussex University. In 1987 he was awarded a Ph.D. in Mathematics from St. Andrews University in astrophysical magneto hydrodynamics. After holding a post-doctoral appointment, he moved to Citibank London, heading research into uses of artificial intelligence for foreign exchange and futures trading. Andrew has worked or consulted for the Commonwealth Bank in Sydney, Zurich Funds Management, Suncorp Metway, Chubb Security, Arthur Andersen, EDS, Alcatel and the Royal Australian Navy. He is the author of* Fixed Income Attribution *(John Wiley & Sons, 2005).*

**Mathieu Cubilié**
*is Product Champion for StatPro's Fixed Income bond performance measurement and attribution system. Mathieu holds a postgraduate degree in Market Finance and Economics from Toulouse University and began his career at Sinopia-Asset-Management, HSBC-group quantitative management subsidiary, in Paris, where he experienced bond portfolio management.*

**Frédéric Bardoux**
*is the Product Strategy Director at StatPro. He manages a team of Product Champions and is responsible for defining the overall product strategy as well as the functional content of the StatPro products. Before his appointment as Product Strategy Director in 2003, Frédéric held the post of Head of Performance Measurement and Managing Director of StatPro France S.A. Prior to joining StatPro in 2000, Frédéric worked for six years at BNP Gestion in Paris where he began his career as a statistical modeler and a fund manager. He then set up and directed the BNP Performance Measurement and Reporting team.*

## INTRODUCTION

The quantitative decomposition of changes in the yield curve into sub-movements that correspond to the terminology used by portfolio managers is a fundamental problem in fixed income attribution. Existing techniques produce results that are highly data- or interval-dependent, with the result that attribution in terms of curve movements is not typically comparable between different managers. More important, the most commonly used curve decompositions do not allow interpretation of results in terms that correspond to the risk decisions that actually drove the portfolio's performance.

*Yield Curve Terminology*

Historically, yield curve movements show a wide variety of forms, such as positively or negatively sloped, flat, humped or inverted (Nelson and Siegel, 1987). Accordingly, changes in the shape of the curve show a correspondingly complex set of behaviors. For instance, Figure 1 shows the month-on-month changes in the Euro yield curve for 2004.

For arbitrage and pricing purposes, market participants use models of the observed term structure that are as accurate as possible. This is usually accomplished by use of smoothing functions, including polynomials, splines, and other special functions.

For fixed income attribution purposes, however, there are a quite different set of requirements: namely, the ability to describe curve movements in commonly used terms. This assumes familiarity with the approaches

**Figure 1: Month-on-Month Changes in Euro Yield Curve During 2004**

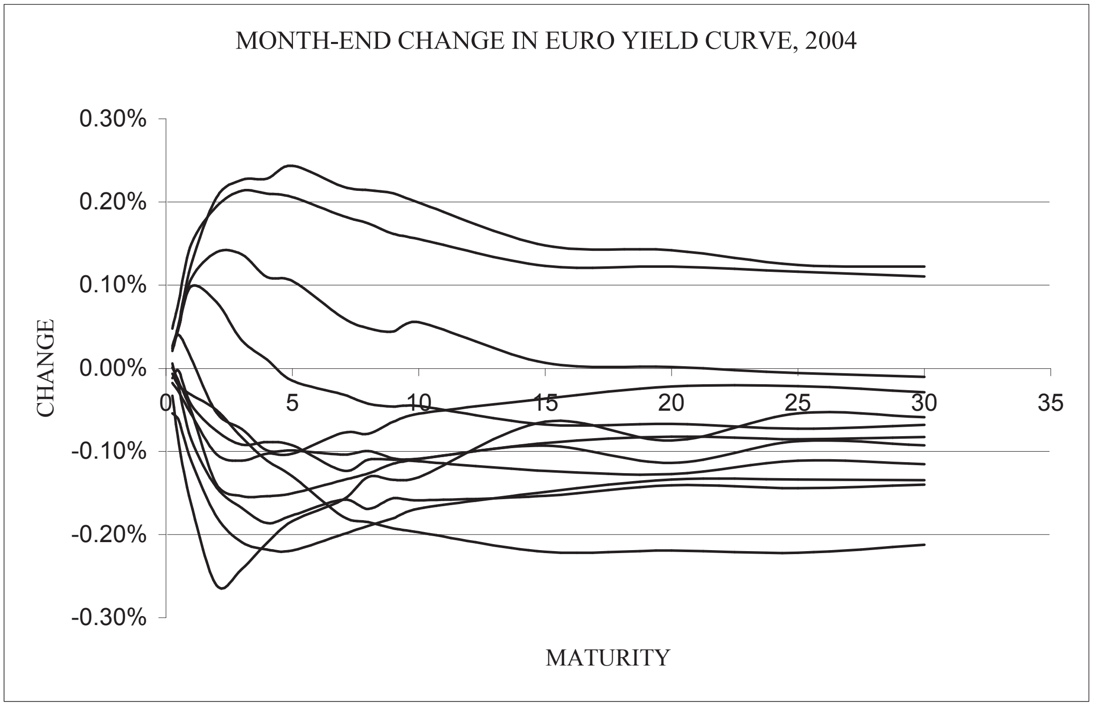


used by traders for hedging risk. For instance, if one of a portfolio manager's trading decisions is to bet on an anticipated steepening of the curve in the 2-5 year region, then an attribution system should be able to describe and measure any such steepening in precisely the same terms so that the returns arising from this decision can be isolated and allocated appropriately.

While there are many ways to describe yield curve movements, the majority of portfolio managers use the terms *shift, twist,* and *curvature* (or butterfly).

- *Shift* is a parallel movement of the curve across all maturities

- *Twist* is a steepening or flattening of the curve

- *Curvature* occurs when a curve becomes more or less humped; in other words, when yields at intermediate maturities change with respect to those at short and long maturities.

Despite the widespread use of these terms, there is no clear definition of what they actually mean, and as a result there are many different ways of assigning numerical measures to them. However calculated, the three effects should add up to the net change in the yield curve, with change types that cannot be conveniently assigned to any of the three categories usually being included in curvature.

## PRINCIPAL COMPONENT ANALYSIS

One commonly used technique to classify such yield movements is the use of principal component analysis (PCA). By examining a large number of historical yield curve changes, a small set of basis functions are determined that can be linearly combined to represent these curve movements in the most economical way. This is accomplished by forming the variance-covariance matrix $V$ from the sample of spot rate changes at the $N$ maturities selected. If we then calculate the $N$ orthogonal eigenvectors of $V$ and rank by order of eigenvalue size, the highest ranked eigenvector forms a basis function that explains as much as possible of the observed curve motion in terms of a single vector. By using a combination of this vector and lower ranked eigenvectors, the underlying data can be approximated to any degree of accuracy required.

The variances of the principal components are given by the magnitudes of the eigenvalues, so that the eigenvec-

**Figure 2: Relative Sizes of Eigenvalues, from Principal Component Analysis on Monthly Changes in Euro Curve, January 2000 to May 2005**

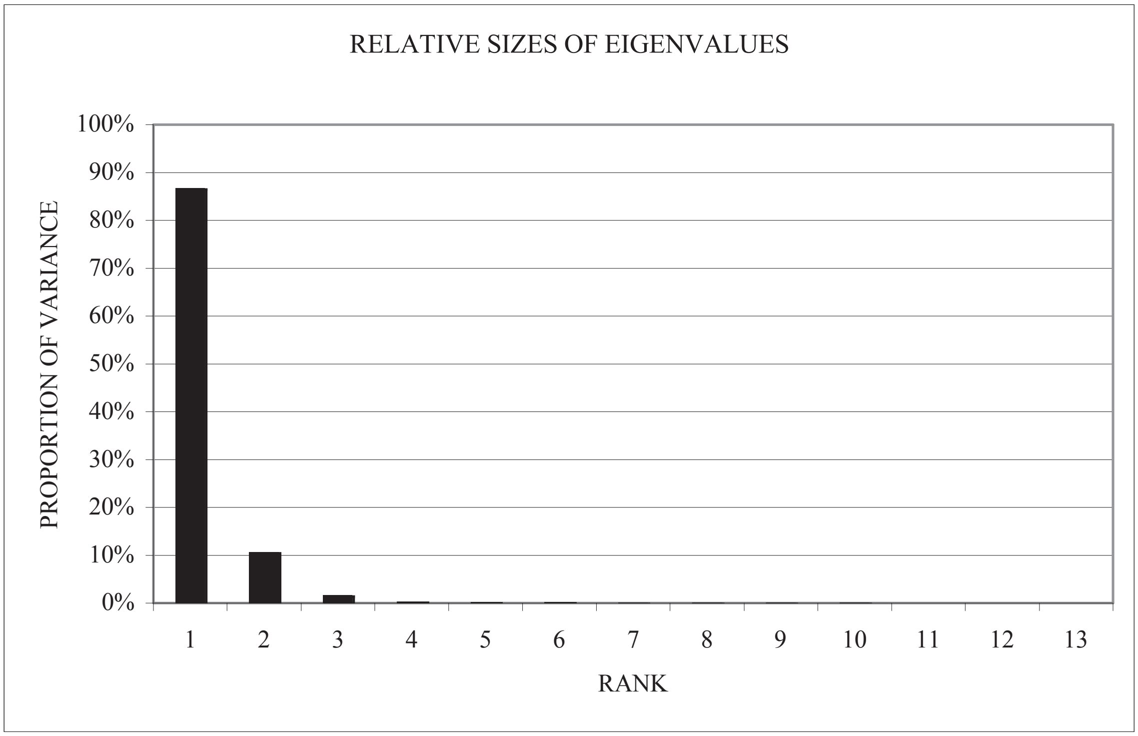


**Figure 3: Shapes of Eigenvector Functions for Three Highest Ranked Eigenvectors from Monthly Changes in Euro Curve, January 2000 to May 2005**

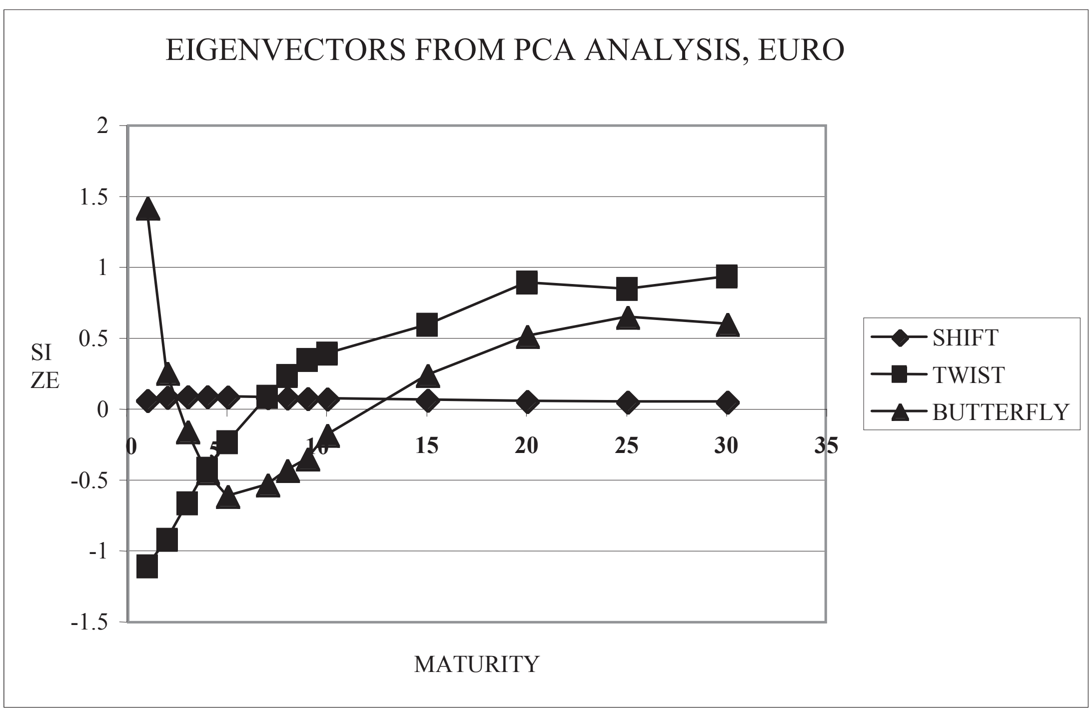

tor with the highest value has the most explanatory power on the underlying data. If the values of the majority of eigenvalues are low, then this indicates that the underlying data can be closely modeled by a small number of functions, which represent some underlying structure in the data. PCA is therefore a useful technique for reducing the dimensionality of a modeling problem. In particular, PCA has been found to work well on yield curve changes (Phoa, 1998; Barber, Copper, 1996), since in practice practically all yield curve changes can be closely approximated using linear combinations of the first three eigenfunctions from a PCA.

PCA on historical yield curve data typically shows that curve movements fall into a number of fairly clearly defined types. Figure 2 shows the relative sizes of the eigenvalues from a PCA for sample yield curve changes for the Euro at monthly intervals from 31st January 2000 to 31st May 2005, where the sample maturities were (1, 2, 3, 4, 5, 7, 8, 9, 10, 15, 20, 25, 30 years). In this case, practically all the movement in these yield curves can be modeled by the first three eigenfunctions, which account for 98.9% of the observed variance. Figure 3 shows the shape of the first three eigenfunctions.

These results are consistent with the PCA studies mentioned above. Specifically, they all show that the first eigenfunction is close to a flat line, that the second rises monotonically (but is seldom a straight line), and that the third imposes some curvature motion. These functions are usually interpreted as shift, twist, and curvature.

To date, most commercially available attribution systems of which the authors are aware have used some form of PCA model to decompose curve movements. Since PCA generally attributes at least 85% of curve changes to the first eigenfunction, this has lead to a general expectation that the bulk of yield curve attribution returns should be generated by shift movements, with only a small fraction from twist and an even smaller fraction from curvature. In view of the subjective nature of these movement types, this is not always easy to achieve but is one of the reasons for the non-PCA

**Figure 4: Shapes of Eigenvector Functions for Three Highest Ranked Eigenvectors from Monthly Changes in Euro Curve, January 2000 to May 2005, with Maturities Including 0.25 and 0.5 Years**

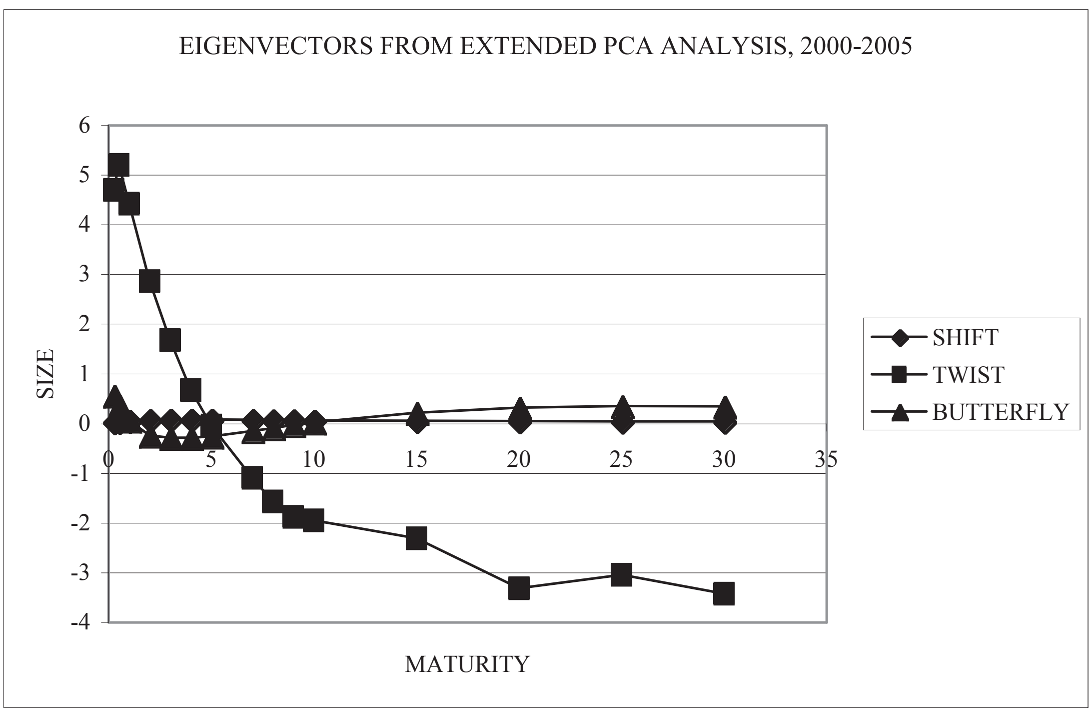

approach described in this paper.

The PCA-based approach is mathematically elegant and assigns returns where the user expects them. However, it has a number of significant drawbacks:

1) The results are dependent on the data set used to construct the basis functions and the interval over which they were constructed. If two different managers use different datasets, they will get different attribution decompositions. This may make it difficult to compare attribution returns. For instance, Figure 4 shows the same PCA study as in Figure 3 but with the maturity interval extended to include 3- and 6-month maturities at the short end of the curve. The shapes of the various functions have changed quite drastically.

2) While the shape of the first two basis functions are approximately zero and first order polynomials respectively, the match is not exact. Figures 3 and 4 show that the shape of the shift function is not a flat line, indicating that the eigenfunctions used to describe curve movements will be different to those in which the risk decisions were expressed. For instance, modified duration (the most widely used measure of interest rate sensitivity) is a response to a parallel movement in a yield curve across all maturities. If the first eigenfunction is *not* a flat line (which is in fact the case in all the above studies), then attribution results will not be presented in terms of the risk decisions made, which is a fundamental requirement for attribution.

3) The results and implications of a PCA are hard to explain to users without a background in mathematical statistics. Few portfolio managers are inclined to make trading decisions in terms of eigenfunctions.

**MODEL FITTING**

Another approach to assigning magnitudes to various curve change types is to fit a function with a low number of parameters to the raw curve data. Typically, these functions include polynomials and Nelson-Siegel functions (Nelson and Siegel, 1987; Diebold and Li, 2005; Colin, 2005).

From an attribution perspective, the major problem with this approach is that the magnitude of the twist and curvature components is often much larger than the user expects, since the simple one or two parameter models used to measure twist are often a poor fit to the observed data. In addition, the use of least-square (LSQ) techniques for function fitting often leads to model parameters that are numerically sensitive to changes in underlying data, which leads in turn to twist and curvature measurements that can vary widely in relative size over time. There are also uniqueness problems in defining shift and twist measurements, requiring that the user set a twist point; this can lead in turn to non intuitive results (see Colin, 2005).

A further difficulty with these approaches is that they attempt to model the dynamics of the entire term structure using a small number of parameters. There is extensive anecdotal evidence to suggest that the dynamics of the cash market (or bill curve), where maturities are less than 2 years, is quite different to that of the longer-term market (or bond curve), with maturities ranging from 5 to 30 years. For instance, the bond curve may steepen but the bill curve may well remain flat, or even flatten. In practice, this suggests that different models should be fitted to different parts of the term structure. We cover this point in more detail below.

**PIECE-WISE INTERPOLATION**

In this paper, we propose representing twist motions in terms of *noncontinuous linear functions*. Equivalently, we represent different portions of the yield curve using different linear functions. The points at which one curve stops being used and is replaced by another is referred to as a *parameter change point*. There is no requirement for these functions to match at such points, which simplifies their calculation quite considerably.

The twist of the curve at a given maturity is then the change in the appropriate linear function at that maturity, after adjusting for any shift movements.

The algorithm below provides for data in each regime being pre-fitted to a Nelson-Siegel function. This provides some elementary smoothing and allows us to isolate the effects of market noise on the term structure. The use of such a smoothing function is not necessary,

**Table 1: Curve Decomposition Algorithm**

| | |
|---|---|
| *Step 1* | Fit a Nelson-Siegel (NS) function to the curve at each interval. For instance, if we have defined a parameter change point at 5 years maturity, we fit two NS curves to the raw data over the intervals [0, 5] and [5, 30].<br><br>An NS function has at most one turning point, which is usually enough to model the broad characteristics of the part of the market we want to model, without including extraneous detail. |
| *Step 2* | At each maturity, calculate the difference between the NS functions over each interval. This set of differences (or curve changes) is what we will model in the remainder of the paper. |
| *Step 3* | Calculate the average difference of the NS functions $\delta NS_i$ over all intervals. This gives the curve shift, denoted by $S$. Curve shift is independent of interval.<br><br>*Note:* Shift must be a constant linear displacement over all intervals for consistency with the widespread use of modified duration as a risk measure, which measures the price sensitivity to curve movements of this type. |
| *Step 4* | Using a least-squares fit, calculate a set of curve parameters $\{a_i, b_i\}$ for each interval so that each function $a_i + b_i t$ models the change in smoothed yields as closely as possible over interval $i$. Then the quantity $a_i + b_i t - S$ gives the twist of the curve at maturity $t$.<br><br>*Note:* In this scheme, twist can be positive in one part of the curve and negative in another, so that the overall contribution to a portfolio's return from twist may cancel out. |
| *Step 5* | At each sample maturity, subtract the shift and twist from the overall smoothed change. This gives the curvature, $\delta NS - (a_i + b_i t - S)$. |

however, and the algorithm can equally well be applied to raw data.

The algorithm assumes that curve data has been supplied at a set of predefined maturities on the yield curve.

This process is shown graphically in Figures 5 and 6.

Figure 6 shows the calculation of shift, twist, and curvature: shift is the average displacement of the current part of the yield curve. Twist is the difference between the shift movement and the line fitted to the curve change. Curvature is the difference between the shift and twist movements and the smoothed curve data, and noise is the difference between the smoothed yield curve data and the true yield curve data.

*Example*

Using end-of-month data for 65 Euro zero-coupon curves between February 2000 and May 2005, with a single parameter change point at 5 years, we fitted a Nelson-Siegel function to the interval 5 - 30 years. The steps in the above algorithm were followed. The results at the 10-year maturity points are shown in Figures 7 and 8, together with the market noise, which is the difference between the Nelson-Siegel function value and the true market yield.

The results show that the bulk of changes are assigned to shift movements. Figure 9 shows the sum of the

**Figure 5: Calculate the Difference Between the Two Smoothed Function; Fit a Straight Line to the Two Portions 0 - 5 years, 5 - 30 years**

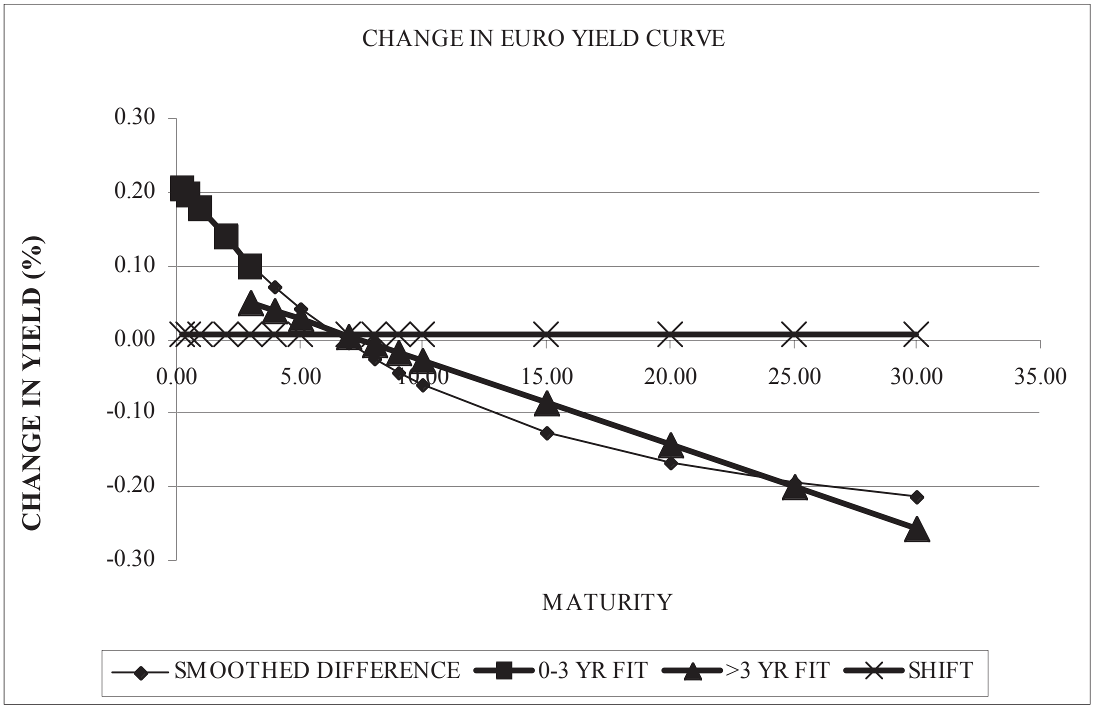


**Figure 6: Calculation of Shift, Twist, and Curvature. (The shift movement has been exaggerated in this diagram for clarity.)**

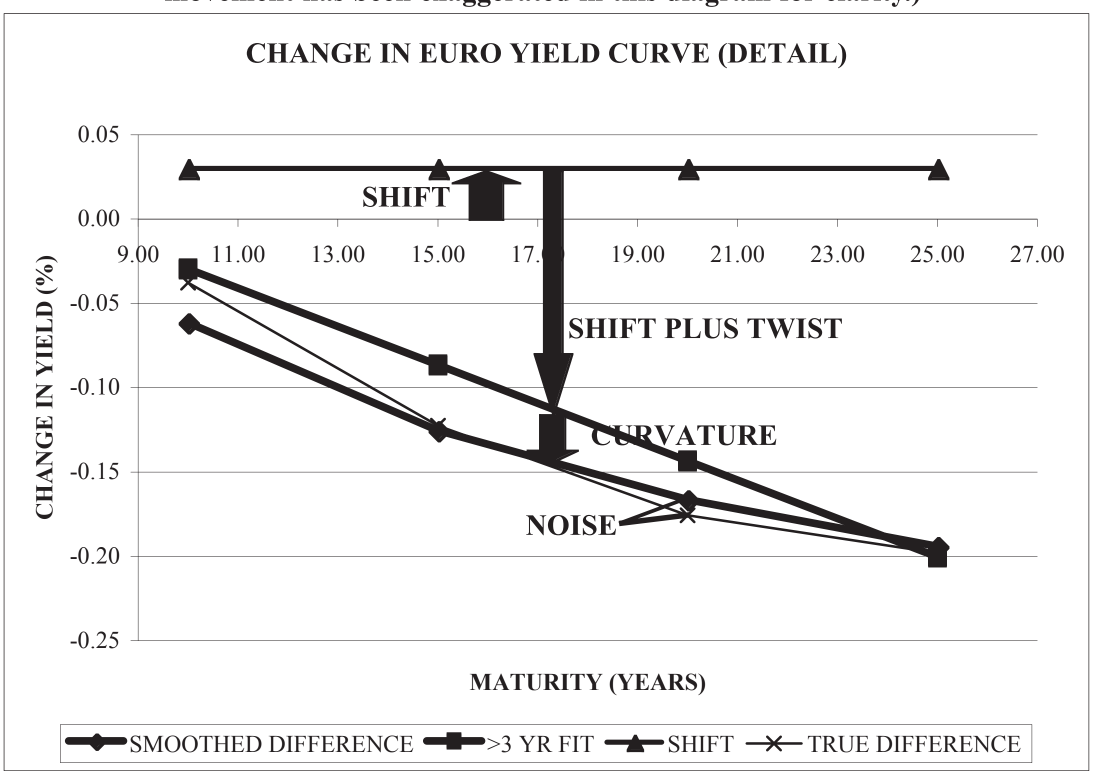

**Figure 7: Euro Yield Curve Decomposition at 10 Years Maturity, January 2002-December 2004**

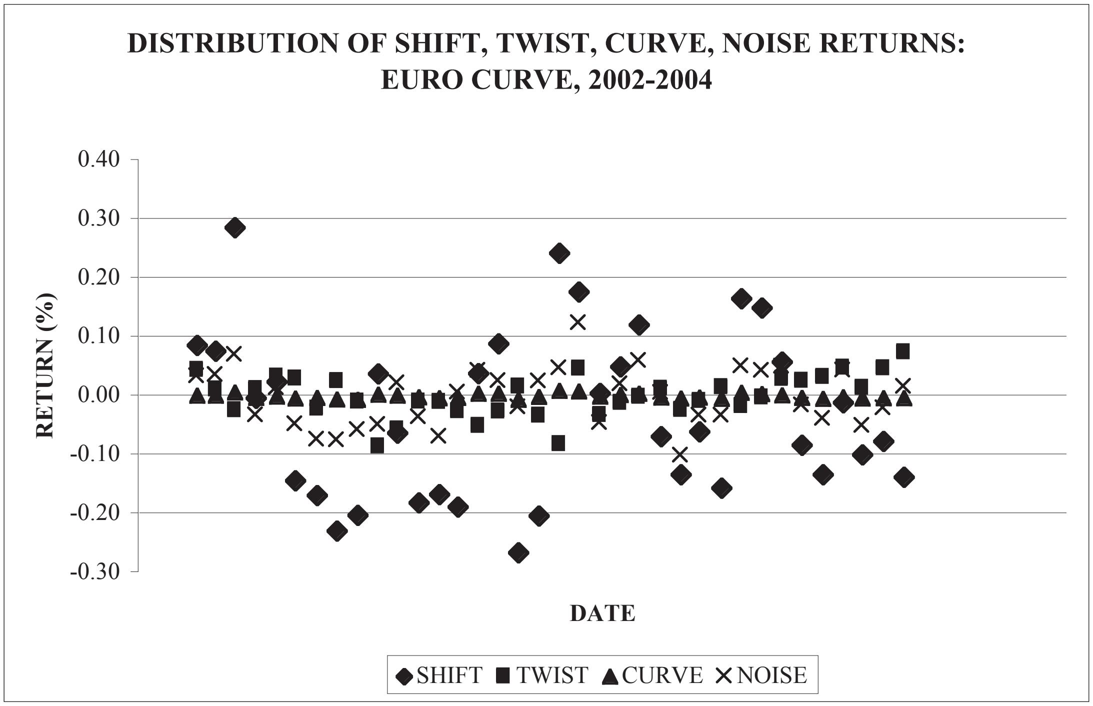


absolute values of the various changes, added over the interval January 2002 through December 2004.

Where there are large non-shift changes, we have verified that the values measured are appropriate. For instance, between 31st January 2000 and 28th February 2000 this calculation gave values of (-0.0572, -0.0871, -0.0195) for shift, twist, and curvature at the 20-year point. Figure 9 shows the graph of the yield curve over this interval, indicating that there was indeed a large curve flattening.

The advantages of this scheme are that

- it is simple to implement and to explain
- it is not data-dependent in that results do not depend on the sample interval used
- the bulk of curve changes are assigned to shift and twist motions. However, when there is a genuine curvature change, this will appear. The scheme therefore bypasses many of the problems of numerical instability that are seen when fitting continuous models
- unlike PCA, the shift motion is always a parallel movement across all maturities, so that the effect of modified duration hedging is shown in the attribution analysis.

The scheme does require the user to specify an appropriate set of parameter change points. Initially, we suggest the use of a single point lying between 2 and 5 years maturity. The number and placement of other points is a subject for future research.

An objection to this scheme is that it has the potential to miss out global curvature movements. For instance, suppose that the yield curve change rises linearly from 0 to 2 years, and then falls linearly from 2 to 30 years, with a parameter change point at 2 years. After removing the upwards-parallel shift component of this movement, we are left with positive twist return in the 0 - 2 year part of the curve, negative twist return in the 2 - 30 year part of the curve, and zero curvature return.

Extending this example, one could construct a curve movement that gave exactly zero net twist and curvature return over the entire term structure, despite having a strongly nonlinear shape.

**Figure 8: Summarized Curve Change Contributions, January 2002-December 2004**

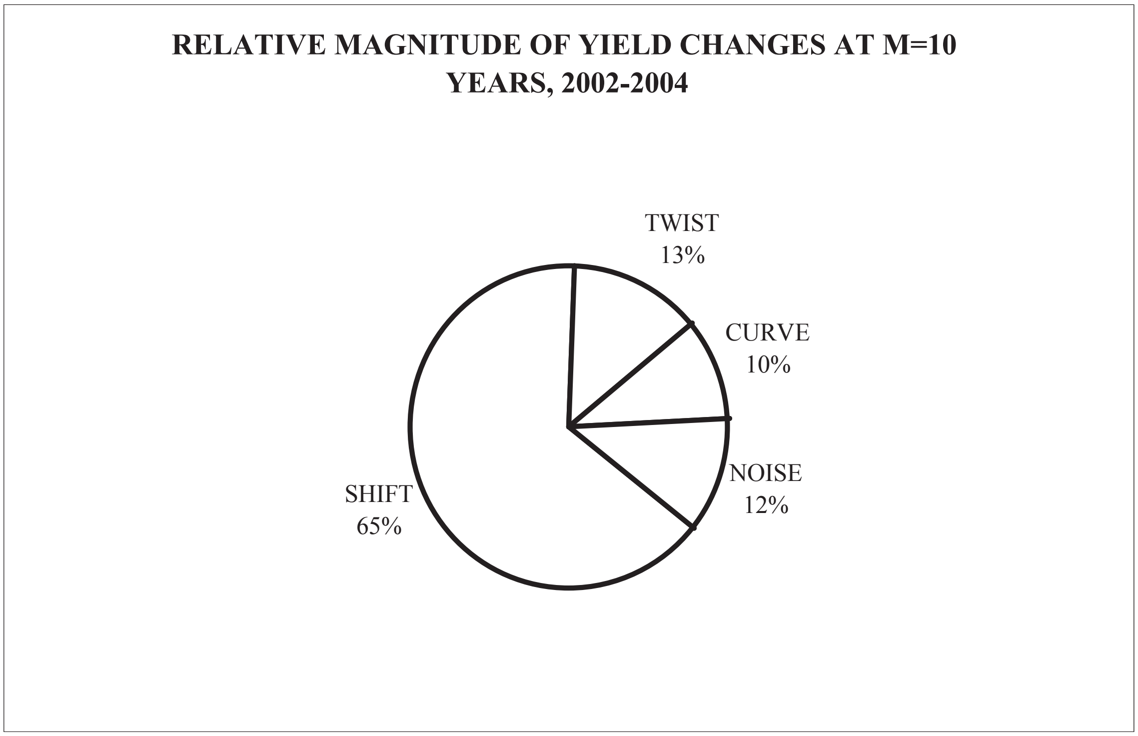


**Figure 9: Euro Yield Curve, 31/01/2000 and 29/02/2000**

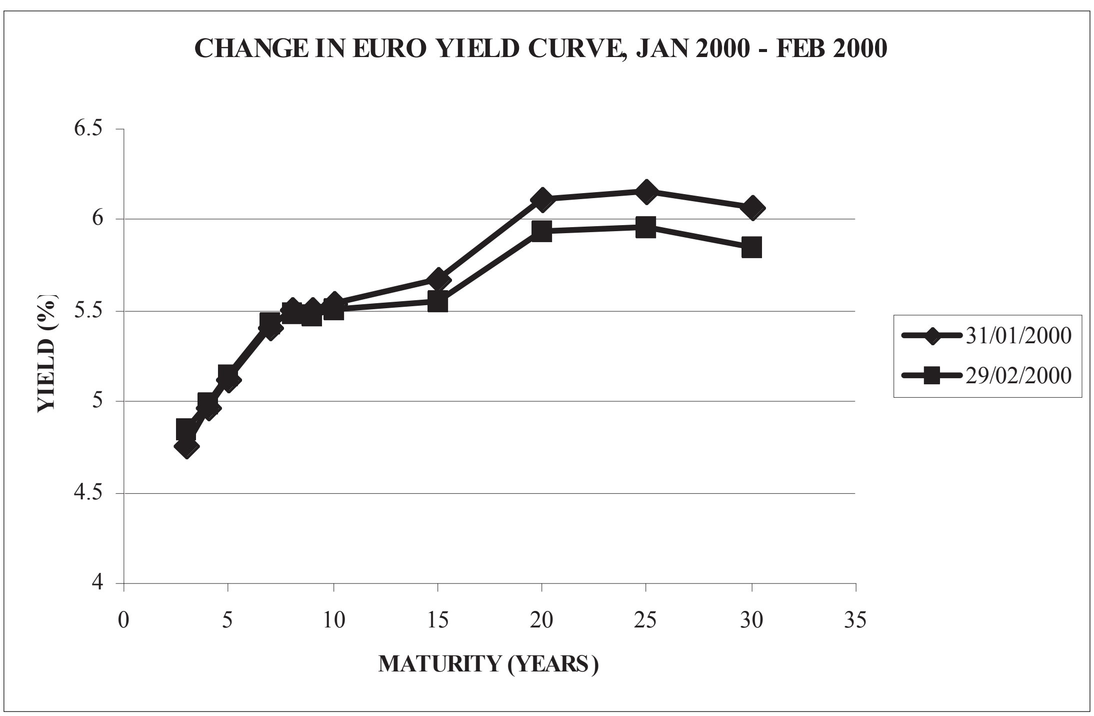

Whether this is a problem depends on how broad a view one takes of the yield curve. For a manager whose risk exposures lie primarily in longer-dated securities, it makes better sense to ignore short-term fluctuations in the cash market curve when positioning a portfolio or running an attribution analysis. Similarly, a trader whose entire exposure is in bank bills and floating-rate notes (FRNs) will form a view of the yield curve that is driven much more by central bank rate movements than by movements in the 30-year bond yield. In both cases, treating different parts of the yield curve as falling in different parameter regimes makes perfect sense.

This result is therefore consistent with our view that different maturities are subject to different parameter regimes, and that the presence of opposite signs of twist between these portions of the curve do not correspond to a genuine curvature movement. If, on the other hand, this behavior were seen at maturities of (say) 15 years, which is a long way from a parameter change point and in the middle of the "bond" portion of the curve, then we would say that there is genuine curvature occurring, and the algorithm will pick this up.

While this can happen, the placement of the parameter change points means that for this to happen the point of global curvature actually lies near the parameter change point, and their placement has been specifically designed so that this does not occur.

It should be noted that the same objection applies to twist functions generated from PCA, which are usually significantly more complex than a simple straight line. In this case the function will encapsulate some curvature as well as twist, leading to the same problem. However, PCA does not allow the user to select parameter change points.

## CONCLUSION

The approach presented in this paper to curve change decompositions circumvents the difficulties inherent in the principal-component and model fitting approaches in a simple, intuitive way that is both transparent and straightforward to communicate. The approach requires determination of points at which the parameter regime for the term structure changes, which may need further research.

## REFERENCES


Barber, Joel R., Mark L. Copper, "Immunization Using Principal Component Analysis," *Journal of Portfolio Management*, Fall 1996.

Colin, Andrew M., Fixed Income Attribution, Wiley, 2005.

Diebold, Francis X., Canlin Li, "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, in press, 2005.

Nelson, Charles R, Andrew F. Siegel, "Parsimonious Modeling of Yield Curves," *Journal of Business 60*, 473-489, 1987.

Phoa, Wesley, "Advanced Fixed Income Analytics," Frank J. Fabozzi Associates, 1998.


## CONTACT INFO


Andrew Colin can be contacted using andrew.colin@statpro.com.

Mathieu Cubilié can be contacted using mathieu.Cubilie@statpro.com.

Frédéric Bardoux can be contacted using frederic.bardoux@statpro.com.