Archetypes of collective yield curve movements

Dominik Robert Dersch

Quantitative Cross Asset Research, Global Markets Research,
Arabellastr 12, Munich D-81925, Germany
E-mail: dominik.dersch@unicreditgroup.de

Abstract: We examine the historical behaviour of interest rate movements in seven major currencies AUD, CAD, CHF, EUR, GBP, JPY and USD. We apply principle components analysis and hierarchical cluster analysis to illustrate, understand and model the past collective movements of yield curves. We show that simple correlations are not able to capture the complex behaviour observed in the data set. In order to model risk factors that are intimately connected, we propose so-called *archetypes* of collective movements as building blocks. Thus, we start from collective movements that are coherent from a historical perspective. A set of risk factor forecasts is then generated by adapting an archetype rather than building single risk factor forecasts from scratch. This approach opens the door to integrated, coherent forecasts created from complex building blocks. The methods may be applied within scenario simulations, forecasting, filtering techniques and technical analysis.

Keywords: interest rate risk management; principle components analysis; PCA; clustering; archetypes; coherent forecasts.

Reference to this paper should be made as follows: Dersch, D.R. (2008) ‘Archetypes of collective yield curve movements’, *Int. J. Services Sciences*, Vol. 1, No. 2, pp.99–114.

Biographical notes: Dominik Robert Dersch is Director, Quantitative Cross Asset Research, Global Markets Research at UniCredit in Munich. His responsibilities cover the development of advanced risk management, asset and liability and portfolio optimisation tools and providing quantitative research services. Previously, he held the position of Head of Research and Development at Crux FE, Australia and Senior Forecasting Analyst in Electricity Trading at Integral Energy, Australia. He holds a Doctorate in Natural Science (LMU Munich). He is founding and acting Regional Director of PRMIA, Munich. His research interests are dynamic portfolio strategies, yield curve modelling, advanced time series analysis and quantitative research.

1 Introduction

The financial services industry is a large driver of innovation with respect to methods and technology. One of the reason is that there are various regulatory requirements imposed on the industry that enforce the implementation of appropriate risk management techniques and procedures. Risk management methodologies that have been initially developed for the banking industry are adopted by non-financial corporations, see for example, Smithson and Simkins (2005) and Glaum (2000). Interest rates (IR) and exchange rates (FX) are the two major market risk factors relevant to corporate treasury. IR plays an important role because corporations have to control and manage their funding and liquidity requirements. FX is relevant whenever the above funds and cash-flows are denominated in different currencies, for example, for cross-border revenues.

Managing IR risk in different currencies requires addressing the problem that future scenarios have to be plausible and consistent with respect to comovements of different markets.

The large dimensionality of the problem requires appropriate techniques to reduce the complexity without the loss of relevant information. The most prominent technique to address this problem is PCA. It has been applied to various modelling and analysis tasks, see for example, Litterman and Scheinkman (1991), Jamshidian and Zhu (1997), Rodrigues (1997) and Lardic et al. (2003).

Considering a set of n risk factors $r_i\ (t)$ we may represent the movement over a period Δt by the absolute change over $\Delta t = t_2 - t_1$ with $t_2 > t_1$

$$\Delta\vec{R}(t) = \begin{pmatrix} \Delta r_1(t) \\ \dots \\ \Delta r_n(t) \end{pmatrix} = \begin{pmatrix} r_1(t+\Delta t) - r_1(t) \\ \dots \\ r_n(t+\Delta t) - r_n(t) \end{pmatrix}. \tag{1}$$

Each historical chance of the n risk factors at a given point in time is therefore represented by a n-dimensional data vector. The set of movements over time form a n-dimensional data cloud. The number of observations defines the number of data vectors in the cloud. Generally, the data cloud may only cover a subspace of the complete space. Areas outside the cloud represent the movements that have not been observed in the past. For the special case of a correlation coefficient of one among all risk factors, the subspace is a one-dimensional manifold, namely the diagonal axis through the origin of the data space. In case the data cloud is embedded in a m-dimensional subspace of lower dimension, for example, $m < n$, PCA is a powerful tool to reduce the complexity of a data set. PCA is a linear base transformation of the orthogonal coordinate systems such that the axes of the new coordinate system are aligned along the directions of largest variation in descending order. The old base A of the n-dimensional space is transformed into the new orthonormal base B:

$$A = A = \{\vec{a}_1, \dots, \vec{a}_n\} \Rightarrow B = \{\vec{e}_1, \dots, e_n\}.$$

The vectors e_i of the new base are the *eigenvectors*. The data vector x is represented in the new base as follows:

$$\vec{x}^b = \sum_{i=1}^{i=n} f_i \vec{e}_i, \tag{2}$$

where the factors f_i are given by the scalar product of the original vector x and eigenvector e_i in the base A according to

$$f_i = \vec{x} \times \vec{e}_i. \tag{3}$$

The set of factors f_i for a component i is called *factor loading*. The standard deviation of the data set along each eigenvector is the corresponding *eigenvalue*. If the eigenvalues quickly drop towards zero for an increasing order of eigenvectors, the sum in Equation (2) may be truncated at $i_{max} < n$ without a significant loss of information.

Thus, PCA may be interpreted as a function approximation by a series of basic functions. Here, the basis functions are derived in a data-driven fashion, rather than predefined, for example, like in a Taylor series expansion. In this paper, a PCA transformation is applied to the movements of IR in seven different currencies.

Representing each yield curve by 13 maturities will result in a 91-dimensional space of collective yield curve movements. The first three PCA components of yield curve movements for a single currency can be interpreted as *level*, *slope* and *curvature*. For a detailed discussion, please see, for example, Littermann and Scheinkman (1991), Jamshidian and Zhu (1997), Rodrigues (1997) and Lardic et al. (2003). Our analysis shows that a similar eigenvector base is obtained, if we perform a PCA on a space that combines movements of all currencies as compared to a PCA on yield movements in a single currency. Here, PCA helps to reduce the dimension of the data space by a factor of over four.

An analysis of the intra- and inter-yield curve factor dynamic shows a significant correlation. However, in a business cycle, we observe regimes of different synchronous and asynchronous risk factor movements. Modelling the co-movements by linear correlations will draw a too much rigid image. Thus, we prefer to focus on co-movements that reappear over time with a given statistical significance and persistence. Such patterns are called *archetypes* (Dersch, 2006). Different archetypes may represent correlated or anti-correlated movements.

In order to derive archetypes from the data set, we apply a second statistical technique namely Hierarchical Cluster Analysis (HCA), Rose et al. (1992). Here, we use a technique that has been thoroughly analysed (Dersch, 1996) and applied in many areas of statistical data analysis, for example, speech and image processing and financial data analysis. The algorithm belongs to the class of hierarchical cluster techniques. It is a powerful tool to find structural information in a high-dimensional data space on multiple length scales by means of a mathematical optimisation process (Dersch and Tavan, 1992).

The main feature of this algorithm is that it maps a set of data vectors $X = \{x_1,\ldots,x_k\}$ onto a smaller set of cluster centres – or as we call them – archetypes. The archetypes represent the local data density on a given length scale. In the course of an annealing process, the data structure is resolved on an increasing finer length scale and areas with a high local data density – become apparent. In this publication, we apply the HCA to identify typical IR movements for different currencies. Besides the static properties of the data set, HCA also provides insight into the dynamic aspects of a process. By assigning a representative label to each archetype, for example, 'rise', 'fall', we are able to generate a sequence of symbolic labels from our data set. This is achieved by assigning each data vector the label of the next neighbour archetype (hard clustering). In a fuzzy clustering case, we assign each data vector with a given probability to a label of a neighbouring archetype. Given the sequence of labels, we may investigate the dynamic behaviour of the process.

The analysis based on HCA provides insight that goes beyond the results obtained from linear correlations. The remainder of this paper is organised as follows: in the next section, we explain the data set and methodology. Sections 3 and 4 show the results for the PCA analysis and HCA, respectively, Section 5 shows the conclusions and an outlook.

2 Data set and methodology

2.1 Data set

We use historical daily data from Bloomberg data services of government zero rates for the period from 30 December 1994 to 2 March 2007 with a total number of 3,165 observations. We consider yield curves in the following currencies: AUD, CAD, CHF, EUR, GBP, JPY and USD on a maturity grid: 1–10 years, 15 years, 20 years and 30 years. For each point in time, we calculate the yield curve movements based on the 60-day absolute changes according to Equation (1) representing approximately three month changes. The collective movements of the seven yield curves are therefore represented by a 91-dimensional data vector (7 currencies multiplied by 13 maturity grid points). In order to reduce the complexity of the data set, we first apply a PCA on the data set.

2.2 PCA

In this publication, we show evidence that a description of the collective movements by a single eigenvector base of the first three components results in a good approximation.

2.2.1 Single PCA space

In order to do this, we first derive eigenvector bases calculated from yield curve movements from each of the 7 currencies. This involves performing individual PCA transformations on seven different data sets of 3,105 13-dimensional data vectors.

2.2.2 Joint PCA space

In order to do this, we first derive eigenvector bases calculated from yield curve movements from each of the 7 currencies. This involves performing individual PCA transformations on seven different data sets of 3,105 13-dimensional data vectors.

We compare the above result with an analysis where all movements are represented by a single eigenvector base. This is achieved by first creating a joint data space of 21,735 ($7 \times 3{,}165$) 13-dimensional vectors and then performing a PCA on the joint data space.

A comparison of the eigenvectors for the two different approaches and of the observed error between the original data and the reduced representation justifies the approach to represent all movements by just one common eigenvector base. A detailed analysis is shown below.

As a result, we may represent multiyield curve movements by the first three components in each of the seven currencies. We therefore achieve a data reduction by a factor of 4.3, that is, from 91-dimensions to just 21-dimensions. In this representation, the movements are approximated by linear combinations of the first three eigenvectors. A joint vector base is a prerequisite for a direct comparison of the factor dynamics between the three factors of each currency as well as between the same factors of different currencies. The analysis of correlations of inter- and intra- factor dynamics gives an insight into the average overall shape and dependencies of yield curve movements.

2.3 HCA

In order to analyse the comovements in different currencies, Jamshidian and Zhu (1997) and Rodrigues (1997) studied the correlation coefficients of the factor dynamics. We focus rather on archetypes that allow us to model movements representing individual correlation patterns.

2 3.1 Static analysis

In the following section, we apply HCA to the 21-dimensional data space of factor dynamics by using 7 archetypes. In the course of the above-described annealing process, the data space is resolved on a successive finer length scale until the set of seven well separated archetypes is obtained. The archetypes are analysed and compared with results derived from simple linear correlations of factor dynamics.

2.3.2 Dynamic analysis

We further investigate the dynamic properties of movements by analysing the time dependent structure of archetype association. This is achieved in the following way. Each data vector is assigned with a given probability to the set of archetypes. The time-series may thus be mapped on a sequence of archetypes labels. The analysis of the label sequence gives an insight into the transition of movements and the appearance of the archetypical patterns in a historic context.

3 Analysis results: PCA

3.1 PCA analysis: single currencies versus joint currency space

In the first analysis, we show the results of single currency PCA compared to a PCA on the joint data space. Figure 1 shows the first eigenvector for the single currency PCA (full lines) and the first eigenvector derived from the joint currency space (broken line).

Figure 1 First eigenvector for individual and joint PCA (see online version for colours)

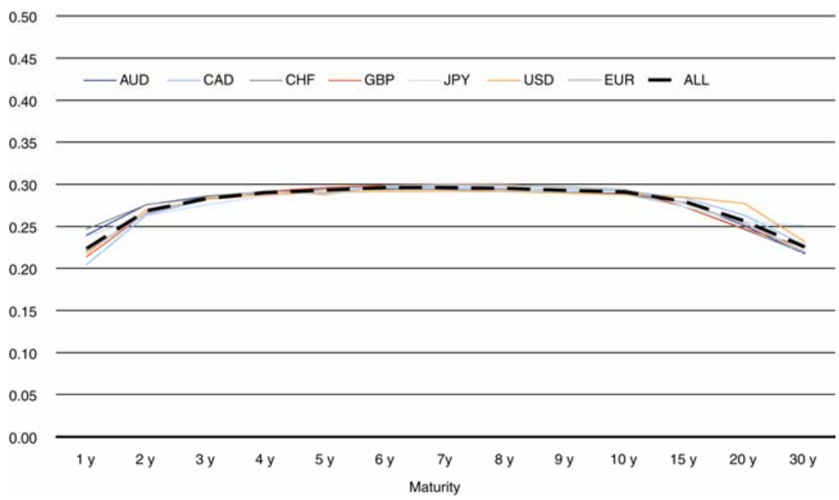

Figures 2 and 3 show the second and third eigenvectors. Please note that, due to the fact that the eigenvectors are orthonormal, all eigenvectors have the same Euclidean length and therefore fall in a similar range.

Figures 1–3 illustrate two important findings:

1. The shape of the eigenvectors follow the well known picture of 'level', 'slope' and 'curvature' already discussed in the literature.
2. The eigenvector of the joint space is a smooth interpolation of the single space eigenvectors.

Figure 2 Second eigenvector for individual and joint PCA (see online version for colours)

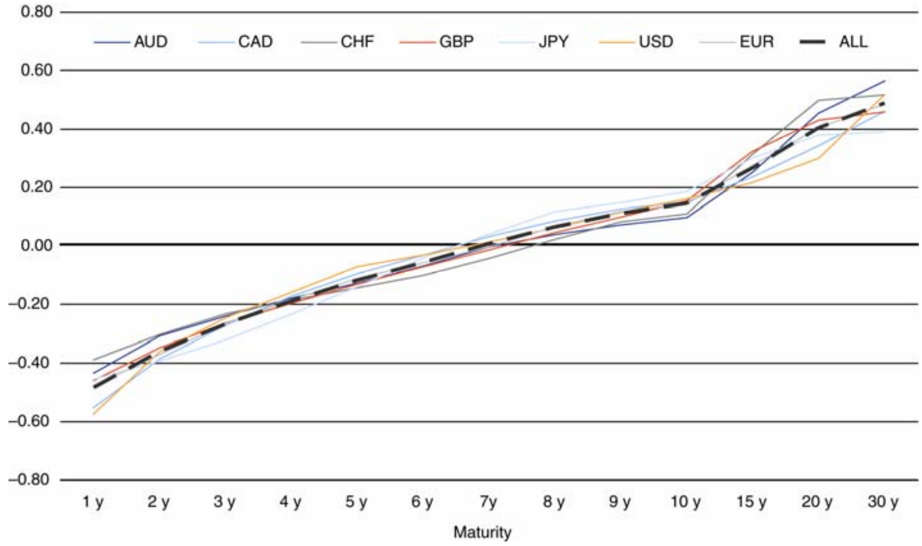

Figure 3 Third eigenvector for individual and joint PCA (see online version for colours)

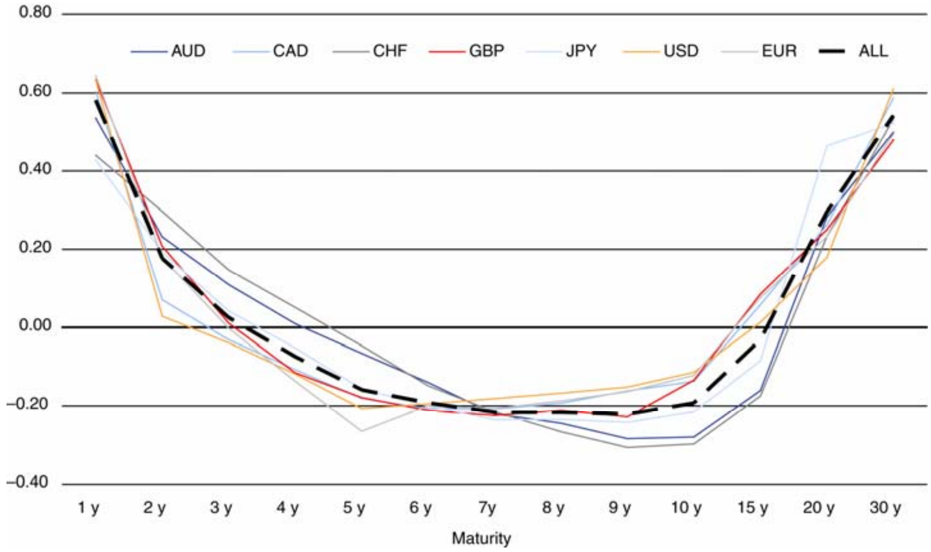

The first result is expected as it has been reported before (see references in the introduction). Nevertheless, it is not obvious as data sets and preprocessing of yield curves described in the literature vary significantly. They cover different phases of business cycles (both pre and post Euro era), different currency sets, different maturity grids and different representations like zero rates or daily rate changes. In our study, we used three month absolute changes on a data set of more than 12 years.

The second finding confirms that the PCA decomposes the movements – independent of the currency into very similar shapes. The third eigenvector shows some variation in the location of the minima. For the AUD and the CHF, the minima are located around a maturity of 10 years, whereas the minimum for the remaining currencies is located at lower maturities. The joint PCA eigenvector smoothly interpolates between these shapes.

We conclude that a common eigenvector base is a good description for the joint movements. This is also confirmed by the analysis of the mean squared error derived from the difference between the original yield curve movement and a reduced representation by the first three components only as described by Equation (2) (i = 1, 2, 3). Table 1 gives the Mean Squared Error (MSE) between the original data set and the reduced representation using the first three PCA components. The first seven columns show the MSE for the PCA of each individual currency, the last column (ALL) shows the PCA on the joint movements.

Table 1 MSE in percent using the first three PCA components for the individual currencies (column 1–7) and the PCA on the joint space (column 8)

Curr	*AUD*	*CAD*	*CHF*	*EUR*	*GBP*	*JPY*	*USD*	*ALL*
MSE	1.55	1.81	2.56	2.17	1.04	1.56	0.53	1.65

The observed error of 1.65% is very close to the average error of the single currencies (1.60%). The standard deviation of the 60-day movements varies for different currencies between 34 bp for the CHF and 53 bp for the AUD. The MSE is therefore in the range of around 1 bp only. The variation of the standard deviation for different maturities within the same currency is relatively low.

As a result of the above analysis, we therefore, use the joint PCA representation for the remainder of this paper.

3.2 Statistical analysis of the joint PCA representation

Table 2 gives the correlation coefficients within the first (upper), second (middle) and third (lower) components of different currencies. The correlations are calculated from the factor loadings f_i (compare Equation (3)) for i = 1, 2, 3. The upper right half of the correlation matrix and the diagonal elements are shaded because they carry redundant information. The dependency between the first, second and third PCA component are discussed in the following.

Table 2 Correlation in percent between first (upper number, $i = 1$), second (middle number, $i = 2$) and third (lower number, $i = 3$) PCA components. The upper right triangle and the diagonal element carry redundant information and are therefore shaded in grey

Curr	*i*	*AUD*	*CAD*	*CHF*	*EUR*	*GBP*	*JPY*	*USD*
AUD	1	100	81.2	53.3	71.1	71.6	48.6	77.8
	2	100	38.5	23.7	37.4	29.7	0.9	41.9
	3	100	25.7	1.7	13.8	27.3	–10.2	25.0
CAD	1	81.2	100	58.6	70.6	62.7	44.6	78.6
	2	38.5	100	44.0	54.0	39.2	6.5	66.5
	3	25.7	100	12.2	43.3	7.1	–2.5	59.7
CHF	1	53.3	58.6	100	77.0	65.0	39.8	61.9
	2	23.7	44.0	100	65.1	35.0	2.6	54.3
	3	1.7	12.2	100	15.1	8.7	5.2	15.3
EUR	1	71.1	70.6	77.0	100	65.0	36.8	70.5
	2	37.4	54.0	65.1	100	35.0	12.7	53.7
	3	13.8	43.3	15.1	100	8.7	9.3	56.2
GBP	1	71.6	62.7	65.0	82.4	100	33.0	71.6
	2	29.7	39.2	35.0	46.3	100	1.7	38.8
	3	27.3	7.1	8.7	28.3	100	–8.3	16.2
JPY	1	48.6	44.6	39.8	36.8	33.0	100	50.2
	2	0.9	6.5	2.6	12.7	1.7	100	–1.4
	3	–10.2	–2.5	5.2	9.3	–8.3	100	7.6
USD	1	77.8	78.6	61.9	70.5	71.6	50.2	100
	2	41.9	66.5	54.3	53.7	38.8	–1.4	100
	3	25.0	59.7	15.3	56.2	16.2	7.6	100

3.2.1 Inter-currency correlation: first PCA component

The first component represents parallel shifts. The top number in each cell describes the correlations between the first PCA components of different currencies. They show a strong positive correlation between all currencies. The largest value of 82.4% is observed between GBP and EUR. This number is the top number in cell GBP/EUR (row/column). Similar large values are observed for the pairs CAD/AUD, USD/CAD, USD/AUD and EUR/CHF. The first PCA component of the JPY shows the lowest correlation to all the other currencies (minimum 33.0% to GBP and maximum 50.2% to USD). EUR and USD show a large overall correlation to the remaining currencies. It is interesting to note that very similar correlations are observed between the three month shifts of the five year zero rates (results not shown).

3.2.2 Inter-currency correlation: second PCA component

The second PCA component – representing twists – shows a similar result, but much smaller in magnitude. The largest values are observed for CAD/USD (66.5%), EUR/CHF, USD/CHF, EUR/CAD and USD/EUR. Some of these pairs also show strong correlations in the first components. Again the JPY shows overall the smallest correlation to the remaining currencies with USD/JPY being the smallest (–1.6%).

3.2.3 Inter-currency correlation: third PCA component

The third components describing curvature are again much smaller in magnitude as compared to the first and second component. The largest values are observed for the pairs CAD/USD, EUR/USD and CAD/EUR. These pairs show also overall a large correlation to the first and second components.

3.2.4 Intra-currency correlation

The correlations between PCA components 1–3, within the same currency, are presented in Table 3. For the sake of a more condensed representation, results for different currencies are combined in single lines in Table 3. The CHF shows a negative correlation of –46.6% between component one and two. This corresponds to cooccurrence of IR rise accompanied by a flattening of the curve and an IR fall accompanied by a steepening of the curve. Similar behaviour, but different in magnitude is observed for the other currencies except the JPY. Here, we observe an IR rise accompanied by a steepening and vice versa. The correlation coefficient indicates a fairly strong relation between shift and steepening (47.6% for factor one and two) for the JPY.

Table 3 Correlation coefficient in percent between different PCA components (1 & 2, 1 & 3 and 2 & 3) each within the same currency

I	*Factor; Curr; factor*	*1*	*2*
2	AUD; CAD	–16.6; –42.2	
	CHF; EUR	–46.6; –22.0	
	GBP; JPY; USD	–31.1; 47.6; –26.2	
3	AUD; CAD	3.6; 4.4	11.7; –53.7
	CHF; EUR	23.8; –12.9	–26.2; –10.6
	GBP; JPY; USD	–20.4; 15.5; –27.3	–13.7; 48.4; –29.1

The above analysis of linear correlations already illustrates the complex relationship of movements between different currencies and different shape components. Furthermore, it is not clear whether these relationships are stable over time. In order to shed further light on the complex relationship, we apply HCA on the data space of factor dynamics. The following section illustrates the results.

4 Analysis results: HCA

We perform a HCA on the data set of 3,105 21-dimensional vectors representing the factor dynamics of seven currencies. In the course of the annealing process, the set of archetypes undergoes a transition process of one, two, four, six and seven archetypes.

A detailed analysis of the complete annealing process goes beyond the scope of this paper. However, we find archetypes that represent combinations of 'fall' and 'rise' categories with increasing finer substructure. In the following, we illustrate the results on the level of seven different archetypes.

4.1 Archetypical joint movements

HCA provides seven archetypes, each coded by a 21-dimensional vector. Using Equation (2) we transform the collective movements back into yield curve shifts for each maturity.

The seven archetypes describe a partitioning of the data set. We derive a statistical weight for each archetype by assigning each data vector to the next archetype. Table 4 summarises the statistical weight for each archetype. The archetypes are placed in descending order with respect to the 5-year movements shown in Figure 4.

Table 4 Statistical weights of each archetype in percent. A weight of 10.3% of archetypes A1 means that 10.3% of all data vectors are next neighbours to archetype A1 in a Euclidian metric

Archetype	*A1*	*A3*	*A2*	*A5*	*A6*	*A4*	*A7*	*Sum*
Stat. weight	10.3	18.5	8.0	25.3	20.2	6.8	10.9	100

Figure 4 The seven archetypes of the EUR yield curve movements (see online version for colours)

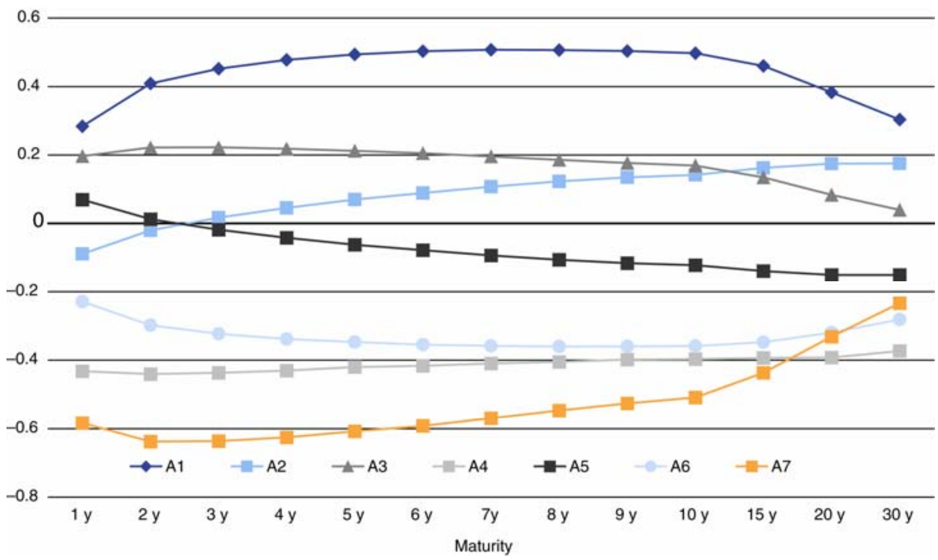

As an example, Figure 4 shows the EUR component of the seven archetypes. We find seven curves that differ in the absolute level and shape. There are three different *rise* and *four* fall movements. The rises differentiate into two fairly flat rises of different magnitude – except a small curvature (A1) and a small flattening (A3). The third *rise* scenario (A2) actually exhibits a slight fall for maturities shorter than two years. It describes an overall steeping of the yield curve of about 30 bp. A2 is almost mirrored by A5 which shows a slight fall and flattening. It is interesting to note that the statistical weight of A5 is about four times higher than the weight of A2 (compare Table 4). Overall, the three *rise* movements account for only 36.8% of all movements. This reflects the fact that the IR fell over the observed historical period.

The remaining three *fall* movements show an additional differentiation in terms of magnitude of shifts and of shape. There are two moderate fall movements that differ only slightly in the levels but show different shapes. The first (A6) exhibits curvature, whereas the other one is almost completely flat (A4). The last and strongest fall (A7) is characterised by a level of about 60 bp and a significant steepening of close to 40 bp. This curve – together with A3 – is consistent with the behaviour expected from the correlation of factor one and two (compare Table 2). However, the symmetric twin scenario (numbers A2 and A5) exhibiting a slight rise and a fall show an opposite picture. Although the levels are not exactly symmetric around zero, A1 and A6 are also almost mirrored images of each other. They are consistent with the anticorrelation between level and curvature (Table 3, EUR, correlation between factors 1 and 2 –12.9%). This may suggest that the correlations between level, slope and curvature are conditioned on the absolute strength and direction of the shift. This behaviour may not be anticipated from the analysis of correlations only!

Figures 5–7 illustrates representative archetypes. Figure 5 shows A3 which may be labelled *moderate rise and flattening*. This archetype is in line with the behaviour anticipated from the linear correlations. The strongest rise is observed for the USD. The JPY shows the smallest rise and is almost flat in the range from 1–10 years. Figure 6 shows A5 which may be labelled *moderate fall and flattening*. This archetype is in line with the behaviour anticipated from the linear correlations of JPY only. The weakest fall is observed for the CHF and looks more like a pure twist (flattening).

Figure 5 The archetype A3 (see online version for colours)

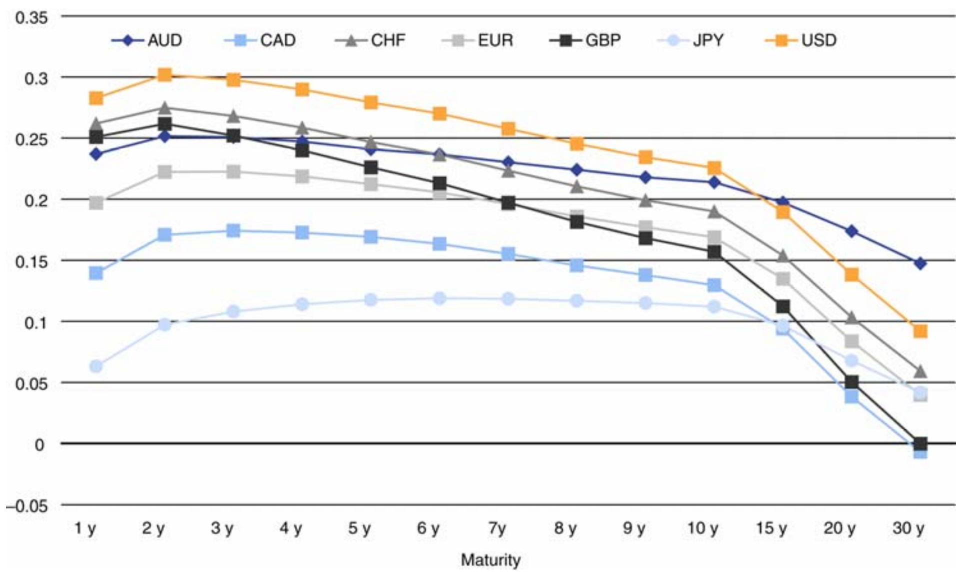

Figure 6 The archetype A5 (see online version for colours)

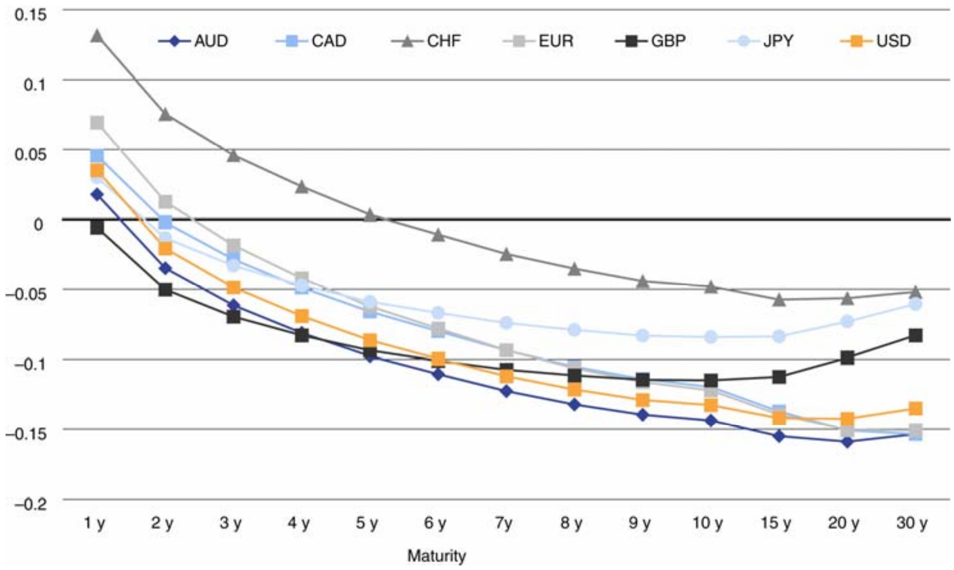

Figure 7 shows A7 which may be labelled *strong fall and steepening*. This archetype is consistent with the behaviour anticipated from the linear correlations.

Figure 7 The archetype A7 (see online version for colours)

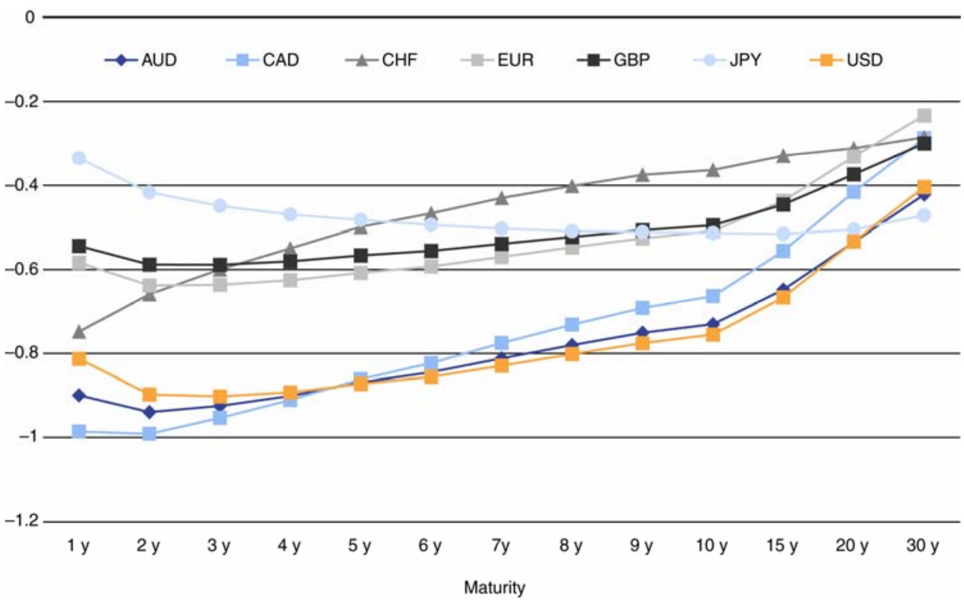

Figure 8 depicts a very interesting behavior. This archetype consists of three different sub sets. The first set is the JPY. It shows a rise and a fairly strong steepening of approximately 40 bp.

Figure 8 The archetype A4 (see online version for colours)

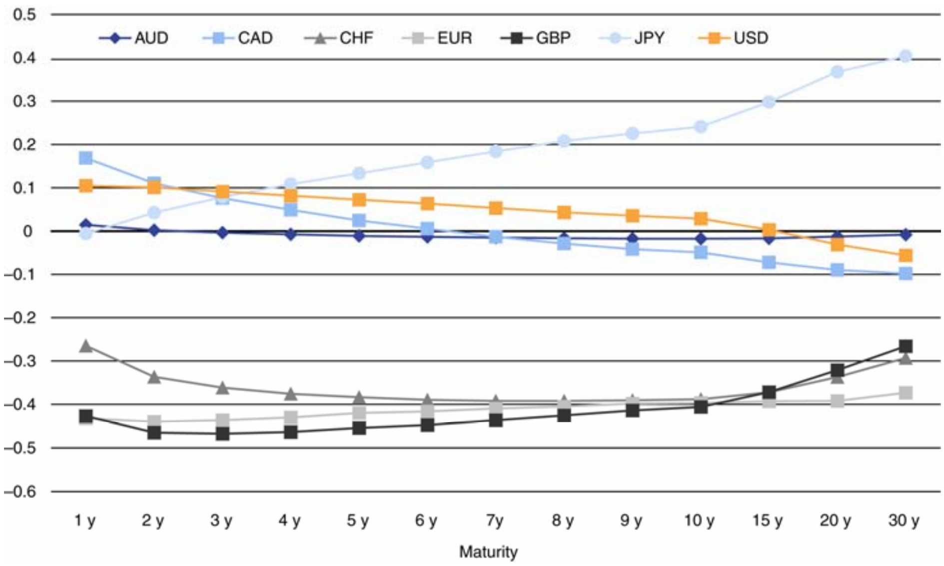

The second set consists of the three dollar currencies: AUD, CAD and USD. They describe a slight flattening around a zero level change (USD and CAD) or stay completely flat at zero level (AUD). The third set consists of the three European currencies EUR, CHF and GBP. They show a moderate fall at a level of 40 bp and slight curvature (CHF) or steepening (EUR, GBP). The archetype A4 describes a decoupled movement of dollar versus European currencies. A4 has a relative low statistical weight of 6.8%. Interestingly, the grouping of currencies is consistent with the results derived from inter-currency linear correlations (please compare Table 2, correlations of level). In the next sub section, we explore the appearance of archetype A4 in the data set.

4.2 Dynamic properties of archetypical movements

Figure 9 shows the projection of the data set (grey diamonds) and the archetypes (red squares) onto the plane defined by the first PCA component of GBP and EUR. The last 100 data vectors in the time series are plotted in red triangles. The first PCA component describes a shift of the yield curves. Data vectors in the lower left corner of Figure 9 therefore illustrate a synchronous downward movement whereas data vectors in the upper right corner a synchronous upward movement of GBP and EUR.

Figure 9 shows an impression about the trajectory of the data set with respect to the position of the archetypes. Depending on the relative distance of a data vector to the archetypes, we may derive an association probability for each data vector. In fact, the HCA already provides these probabilities (Dersch, 1996). In this example, we get a seven dimensional hypothesis vector of association probabilities. The probabilities of association for archetype A4 as a function of time is shown in Figure 10. A4 showed a very interesting behaviour because of the differentiation between currency sets (compare Figure 8).

Figure 9 A two-dimensional projection of the data set (black diamonds) and the archetypes (pink circles) onto the first PCA component of GBP and EUR. The red triangles show the last 100 data vectors (13 October 2006 to 2 March 2007) (see online version for colours)

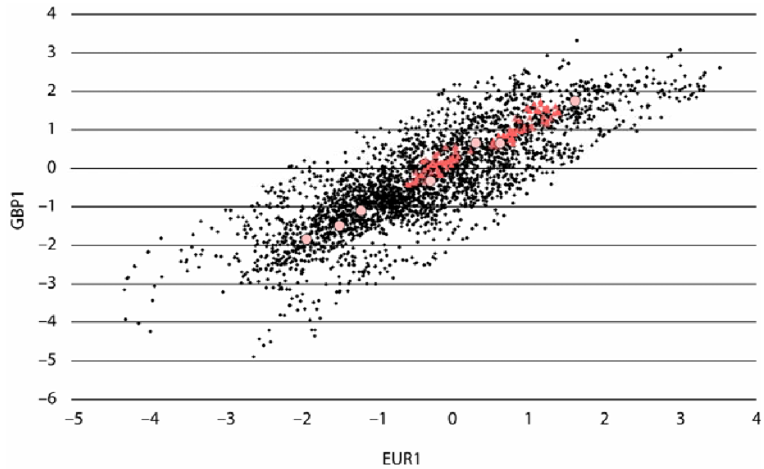

Figure 10 The historical association probability of A4 to a data vector. The dashed-line indicates the 75% probability line (see online version for colours)

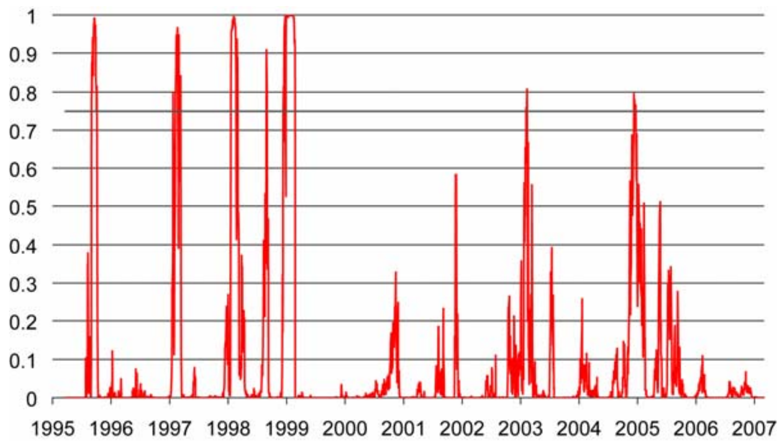

Figure 10 shows the periods where A4 is closest to the trajectory. From Table 4, we already know that this is the case in only 6.8% of the overall time. Time steps where A4 possesses an association probability above 75% (above dashed line in Figure 10) are: Sept. 1995, Feb. 1997, Feb. and Aug. 1998, Jan. 1999, Feb. 2003 and Dec. 2004. These findings confirm an important necessary condition required for archetypes: they reappear over time. As shown above, the archetype A4 describes a decoupling of the dollar (USD, CAD and AUD) versus the European (EUR, CHF and GBP) market. Figure 10 shows the historical periods when this happened and therefore allows segmenting the time series in a complete data driven fashion.

Further work is required to link these findings to economic factors in a business cycle. In a previous paper, these have been done for combined IR/FX movements (Dersch, 2006).

5 Conclusion and outlook

In this paper, we applied the novel concept of archetypes to analyse the complex movements of 91 risk factors. We showed that our findings are consistent with results derived from linear correlations, but draw a much more detailed picture. In the remainder of this section, we would like to provide an outlook on how this scheme may assist in various areas like scenario simulation, forecasting, segmentation or novelty detection. The foundation of these applications is a set of archetypes calibrated on a historical data set.

5.1 Scenario simulations

Jamshidian et al. (1997) applied a PCA to speed-up VaR calculations. Scenarios are generated using linear correlations and a discretisation of states based on a binomial distribution. As a promising alternative, we propose a set of archetypes as a basis for scenario simulations. Scenarios may be generated from single archetypes, their linear combinations or even by stretching given archetypes. The latter allows generating stress scenarios not observed in the past. Different statistical weights and transition probabilities may also be incorporated in this approach. It allows generating scenarios that are consistent and plausible with respect to the historical behaviour and may represent different correlation regimes.

5.2 Risk factor forecasting

In a similar way as described above, we may apply the concept of archetype to generate historical plausible forecasts. The big advantage as compared to bottom-up built forecasts is that we do not have to come up with an estimate for all individual risk factors. There is no need to explicitly make sure that individual forecasts are coherent from a historical point of view, as this is already built into the construction of archetypes. We want to illustrate this with a very simple example. Assume an analyst expects the EUR curve to rise by 20 bp with a flattening on the long end of the curve in the next three months. He may choose A3 from the set of archetypes (compare Figure 4) to define the movement of the EUR curve. Without any further opinion on the other markets, he may just choose A3 as forecasts for the remaining six markets. A further extension of this approach is to use linear combinations of archetypes rather than just one archetype or sets or archetypes each assigned with a given probability.

5.3 Financial time series segmentation

Time series segmentation is a special case of filtering. As outlined in Section 4.2 (compare Figure 10), our approach creates the opportunity to label and segment a financial time series. The mapping of a time series onto a stream of symbols is a first step toward pattern recognition or time series analysis (Weigend and Gershenfeld, 1993). In a previous study, we investigated archetypes of FX and IR movements and found that archetypes link to different stages in business cycles (Dersch, 2006). In addition, the segmentation might be a first step towards technical analysis applications.

5.4 Novelty detection

The analysis of A4 illustrates that the concept of archetypes may help to find relationships in data sets that are not obvious at first. A4 teaches us that temporary decoupling of different markets reoccurs over time. Thus, different points in time are put into a relation. An analysis of these periods may provide a deeper understanding of the mechanics of IR markets.

Acknowledgement

The author is grateful to Thorsten Weinelt and UniCredit Markets and Investment Banking to support this work, to Christina Yassouridis for research assistance and David Dakshaw for proof reading the manuscript. Views expressed in this paper are those of the author and do not necessarily reflect positions of UniCredit Markets and Investment Banking.

References

Dersch, D. and Tavan, P. (1992) 'Load balanced vector quantization', *Proceedings of ICANN'94*, Italy, May, pp.1067–1070.

Dersch, D.R. (1996) *Eigenschaften Neuronaler Vektorquantisierer und Ihre Anwendung in der Sprachverarbeitung*, Reihe Physik, Harri Deutsch, ISBN 3-8171-1492-3.

Dersch, D.R. (2006) *Archetypes of IR and FX movements, Quantitative Cross Asset Strategy, UniCredit Markets and Investment Banking*, November.

Glaum, M. (2000) 'Foreign exchange risk management in German non-financial corporations: an empirical analysis', *Risk Management – Challenge and Opportunity,* Springer Berlin, pp.373–379.

Jamshidian, F. and Zhu, Y. (1997) 'Scenario Simulation: theory and methodology', *Finance and Stochastics*, Vol. 1, pp.43–67.

Lardic, S., Priaulet, P. and Priaulet, S. (2003) 'PCA of the yield curve dynamics: questions of the methodologies', *Journal of Bond Trading and Management*, pp.327–349.

Litterman, R. and Scheinkman, J. (1991) 'Common factors affecting bond returns', *Journal of Fixed Income*, Vol. 1, pp.54–61.

Rodrigues, A.P. (1997) *Term Structure and Volatility Shocks*, Federal Reserve Bank of New York.

Rose, K., Gurewitz, E. and Fox, G.C. (1992) 'Vector quantization by deterministic annealing', *IEEE Transaction on Information Theory*, Vol. 38, pp.1249–1257.

Smithson, C. and Simkins, B.J. (2005) 'Does risk management add value? A survey of the evidence', *Journal of Applied Corporate Finance*, Vol. 17, No. 3, pp.8–17.

Weigend, A.S. and Gershenfeld, N.A. (1993) 'Ti*me Series Prediction: Forecasting the Future and Understanding the Past*, Addison-Wesley.