Yield Curve Construction Models – Tools & Techniques

Summary

Yield curve models are used to price a wide variety of interest rate-contingent claims. The existence of several different competing methods of curve construction available and there is no single standard method for constructing yield curves and alternate procedures are adopted in different business areas to suit local requirements and market conditions. This fragmentation has often led to confusion amongst some users of the models as to their precise functionality and uncertainty as to which is the most appropriate modeling technique. In addition, recent market conditions, which inter-alia have seen elevated levels of LIBOR basis volatility, have served to heighten concerns amongst some risk managers and other model users about the output of the models and the validity of the underlying modeling methods.

The purpose of this review, which was carried out in conjunction with research analyst Xu Bai, now at Morgan Stanley, was to gain a thorough understanding of current methodologies, to validate their theoretical frameworks and implementation, identify any weaknesses in the current modeling methodologies, and to suggest improvements or alternative approaches that may enhance the accuracy, generality and robustness of modeling procedures.

Model Description

In this section, we describe in detail the three methodologies used in the curve construction. We choose the curve inputs as of 01/22/2008 as the typical set-up so that we have a concrete example based on which we can write the documentation and do testing and analyses. As shown in the chart below, this typical set-up includes two cash points, sixteen futures points and fifteen swap points. The basic requirement for a curve construction methodology is to correctly price the input instruments. However the interpolation methods will differ as we shall see in the following sections.

Cash			Futures				Swaps					
Term	Shift	Rate	Instr	Conv	Shift	Price	Term	Shift	Rate	StartDate	22-Jan-08	
2D	0	3.700	MAR08	0.07	0	9696.5	5Y	0	3.313	SpotDate	24-Jan-08	
3M	0	3.418	JUN08	0.12	0	9751.5	6Y	0	3.518	Method	RAW	
			SEP08	0.26	0	9763.0	7Y	0	3.694			
			DEC08	0.49	0	9762.0	8Y	0	3.845	Generate Curve		
			MAR09	0.8	0	9755.5	9Y	0	3.977			
			JUN09	1.2	0	9743.0	10Y	0	4.089	Time		3.500
			SEP09	1.69	0	9728.0	11Y	0	4.185			
			DEC09	2.26	0	9712.0	12Y	0	4.269			
			MAR10	2.9	0	9696.5	15Y	0	4.459			
			JUN10	3.63	0	9678.5	20Y	0	4.613			
			SEP10	4.43	0	9661.5	25Y	0	4.667			
			DEC10	5.3	0	9646.5	30Y	0	4.691			
			MAR11	6.25	0	9633.5	35Y	0	4.691			
			JUN11	7.26	0	9619.5	40Y	0	4.691			
			SEP11	8.34	0	9606.5	45Y	0	4.681			
			DEC11	9.49	0	9594.0						

Table 1: Cash, Futures and Swap Data Grid

RAW METHODOLOGY

The RAW is the simplest among the four methodologies. The interpolation in RAW is basically log-linear, or linear interpolation of logarithm of the discount factors, which results in piecewise linear forward rate curves. The procedure to build up a curve using RAW can be decomposed into four steps: turn-adjustment, add-cash-points, add-futures-points and add-swap-points.

Turn-adjustment

The year-end turn of the yield curve is defined as the sudden jump in yields during the change of the year. This usually happens at the end of the calendar year, reflecting increased market activity related to year-end portfolio adjustments and hedging activity. However the turn can be at other times of the year, for instance around common

financial year end dates such as the March-April period in the Japanese market. Currently methodology permits the yield curve to have turns only in the futures sector. When there is a year turn(s), two discount curves are constructed: one for turn discount factors and one for the discount factors calculated from the input instruments after adjustments and the discount factor at any time is the multiplication of two.

The very first turn is a user input: the user can specify the turn start date, for example the last business day of the year, the turn end date, usually the first business day of the following year and the turn spread. We use the following notation:

$TotalPeriod$=number of days in the futures period that comprises the year turn
$TurnPeriod$=number of days in the turn period
$InputTurnSpread$=user defined turn spread in bps
$AdjustedTurnSpread$=turn spread adjusted after lowering the corresponding futures implied rate

Then the turn discount factor can be calculated as follows:

$TurnPeriod = TurnEndDate - TurnStartDate$
$TotalPeriod = FuturesEndDate - FutureStartDate$
$AdjustedTurnSpread = InputTurnSpread * TotalPeriod / (TotalPeriod - TurnPeriod)$
$TurnDiscountFactor = exp(-AdjustedTurnSpread * TotalPeriod / 3650000)$

To lower the futures implied rate accordingly, we have:

$FutDcf$=futures daycount fraction=$TotalPeriod / 360$
$FutDiscFactor$=discount factor implied by the futures price before adustment
$\qquad = 1/(1+FutRate*FutDcf)$
$FutRate$=futures implied rate=$1-(FuturesPrice+Convexity)/10000$
$AdjustedFutRate=(TurnDiscountFactor / FutDiscFactor-1)/FutDcf$

The reason for adjusting the turn spread and the futures implied rate is to ensure the equality of the integral of the forward rates before and after the adjustments.

After the first turn inputted by user, the far turns are implied by the futures prices. For instance, there are a turn in Dec09-Jan10 and another turn in Dec10-Jan11. The implied spread is calculated using three futures that are close to the turn:

$ImpledSpread = -10000 * \left[(FutRate_{i-1} + FutRate_{i+1}) / 2 - FutRate_i \right]$
$FutRate_i$: rate implied by the futures that straddles the turn
$FutRate_{i-1}$: rate implied by the futures before
$FutRate_{i+1}$: rate implied by the futures after

From the formulae we can see that nothing a priori excludes a negative implied spread. Then the turn discount factor and the adjusted futures implied rate can be calculated similarly as for the first turn:

$TurnDiscountFactor=exp(-ImpliedSpread*TotalPeriod*TotalPeriod/(TotalPeriod - TurnPeriod)/3650000)$
$AdjustedFutRate=(TurnDiscountFactor / FutDiscFactor-1)/FutDcf$

For each turn point, the cumulative value of logarithm of turn discount factor is stored along with the corresponding turn date.

Add-cash-points

Once the turn discount curve has been constructed, we can proceed to add discount factors to the regular discount curve by starting with the cash instruments. The cash rates are simple rates with act/360 daycount convention. The discount factors can be calculated from the cash rates as follows:

$$DF = 1/(1 + dcf * CashRate)$$
$$dcf = daycount\ fraction = (SpotDate\text{-}StartDate)/360$$
$$StartDate:\ curve\ start\ date$$
$$SpotDate:\ settlement\ date,\ usually\ 2\ days\ from\ the\ start\ date$$
$$CashPointValue = log(DF)$$
$$CashPointDate = CashRateDate$$

So far we have added three points to the curve: DF(StartDate), DF(2D) and DF(3M). The very first point DF(StartDate) is equal to one of course.

Add-futures-points

There are 16 futures in our example and therefore we have to add 16 discount factors to the curve. As the first futures point (MAR08) is overlapped with the last cash point (3M), we have to replace the last cash point by the first futures point. Therefore the date for the last cash point is not 3M anymore and is replaced by 03-19-08. The corresponding curve point value is given by log-linear interpolation as follows:

$$dcf = (FutEndDate - FutStartDate)/360$$
$$LogDF = -log(1 + dcf * FutRate)$$
$$LastCashPointValue = LastYiledCurvePointValue - LogDF * (LastCashPoint\ Date - FutStartDate)\dots$$
$$/(FutEndDate - FutStartDate)$$

For other futures points, the following formulas are used:

$$dcf = (FutEndDate - FutStartDate)/360$$
$$LogDF = -log(1 + dcf * FutRate)$$
$$FutPointValue = LastFutPointValue + LogDF$$
$$FutPoint\ Date = FutEndDate$$

Add-swap-points

The discount factors on the swap dates are calculated using the traditional bootstrapping technique: given the previous found discount factors and log-linear interpolation, one find the discount factor or precisely the increment of logarithm of discount factor at the swap maturity date that makes equal the fixed leg and the floating leg of the swap. The fixed leg is usually semi-annual with 30/360 day count convention while the floating leg is quarterly with act/360 day count convention. Mathematically we have:

$$V_{fixed} = C_i \times \left(\sum_{j=1}^{M} dcf_j \times DF_j + \sum_{j=M+1}^{N} dcf_j \times DF_j(fwd) \right)$$

C:ith swap coupon

dcf_j:day count fraction for jth coupon payment

DF_j: discount factor at jth coupon payment date

M: number of payments below the curve end date

N: total number of payments

fwd: the incremental logDF to be found

$DF_j(fwd)$: discount factor at jth coupon payment date as a function of fwd

$DF_j(fwd)=DF_{end} \times exp(-fwd \times TS_j)$

TS_j:jth fixed leg time span=(jth pay date-curve end date)/FixedTimeSpan

$FixedTimeSpan$=ith swap maturity date-curve end date

DF_{end}: discount factor at the curve end date

$$V_{floating} = \sum_{j=1}^{M} (DF_{j-1} - DF_j) + \sum_{j=M+1}^{N} DF_j(fwd) \times (exp(fwd \times dT_j) - 1)$$

DF_j: discount factor at jth coupon payment date

M: number of payments below the curve end date

N: total number of payments

fwd: the incremental logDF to be found

$DF_j(fwd)$: discount factor at jth coupon payment date as a function of fwd

$DF_j(fwd)=DF_{end} \times exp(-fwd \times TS_j)$

TS_j:jth floating leg time span=(jth pay date-curve end date)/FloatingTimeSpan

$FloatingTimeSpan$=ith swap maturity date-curve end date

DF_{end}: discount factor at the curve end date

$dT_j = (PayDate(j) - PayDate(j-1))/FloatingTimeSpan$

with $dT_j = (PayDate(j) - CurveEndDate)/FloatingTimeSpan$ if $j=M+1$

Then

$$SwapValue(fwd) = V_{fixed}(fwd) - V_{floating}(fwd)$$

The objective now is to find a value of fwd that makes the swap value equal to zero. Some rooting finding routines such as *Brent's method*[1] can be used to quickly find the appropriate forward rate value. Once the fwd value is found, the point for the current swap can be added to the curve:

$$CurrentPointValue = LastPointValue - fwd$$
$$CurentPoint Date = SwapMaturityDate$$

The above process can be repeated until we find the point values for all 15 swaps in our case.

Log-linear interpolation
The key assumption in RAW methodology is the linear interpolation of the logarithm of discount factors. This will lead to the constant forward rates between any two input points as we show below. Let's assume we have two discount factors DF_1 and DF_2 at t_1 and t_2. For any time t between t_1 and t_2, the instantaneous forward rate is calculated as:

$$log(DF_t) = log(DF_1) + (t - t_1)(log(DF_2) - log(DF_1))/ (t_2 - t_1)$$

$$Forward = \frac{-\partial log(DF_t)}{\partial t} = -\frac{log(DF_2) - log(DF_1)}{(t_2 - t_1)}$$

The instantaneous forward rate clearly only depends on the DF_1 and DF_2. Consequently the forward rates are constant if they fall into the time interval t_1 and t_2. The chart below illustrates this property of the RAW methodology.
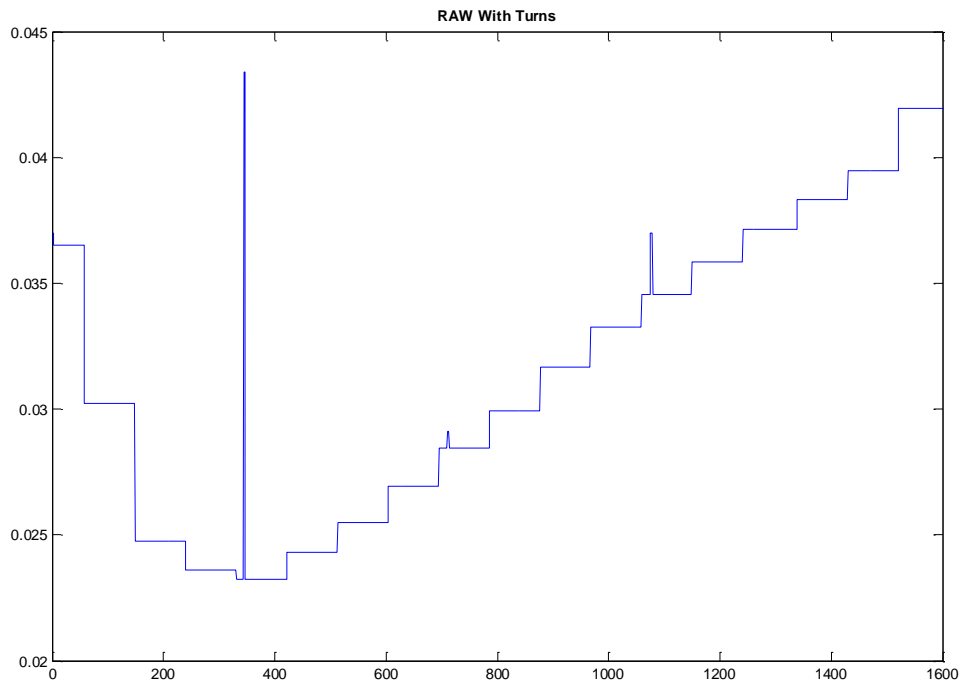

Fig. 1 - RAW piecewise-linear forward rate curve with year-end turns

SPLINE METHODOLOGY

The second methodology is named Tension Spline or Spline for short. Spline techniques have become widely adopted as a means of ensuring that forward rate curves are smooth, continuous and differentiable.

The Spline methodology also includes turn-adjustment, add-cash-points, add-futures-points and add-swap-points. The turn-adjustment in Spline is slightly different from that in RAW. For the first turn, we have:

$$TurnDiscountFactor = exp(-InputTurnSpread * TotalPeriod / 3600000)$$
$$TurnPeriod = TurnEndDate - TurnStartDate$$

To adjust the futures rate:

$$FutDcf=(FutEndDate-FutStartDate) / 360$$
$$FutDiscFactor=1/(1+FutRate*FutDcf)$$
$$FutRate=1-(FuturesPrice+Convexity)/10000$$
$$AdjustedFutRate=(TurnDiscountFactor / FutDiscFactor-1)/FutDcf$$

For far turns:

$$Im\,pledSpread = -10000 * \left[\left(FutRate_{i-1} + FutRate_{i+1} \right) / \ 2 - FutRate_i \right]$$

$FutRate_i$: rate implied by the futures that straddles the turn

$FutRate_{i-1}$: rate implied by the futures before

$FutRate_{i+1}$: rate implied by the futures after

$TurnDiscountFactor = exp(-ImpliedSpread * TotalPeriod / 3600000)$

$AdjustedFutRate = (TurnDiscountFactor / \ FutDiscFactor-1) / FutDcf$

Add-cash-points[1] and add-futures-points remain the same as in RAW. The big differences reside in the interpolation method and the add-swap-points.

Tension-Spline Interpolator

The Spline methodology[2] was developed by FAST about ten years ago. Although the original documentation describes that the instantaneous forward rate is modeled by a tension spline, an examination of the code revealed that it is actually the (negative) logarithm of the discount factor (with some transformations) that is modeled.

A tension spline can have the following form on all intervals $\left[t_j, t_{j+1} \right]$:

$$f(t) = a_0 + a_1 t + a_2 exp(\sigma t) + a_3 exp(-\sigma t)$$

$$t \in \left[t_j, t_{j+1} \right]$$

σ is the tension parameter

When the sigma tends to zero, the tension spline degenerates into a regular cubic spline. On the other hand, when the sigma increases, the tension spline asymptotically reduces to piecewise linear interpolation.

We describe briefly the curve modeling scheme using the tension spline. We have as inputs n discount factors DF_i and the associated dates D_i and the following transformations are performed first:

$$TimeScaling = (t_n - t_1) / (n - 1)$$

$$TimeShift = M_1 - TimeScaling$$

$$t_i = D_i - AsOfDate$$

n: number of points

t_i: ith discount factor maturity in days

D_i: ith discount factor date

DF_i: ith discount factor

AsOfDate: curve date

$$Z_i = (-log(DF_i)) / \ t_i$$

$$t_i = (t_i - TimeShift) / \ TimeScaling$$

$$Z_i = Z_i \times t_i$$

With these transformations, the maturities of the discount factors will range from 1 to n and Z can be considered as the negative logarithm of the discount factor on the new time axis. It is ultimately the Z and t that are modeled using tension spline.

---

[1] As we shall see soon we detected an inconsistency in overlapped cash-futures region which caused some excessive swings for Spine interpolation.

For each transformed time interval $[t_j, t_{j+1}]$, a relative time scale is used again so that the time interval becomes $[t_0, t_1]$ with $t_0 = 0$ and $t_1 = t_{j+1} - t_j$. Then the tension spline can be rewritten in the following form:

$$Z(t) = Z_0 + Z_0't + Z_0'' \frac{COSH - 1}{\sigma^2} + Z_0''' \frac{SINH - \sigma t}{\sigma^3}$$

$t \in [t_0, t_1]$

$COSH = 0.5 \times [exp(\sigma t) + exp(-\sigma t)]$

$SINH = 0.5 \times [exp(\sigma t) - exp(-\sigma t)]$

$Z_0 = Z(t_0)$: value of Z at t_0

$Z_0' = Z'(t_0)$: first derivative of Z at t_0

$Z_0'' = Z''(t_0)$: second derivative of Z at t_0

$Z_0''' = Z'''(t_0)$: third derivative of Z at t_0

We can easily verify that this form is equivalent to the form we defined earlier and the conditions are satisfied.

The key to have a smooth curve is not only make the function $Z(t)$ go through the transformed input points, i.e., $Z(t_0)$ and $Z(t_1)$ equal to the tabulated values at t_0 and t_1 for each interval, but also make the $Z'(t)$ and $Z''(t)$ equal at the joint points. Along with some boundary conditions, this is equivalent to solving a matrix equation whose unique solution determines for each interval the values of Z_0, Z_0', Z_0'' and Z_0''' which totally define the curve through each time interval.

Add-swap-points
As the spline curve depends on neighboring points, one cannot use the usual bootstrapping technique to find the discount factor that matches the input swap rate one after another. Instead a trial-and-error approach is used to simultaneously find the discount factors that matched all input swap rates. Mathematically we have for the i[th] swap:

$$V_{fixed}^i = \sum_{j=1}^{M} dcf_j \times DF_j$$

dcf_j: day count fraction for jth fixed coupon payment

DF_j: discount factor at jth coupon payment date

M: total number of fixed payments

$$V_{floating}^i = \sum_{j=1}^{N} (DF_{j-1} - DF_j)$$

DF_j: discount factor at jth floating coupon payment date

N: total number of floating payments

Then the i[th] swap rate is calculated as:

$$SwapRate_i = V_{floating}^i / V_{fixed}^i$$

If we have K discount factor to be found in order to match K input swap rates (K=15 in our case), we can use a trial vector $[df_1, df_2, ..., df_K]$ along with the discount factors we already found in cash and futures region to build up the spline curve. Then we have to zero the following vector:

$$[(SwapRate_1 - InputSwapRate_1), (SwapRate_2 - InputSwapRate_2), ..., (SwapRate_K - InputSwapRate_K)]$$

We can use multidimensional secant methods such as *Broyden's method* [1] to find the discount factor vector $[df_1, df_2, \ldots, df_K]$ that makes the output swap rates equal to input swap rates.
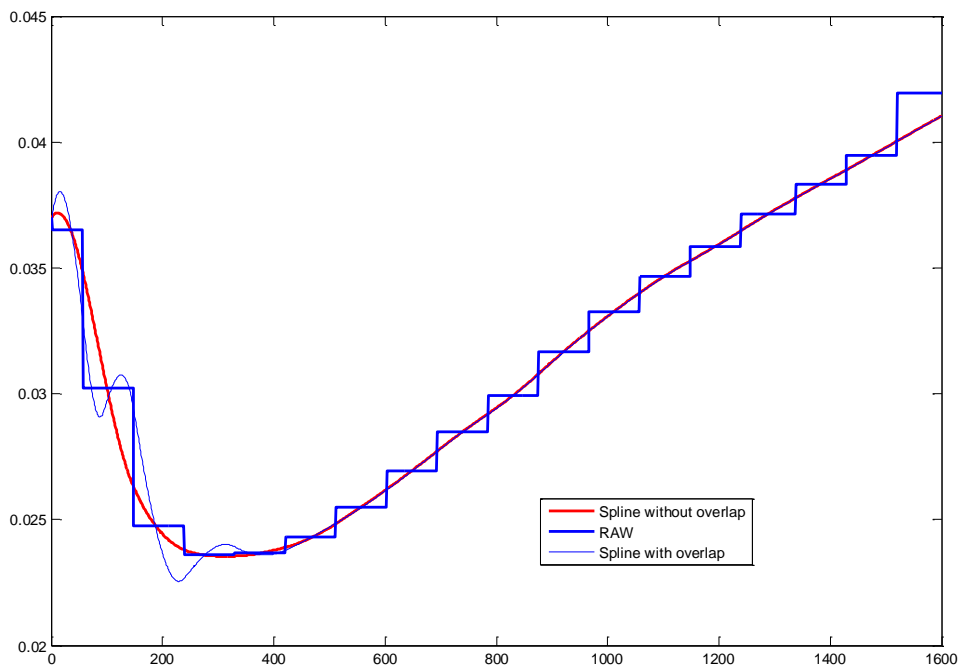

Fig 2 - RAW and Spline-interpolated forward rate curves (with and without overlap)

As we mentioned earlier, we found an inconsistency in overlapped cash-futures region where two discount factors calculated from a cash point and a futures point overlaps and causes excessive oscillation at the short-end of the curve as shown in the figure above. The solution is simply to make a choice between the two competing rates, rather than attempting to use both, especially when they are very different. As illustrated in the figure above, when we replace the last cash point by the first futures point (as in RAW), the spline curve becomes much smoother.

In figure 3 below, we show the spline curve with different tension values: when the tension increases, the interpolation becomes more linear. This effect is more obvious at the long-end of the curve due to the longer time interval.
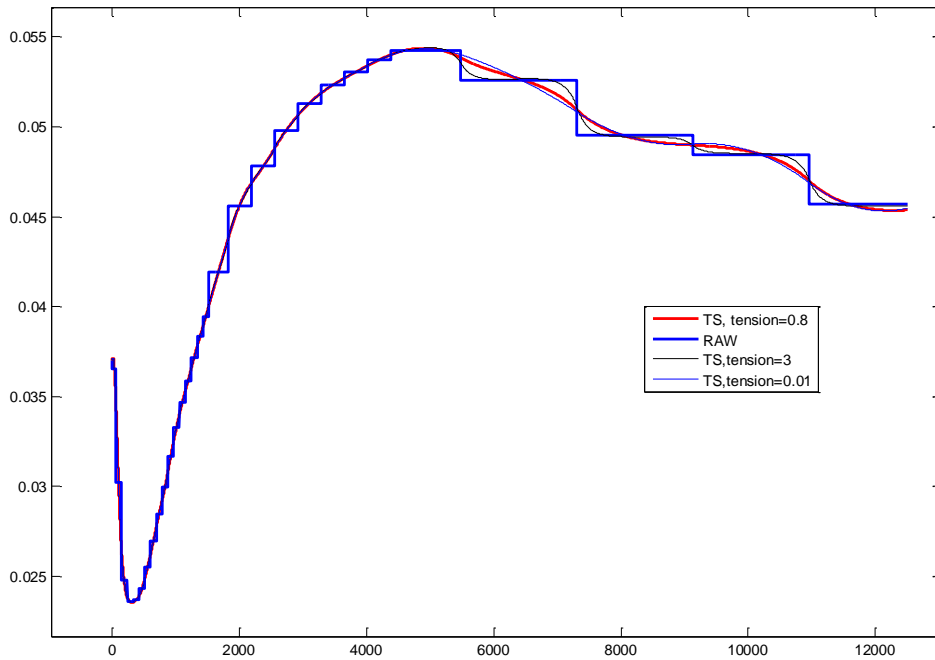
Fig. 3 - Forward rate curves with different tension parameter values

DENSE METHODOLOGY

The DENSE curve uses the same RAW methodology for cash and futures region, i.e., log-linear interpolation which results in piecewise constant forwards. For the swaps, a spline curve is first built. From this spline curve, quarterly spaced swap rates are generated. Then we revert to the RAW method to generate a forward curve based on these quarterly apart swaps. As shown in the figure below, this allows the curve to follow the general shape of a smooth spline curve and at the same time to have constant forward rates between quarterly grid points in the swaps region.
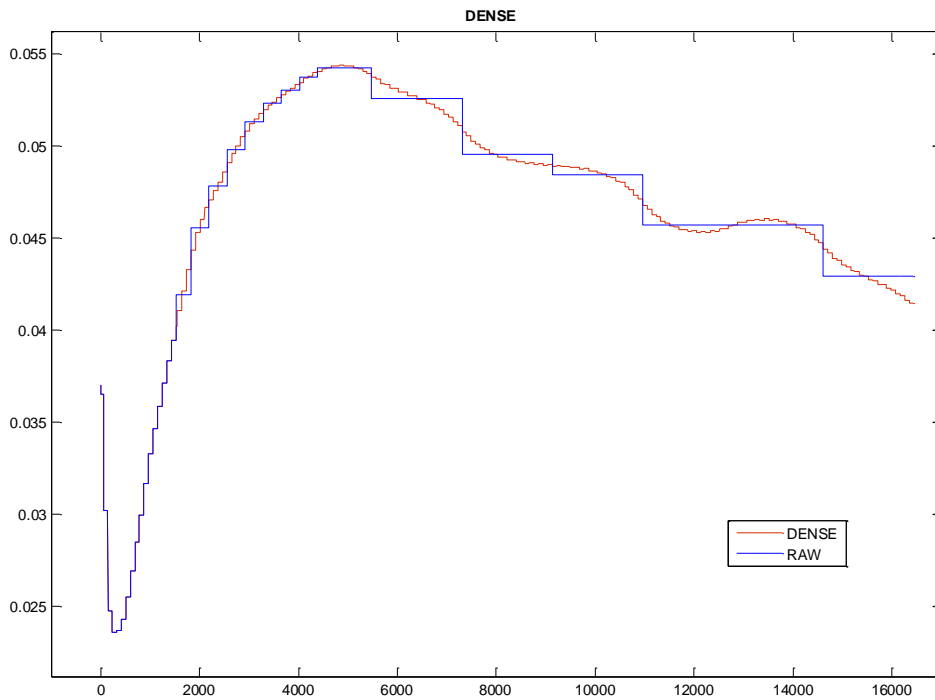
Fig 4 - DENSE and RAW forward rate curves

DUALDENSE METHODOLOGY

For the short-end of the curve, daily grid points are usually generated from neighboring input points. From these daily points, one can use one of five interpolation methods: linear cash, exponential cash, linear DF, exponential DF and log-linear DF to generate forward curve and at the same to reprice correctly the input instruments. For the long-end of the curve, the procedure used by DUALDENSE is the same as for DENSE. Furthermore, rather than applying a single tension parameter for the entire curve, one can use different tensions for each curve segment in order to optimize the curvature and fit to target instruments with corresponding weights[2]. The following figure illustrates this curve methodology.

---

[2] For more details see [4].
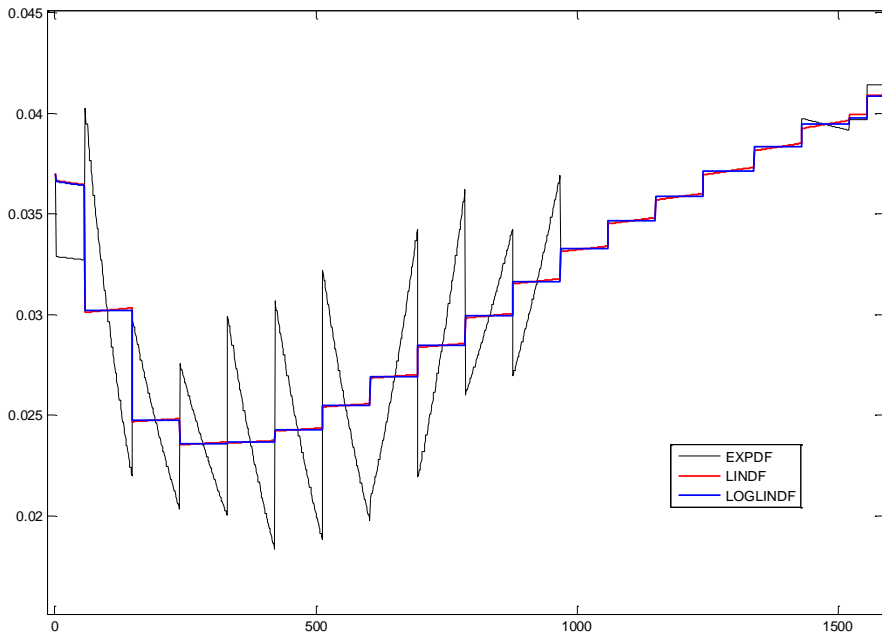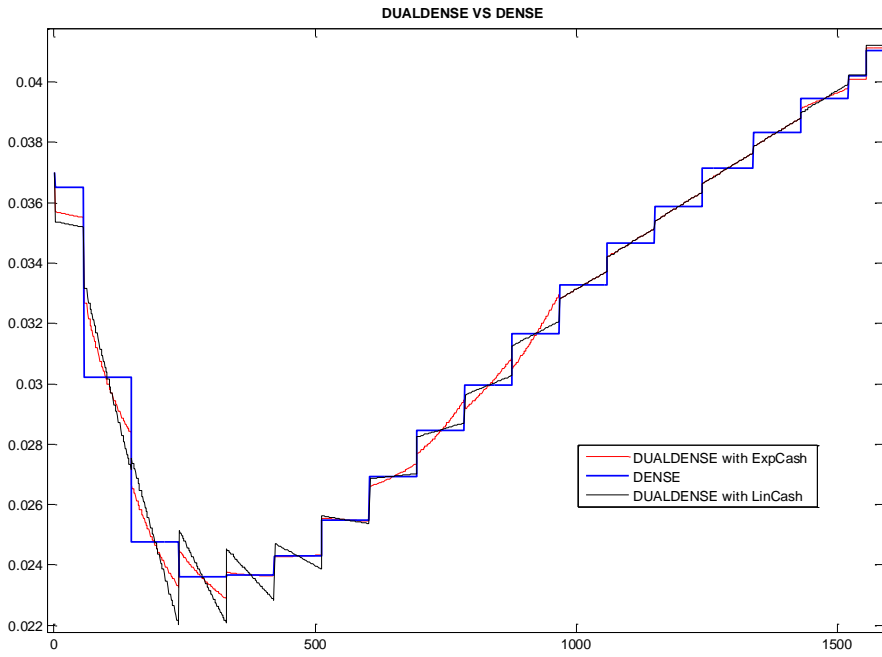
Fig 5 - DENSE and DUALDENSE methodologies compared for the front end of the forward curve

Extensions and Alternative Approaches

Tension / Smoothed Splines

The tension spline is the interpolation method employed by the SPLINE model and is the base of the DENSE and DUALDENSE models.

Consider a standard cubic spline g(t) interpolating a set of data points (t_j, g_j), $j = 1,…,M$, where the t_j are said to be knot points. The spline function is piecewise linear in its second derivative:

$$g''(t) = \frac{t_{j+1} - t}{t_{j+1} - t_j} g''_j + \frac{t - t_j}{t_{j+1} - t_j} g''_{j+1}, \quad t \in [t_j, t_{j+1}]$$

Explicit equations for the interpolating spline can be recovered from the above by integrating the equation and requiring the curve to pass through the given grid points, as well as having continuous first derivatives across knots. To uniquely specify the cubic spline, derivative boundary conditions must be specified at $t = t_1$ and $t = t_M$. The so-called natural spline has boundary conditions setting the second derivative to zero at these end points.

Cubic splines suffer from the defect that large values of the second derivative are hard to achieve, making tight turns difficult, and exhibit non-local behavior in the sense that perturbations at one point in the curve will have secondary effects in regions of the curve far from the point of disturbance. Tension splines seek to remedy this by applying a tensile force at the end-points of the spline, formally:

$$g''(t) - \sigma^2 g(t) = \frac{t_{j+1} - t}{t_{j+1} - t_j} \left(g''_j - \sigma^2 g_j \right) + \frac{t - t_j}{t_{j+1} - t_j} \left(g''_{j+1} - \sigma^2 g_{j+1} \right), \quad t \in [t_j, t_{j+1}]$$

For σ = 0, this reduces to the standard cubic spline form.

Let φ be a representation of the yield curve such that

$$P(t) = P(t, \varphi(t))$$

And let

$$\varphi = \left[\varphi(t_1), \dots, \varphi(t_M) \right]^T$$

be the curve we have used as a representation of the discount curve.

Now let $\quad V = cP$

Where

$V = (V_1, \dots, V_N)^T$ is a vector of observable securities prices,
$C = \{c_{ij}\}$ is an N x M matrix of cash flows and
$P = (P(t_1), \dots, P(t_M))^T$ is a vector of discount factors

Then we aim to minimize the norm

$$J(\varphi) = \frac{1}{N} (V - cP(\varphi))^T W^2 (V - cP(\varphi)) + \lambda \left(\int_{t_1}^{t_M} \left[\varphi''(t)^2 + \sigma^2 \varphi'(t)^2 \right] dt \right)$$

Where **W** is a diagonal N x M matrix of weights which can be used to assign different importance to various of the benchmark securities and λ and σ are positive constants relating the smoothness and tension of the spline curve.

The first of the two terms in the norm consists of the standard least-squares penalty term.

The second term, $\int_{t_1}^{t_M} |\phi''(t)|^2$ is a weighted smoothness term which penalizes high second order gradients of ϕ to avoid kinks an discontinuities.

The optional final term, $\sigma^2 \int_{t_1}^{t_M} |\phi'(t)|^2$ penalizes oscillations and excess convexity/concavity.

The norm is typically minimized using Gauss-Newton iteration.

DUALDENSE also allows for non-uniform tension in the spline, with local tension parameters σ_j. In this case the norm to be minimized is:

$$J(\phi) : \frac{1}{N}(V - \mathbf{cP}(\phi))^T W^2(V - \mathbf{cP}(\phi)) + \lambda \left(\sum_{j=1}^{M-1} \int_{t_1}^{t_{j+1}} \left[|\phi''(t)|^2 + \sigma_j^2 |\phi'(t)|^2 \right] dt \right)$$

One problem not addressed by the current methodology is the choice of smoothing constant λ. Wahba (1990) provides a technique, referred to as generalized cross validation (GCV), for estimating λ by non-linear least squares minimization of an expression in the form:

$$\min_{\phi(\lambda)} \left[(P - \phi(\phi(\lambda)))^T (P - \phi(\phi(\lambda))) + \lambda \phi(\lambda)^T \mathbf{H} \phi(\lambda) \right]$$

Where **H** is a band-diagonal matrix determined by the knot points of a B-spline basis.

We return to this issue in the section on Basis Splines methodology.

Implementation of Tension / Smoothed Spline Model

Figure 6 shows the forward rate curve for the test data set estimated using a standard cubic spline, compared to the constant forward rate approach of the RAW model.
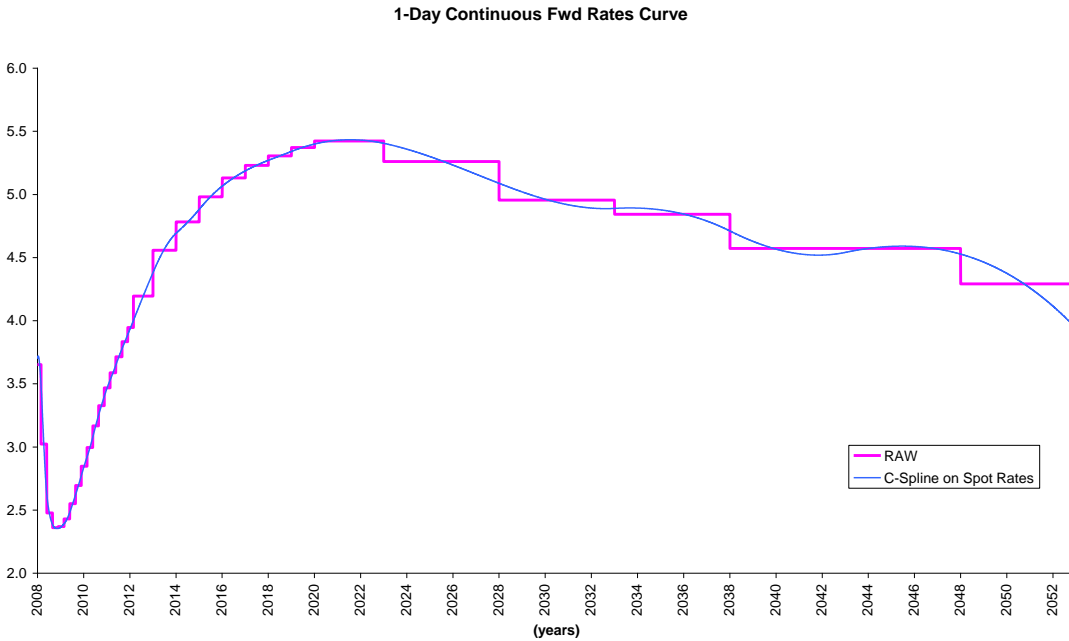

Fig 6 - RAW and cubic spline forward rate curves

The attractions of the C-spline model are immediately obvious in terms of producing of a smooth, continuous forward rate curve. However, the drawbacks of the model are equally apparent in the excess fluctuations at the long end of the curve, demonstrating the need for a cost function which penalizes such oscillations. As noted above, these shortcomings can be addressed by introducing a smoothing term, the effects of which are described in figure 7. The simplest approach, in which a single smoothing parameter is used for the entire curve, is clearly inadequate – the problems of excess oscillation at the far end of the curve are addressed at the expense of over-smoothing the front-end. The preferred approach is to allow the smoothing parameter (or tension) to vary from segment to segment, producing a spline of variable smoothness as illustrated in the figure.
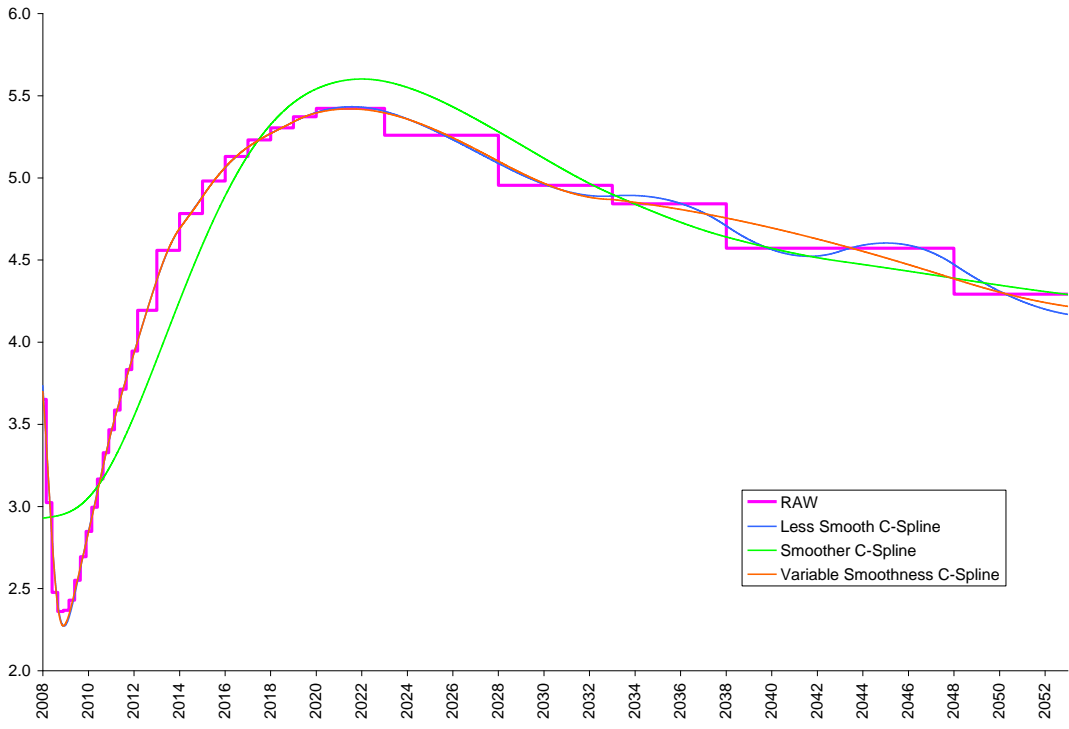
Fig. 7 – Splines fitted with constant and variable smoothness

Problems Caused by Perturbations

Recent market volatility and concerns over liquidity have accentuated basis fluctuations in LIBOR markets as the trading volume in near-term futures contracts has expanded to meet hedging and trading demands. This has created difficulties for curve estimation using standard models, which rely on spline techniques. The difficulties are illustrated in the figures below, which show how the spline smoothing conditions create significant oscillation in the curve outside the local area of perturbation.
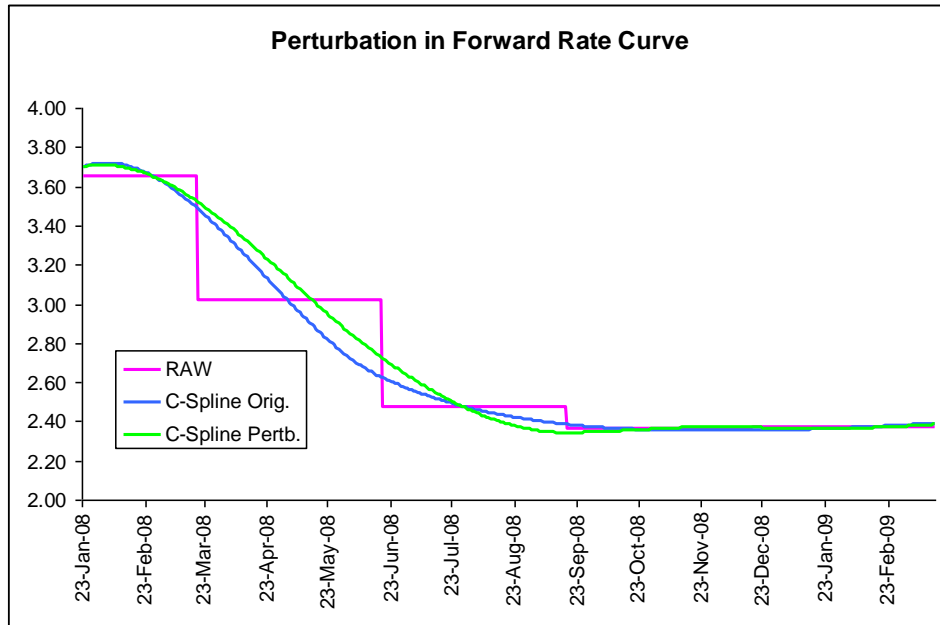
Fig. 8 – Effect of perturbation in the forward grid rates
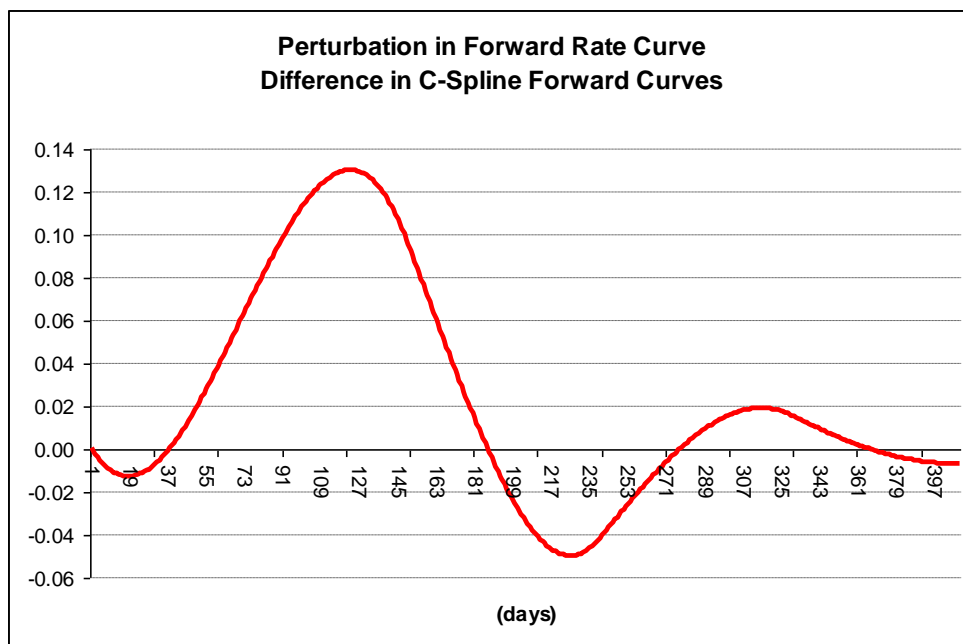
Fig. 9 – Difference in cubic spline forward curves

As foreshadowed in the initial discussion of the spline methodology, excessive oscillatory behavior can be corrected by applying appropriate weightings to the grid relevant points that the modeler judges to be adversely affected. We illustrate this with the following example, in which we introduce a very substantial uplift in the yield on the 4[th] futures contract. This produces the expected step-change in the forward curve derived from the RAW model and would be expected to induce severe oscillation in the spline curve, not only in the region of local disturbance, but further afield in sectors of the curve around the third and fifth grid points. However, by providing an appropriate weighting adjustment for the outlier, the modeler can readily control its influence, as illustrated in the figures following.

The appropriate weighting to apply to securities prices that are potential outliers is very difficult to calibrate. Statistical techniques exist, but these represent a rather blunt instrument in the context of yield curve modeling and, typically, weightings are a matter best left to the expert judgment of market practitioners. This methodology represents a more justifiable approach to making appropriate corrections for local disturbances than manipulating convexity adjustment factors for the futures contracts themselves, which appears to be a practice adopted by some model users. Instead, outliers should be treated by adjustments to the weighting matrix, while convexity adjustment factors should be calibrated to an appropriate interest rate model, such as, for example, Hull & White, in conformance with standard industry practice.
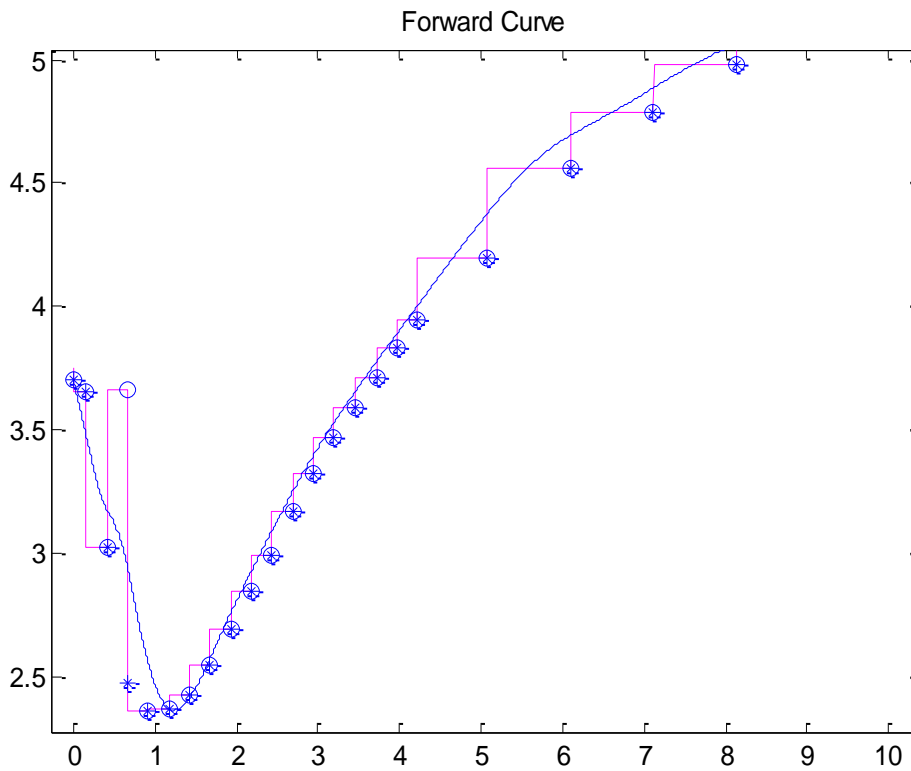

Fig.10 – Original (*) and perturbed (O) grid points and RAW (red) and weight-adjusted spline (blue) forward curves.
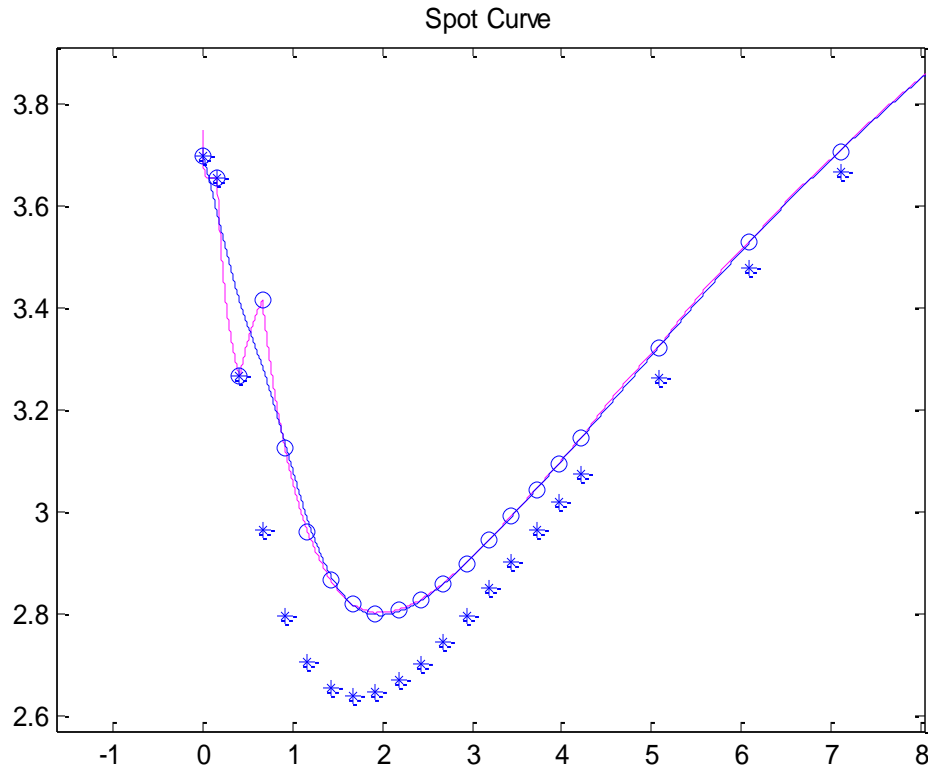
Fig 11 – Original (*) and perturbed (O) spot curve grid points and RAW (red) and spline (blue) spot rate curves

Basis Spline Modelling

Under the B-spline scheme, a spline function f is defined as a linear combination of underlying spline functions known as basis splines, each of which is positive on some interval defined by knot points. More formally,

$$f = \sum_{i=1}^{n} B_{ik} a(:, i)$$

Where B_{ik} is the i[th] order spline of order k, with knot points t(i), …, t(k), over which interval it is positive in value. It is a piecewise polynomial of order k with breaks at the knot points t(i), …, t(k). These knots may coincide, and the precise multiplicity governs the smoothness with which the two polynomial pieces join there.

For example, Fig 12 below shows a picture of a B-spline with knot sequence [0, 1.5, 2.3, 4, 5] which is of order 4, together with the underlying polynomials that make up the spline.
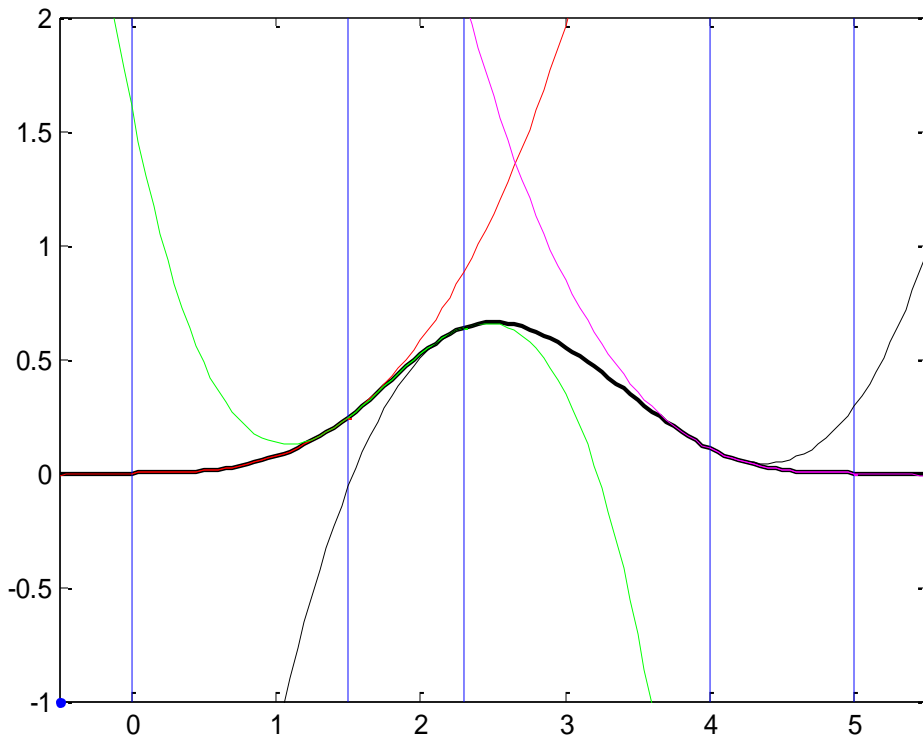
Fig. 12 – a B-spline of order 4 and the cubic polynomials of which it is composed

We can use B-splines to represent the discount function, the spot rate curve, the forward rate curve, or some function of these and later in this report we discuss some of the pros and cons of the alternative approaches. For now, we focus on modeling the spot rate grid.

As before, we employ norm minimization to identify the B-spine form of the smoothest function that lies within a specified tolerance of the specified spot rate grid (t(j), s(j), j = 1, . . ., M.

The penalty function contains both a distance metric

$$E(f) = \sum_{j=1}^{n} w(j)|(y(:,j) - f(x(j)))|^2$$

and a roughness measure
$$F(D^m f) = \int_{\min(x)}^{\max(x)} \lambda(t)|D^m f(t)|^2 dt$$

where $D^m f$ denotes the mth derivative of **f** and **W** is a N x M matrix of weights which can be used to assign different importance to various of the grid inputs.

The spline f is determined as the unique minimizer of the expression $\rho E(f) + F(D^m f)$ where ρ is a smoothing parameter.

This methodology provides an automated procedure for estimating parameters that result in a curve which closely fits the grid data with a smooth, continuous forward rate curve with a B-Spline function for which the number and

location of knot points is optimized. Applying this modeling procedure to our test data set results in a quartic B-spline with 32 knots, shown in figure 13.
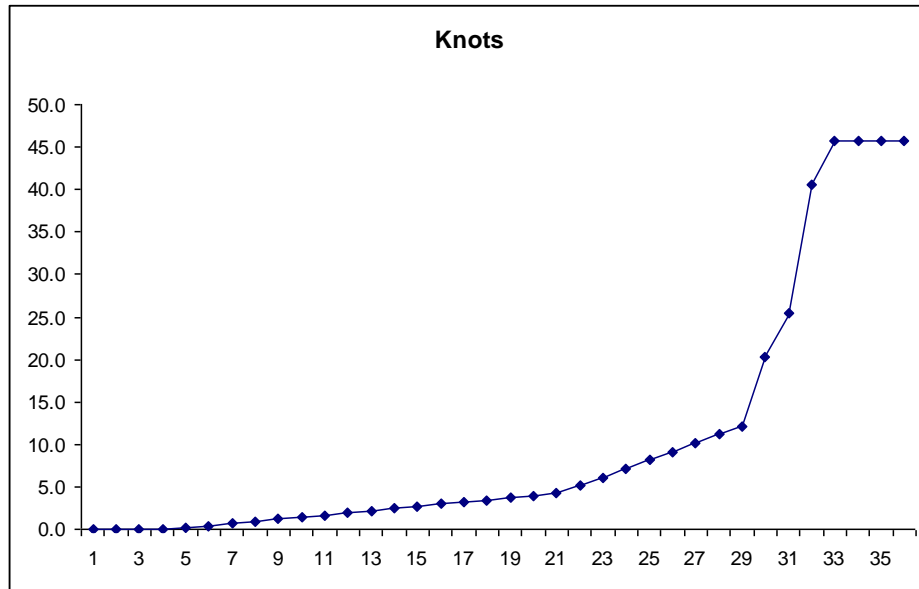

Fig. 13 – Knot points for the fitted B-spline function

Results from the optimized B-spline are contrasted with the estimated C-spline forward curves in figures 14 and 15 below. The fits are very similar both in the belly of the curve and at the long end, but the B-Spline makes a tighter turn around the main inflection point in the front end of the curve, as can be seen from figure 15.
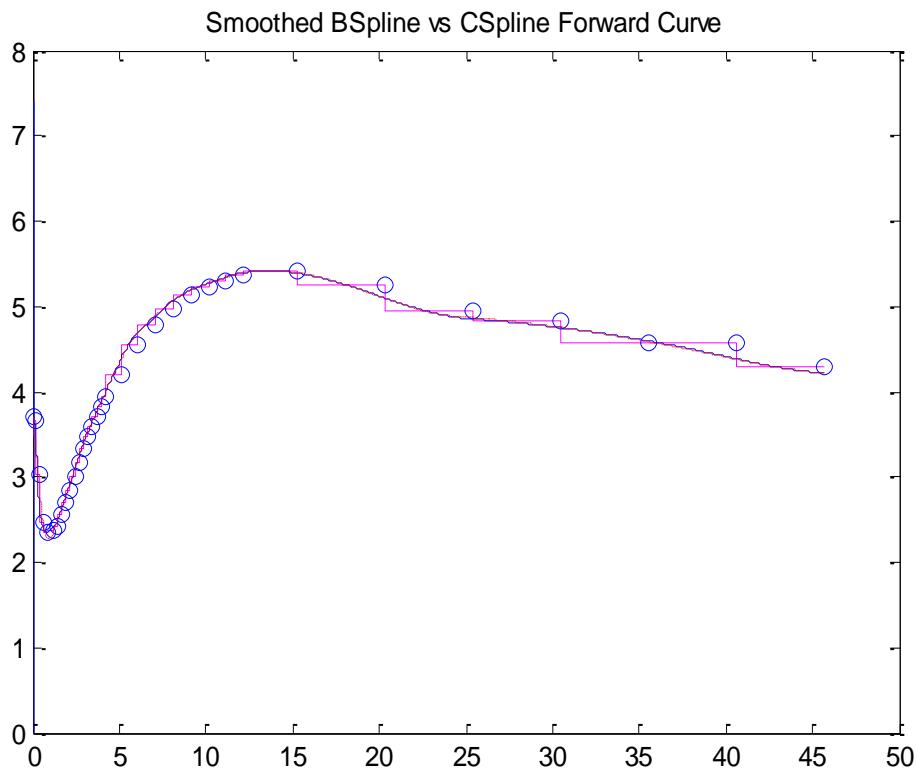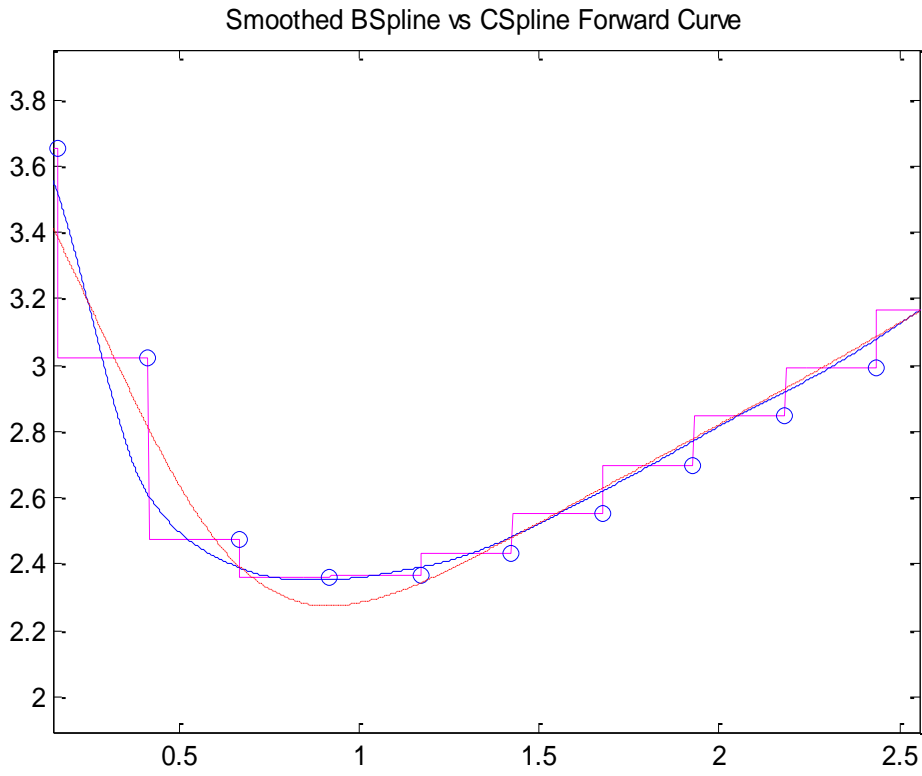
Fig 15 –Comparison of B-spline (blue) and C-spline (red) forward rate curves

Perturbation and Smoothing of the B-Spline curve

As with the C-spline model, weights can be applied to the spot rates at the each grid point to reduce the impact of outliers. Applying the same shock to the 4th futures contract as before, and reducing the weighting of the 4th grid point relative to its near neighbors, results in a curve that smoothes out the effects of the outlier, as illustrated in the figure 16.
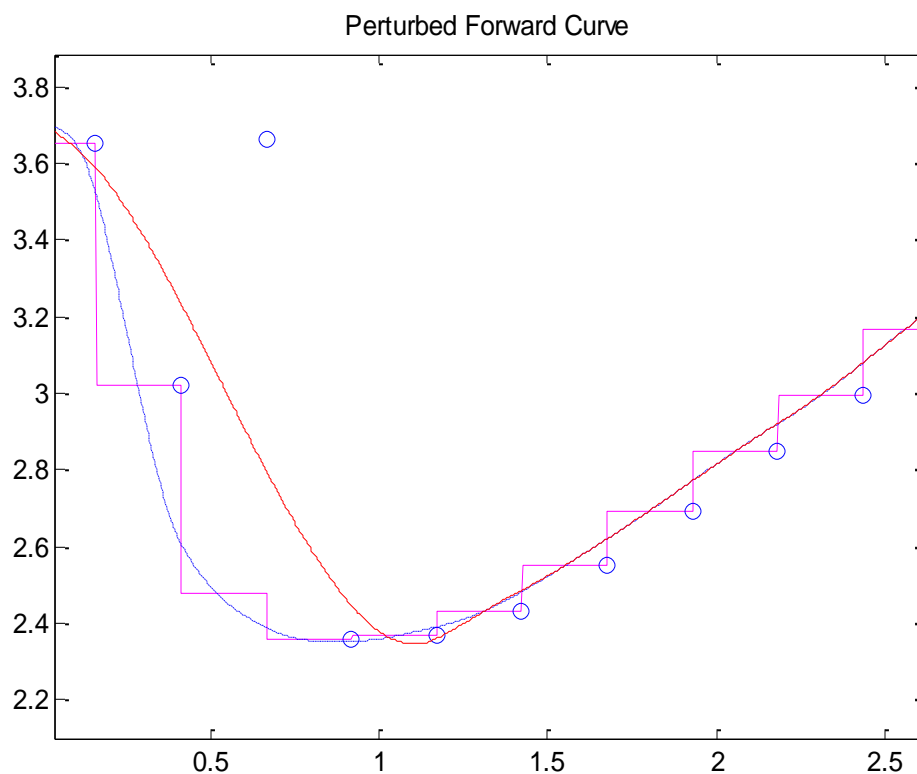

Fig 16 –Comparison of original B-spline (blue) and smoothed B-spline (red) forward rate curves after perturbation to the futures grid data

Procedures for Handling Outliers

Procedures for handling outliers have taken on additional importance given recent market turbulence, which has produced unusually high levels of volatility in the LIBOR basis as investors seek to adjust position hedges and/or favor safe-haven investments. The greatest effects are seen in the near-term Eurodollar futures contracts, and these price "aberrations" are transmitted onwards through the strip as investors roll futures positions. To the extent that analyst believes that, under conditions of market stress, certain individual futures prices are outliers with respect to the true underlying forward rate curve, it makes sense to consider methods by which adjustments can be made to the estimation methodology that will produce more reliable curve estimates.

Weighting

Making adjustments to the weight matrix, as described earlier in this report, is an appropriate and effective method of handling outliers, but the procedure is somewhat arbitrary since it relies exclusively on the judgment of the analyst. A more systematic approach that reflects current concerns about liquidity effects at the front of the futures strip would be to derive weights based on (changes in) daily trading volumes and open interest. For example, to take account of the forward-roll effect described above, weights for near-term futures contracts could be adjusted by a factor representing a combination of the remaining term to maturity and changes in daily trading volume (relative

to open interest). Although we have not formulated a specific procedure in this analysis, it is likely that a feasible result could be attained by a procedure of this kind.

FRAs

An alternative approach that deserves consideration would be to use FRA rates in the construction of the front of the curve, supplementing and/or replacing some of the near-term futures contracts as the instruments of choice for certain grid dates. This would serve as a consistency check for short term forward rates and help to identify outliers in the futures strip. FRA's have the added advantage of avoiding the need for convexity adjustments, which can cause problems for an estimation procedure that is reliant on using futures prices exclusively. Liquidity in FRA contracts is rapidly improving and this approach is probably viable as an adjunct to existing modelling procedures for US$ and CAN$ curves. We understand that F.A.S.T. is currently developing new functionality to enable FRA rates to be used in curve construction.

Which Grid Data to Fit?

One unanswered question is whether there are any relative advantages of fitting the curve to securities prices, discount factors, spot rates, or to the forward rate grid.

To address this question we follow a procedure employed by Fisher, Nychka and Zervos (1994), using Monte Carlo simulation to gauge the ability of the various estimation techniques to price securities accurately and to uncover the true spot and forward rate curves. We begin by assuming a particular functional form for the forward term structure, from which we generate "true" spot rates and zero coupon discount factors. Adding noise to these discount factors we produce "observed" ZCB prices which include a Normal random term distributed with zero mean and deterministic standard error σ. We then fit the term structure to each of these sets of generated data using various functional forms and parameterizations, estimating an optimized B-spline curve using the generalized cross validation (GCV) procedure described in the foregoing section of this report. This procedure is repeated in R simulation runs.

We look at two sorts of criteria for goodness of fit, one relating to how well we uncover true zero coupon bond prices and the other relating to how well we uncover spot and forward rate curves.

For bond prices we examine the mean absolute pricing error (MAPE):

$$\mathbf{MAPE}: \quad \frac{1}{Rn} \sum_{r:\,1}^{R} \sum_{i:\,1}^{n} \left| \hat{p}_r(t_i) \cdot p_r(t_i) \right|$$

Where the $\hat{p}_r(t_i)$ are the observed bond prices, and $p_r(t_i)$ are the true bond prices at time t_i.

For spot and forward rates we consider both the bias, which for a given maturity t_i is defined as:

$$B(t_i): \quad \frac{1}{R} \sum_{r:\,1}^{R} \left(\hat{\zeta}_r(t_i) \cdot \zeta_r(t_i) \right)$$

and the standard error, which for a given maturity t_i is defined as:

$$S(t_i): \quad \left(\frac{1}{R} \sum_{r:\,1}^{R} \left(\hat{\zeta}_r(t_i) \cdot \frac{1}{R} \sum_{r:\,1}^{R} \hat{\zeta}_r(t_i) \right)^2 \right)^{1/2}$$

Where ζ can refer to either the spot or forward rate.

The function forms chosen for the forward rate curve for test purposes are as follows:

$f_1(t) = 0.05 + 0.000584t$

$f_2(t) = 0.05 + 0.0004t - 0.0000133t^2$

$f_3(t) = 0.05 + 0.000266t + 0.000044t^2 - 2.3 \times 10^{-6}t^3 + 2.4 \times 10^{-8} t^4 + 0.002$
$\quad\quad Cos[0.566t]$

In this test we assume that price volatility at time t_i is function of the square root of time, i.e. $\sigma_{t_i} = \sigma \sqrt{t_i}$

where $\sigma = 10^{-5}$

Bias and standard errors are reported in basis points, while MAPE is reported in 1/10 basis point.
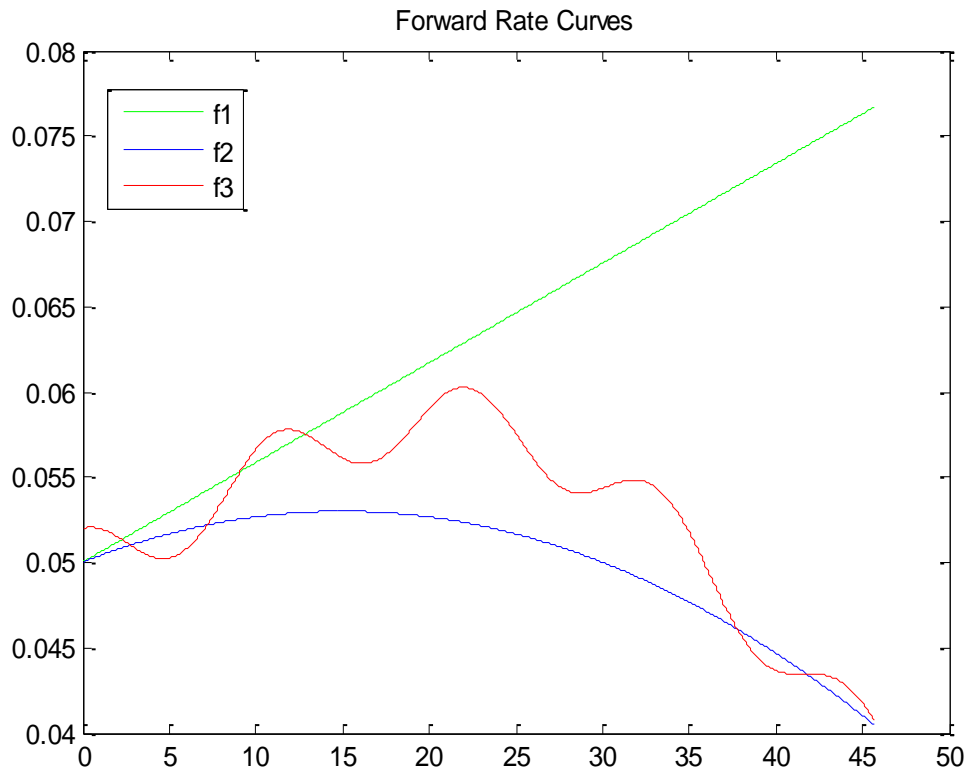
Fig. 17 - Synthetic forward rate curves

Test Results Summary

Results from the testing process, using 100 runs of the Monte-Carlo simulation model, are shown in the charts and tables following. In summary, in line with the results obtained by other researchers, we find that using forward rates for curve estimation generally improves the ability to discover bond prices, as well as spot and forward rates, by a significant margin. The pricing error in bond prices, as well as bias in estimated spot and forward rates, is typically at least one order of magnitude smaller for curves fitted to forward rates, compared to other grid data, as can be seen from Tables 2.1 – 2.3. Conversely, curves fitted to a grid of discount factors, the traditional approach originally adopted by McCulloch, performs the worst of all three methods on most criteria, across all of the functional forms considered here.

Turning to the detail of the results, for the linear functional form F1, all three techniques perform reasonably well in fitting the spot and forward curves out to the 10-year horizon. Thereafter, however, both DF-fitted and spot rate-fitted curves have an upward bias in the range of between 6 and 11 basis points in both spot and forward curves, compared with a bias of less than 1/10 of a basis point for the forward rate-fitted curves (see charts in figure 18). While the standard error of the spot rates is comparable amongst all three methods, the forward-rate fitted curve has an average standard error in fitted forward rates that is significantly lower than for curves estimated by other techniques.

The same pattern of results holds for the second and third functional forms, as can be seen from figures 19 and 20 and the results in Tables 2.2. and 2.3. The error in the forward rate-fitted curve is an order of magnitude lower than for other curves, on most goodness-of-fit criteria.

It is noteworthy that the B-spline GVC approach performs particularly well when applied to a forward rate grid representing the most complex functional form F3. The curve fits the true forward curve closely at the front of the curve where grid data is plentiful and then smoothes the curve in an appropriate manner (i.e. with increasing linearity) from the 20-year maturity, when grid data becomes relatively scarce. Here, too, the average biases in estimated spot and forward rates are an order of magnitude lower than for curves fitted to discount factors and spot rates.

Goodness of Fit Test Results

Functional Form F1

ESTIMATION	MAPE	SPOT BIAS	SPOT SE	FWD BIAS	FWD SE
DF	30.03	10.41	0.0013	11.39	0.0126
SPOT	29.71	6.27	0.0011	11.06	0.0136
FORWARD	0.42	0.04	0.0012	0.02	0.0095

Table 2.1 - Goodness of fit test results – functional form F1

Functional Form F2

ESTIMATION	MAPE	SPOT BIAS	SPOT SE	FWD BIAS	FWD SE
DF	12.88	2.58	0.0012	-7.28	0.0078
SPOT	12.07	-1.65	0.0008	-7.08	0.0089
FORWARD	0.43	0.00	0.0011	-0.04	0.0078

Table 2.2 - Goodness of fit test results – functional form F2

Functional Form F3

ESTIMATION	MAPE	SPOT BIAS	SPOT SE	FWD BIAS	FWD SE
DF	18.18	3.05	0.0021	-9.29	0.0088
SPOT	16.73	-0.68	0.0009	-8.94	0.0096
FORWARD	2.05	0.17	0.0012	-0.61	0.0079

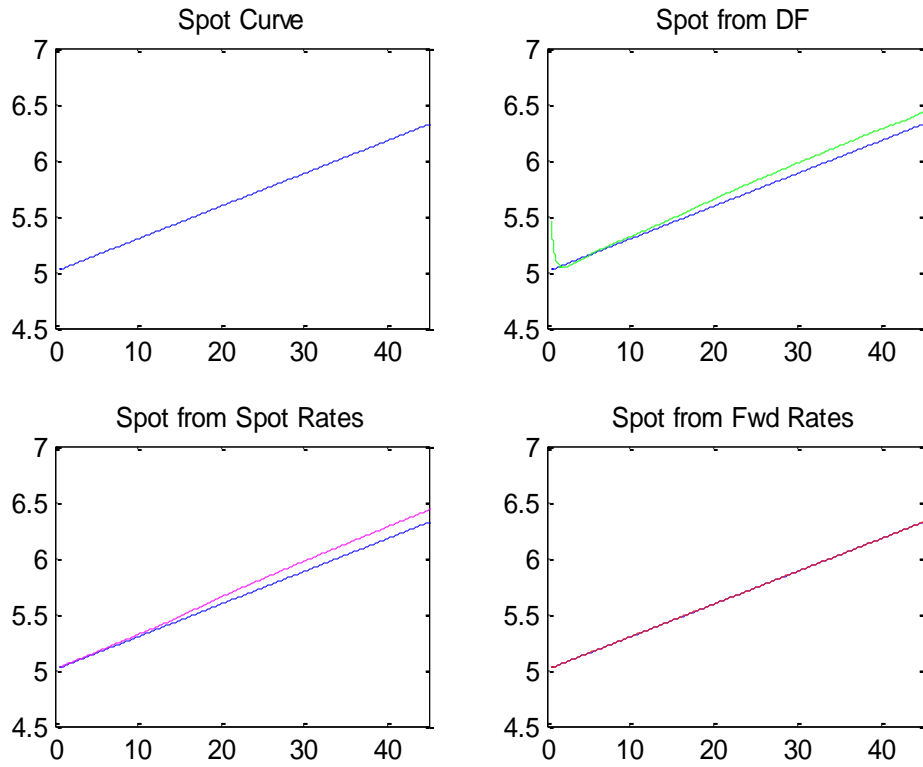Table 2.3 - Goodness of fit test results – functional form F3

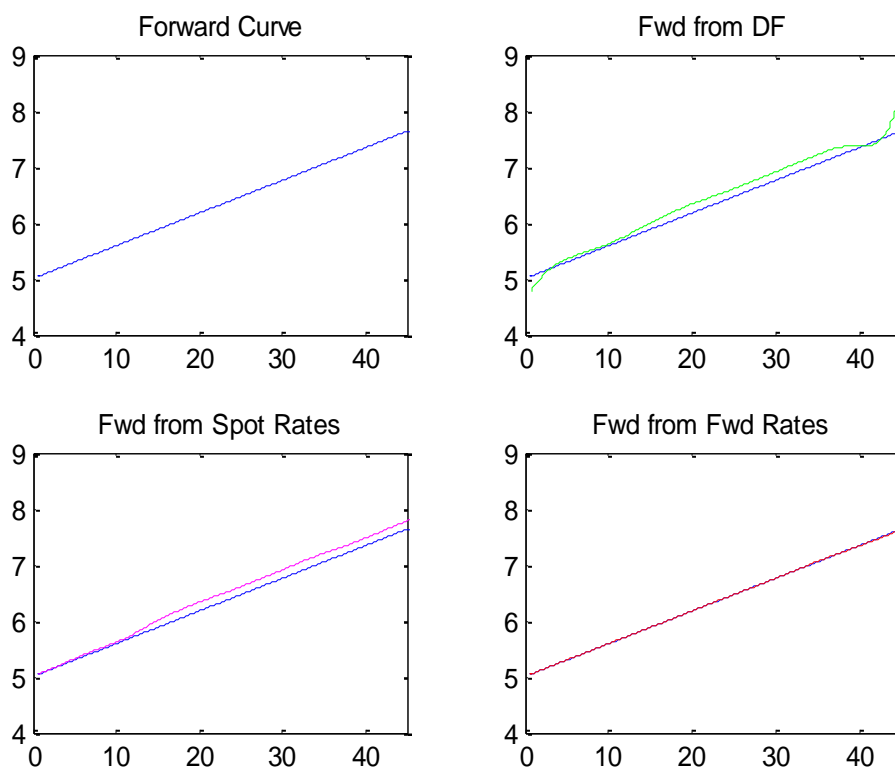Fig. 18 – True and estimated forward curves – functional form F1

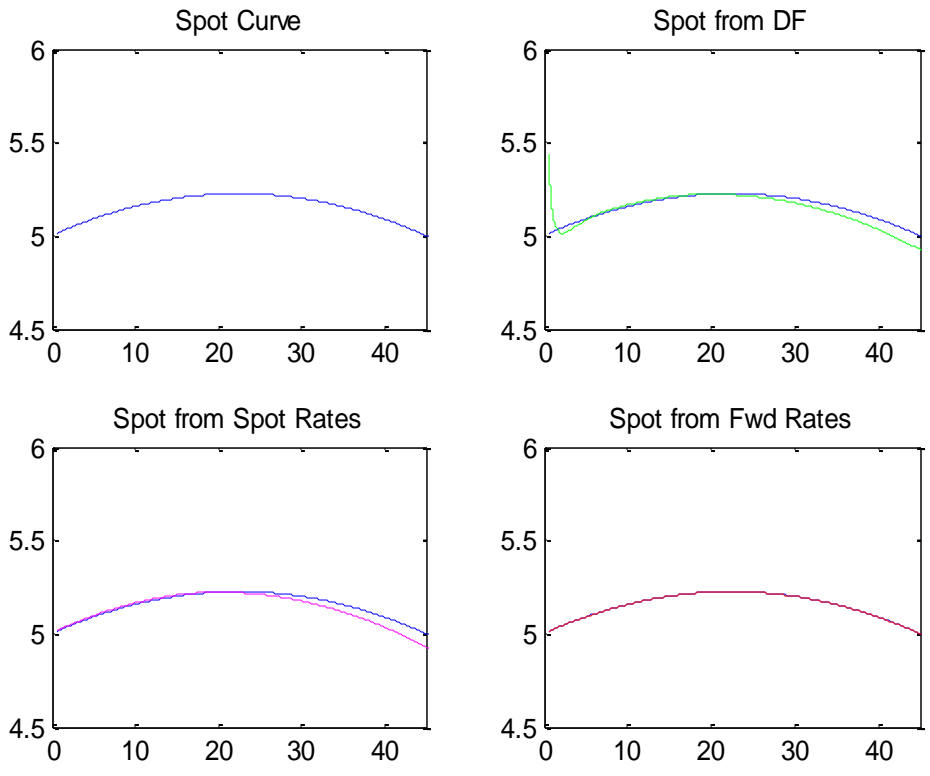Fig. 19 – True and estimated forward curves – functional form F1

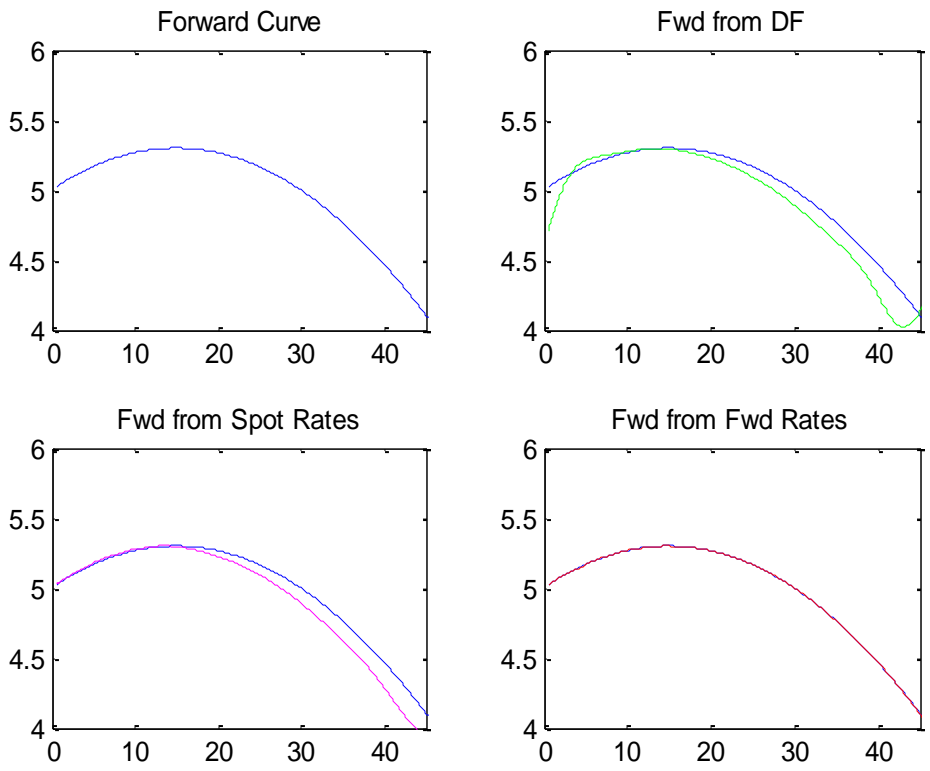Fig. 20 – True and estimated spot curves – functional form F1

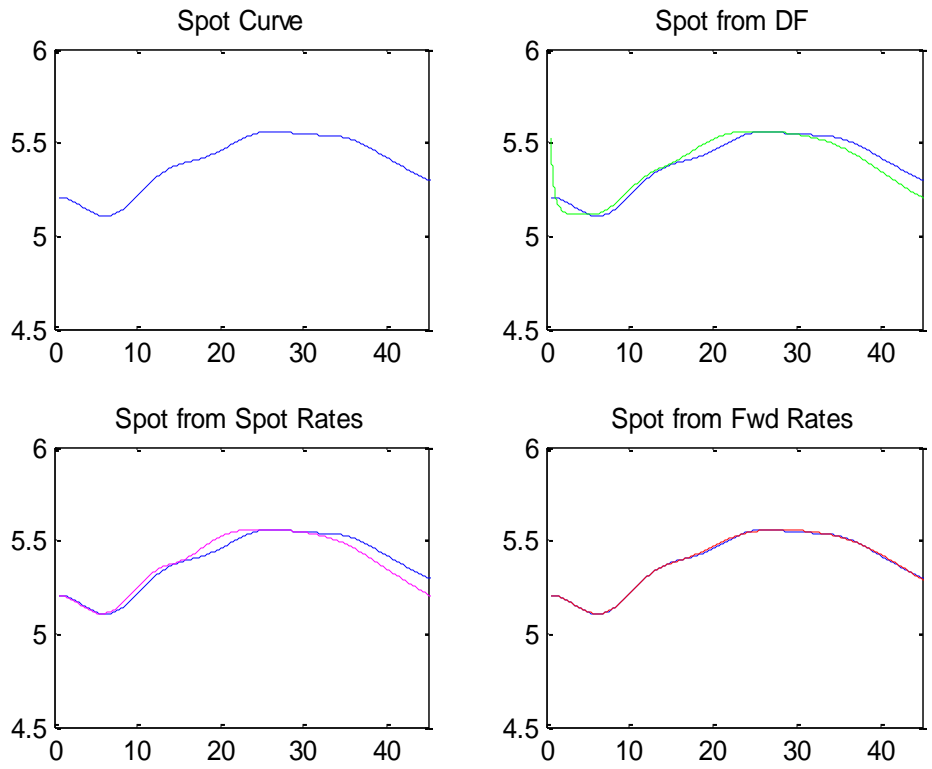Fig. 21 – True and estimated forward curves – functional form F2

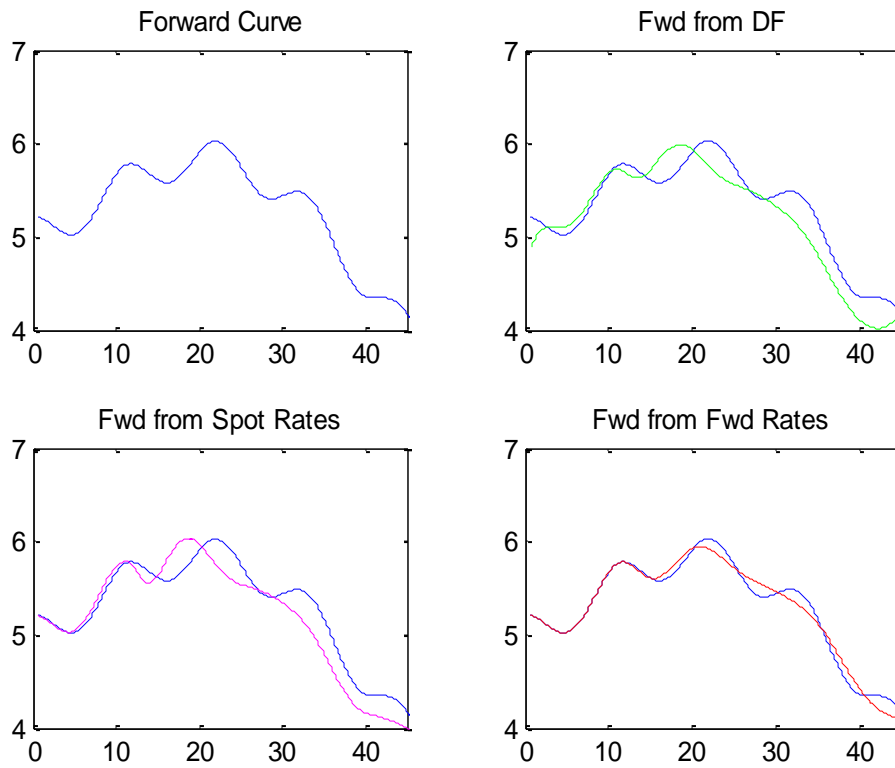Fig. 22 – True and estimated spot curves – functional form F3

Fig. 23 – True and estimated forward curves – functional form F3

Recommendations

1. The SPLINE function should be changed to enable it to handle situations where there is inconsistency in overlapped cash-futures grid rates.

2. Consideration should be given to developing a methodology to automatically adjust grid-point weights, taking account of liquidity in the futures strip.

3. Analysts should be discouraged from using convexity adjustment factors to manipulate outlier futures prices. Convexity adjustment factors should be derived from a suitable interest rate model, such as Hull-White. Outliers should be handled by the weight-adjustment procedure.

4. Consideration should be given to incorporating FRA data in the construction of US$ and CAN$ curves.

5. Curves should be fitted to forward rate grid data rather than to spot rate or discount factor grid data.

6. The GCV B-Spline approach should be considered for implementation in a future release.

Acknowledgements

References

[1] Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vettening, W. T., "Numerical Recipes in C++", Cambridge University Press, 2002.

[2] Sankar, L., "OFUTS – An Alternative Yield Curve Interpolator", F.A.S.T. research documentation, Bear Stearns, June 1997.

[3] Ezra Nahum, "Changes to Yield Curve Construction, Linear Stripping of Short End of the curve", F.A.S.T. research documentation, Bear Stearns, July 2004.

[4] "Dual Dense Yield Curve Methodology", F.A.S.T. research documentation, Bear Stearns.

[5] Adams, K.J., and Van Deventer, D.R., "Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness", Journal of Fixed Income, June (1994): 52-62

[6] Andersen, L., "Discount Curve Construction with Tension Splines", Banc of America Securities (2005)

[7] De Boor, C., "A Practical Guide to Splines", Springer-Verlag, (1978)

[8] Fisher, M., Nychka, D., and Zervos, D., "Fitting the Term Structure of Interest rates with Smoothing Splines", (1994)

[9] McCulloch, J.H., "Measuring the Term Structure of Interest Rates", Journal of Business, 44 (1971): 19-31

[10] Wahba, G., "Spline Models for Observational Data", SIAM; Philadelphia (1990)