Evolving Yield Curves in the Real-World Measure: a Semi-Parametric Approach

Riccardo Rebonato, Sukhdeep Mahal, Mark Joshi, Lars-Dierk Bucholz
Quantitative Research Centre (QUARC) - Royal Bank of Scotland

November 25, 2003

Abstract

In this paper we show how to evolve a yield curve over time horizons of the order of years using a simple but effective semiparametric method. The proposed technique preserves in the limit all the eigenvalues and eigenvectors of the observed changes in yields. It also recovers in a satisfactory way several important statistical features (unconditional variance, serial autocorrelation, distrubution of curvatures) of the real-world data. A simple financial explanation can be provided for the methodology. The possible financial applications are discussed.

1 Introduction and Motivation

1.1 The Setting

The literature on models to evolve the yield curve is vast, and entire books have been devoted to the topic. The evolution from the early short-rate-based models to the modern pricing approach has been highlighted, for instance, in Morton (1996) [22], Brigo and Mercurio (2001)[6], Rebonato (2002)[24] etc. These models, however, prescribe how a yield curve should evolve if a trader wanted to price a replicable interest-rate derivative product and avoid arbitrage. At the root of the approach is the (Girsanov) transformation between the real-world (econometric) measure and the pricing measure. The constraints on this transformation are very weak, and amount to the principle of absolute continuity and to the equivalence of the real-world and pricing measures. Even when markets can be assumed to be complete, this pricing measure is in general not unique, depending as it does on the numeraire chosen for present-valuing future payoffs. Depending on the chosen numeraire, the resulting measure is described by different names (risk-neutral measure, forward measure, terminal measure, etc).

More generally, if perfect replication is not possible, there exist an infinity of pricing measures consistent with absence of arbitrage. To each of these pricing measures and numeraires there corresponds a different set of drifts to be applied to the relevant financial quantities (eg, forward rates) to avoid arbitrage and therefore different evolutions of the universe of interest-rate based assets. This causes no concern for relative pricing purposes, but creates an ambiguity for applications that require the real-world evolution of the yield curve. Yet, in many applications what it required is a simulation of *the* real-world evolution, as opposed to one of the arbitrage-free evolution*s*, of the yield curve. This is the problem addressed in this paper.

1.2 Possible Applications

These applications are numerous and important. Some examples are the following.

1. *Evaluation of Potential Future Exposure (PFE) for counterparty credit risk assessment.* This case is typically associated with off-balance sheet transactions (such as swaps) which require zero or minimal outlay of cash upfront, but that expose the counterparty to a potential future credit exposure according to whether the deal will be in the money at the time of a possible future default In order to evaluate the PFE the relevant yield curve will have to be evolved, typically using a Monte Carlo procedure, and the conditional exposure in the various states of the world computed.

2. *Assessment of the hedging performance of interest-rate option models.* Judging the quality of a model from the plausibility of its assumptions is notoriously a dangerous task. The strongest test to which a new interest-rate model or calibration methodology can be put is a hedging 'beauty context' against a rival approach. Clearly, the result will strongly depend on the assumed dynamics of the yield curve. Not surprisingly, if, as it customary, the yield curve is assumed to be shocked by a few eigenvalues obtained from the orthogonalization of the covariance matrix, a simple diffusive model with deterministic volatilities will perform rather well. This is likely to be due, however, to the very simplified nature of the yield curve evolution, rather than to any intrinsic virtues of the modelling approach. It should be remembered, in fact, that even if all the eigenvalues are retained, the resulting yield-curve dynamics will reproduce the real-world one only if the original process were in fact a diffusion. This, in general, is far from being the case.

3. *Assessment of different investment strategies in interest-rate sensitive investment portfolios (Asset/Liability Management).* Investment portfolios are typical not managed and/or assessed on the basis of a daily mark-to-market, and their performance is evaluated with reference to the net interest income (NII) they generate. Since it is not difficult to 'engineer' a favourable NII over a short period of time it is essential that the evaluation of the relative NII performance of competing interest-rate-based portfolios should be assessed over a suitably long time horizon.

4. *Economic-capital calculations*. Typically with these techniques the marginal contribution of an investment or a business activity to a total loss profile is estimated. The desirability of the investment or activity is assessed on the basis of the trade-off between the incremental expected return and some measure of the change in the portfolio risk. These applications can be used to allocate scarce capital, or to evaluate the performance of business lines, and require the construction of a profit-and-loss profile over horizons of at least one year. In these applications the evolution of yield curves plays an important role, because the loss profile is affected by interest rates not only directly, (eg, via traded interest-rate sensitive instruments), but also indirectly via the behaviouralization of, say, depositors' behaviour or mortgage pre-payment patterns.

All these applications have some important features in common:

- the evolution of the yield curve must be carried out over a potentially very long time horizon (many years);
- the resulting yield curves must be accurately representative of the population of the future yield curves, since misleading results can easily be obtained if an over-stylized model evolution is chosen;
- typically, contingent on a state of the world being reached in the future, the evaluation of the fair future conditional value of trading instruments will have to be carried out: for PFE calculations, the future value of, say, a swap will have to be calculated; for the assessment of hedging strategies option positions will have to be re-valued, re-hedged and rebalanced; for NII calculations a re-investment strategy will call for the estimate of a future 'par coupon'. All these applications will require switching from the real-world to one of the future pricing measures. Therefore the real-world evolution of the yield curve complements but does not substitute the sampling of the risk-adjusted measure.

The most important difference between these and related applications on the one hand, and the more typical Value-at-Risk-based (VaR-based) calculations on the other is the length of the time horizon. For trading-book applications the time interval over which the change in risk factors gives rise to P&L variations ranges from one to a few days. While it is difficult to make totally general and process-independent statements, it is in general safe to say that, over these short time periods, the stochastic term will dominate the deterministic component. See, eg, Jorion (1997). In particular, if the underlying real-world processes were simple diffusions, the drift term would scale as Δt, but the innovation component as $\sqrt{\Delta t}$. To give an order-of-magnitude estimate, for $\Delta t=$ 1 day, the drift term would have to be approximately twenty times as large as the diffusive coefficient to make a comparable impact on the total price move.

One can no longer rely on this dominance of the stochastic term, however, when time horizons are as long as many years. We shall show below that a careful treatment of the deterministic term becomes necessary in these situations, and is actually at the heart of the approach we suggest.

1.3 Features of the Proposed Approach

We propose in this paper a simple methodology to evolve in the real-world measure yield curves over long periods of time (of the order of years). These yield curves can then be used for the applications briefly discussed above, or for similar uses. The approach enjoys the following features:

- asymptotically, it reproduces exactly the eigenvalues and the eigenvectors of the real-world one-day changes;
- asymptotically, it reproduces exactly the one-day unconditional variance of the changes in all the yields used to describe the yield curve (the benchmark yields);
- it recovers approximately the distribution of yield curve curvatures;
- it recovers approximately the unconditional variance for many-day changes in the benchmark yields;
- it recovers approximately the serial auto-correlation function for one- and many-day changes in the benchmark yields;
- it is very simple to implement;
- it does not make strong statistical assumptions about the generating joint process, and requires very few parameters: for n benchmark yields, approximately $2n$ parameters will be required for the evolution of the whole yield curve over several years;
- the most important parameters introduced (the 'spring constants') are financially motivated.

The method will be presented in stages, starting from the a totally non-parametric approach, because this style of presentation, although slightly more cumbersome, allows a clear understanding of the origin of the shortcomings of the various methodologies.

2 Links with the Existing Literature

The yield curve has been studied extensively both parametrically and non-parametrically. The sheer number of the publications prevents us from providing a comprehensive survey. Therefore, only a few relevant studies will be mentioned.

One strand has been linked with asset pricing (often derivative pricing). In this context, models for the evolution of the yield curve have been proposed (in approximate chronological order) for instance by Vasicek (1977)[26], Cox Ingersoll and Ross (1985)[8], Ho and Lee[15], Hull and White(1990)[16], Longstaff and Schwartz(1992)[18], Heath Jarrow and Morton (1989)[14]. All these works assume a diffusive behaviour for the yield curve, and the only part of their description of relevance to the present study is the stochastic innovation (this is because the 'drift' is obtained by no arbitrage, and cannot be simply related to the real-world drift). The relatively recent appearance of 'smiles' in the implied volatility curve has suggested that jumps (see, eg, Glasserman and Kou (2000)[11]), stochastic volatility (see, eg, Joshi and Rebonato (2003)[9]) or a CEV-type behaviour (see, eg, Andersen and Andreasen (1999)[3]) might be important.

Another stand of research has focussed on the real-world evolution of the term structure. Although still numerous, the works in this area have been fewer in comparison. One of the best known is the study by Chan, Karolyi, Longstaff and Sanders (1992)[7], who carry out a parametric estimation of some of the parameters of the process (in particular, the volatility). A related but different (non-parametric) approach is followed by Stanton (1997)[25], who focusses on the US$ three-month Treasury Bill rate and obtains a similar estimate as Chan et al. for the volatility, but discovers strong evidence for non-linearity in the drift. Stanton also estimates from the data the market price of risk. The approach by Stanton inspired the work by Prigent, Renault and Scaillet (2000)[23], who study the dynamics of the 10-year maturity yield of corporate bonds.

There are several studies of specific portions of the yield curve, especially of the very short end (the 'short rate'). Ait-Sahalia (1996)[2], for instance, argues that the short rate should be close to a random walk in the middle of its historical range (approximately between 4% and 17%), but strongly mean-reverting outside this range. Hamilton (1989)[13], Garcia and Perron (1996)[10], Gray (1996)[12], Ang and Bakaert (1998)[4], Naik and Lee (1997)[20] and Bansal and Zhou (2002)[5], among others, argue that a switch behaviour, with a reversion level and reversion speed for each regime, should be more appropriate.

None of these approaches, however, tackles the problem of the joint evolution of a whole yield curve as described by several reference benchmark yields. This type of investigation is more frequently carried out in a Principal-Component-Analysis context. Again, there have been a very large number of studies, many of which have been recently reviewed in Martellini and Priaulet (2001)[21]. Whatever the true process driving the yield curve, the eigenvectors and eigenvalues always convey interesting summary information about the underlying dynamics. Unfortunately, while informative from a descriptive point of view, these investigations are only valid for the purpose of reproducing the dynamics of the yield curve if the underlying joint process were indeed a diffusion, with iid increments. We find in our investigation strong evidence to reject this hypothesis.

The approach we propose does not posit any particular form for the joint processes that drive the yield curve, and attempts to extract as much information as possible from the empirical data in an (almost) model-free way. In particular, since we will use as the fundamental building blocks of our algorithm *synchronous* vectors of yield changes, the degree and the nature of the co-dependence between changes for different-maturity yields will be recovered with accuracy.

Having said this, our approach is not totally model-free. This is not just because we will have to introduce some constraints on the deterministic innovation part of the implicit (and unknown) process; but also because we have to choose whether the observed changes should be interpreted as proportional or absolute (or otherwise). The results we present below are obtained using the proportional assumption (that is guaranteed to ensure positivity to all the future rates). We revisit this topic, however, in the second-to-last last section of the paper.

3 The Empirical Data

3.1 Description

We have data for USD LIBOR curves from 29-Sep-1993 to 4-Dec-2001 (2,135 days). The yield curve is assumed to be described by 8 points (benchmark maturities), located at 3 and 6 months, and 1, 2, 5, 10, 20 and 30 years. The associated yields are labelled in the following 3m, 6m 1y, 2y, 5y, 10y 20y and 30y. The whole data set therefore consists of 2,135 x 8 = 17,080 points. The time behaviour of each rate is presented in Figures 1a to 1h. The data set was carefully checked for outliers, and, wherever dubious points were encountered, corroboration was sought from independent sources. We chose to work with yields rather than, say, with forward rates, because the former can be obtained from observed prices in a model-independent way, while the specific technique used to distill forward rates could easily produce spurious statistical results.

Fig 1a: The time series of the 3m rate from 29 September 1993 to 4 December 2001

Fig 1b: The time series of the 6m rate from 29 September 1993 to 4 December 2001

Fig 1c: The time series of the 1y rate from 29 September 1993 to 4 December 2001

Fig 1d: The time series of the 2y rate from 29 September 1993 to 4 December 2001

Fig 1e: The time series of the 5y rate from 29 September 1993 to 4 December 2001

Fig 1f: The time series of the 10y rate from 29 September 1993 to 4 December 2001

Fig 1g: The time series of the 20y rate from 29 September 1993 to 4 December 2001

Fig 1h: The time series of the 30y rate from 29 September 1993 to 4 December 2001

Tab 1: Descriptive statistics of the reference yields

Tab 2: Descriptive statistics of the changes in the reference yields

We present statistics of the data values and of the daily percentage changes in the data values and in Tabs 1 and 2, respectively. More precisely, Tabs 1 and 2 shows the mean, standard deviation, skew and kurtosis of the rates and of their relative changes. We observe that the volatility as a function of the maturity has a maximum around the 1 year rate. We also note that the mean of the changes is very close to zero.

We performed a principal component analysis of the historic data, obtained by orthogonalizing the matrix of the percentage changes in rates. We observed that the first component explains more than 90% of the variance. We also see that the first eigenvector has loadings of the same sign for all the yields, the second eigenvector displays one sign change, and the third eigenvector has a 'V' shape, with two crossings of the x-axis. See Figs 7, 8 and 9. Therefore we find again the well-established interpretation of the first modes of deformation as the level, slope and curvature changes of the yield curve.

In the following the curvature of the yield curve will become a central quantity of our analysis. For future reference we therefore report in Section 3.3 below some descriptive statistics about the curvatures and the changes of curvature.

As for the qualitative behaviour of the yield curves, displayed in Figs 1a to 1h, one can distinguish short-, long- and intermediate-maturity rates. The 3m rate shows 3 regimes: rates rising from ~3% in 1993 to ~6% in 1995, followed by a period of relatively stable rates, followed by falling rates from January 2001 onwards. The 6m and 1y rates behave in a manner similar to the 3m rate. Mid-maturity rates are more volatile during the short rate's stable period. Long term rates show a steady but volatile decline. The Russia crisis period of 1998 is clearly evident in rates of 1y and longer.

3.2 Unconditional Variances

We present the variance of non-overlapping changes of the data values in Figures 2a to 2g. Call N the total number of days (ie 2,135), m the length (in days) of each non-overlapping period for rate i, $i = 1, 2, ..., 8$, and Var_m^i the variance of the m-day increments for rate i. ($i = 1$ corresponds to the shortest maturity (3 months) and $i = 8$ corresponds to the longest, 30 years). The straight line in Figs 2a to 2g is the extension of the variance of the 1-day changes, Var_1^i. So, the quantity

$$Var_m^i = m\ Var_1^i \tag{1}$$

would be the variance of the m-day changes that would be observed if the increments were iid[1]. We clearly see a deviation from linearity in the behaviour of the true one-day variance. As the length of the non-overlapping periods increases,

[1] It could be argued that, by the way this striaght line has been constructed, its slope is very sensitive to small estimation error. While this is true, the 1-day variance is the quantity estimated with the largets number of non-overlapping data points, and therefore the most reliable. We checked that very similar results were obtained if the 2-day variance, or a combination of the 1-day and the 2-day variances had been used.

rates of short maturities show *higher* variance than would be obtained if the increments were iid. We remark at this point that a less-then-linear increase in the m-day variance could be compatible with positive auto-correlation.

Fig 2a: The empirical variance of the m-day changes in the 3m yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 6m yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 1y yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 2y yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 10y yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 20y yield (m on the x-axis), with the tangent to the curve through the origin.

Fig 2a: The empirical variance of the m-day changes in the 30y yield (m on the x-axis), with the tangent to the curve through the origin.

When we move to longer maturities, as the length of the non-overlapping periods increases, these longer-maturities rates display a variance that increases less than linearly with the length of the non-overlapping period. We remark at this stage that this observation could be compatible with mean reversion.

As the number of days, m, in each overlapping period increases we have fewer and fewer observations. As a consequence our estimates of the m-day variance become noisier and noisier. The effect, however, is very clear and consistent and displays a smooth behaviour across maturities.

3.3 Distribution of Curvatures

An important statistic for the analysis we present in Section 4 is the distribution of yield curve curvatures. We define the curvature, ξ_i, at point $\tau_i = (T_{i+1} - T_{1-1})/2$, $i = 2, 3, ..., 7$ between 3 successive grid points (points i, $i+1$ and $i-1$) on the curve as

$$\xi_i = \frac{\frac{y_{i+1}-y_i}{T_{i+1}-T_i} - \frac{y_i - y_{i-1}}{T_i - T_{i-1}}}{\frac{T_{i+1}+T_i}{2} - \frac{T_i + T_{i-1}}{2}} \qquad (2)$$

where y_i is the yield at grid point (maturity) T_i. (We have to make use of this rather awkward definition of the second derivative because our grid points are not evenly spaced, and we prefer not to 'pollute' our data with any form of interpolation).

We have already presented statistics of the curvature values and of the daily changes in curvature values in Table 3 and Table 4 respectively. Plots of selected densities of the curvatures for several maturities are shown in Figs 3a to 3d. For future discussions, it is important to point out that the distribution of curvatures is much wider at the short end of the curve. Figure 4 shows the standard deviation of the curvatures as a function of the approximate maturity along the yield curve (defined as $(T_{i+1} - T_{i-1})/2$).

Fig 3a: The density of curvatures around the 6m yield point
Fig 3b: The density of curvatures around the 1y yield point
Fig 3c: The density of curvatures around the 10y yield point
Fig 3d: The density of curvatures around the 20y yield point
Fig 4: Standard deviation of the curvature as a function of the yield maturity

Tab 3: Descriptive statistics for the curvatures
Tab 4: Descriptive statistics for the changes in the curvatures

One can conclude from the data that an important feature of real-world yield curve shapes is that they should be much less 'twisted' at the long end than for short maturities. This feature should be reproduced by the simulated yield curves, and we shall show that some naive approaches fail to capture it. The ability to capture this empirical feature is one of the motivators of the more complex simulation methodology proposed in Section 5.

3.4 Serial Auto-Correlations

We present in Figure 5 the lag-1 auto-correlation of non-overlapping n-day changes of all the yields. For all maturities, except, possibly, for the 30-year rate, one readily observes a positive serial autocorrelation that increases as a function of n. The effect is more pronounced for short maturities, and appears to decrease monotonically in magnitude as the rate maturity increases. Somewhat surprising is the fact that the auto-correlation is significantly non-zero for n-day non-overlapping changes even when n takes values as high as 50.

Fig 5: The empirical lag-1 autocorrelation for non-overlapping n-day changes in yields (n on the x-axis)

When discussing the sub- or super-linearity of the variance as a function of number of days change, we remarked that the super-linearity could be due to positive serial correlation. Indeed, the auto-correlation graphs show that the maturities that display strongest super-linearity are also the maturities that have most pronounced positive serial autocorrelation. The transition from super- to sub-linearity could be explained by an additional mean-reverting component (possibly of different strength) for all rates. This would produce sub-linear variance, and a positive serial auto-correlation of decreasing strength as the maturity increases producing a transition from sub- to super-linearity (the more so where the seral correlation is more positive).

These strong positive auto-correlations, *prima facie* somewhat surprising, are actually just a reflection of the very 'trendy' nature of rates during the period of observation. See the discussion in Section 2. It should be emphasized that, by themselves, these results do not imply anything about the efficiency of the market. In particular we are not analyzing the auto-correlation or any other statistic of *forward* rates, or of the difference between 'projected' and realized rates. The observation of a positive serial auto-correlation *per se* does not create any trading opportunities, but simply reflects the fact that the monetary authorities effect their monetary tightening/easing gradually and over long periods of time. So, contingent on the change in, say, the 30-day rate having been positive (negative), it is likely that we are in a tightening (loosening) cycle, and that therefore the next 30-day change will also be positive (negative). Forward rates typically reflect this, since implied forward-rate curves remain positively or negatively sloped for extended periods of time. A trading advantage would come from knowing whether the forward curve is too steeply (or not steeply enough) sloped with respect to the realized future rates (after adjusting for risk). This aspect is not captured by the statistical observations reported in this section.

4 Modelling Approaches

The real-world data is a single realization from the universe of possible yield curve evolutions. We therefore calibrate the models using path statistics, ie, we calculate the statistics along each model path in an evolution and average over many evolutions.

In order to assess the quality of the output from our model we take into account the following characteristics of real-world curves:

- Variance of rates
- Variance of curvatures
- Variance of n-day non-overlapping changes
- Lag-1 auto-correlation of n-day non-overlapping changes

4.1 A Totally Non-Parametric Approach: Description

The simplest approach to evolve the current yield curve is to assume that the each vector of rate changes is drawn at random (with equal probability) from the 2,134 observed vector changes. More precisely, let $\{y_i^k\}$ denote the vector of rates on day k, $k = 1, 2, ..., N$, and $\{y_i^N\}, i = 1, 2, ..., 8$ be the vector of rates describing today's yield curve. Let $\Delta y_i^k, k = 2, 3, ..., N$ be the change in the i-th rate between day k and day $k-1$:

$$\Delta y_i^k = y_i^k - y_i^{k-1} \tag{3}$$

Suppose that we want to evolve the yield curve M days forward. Draw M independent variates, U_r, $r = 1, 2, ..., M$, from the uniform discrete $[1 \quad N]$ distribution[2]. In the naive approach, the future yield curve that will prevail M

[2] In the uniform discrete $[1 \quad N]$ distribution all integers between 1 and N have the identical mass probability of occurrence of $1/N$.

days after today is simply given by

$$y_i^{N+M} = y_i^N + \sum_{r=1,M} \Delta y_i^{U_r} \tag{4}$$

(A remark on notation: to avoid the proliferation of symbols, we use the use the same symbol, y, to denote past and present observed rates and future simulated rates. Ambiguity is avoided thanks to the range of the superscript: the first historical vector of rates in our data base is labelled by the superscript 1. Since we have N distinct days of data the current curve has the superscript N. Future rates are identified by having a superscript greater than N). By drawing a whole synchronous vector the procedure attempts to preserve whatever co-dependence structure the original data might show across maturities. If the increments were iid, the procedure would asymptotically fully capture both the serial auto-correlation structure (or rather, the lack thereof) and the cross-sectional co-dependence. For general increments, however, this is no longer true. The analysis below show that the deviation form iid behaviour (which could already be expected in the light of the data description in Section 2) is sufficiently strong to make the procedure unsatisfactory. Nonetheless, it is interesting to point out how many features even this naive approach *does* recover. In particular, it is straightforward to show that *all* the eigenvectors and eigenvalues of the sample data are fully recovered by the naive procedure as are all the unconditional 1-day variances. This inserts a strong *caveat* regarding the validity of those procedures that invoke orthogonalization of the covariance matrix (and retention of a small number of eigenvectors).

In the discussion presented below we refer to this method as the random-sampling approach.

4.1.1 Positive Features and Shortcomings of the Random-Sampling Approach

Given a time series of changes of a set of time-indexed data, it is obviously always possible to construct and orthogonalize a covariance matrix. In many fields (eg, interest-rate modelling) it has become customary to work with the resulting eigenvectors and eigenvalues (for instance, in the LIBOR-market-model context, this is often done by calibrating the model-implied correlation matrix among forward rates to the corresponding statistical quantity). Working with the full eigen-system has in practice often proven to be too computationally expensive. The computational difficulties associated with multi-factor no-arbitrage pricing have therefore forced the modeller to drop many of the eigenvectors, and to rescale the few preserved eigenvalues in such a way as to recover some important quantities (eg, the prices of caplets). It has been claimed Longstaff et al (2000)[19] that the inability to work with a full-factor model can have very adverse financial consequences (their controversial paper on this topic in the context of Bermudan swaptions was titled 'Throwing away a billion dollars"). It has therefore become customary to consider the ability to work with the full dimensionality of the eigenvectors almost like a Holy Grail. The analysis presented below shows how little is actually achieved in terms of realistic modelling of the yield curve even when all the eigenvectors and eigenvalues are correctly recovered. In particular, Figures 6a to 6d, which display a few typical yield curves evolved by the random-sampling method over a period of 5 years, show the qualitatively unsatisfactory nature of the procedure: it is difficult, for instance, to find 'convincing' a yield curve which drops from 10% to 7% for maturities out to one year, and then rises to above 13% for two-year yields (see Fig 6a). Yet, the naive procedure proposed above recovers (asymptotically) exactly all the eigenvalues and eigenvectors. See Figs 7 to 10. This is because the vectors of changes are exactly the same as in the real world, and, if the evolution of the yield curve is sufficiently long, they will all be represented with the same frequency as in the real-world data set. Therefore the model covariance matrix will asymptotically be the same as the empirical one, and so will therefore be its eigenvalues and eigenvectors. Clearly, any single realization of the yield curve evolution will produce different eigenvectors and eigenvalues, but this is nothing more than the effect of a finite-size sampling.

Fig 6a to 6d: A few 'unsatisfactory' yield curves obtained with the random-sampling method

Fig 7: The magnitudes of the eigenvalues obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

Fig 8: The first eigenvector obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

Fig 9: The second eigenvector obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

Fig 10: The third eigenvector obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

Fig 11: Lag-1 autocorrelation for non-overlapping n-day changes in yields obtained with the random-sampling method (n on the x-axis)

By the same token, every statistical quantity that does not depend on the order along the time axis of the quantities used for its computation will be recovered exactly by the non-parametric approach. Therefore, also the one-day variance and the correlation matrix will be exactly recovered. Crucially, the same is not true, however, for the many-day variance, where serial correlation becomes important.

Besides the 'implausible' look of the yield curves generated by the random-sampling method, the shortcomings of this approach just introduced were that the distribution of curvatures, the serial auto-correlation and the many-day variances turned out to be at odds with the corresponding real-world quantities. See Fig 11. In particular, the serial auto-correlation was, within statistical noise, zero (not surprisingly, given the method of sampling), and the many-day variance was observed to grow almost linearly with the number, m, of days in the change. Both these features point to the fact that the most important mis-specification of the approach is to be found in the fact that it implies, by construction, independent increments.

5 A Financial and Computational Model

5.1 Arbitrageurs and the Curvature of the Yield Curve: the Barbell Mechanism

In order to obviate the shortcomings associated with the naive approach described above a simple mechanism can be introduced. Looking at the shapes of the yield curves produced using the simple approach, one of the unsatisfactory features is the presence of 'kinks' and pronounced changes in slope for long maturities. In reality, the width of the statistical distribution of curvatures, as discussed above, shows a strong dependence on the portion of the yield curve. This can be explained and modelled as follows. Let us assume that, as the yield curve evolves and changes shape, a 'kink' appears. (A 'kink' in this context, is defined as a triplet of yields, y_{i-1}, y_i, y_{i+1}, associated with contiguous maturities T_{i-1}, T_i, T_{i+1}, such that $y_{i-1} < y_i$ and $y_i > y_{i+1}$, or $y_{i-1} > y_i$ and $y_i < y_{i+1}$). When this happens relative-value traders ('arbitrageurs') might have an incentive to put in place a barbell strategy, whereby the receive the 'high' yields y_{i-1} and y_{i+1}, and pay the low yield y_i (or vice versa). If the appearance of the kink were due, say, to a momentary fluctuation in supply and demand, the action of the 'arbitrageurs' would be likely to be successful, and would contribute to flattening out the kink itself. In some circumstances, however, the 'arbitrageurs' can be reluctant to enter the barbell strategy. This might occur, for instance, if they were afraid that the sharp change in the shape of the yield curve associated with the kink might be due to rapidly changing expectations about the future behaviour of rates[3]. Note that the arbitrageurs can expect to profit from such kinks only if they do not reflect expectations about future rates. Observation of the forward rates *per se* gives no direct information about expectations of future rates because of the existence of the risk premium. The impact of a risk-premium on the shape of the yield curve, however, becomes more pronounced as the maturity of the associated yield increases (see, eg, Rebonato (1998)). Therefore, the observation of a rapidly-changing yield curve at the short end is more likely to be directly associated with changing expectations. As a consequence, 'arbitrageurs' are plausibly more reluctant to put on their barbell trades at the short end of the yield curve, because the observed kink might simply reflect expectations about a complex future rates path. In other words, a kink in the 8-, 9- and 10-year area is more difficult to explain in terms of expectations about the future short rate in 8, 9 and 10 years' time than the same kink appearing, say, in the 3-, 6- and 9-month area.

If this explanation is correct, the arbitrageurs would be more and more active

[3] As an example, the rapidly-changing shape of the US yield curve in the spring/summer of 2001 was commonly rationlaized as embedding complex expectations about a V-shaped recovery, and the concomitant actions of the monetary authorities.

the longer the maturity of the rates in question. Their actions would force the yield curve to be 'straighter' for long maturities, but would allow it to remain 'kinky' at the short end.

How can the buy-sell-buy (or sell-buy-sell) action of the arbitrageurs be modelled? Their action would be equivalent to the presence of springs with a spring constants, k_i, dependent on the yield maturity. Therefore, we choose to simulate the value of the i-th yield M days after today as

$$y_i^{N+M} = y_i^N + \sum_{r=1,M} \Delta y_i^{U_r} + \sum_{r=1,M} k_i \xi_i^{N+r} \qquad i = 2, 3, ..., 7 \tag{5}$$

where the first two terms on the RHS are the same as for the naive updating of the yield curve described in Section 4, ξ_i^{N+r} is the curvature prevailing at the i-th point of the yield curve r days after today and k_i is a spring constant for the yield i to be determined as described later on.

In leaving this section it is important to point out that the proposed spring mechanism will introduce a non-zero auto-correlation for the various rates. This will also affect the many-day variances. The implications of this observation is discussed in the sections below.

By the way they have been defined the spring constants will have a direct effect on the distribution of curvatures. We therefore determine their values by matching (approximately) the model-produced and real-world variances of the distribution of curvatures.

Finally, the proposed approach does not affect the two ends of the yield curve. (Note that the summation in Equation 5 runs from 2 to 7). In order to deal with the short and long end, we introduced a simple mean-reverting term, with the mean-reversion level equal to the unconditional average of the 3m and 30y rates. The reversion speed was chosen in conjunction with the spring constants so as to obtain a good pattern for the variance of the curvatures, as discussed below.

In the following discussion we refer to this method as the barbell or the spring mechanism.

5.1.1 Positive Features of the Barbell Mechanism

The spring constants, the reversion levels and the reversion speeds required to obtain a good match of the variances of the curvatures were found to have the values in the table below.

k_1	k_2	k_3	k_4	k_5	k_6
0.004	0.0013	0.01	0.02	0.03	0.03

As for the reversion levels, they were calculated to be 5.365% and 6.662% for the 3m and the 30y yields, respectively. The reversion speed was found to be 0.4.

The spring-constant mechanism is equivalent to introducing a drift term in the real-world dynamics of the yield curve. In the limit as Δt becomes very small, a drift term in Δt gives a negligible contribution to the covariance matrix of daily changes. Therefore the orthogonalization of this matrix will produce approximately the same results as in the naive (totally non-parametric) case described in Section 4. All the eigenvectors and eigenvalues should therefore remain essentially identical to the empirical ones. Figs 7 to 10 show that this is indeed the case. We were also successful in obtaining the desired pattern for the variances of the curvatures. These are shown in Fig 12 and in Tab 5.

Despite these positive features, two important shortcomings remained. These are discussed below.

Fig 12: Variances of the curvatures obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

Tab 5: Variances of the curvatures obtained with the random-sampling, box and spring-and-box methods, and the corresponding empirical quantity

5.1.2 Shortcomings of the Barbell Mechanism

The serial auto-correlation and the many-day variances produced by the barbell mechanism were still unsatisfactory, and, within statistical noise, not different from zero. See Fig 13. As for the many-day-variances, they were observed to grow less-than-linearly.

Fig 13: The lag-1 autocorrelation for non-overlapping n-day changes in yields obtained using the spring mechanism (n on the x-axis)

Understanding reasons for the success and shortcomings of the barbell mechanism is instructive. The variances of the curvatures were successfully recovered by introducing spring constants. These, however, will introduce negative serial autocorrelation, and a many-day variance growing sub-linearly in the number of days in the interval. The empirical data show that the many-day variances grow *super*-linearly for maturities under approximately two years, and *sub*-linearly for longer maturities. It is therefore plausible to expect that, in addition to the springs, a second mechanism with an opposite effect on the auto-correlation could be at play. If this mechanism were of constant strength across the yield curve, it could produce the desired effect, because we know that the springs are 'stronger' for the longer maturities.

5.2 The Modified Barbell Mechanism with Serial Sampling

In order to capture the auto-correlation features of the real-world yield curves that are still missing in the model data, we propose the last variation on the theme. We define two additional quantities, the window length, an integer, w_l, and the jump frequency, λ. We sample again using Equation 5 using a starting point determined by the draw U_r. We carry out the successive draws, however, using *consecutive* vectors for at most w_l draws. After each draw we allow the possibility of jumping out of the window with probability λ. So, at most w_l consecutive draws will be carried out. Once a window has been exhausted, either because its end has been reached or because a jump has occurred, we move to a different window, of the same length, and with an independently drawn starting point.

The following section discusses how well this modified mechanism (that we call 'modified barbell' or 'box and spring') works in practice.

5.2.1 Positive Features of the Modified Barbell Mechanism

The results reported were obtained with the same curve springs determined above, with a window length, w_l, of 40 and a jump intensity, λ, of 0.05. The result were stable with respect to reasonable variations in the parameters. For instance, similar results (not reported for the sake of brevity) were obtained for a window length of 60 and a jump intensity of 0.1 (jumps/year).

Also this algorithm asymptotically (ie, as Δt becomes small) does not disturb the covariance matrix, and the one-day variances. Indeed all the eigenvectors and eigenvalues are correctly recovered within numerical noise. This is shown in Figs 7 to 10.

Given the persistence of the positive lag-1 serial auto-correlation as a function of the number days, the relatively small window lengths do not allow to recover accurately the empirical pattern for this quantity. The algorithm proposed, in particular, will by construction not be able to produce a non-zero autocorrelation for changes of more than w_l days. Nonetheless Fig 14 show that the modified barbell mechanism is certainly an improvement over the random or the simple spring procedures. For greater clarity the empirical and box-and-spring serial auto-correlation are also shown in Figs 15a and 15b.

Fig 14: The lag-1 autocorrelation for non-overlapping n-day changes in yields obtained using the spring-and-box mechanism (n on the x-axis)

Fig 15a: The lag-1 autocorrelation for non-overlapping n-day changes in yields: empirical data and the model data obtained using the spring-and-box mechanism (n on the x-axis, yield maturities 3m to 5y)

Fig 15b: The lag-1 autocorrelation for non-overlapping n-day changes in yields: empirical data and the model data obtained using the spring-and-box mechanism (n on the x-axis, yield maturities 30y and 30y)

The most significant improvement, however, is in the variances for several-day changes. In Figs 16a to 16h we show the real-world variances, the tangent through the origin of the same curve (this would be the multi-day variance function if the draws were iid), and the model multi-day variance obtained with the spring and box mechanism. Note how the model variances switch from growing super-linearly to growing sub-linearly, in agreement with the real-world data.

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 3m rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 6m rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 1y rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 2y rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 5y rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 10y rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 20y rate

Fig 16a: The emprical variance of the m-day yields, the tangent to this curve through the origin, and the variance of the m-day yields obtained with the spring-and-box mechanism for the 30y rate

Qualitatively, the consecutive sampling within each box is sufficient to create a sufficient degree of serial autocorrelation to more than compensate for the negative autocorrelation introduced by the spring mechanism. These two features used together with the random sampling appear to describe to a satisfactory degree the dynamics of the real-world yield curve.

Our initial search for a methodology better than random sampling was motivated by the 'implausibility' of the synthetic yield curves. It is therefore important to check whether the model real-world yield curves produced by the spring-and-box technique satisfy this optical test. Three typical evolutions of the yield curve over a seven-year period are presented in Figs 17 to 19. Despite the fact that plausibility, much as beauty, is in the eye of the beholder, we find these synthetic curves 'believable' and 'desirable'.

Figs 17 to 19: 'Desirable' yield curves produced by the spring-and-box mechanism (the yield curves are equally spaced in time, and the yield curve labelled 'Series 5' corresponds to 1800 business days in the future).

6 Level Dependence of the Changes in Rates

In the analysis presented above we have worked with percentage changes in rates. This had the obvious advantage of guaranteeing the positivity of all future rates. However, it should not be taken for granted that changes in rates in the real world do follow such a proportional law.

In Fig 20 and Tab 6 we show the correlation between the absolute magnitude of the (absolute or percentage) empirical changes and the level of the rates for the eight reference yield. It is clear that there is a strong negative correlation for percentage yields, which is significantly reduced if one worked with absolute changes. We know, however, that when rates become *very* low, changes cannot be absolute (when rates, as they have been in Japan, are below 25 bp, the 'classic' quarter-point cut is no longer possible). Some researchers (see, eg, Andersen and Andreasen(2000)[3]) have therefore recently proposed that rates might become log-normal when extremely low. A smooth transition from a normal behaviour for 'high' rates to a log-normal regime for 'very low' rates is therefore intuitively appealing. Such a behaviour could be provided by a modified CEV process, of the type

$$dy = y^{\beta(y)}\sigma_{\beta}dz \tag{6}$$

with

$$\beta(y) = 1 \qquad \text{for } y \le y_{\min} \tag{7}$$

$$\beta(y) = 0 \qquad \text{for } y \ge y_{\max} \tag{8}$$

In principle one could try to in interpret the vector of changes drawn for a particular future day as being of this extended CEV type. However, for a steep yield curve (such as the US$ market yield curve in late 2003) this could give rise to a quasi-lognormal behaviour at the short end, and to an almost-normal behaviour at the long end. Incidentally, to give an idea of the complexity of the empirical studies, Ait-Sahalia (1996) [2] and Chan et al (1992) [7] find evidence that the exponent should be *greater than* 1. Perhaps it would be interesting to repeat some of the econometric investigations carried out in the past with recent market data, that would contain the exceptionally low rates encountered in US$, EUR, JPY and, to some, extent, GBP.

These complications could be handled with some care, but, given the non-process-based approach that we have chosen to follow in this paper, we prefer to leave the analysis at this stage.

Fig 20: Correlation between the absoulte magnitude of the (absolute or percentage) yield change and the level of the yield at the time of the change

Tab 6: The data in Fig 20 in tabular form

7 Discussion and Conclusions

We have shown in this paper how the real-world evolution of the yield curve can be simulated in a simple and effective semi-parametric way. The proposed procedure enjoys several positive features: it recovers almost exactly all the eigenvectors, eigenvectors and variances associated with one-day changes observed in the statistical data. It also reproduces very well the observed distribution of yield curve curvatures and the many-day-return variances, and fairly well the lag-1 serial auto-correlation. Last but not least, it produces future yield curves that are 'optically' acceptable.

The procedure has been constructed algorithmically, and it makes as few statistical assumptions as possible. In particular no generating stochastic process is implicitly or explicitly assumed. The model is very parsimonious: the only free parameters are the spring constants (equal in number to the number of independent yields minus 2), the window length, the out-of-window jump frequency, two reversion levels and two reversion speeds.

As a by-product of the analysis we have observed how weak the requirement is that the model yield curves should reproduce the observed eigenvectors and eigenvalues. This observation cautions strongly against the common implicit assumption of independence of the yield changes. The lack of 'plausibility' of the resulting yield curves obtained using the naive procedure was quantified by analyzing the distribution of yield curve curvatures. A simple, financially justifiable, mechanism was introduced to produce more realistic curves.

We have empirically observed (and reproduced with our procedure) a strong positive serial auto-correlation on n-day returns. The dependence on n of the strength of the auto-correlation was compatible with the financial quasi-arbitrage mechanism posited to account for the curvatures. More precisely, the observed data could be explained by a positive serial correlation for all maturities, with a superimposed quasi-arbitrage activity more pronounced at the long end producing a maturity dependence both to the n-day unconditional variances and to the serial auto-correlations.

Work is in progress to extend the analysis to the joint evolutions of several yield curves (eg, Treasury, swaps and mortgage-backed securities) with stochastic, maturity-dependent spreads.

References

[1] Black F, Derman E, Toy W, (1990), "A one-factor model of interest rates and its application to Treasury bond options", Financial Analyst Journal, 1990, 33-339

[2] Ait-Sahalia Y, 1996, 'Testing Continuous-Time Models of the Spot Interest Rate', Review of Financial Studies, 9, 385-426

[3] Andersen L, Andreasen J, (2000), 'Volatility Skews and Extensions of the LIBOR Market Model', Applied Mathematical Finance, 7, March, 1-32

[4] Ang A, Bekaert G (1998), 'Regime Swirches in Interest Rates', NBER working paper 6508, 91, 61-87

[5] Bansal R, Zhou H, (2002), 'Term Structure of Interest Rates with Regime Shifts', Journal of Finance, LVII, 5, October, 1997-2043

[6] Brigo F, Mercurio D, (2001) 'Interest-Rate Models - Theory and Practice' Springer Verlag, Berlin, Heidelberg

[7] Chan KC, Karolyi GA, Longstaff FA, Sanders AB, (1992) 'An Empirical Comparison of Alternative Models of the Short-Term Interest Rate', Journal of Finance, 57, May 1992.

[8] Cox J., Ingersoll J.E., Ross S.A., (1985) 'A Theory of the Term Structure of Interest Rates' Econometrica, 53, 385-407

[9] Joshi M, Rebonato R, (2003), 'A Stochastic-Volatility Extension of the LIBOR MArket Model', Quantitative Finance, accepted for publication

[10] Garcia P, Perron P, (1996) 'An Analysis of the Real Ineterest Rate Under Regime Shifts', Review of Economics and Statistics, 78, 111-125

[11] Glasserman P, Kou, S.G (2000) 'The Term Structure of Simple Forward Rates with Jump Risk' – Working paper - Columbia University

[12] Gray, S F, (1996) 'Modeling the Distribution of Interest Rates as a Regime-Switching Process', Journal of Financial Economics, 42, 27-62

[13] Hamilton J (1989) 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle', Econometrica, 57, 357-384

[14] Heath D, Jarrow R A, Morton A, (1989) "Bond Pricing and the Term Structure of Interest Rates: a New Methodology", working paper, (revised edition) , Cornell University

[15] Ho and Lee

[16] Hull J, White A, (1990) 'Pricing Interest Rate Securities', Review of Financial Studies, 3, 573-592

[17] Jorion P, (1997) 'Value at Risk', Irving Professional Publishing, Chicago, London, Singapore

[18] Longstaff, F A, Schwartz, E S (1992) 'Interest Rate Volatility and the Term Structure: a Two-Factor General Equilibrium Model', Journal of Finance, 47, 1259-1282

[19] Longstaff F.A., Santa-Clara P, Schwartz, ES, (2000) 'Throwing Away a Billion Dollars: The Cost of Sub-Optimal Exercise Strategies in the Swaptions Markets', Working Paper, UCLA, presented at the ICBI Risk Conference – Geneva (2000)

[20] Naik V, Lee M H, (1997), 'Yield Curve Dynamics with Discrete Shifts in Economic Regimes: Theory and Estimation', working paper, Faculty of Commerce, University of British Columbia

[21] Martellini Priaulet (2001)

[22] Morton A, (1996) PhD Thesis Cornell University

[23] Prigent J-L, Renault O, Scaillet O, (2000) 'An Empirical Investigation in Credit Spread Indices', London School of Economics, Discussion Paper 363

[24] Rebonato R, (2002) 'Modern Pricing of Interest-Rate Derivatives: the LIBOR Market Model and Beyond', Chapter 1, Princeton University Press, Princeton, New Jersey

[25] Stanton, R, (1997) 'A Non-Parametric Model of Term Structure Dynamics and the Market Price of Interest rate Risk', Journal of Finance, 52, 1973-2002

[26] Vasicek O, (1977) 'An Equilibrium Characterization of the Term Structure' Journal of Financial Economics, 5, 177-188