SERIOUS ZERO-CURVES

Roger J-B Wets

Department of Mathematics
University of California, Davis
rjbwets@ucdavis.edu

Stephen Bianchi

EpiSolutions Inc
El Cerrito, CA
sbianchi@episolutionsinc.com

Liming Yang

KMV Corporation
San Francisco, CA
liming.yang@kmv.com

Abstract. All valuations (discounted cash flow, instrument pricing, option pricing) and other financial calculations require an estimate of the evolution of the risk-free rates as implied by the term structure. This presumes that one has, if not complete knowledge, at least a very good estimate of the term structure, for the so-called zero-curves (spot and forward rate curves, discount factor curve, etc.). This paper is concerned with the methodology of deriving these zero-curves. It reviews the methodology and the limitations of standard BootStrapping and proposes a quite different approach based on Approximation Theory that has been implemented at EpiSolutions Inc.

Keywords: spot and forward rate curve, discount factors curve, term structures, best fit, max-error criterion, error-sum criterion,

Date: January 21, 2002

1 An example

We deal with the following simple, but fundamental, issue: Find the zero-curves associated with a given portfolio. The term *zero-curve* is used here in a generic sense to designate any one of the following financial curves, *spot rates, forward rates, discount factors and discount rates*, since any one of these curves completely determines the others.

Zero-curves are the corner stone to practically all valuations of financial instruments. So, it may come as a surprise that the zero-curves associated with a well-defined portfolio, derived using different methodologies, might vary significantly? This can best be illustrated by an example: On their web-site, TechHackers describes a portfolio that includes Eurodollar Deposits, Eurodollar Futures and Swaps from June 10, 1997. The four pairs of Spot and Forward Rate curves in Figure 1 are derived using three different implementations of the BootStrapping approach and a novel approach, *EpiCurves*, laid out in this paper and implemented by EpiSolutions Inc. (`www.episolutionsinc.com`)

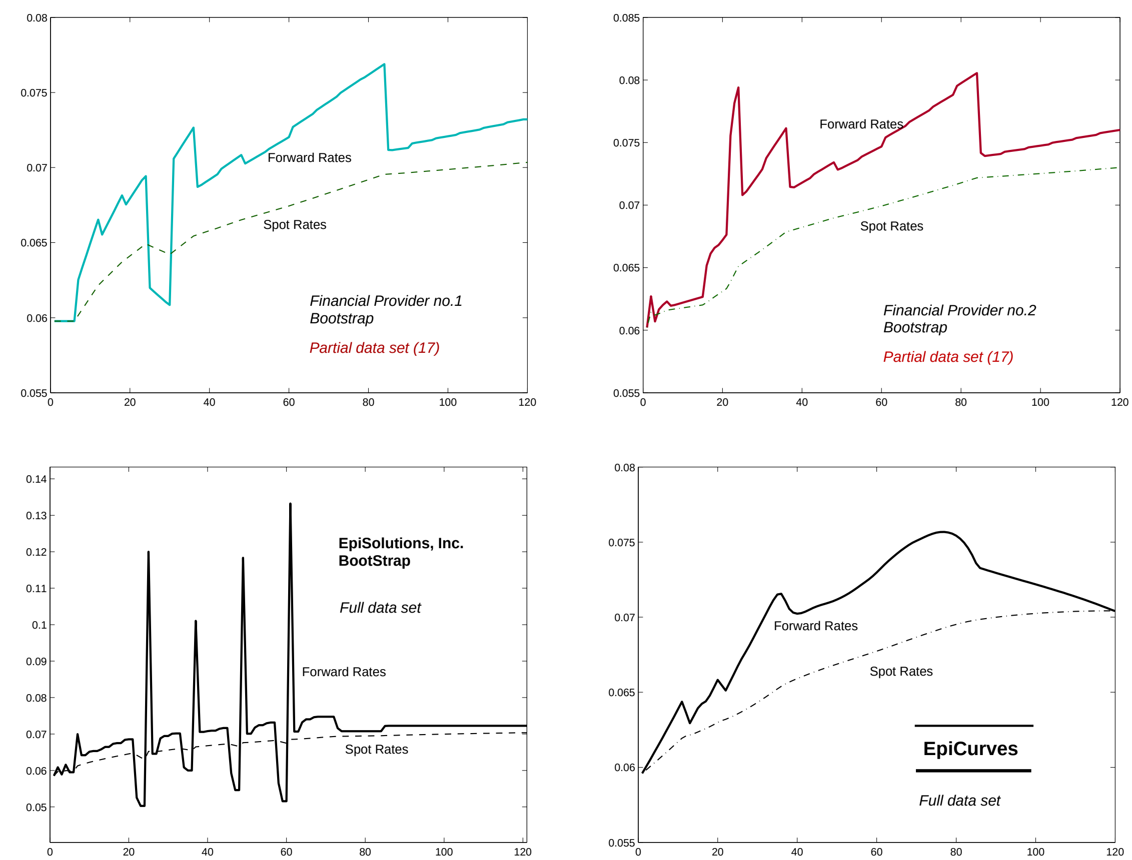

Figure 1: Spot and Forward Rate curves: Four different functionalities

At this point one might wonder if one should have any confidence in any one of these pairs! There are even noticeable differences between the Spot Rate curves. To understand the underlying reasons for these differences, one needs to examine the hypotheses under which these zero-curves were obtained. To do so, let's begin with an overview of the BootStrapping methodology.

2 BootStrapping

The valuation of fixed-income securities, and derivatives written on fixed-income securities, requires an estimation of the underlying risk-free term structure of interest rates. In principle, the term structure of interest rates is defined by a collection of zero-coupon bond prices (and their respective maturities), spanning the horizon over which a fixed-income security is to be valued. However, unless a zero-coupon bond exists for *every* maturity for which a discount factor is desired, some form of estimation will be required to produce a discount factor for any 'off-maturity' time. In practice, zero-coupon bond prices are available for a limited number of maturities (typically ≤ 1 year). If zero-coupon bonds for other maturities are available, a lack of liquidity may prevent the determination of an accurate or reliable price. As a result, the zero-curve is typically 'built' from a combination of liquid securities, both zero-coupon and coupon-bearing, for which prices are readily available. This can include Treasury Bills, Deposits, Futures, Forward-Rate Agreements, Swaps, Treasury Notes and Treasury Bonds. The list need not be limited to the securities just noted, though current vendor solutions *are* limited to these securities. Given a spanning set of securities, the zero-curve is then built using one of two forms of *BootStrapping*.

Under one BootStrapping method, the first step is to construct a larger set of spanning securities, by creating an 'artificial' security that matures on every date for which a cash flow is expected, and on which no security in the original set matures. For example, given a 5-year and a 6-year swap (paying semi-annually), a 5.5-year swap would be constructed, with a fixed-rate somewhere in between (based on some sort of interpolation) the fixed-rates for the 5-year and the 6-year swaps. Then, 'standard' BootStrapping may be applied to the expanded set of securities, giving discount factors for each maturity and cash flow date in the security set.

Another BootStrapping approach, is to make an assumption about how the instantaneous (or periodic) forward rates evolve between maturities in the security set. One assumption might be that forward rates stay constant between maturities, another might be that they increase or decrease in a linear fashion. Whatever the form of the forward rate evolution, *some assumption must be made.* Under this approach, instead of solving for a single discount factor for each successive security, a forward rate (or a parameter governing forward rate evolution) is solved for that will give the appropriate discount factor(s) between two maturity dates. For the example of the 5-year and 6-year swaps, given that the discount factors through the 5-year maturity have already been calculated, a forward rate is determined for the period between 5 and 6 years, which gives a 5.5-year discount factor and a 6-year discount factor that (when combined with the previous discount factors) will value the fixed side of the 6-year swap at par.

The results of either approach 'look' somewhat similar, a set of discount factors and corresponding dates spanning the horizon from today to the last maturity date in the security set. However, when using this set of discount factors as a basis for the valuation of other fixed-income securities, it will rarely be true that the cash flows of these securities will fall directly on the discount factor dates of the newly created zero-curve. In this case, a discount factor or a zero-coupon rate must be interpolated from the zero-curve. Typically available interpolation methodologies for this include linear, log-linear, exponential, cubic-spline, or any of a number of variations on fitting the zero-curve with a polynomial.

As long as two securities do not share a maturity date, any combination of securities may theoretically be used in constructing a zero-curve, even with the BootStrapping methodology. A limitation of currently available technology, is that the user must typically 'switch' from one security type to another during the BootStrapping process. For example, given a set of deposits, futures, and swaps, the currently available methods will not allow for inclusion of a deposit and a futures contract, where the underlying deposit maturity date of the futures contract is *prior* to the maturity date of the deposit (similarly, no 'overlap' is allowed between futures contracts and swaps). In practice, this may not be a severe limitation, as users may very well wish to describe different 'sections' of the zero-curve using certain types of securities. Almost all commercially available zero-curve construction technologies rely on a form of *BootStrapping*, in combination with a variety of interpolation methods. Furthermore, they are limited to using certain types of securities, and are also limited in the ways in which these securities may be combined.

DFS-Portfolio (Deposits, Futures and Swaps): Let's now return to the example found on the TechHackers web-site. A detailed description of the portfolio is given in Tables 1–3 in the Appendix, but at this point at brief description will suffice. There are 36 instruments in the full data set, broken down as follows: 6 Eurodollar Deposits with term-to-maturities ranging from overnight to 12 months, 24 Eurodollar Futures with 90-day deposit maturities ranging from 3 months to 6 years, and 6 Swaps with term-to-maturities ranging from 2 to 10 years. The yields of the deposits vary from 5.55% to 6.27% with the higher yields corresponding to those with larger maturity. Similarly, the 24 futures have yields that vary from 5.89% to 7.27%. And the yields for the swaps vary from 6.20% to 6.86%, again with the yields increasing as maturities get larger.

BootStrapping Sub-portfolio. In this example, there is overlap between the maturity dates for the EuroDollars Deposits (1 day, 1 month, 2 m., 3 m., 6 m., 12 m.) and the EuroDollars Futures (3 months, 6+ m., 9 m., 12 m., 15 m., . . .). Thus *not all instruments can be included* in a 'BootStrapping sub-portfolio.' One possible choice, the default setting for one of the Financial Providers, is to switch from Eurodollar Deposits to Eurodollar Futures at the earliest possible time (which simply means using the first available Futures contract). The same idea is employed when switching from Eurodollar Futures to Swaps – the first available Swap is used. Of course the user may choose any time they wish for making these switches, the point is that they *must* be made somewhere along the line. The data set used to generate the 'BootStrapping portfolio' is a subset of the full DFS-Portfolio data set with switches set to include the first 5 Eurodollar Deposits, the 2nd through the 7th Eurodollar Futures contracts and all 6 Swaps – for a total of 17 instruments.

BootStrapping Results. The first two pairs of Spot and Forward Rate curves in Figure 1 are those derived from functionalities made available by two Financial Providers. Although all the Spot Rate curves appear to be relatively similar (except for the time span from month 20 to month 25), the Forward Rate curves are quite dissimilar. Of course, this can be traced back to the different BootStrapping implementations that rely on different strategies described earlier.

The implementation of the BootStrapping technique at EpiSolutions Inc. –based on the simple precept that the instantaneous forward rates are constant between (adjacent) maturity dates— for the DFS-Portfolio yields the pair of zero-curves in Figure 2.

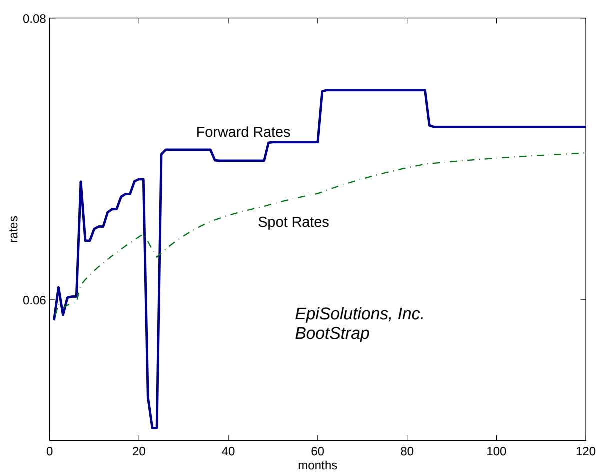

Figure 2: EpiSolutions Inc.-BootStrapping for BootStrapping Portfolio

3 Guiding principles

These 'variations' between these zero-curves are bound to be somewhat disturbing, but more unsettling is the fact that only 17 of the 36 instruments were used in deriving them. As already indicated earlier, the standard BootStrapping methodology shackles us with a portfolio with 'no overlapping maturities'. A sub-portfolio must be selected, and the choice of instruments to include in this sub-portfolio is not necessarily unique; in the DFS-Portfolio, it's not unique. It's not difficult to create a hypothetical example of a portfolio where the 'not-included' instruments would yield dramatically different zero-curves than those generated by the selected BootStrapping portfolio!

From these observations, it becomes self-evident that the only acceptable methods that will generate 'serious' zero-curves must be based on the *full portfolio*. That's going to be our *first guiding principle*.

Because of the 'non-overlapping condition,' the BootStrapping implementations of Financial Providers no.1 and no.2 can't deal with the full data set. The implementation of the BootStrapping by EpiSolutions Inc. does allow us to deal with the full collection of instruments included in the DFS-Portfolio. The resulting Spot and Forward Rate curves are those found in Figure 1 (lower left hand corner). Unfortunately, although the Spot Rate curve might appear reasonable, the Forward Rates curve come with spikes that go from about 5% to nearly 14% almost instanteneously! This is certainly not something that we expect in practice. What we expect to see is an evolution of the rates that doesn't come with such abrupt changes. In other words, a *certain level of smoothness* should be required from zero-curves. That's going to be our *second guiding principle*.

The methodology of *EpiCurves*, to be detailed in the following sections, is guided by these two principles (complete portfolio, smooth curves). If we run *EpiCurves* on the full portfolio of the DFS-Portfolio, the resulting Spot and Forward Rate curves are in the lower right hand corner of Figure 1.

Figure 3 allows us to make a comparison between the Spot and the Forward Rate curves obtained via BootStrapping and EpiCurves, based on the *given* portfolio, not a sub-portfolio chosen to accommodate the methodology. And because of the scaling, it's now possible to see that the Sport Rate curve generated by the BootStrapping method isn't really acceptable, either. It's also useful to compare these results with those coming from the BootStrapping method applied to the restricted 17-instruments BootStrapping sub-portfolio.

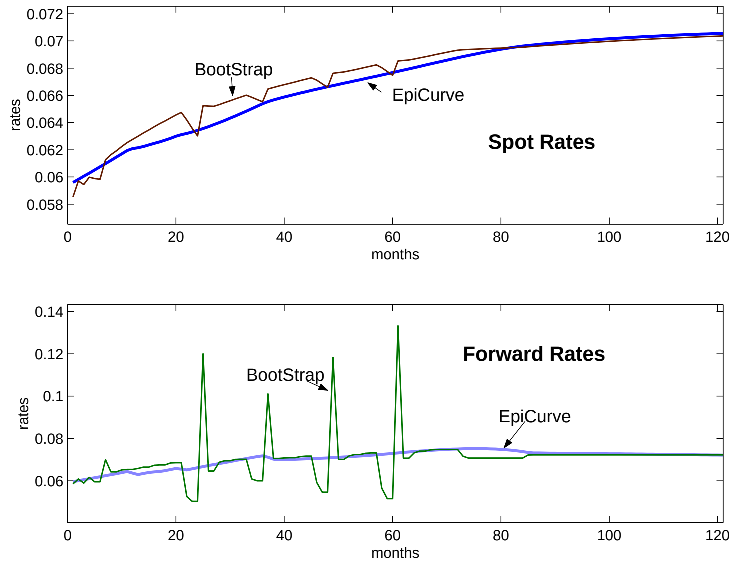

Figure 3: Comparison of Spot and Forward Rate curves.

To render the problem more manageable via BootStrapping, let's restrict the DFS-Portfolio to Deposits and Futures only, i.e., without the Swaps. Because of the overlap of maturity dates, we have to use the EpiSolutions Inc. BootStrapping to make any comparison with the *EpiCurves* results.

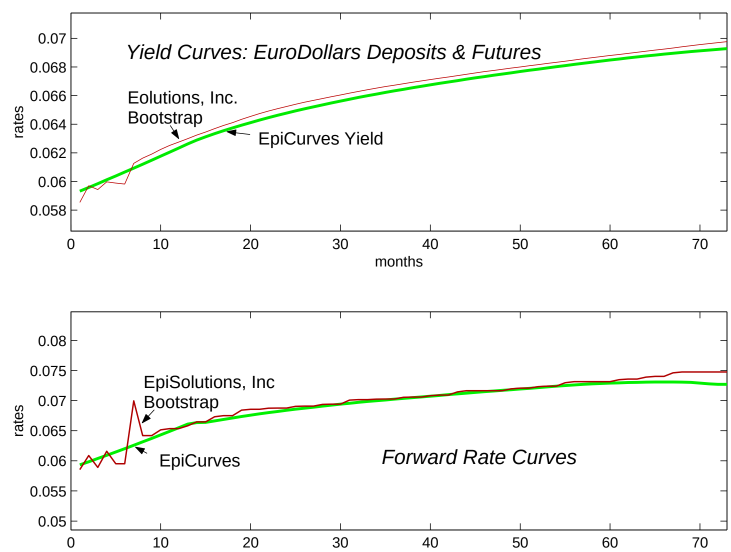

Figure 4: Comparison of Spot and Forward Rate curves: Deposits & Futures.

4 Smoothness

The definition of smooth for a curve might almost be subjective, but it's certainly application dependent. In the world of zero-curves, one is certainly not going to be satisfied with a 'smooth' Spot Rate curve, or a 'smooth' discount factors curve, when an associated zero-curve, for example the Forward Rate curve, is seesawing; that this can actually occur was already pointed out in [4]. In fact, it suffices return to Figure 1 and look at the Forward Rate curves generated by BootStrapping, even those derived from the 17-instruments sub-portfolio.

In [1], Adams and Van Deventer rely on a criterion used in engineering applications, cf. [3], in their derivations of Spot and Forward Rate curves with 'maximum smoothness.' They propose finding a forward rates curve, fw, such that for each instrument in the portfolio:

$$P_i = \exp\left(- \int_0^{t_i} fw(s)\, ds \right), \quad i = 1, \ldots, I,$$

where P_i is the (today's) price and t_i the maturity date of instrument i, and

$$\int_0^T fw''(s)\, ds \qquad \text{is minimized,}$$

$[0, T]$ is the time span in which we are interested; usually T is the largest maturity of the instruments in the portfolio. It's shown that the solution is a 4th order spline, certainly smooth, whose coefficients can be easily computed; 'maximum' smoothness is achieved in terms of a criterion attributed to Vasicek.

The Achilles' heel of this approach, at least as laid out in [1], is that the only instruments that can be included in the 'maximum smoothness' portfolio are zero-coupon bonds, and zero-coupon bonds with maturities exceeding one year are extremely rare. In order to obtain zero-curves that span more than a few months, one possibility is to fabricate artificial (long term) zero-coupon bonds that have similar financial characteristics to those instruments found in the portfolio; this requires interpolations of some type. Moreover, prices P_i are present day prices and so no future contracts can be included in the 'maximum smoothness' portfolio. Presumably, this can also be skirted by some adjustments. In the final analysis, like for 'standard' BootStrapping, we have to create a sub-portfolio and then enrich it by artificial instruments in order to be able to apply the suggested method.

From our previous examples and analysis, it's clear that if one is going to derive zero-curves by taking into account more than just a few well chosen instruments, and one is going to aim at an acceptable level of smoothness, there is going to be a 'price' to pay for this. In Function Theory, the smoothness of a curve is identified in terms of the number of times it's continuously differentiable. A curve $z : [0,T] \to I\!R$ is said to be of class $\mathcal{C}^q$ if its qth derivative is continuous. So, if z is of class $\mathcal{C}^2$ it means that it can be differentiated twice and the second derivative is continuous. If it's of class $\mathcal{C}^0$ then z is just continuous, and if it's of class $\mathcal{C}^\infty$ then all derivatives, of any order, exist and are continuous. It's evident that $\mathcal{C}^0 \supset \mathcal{C}^1 \supset \cdots \supset \mathcal{C}^\infty$. In those terms, one might wish the zero-curves to be of class $\mathcal{C}^\infty$, but it's evident that this is a much smaller family of curves than those that are just continuous, or just continuously differentiable ($\mathcal{C}^1$). And consequently, by insisting that our zero-curve be of class $\mathcal{C}^\infty$, we might very well have excluded those curves that have the 'accuracy' properties we are looking for. Consequently, usually, we have to be content with smooth curves that are less than infinitely smooth.

5 EpiCurves

EpiCurves come from a large, but specific, sub-family of curves that are of class $\mathcal{C}^q$ for some $q = 1, 2, \ldots$. To simplify the presentation, let's suppose that we are interested in finding a $\mathcal{C}^2$-curve. Every $\mathcal{C}^2$-curve on $[0,T]$ can be written as

$$z(t) = z_0 + v_0 t + \frac{1}{2} a_0 t^2 + \int_0^t \int_0^\tau x(s)\, ds\, d\tau, \quad t \in [0,T],$$

where

- $x : (0,T) \to I\!R$ is an arbitrary piecewise continuous function that corresponds to the 3rd derivative of z;
- a_0, v_0, z_0 are constants that can be viewed as integration constants

Once the function x (3rd derivative) and the constants a_0, v_0, z_0 have been chosen, the function z is completely determined.

Now, let's go one step further. Instead of allowing for any choice for x, let's restrict the choice of x to piecewise constant functions of the following type: split $[0,T]$ into N sub-intervals of length $1/N$ and let the function x be constant on each one of these intervals, and defined as follows: for $k = 1, \ldots, N$,

$$x(t) = x_k, \qquad \text{when} \qquad t \in (t_{k-1}, t_k],$$

where $t_0, t_1, \ldots, t_L$ are the end points of the N sub-intervals. The corresponding curve z on $[0,T]$ is completely determined by the choice of

$$a_0, v_0, z_0 \quad \text{and} \quad x_1, x_2, \ldots, x_N,$$

i.e., by the choice of a finite number of parameters, exactly $N+3$.

Namely, for $k=1,\dots,N$, $t\in(t_{k-1},t_k]$ with $\delta=1/N$ and $\tau=t-t_{k-1}$,

$$\begin{aligned}
z(t) &= z_0 + v_0 t + \frac{1}{2}a_0 t^2 + \frac{\tau^3}{6}x_k \\
&\qquad + \sum_{j=1}^{k-1} x_j \left[\frac{\delta^3}{2}\Big(\frac{1}{3} + (k-j-1)(k-j)\Big) + \delta^2\tau(k-j-0.5) + \frac{\delta\tau^2}{2} \right], \\
z'(t) &= v_0 + a_0 t + \sum_{j=1}^{k-1} x_j \Big(\delta^2(k-j-0.5) + \delta\tau\Big) + \frac{\tau^2}{2}x_k, \\
z''(t) &= a_0 + \delta\sum_{j=1}^{k-1} x_j + \tau x_k, \\
z'''(t) &= x_k.
\end{aligned}$$

By restricting the choice of x to piecewise constant functions, the resulting z-curves are restricted to those curves in $\mathcal{C}^2$ that have (continuous) piecewise linear second derivatives. Let's designate this family of curves by $\mathcal{C}^{2,pl}$ where *pl* stands for piecewise linear; whenever appropriate we will use the more complete designation $\mathcal{C}^{2,pl}([0.T],N)$ with $[0,T]$ the range on which these curves are defined and N the number of pieces, but usually the context will make it evident on which interval these curves are defined. Clearly, not all $\mathcal{C}^2$-curves are of this type. However, Approximation Theory for functions, tells us that any $\mathcal{C}^2$-curve can be approximated *arbitrarily closely* by one whose second derivative is a continuous piecewise linear function, i.e., a curve in $\mathcal{C}^{2,pl}([0,T],N)$, by letting $N\to\infty$. This provides us with the justification one needs to restrict the search for 'serious' zero-curves to those in this particular sub-family of $\mathcal{C}^2$-curves.

In summary, the building of *EpiCurves* starts by selecting the level of smoothness desired ($z\in\mathcal{C}^q$), and then a zero-curve is built whose qth derivative is a continuous piecewise linear function. This requires fixing a finite number of parameters; actually $N+q+1$ parameters. If the resulting curve doesn't meet certain accuracy criteria, the step size $(1/N)$ is decreased by letting $N\to\infty$.

6 Zero-curves from spot rates

To set the stage for finding the zero-curves associated with a collection of instruments generating cash flow streams, let's consider an *EpiCurves* approach to fitting spot rates to obtain a Spot Rate curve. The data come in a pair of arrays,

$$s=(s_1,s_2,\dots,s_L), \qquad m=(m_1,m_2,\dots,m_L),$$

that give us the spot rates for a collection of instruments of different maturities, for example, Treasury Notes. The task is to find a Spot Rate curve that 'fits' these data points. That's easy enough. Assuming that $m_1 < m_2 < \cdots < m_L$, one could simply derive a Spot Rate curve by linear interpolation between adjacent pairs. That's actually a perfect fit. Generally, this is not a 'smooth' curve. And, in turn, this will usually generate a forward rates curves that can be quite jagged. So this 'simple' solution almost never produces zero-curves that practitioners would consider acceptable. Of course, one can use another interpolation method, such as via splines, that generates significantly better results. Another possibility is to set-up an artificial portfolio with coupon-bonds whose yields would match the given spot rates. The problem is then reduced to one of finding the zero-curves associated with the cash flow stream of this (artificial) portfolio. This is dealt with in the next section. But, this latter approach, found in the packages of some financial technology providers, circles around the problem, at least one too many times, before dealing with it.

The use of the *EpiCurves* technology provides an elegant solution that generates with smooth zero-curves. The strategy is to find a Spot Rate curve of the type described in the previous section, say again a $\mathcal{C}^{2,pl}$-curve that will match the given spot rates. One must accept the possibility that there we won't be able to find, for a fixed N, a $\mathcal{C}^{2,pl}([0,T],N)$-curve that fits perfectly the given data. So, the problem becomes one of finding the 'best' possible fit. Best possible can be defined in a variety of ways but it always comes down to minimizing the 'error', i.e., the distance between the *EpiCurves* result and the given spot rates. Mathematically, the problem can be formulated as follows:

$$\text{find} \quad z \in \mathcal{C}^{2,pl}([0,T],N) \text{ so that } \|s - z(m_1 : m_L)\|_p \quad \text{is minimized} \;,$$

where $z(m_1 : m_L) = (z(m_1), z(m_2), \ldots, z(m_L))$ and $\|a\|_p$ is the ℓ^p-norm of the vector a. With $p = 1$, one would be minimizing the sum of the (absolute) errors, with $p = 2$ one minimize the sum of the squares of the errors, and with $p = \infty$, it would be the maximum (absolute) error that would be minimized. The present implementation by EpiSolutions Inc., has $p = 1$ and thus minimizes the sum of the errors, since

$$\|s - z(m_1 : m_L)\|_1 = \sum_{l=1}^{L} |s_l - z(m_l)|.$$

The resulting optimization problem can then be reduced to a linear programming problem, since, as explained in the previous section, the functions z in $\mathcal{C}^{2,pl}$ are completely determined by a finite number of parameters.

To illustrate the results, let's apply both linear interpolation and the *EpiCurves* technology to obtain a Spot Rate curve that fits the spot rates (for T-bills and Treasury notes) of October 1982:

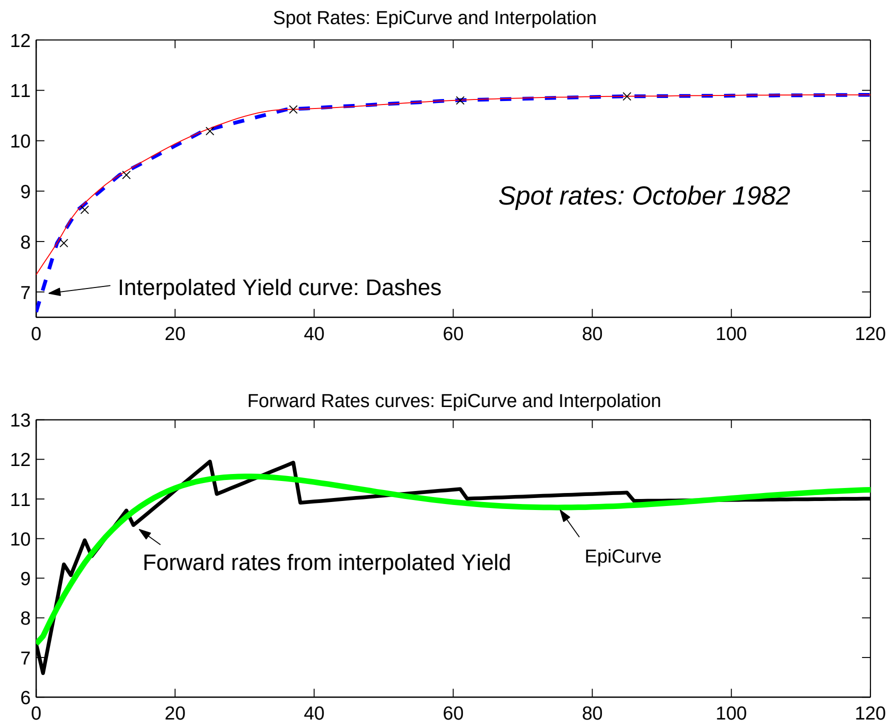

Figure 5: Linear interpolation vs. EpiCurves.

$$m = (3, 6, 12, 24, 36, 60, 84, 120, 240, 360),$$
$$s = (7.97, 8.63, 9.32, 10.19, 10.62, 10.8, 10.88, 10.91, 10.97, 11.17);$$

the time unit is 1 month. The Spot and Forward Rate curves can be found in Figure 5. It's barely possible to see the difference between the Spot Rate curves, but the difference between the Forward rates curve is more than noticeable. The difference, of course, can be traced back to the intrinsic smoothness of the Spot Rate curve when it's generated from *EpiCurves*.

Let also consider the case spot rates for January 1982, the maturities-array m is the same, but now

$$s = (12.92, 13.90, 14.32, 14.57, 14.64, 14.65, 14.67, 14.59, 14.57, 14.22).$$

Running *EpiCurves* yields the result in Figure 6. Note that, the Forward Rates curve is rather unsettled up to the end of year 1, it actually reflects almost perfectly the 'unsettled' market situation at that time (begin 1982).

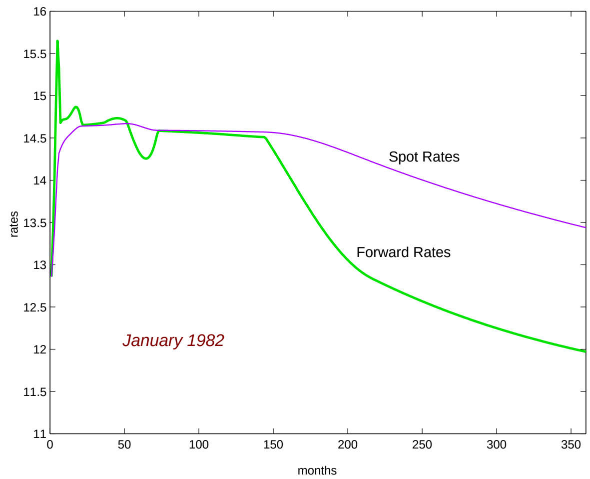

Figure 6: Spot and Forward Rate curves from spot rates.

7 Zero-curves from cash flow streams

Let's review briefly our goals and guiding principles. Given the increased complexity of the instruments being traded, it certainly is no longer sufficient to be able to build zero-curves based on just zero-coupon bonds. It should be possible to build the Spot Rate curve associated with *any* collection of instruments, for example, AAA- or AA-rated corporate bonds, any mixture of swaps, futures and bonds, etc. Notwithstanding a relatively large literature devoted to zero-curves, cf. Buono, Gregory-Allen and Yaari [2], there has never been any serious attempt at dealing with the building of zero-curves at this more comprehensive level.

Of course, given an arbitrary collection of instruments, each one generating its own cash flow stream, it might be possible* to generate, via BootStrapping for example, any one of the zero-curves. However, as every practitioner knows all too well and as was reviewed in §2, some of the resulting curves will be, to say the least, unwieldy, and have every characteristic except 'believable.' The insistence on 'smoothness', cf. Vasicek and Fong [5], Shea [4], Adams and Van Deventer [1], is motivated by the strongly held belief, that's also supported by historical data, that zero-curves don't come with kinks, and spikes i.e., extremely abrupt changes in the rates.

Keeping this in mind, the problem of generating zero-curves could be roughly formulated as follows: Given a collection of instruments, each one generating a given cash flow stream, find smooth zero-curves so that for each instrument (in the collection), the *net present value (NPV)* of the associated cash flow matches its present price.

*assuming that maturities occur at different dates

Although this formulation allows us to include zero-coupon bonds, coupon bonds, swaps, etc., in our collection of instruments, it doesn't allow for futures, future swaps, etc. To do so, let's reformulate the problem in the following more general terms: With each instrument i in our collection, we associate a

$$\text{Time Array:} \quad (t_{i1}, t_{i2}, \dots, t_{i,L_i})$$

the dates, or time span, at which cash payments will occur, and a

$$\text{Payments Array:} \quad (p_{i1}, p_{i2}, \dots, p_{i,L_i})$$

with cash flow p_{il} received at time t_{il}, L_i is the maturity date. One then interprets $p_{il} > 0$ as cash received and $p_{il} < 0$ as cash disbursed. For example, in the case of a coupon bond, bought today for \$100, with semi-annual \$3 coupons and a two year-maturity, one would have:

$$\text{Time Array:} \quad (0, 6, 12, 18, 24),$$

assuming the time unit is '1 month', and

$$\text{Payments Array:} \quad (-100, 3, 3, 3, 103).$$

This allows us to include almost any conceivable instrument in our collection, as long as it comes with an explicit cash flow stream. For example, in the case of the following T-bill forward: A bank will deliver in 3 months from now, a 6-month Treasury bill of face value \$100, and 10% the annual forward rate for that 6 months period. The value of such a contract would be \$95.24 that would have to be paid in 3 months. This contract would then come with the following arrays:

$$\text{Time Array:} \quad (3, 9), \qquad \text{Payments Array:} \quad (-95.24, 100).$$

Now, in this frame of reference, the zero-curve problem could be formulated in the following terms: Given a (finite) collection of instruments that generate cash flow streams, find a *discount factor curve* such that

- the net present value (NPV) of each individual instrument (contract) turns out to be 0 when all cash payments received and all disbursements are accounted for;
- all associated zero-curves (forward, spot, discount rates) are 'smooth'.

When formulated at this level of generality, the zero-curve problem is usually not feasible. In fact, it's not difficult to fabricate an 'infeasible' problem. Simply, let the collection consist of two one-coupon bonds that have the same nominal value, the same maturity and the same price (today). Both coupons are to be collected at maturity but have different face value. Clearly, there is no discount factor curve so that the net present value (NPV) of both of these cash flows turns out to be 0! Of course, this is an unrealistic example, the financial markets wouldn't have assigned the same price to these two instruments; arbitrage would be a distinct possibility in such a situation. But since we allow for *any* collection of instruments, there is the distinct possibility that there are practical instances when one can't find a 'smooth' discount factors curve so that the NPV of all cash flow streams factors out to 0. So, given that we want to be able to deal with any eclectic collection of instruments, as well as the 'standard' ones, instead of asking for the NPV of all cash flow streams to be 0, we are going to ask that they be as close to 0 as possible.

Smoothness is going to be achieved by restricting the choice of the discount factors curve to $\mathcal{C}^{q,pl}$, i.e., curves whose qth derivative is continuous and piecewise linear; cf. §5 for a more detailed description of such curves. To render our presentation more concrete, and easier to follow, we are going to proceed with $q = 2$.

The problem is now well defined mathematically:

$$\text{find a discount factors curve:} \quad df \in \mathcal{C}^{2,pl}([0,T],N) \text{ so that } \|v\|_p \quad \text{is minimized} \ .$$

where $\|v\|$ is the ℓ^p-norm of v,

$$v = (v_1, v_2, \dots, v_I), \qquad v_i = \sum_{l=1}^{L_i} df(t_{il}) p_{il};$$

v_i is the net present value of instrument 'i' given that the cash flow is discounted using the discount factors $df(t_{il})$. The EpiSolutions Inc., implementation relies on the ℓ^∞-norm,

$$\|v\|_\infty = \max\left[\,|v_1|, |v_2|, \dots, |v_I|\,\right],$$

so let's proceed with this criterion but it should be noted that one can choose any $p \in [1, \infty)$ that might better represent the decision maker's preferences or concerns. In fact, except for extremely unusual portfolio, the differences between the solutions should be insignificant.

Since df belongs to $\mathcal{C}^{2,pl}([0,T],N$, it's of the form: for $k = 1, \dots, N$, $\delta = 1/N$, $t \in (\delta(k-1), \delta k]$ and $\tau = t - \delta(k-1)$

$$\begin{aligned} df(t) = 1 + v_0 t + \frac{1}{2} a_0 t^2 + \frac{\tau^3}{6} x_k \\ + \sum_{j=1}^{k-1} x_j \left[\frac{\delta^3}{2}\Big(\frac{1}{3} + (k-j-1)(k-j)\Big) + \delta^2 \tau (k-j-0.5) + \frac{\delta \tau^2}{2} \right]; \end{aligned}$$

where $a_0, v_0, x_1, x_2, \ldots, x_N$ are parameters to be determined; note that the discount factor at time $t = 0$ is 1, so we can fix this 'constant' (z_0). But simply being of this form doesn't make *df* a discount factors curve. We already have that $df(0) = 1$, we need to add two conditions:

- *df* should remain non-negative, thus we have to introduce the constraints: $df(t) \geq 0$ for all $t \in [0, T]$;
- *df* should be decreasing, at least non-increasing, this means that $df'(t) \leq 0$ for all $t \in [0, T]$, a condition that translates into the constraints:

$$df'(t) = v_0 + a_0 t + \sum_{j=1}^{k-1} x_j \left(\delta^2 (k - j - 0.5) + \delta\tau \right) + \frac{\tau^2}{2} x_k \leq 0, \quad \forall\, t \in (0, T].$$

Putting this all together with *df* as defined above, we end up with the following optimization problem:

$$\begin{aligned} \min\ & \theta \\ \text{so that}\ & \theta \geq \sum_{l=1}^{L_i} df(t_{il}) p_{il}, \qquad i = 1, \ldots, I, \\ & \theta \geq -\sum_{l=1}^{L_i} df(t_{il}) p_{il}, \qquad i = 1, \ldots, I, \\ & df(t) \geq 0, \qquad t \in [0, T], \\ & v_0 + a_0 t + \sum_{j=1}^{k-1} x_j \left(\delta^2 (k - j - 0.5) + \delta\tau \right) + \frac{\tau^2}{2} x_k \leq 0, \quad t \in (0, T], \\ & v_0 \leq 0,\ a_0 \geq 0,\ x_k \in I\!R, k = 1, \ldots, N; \end{aligned}$$

the restriction $v_0 \leq 0$ means that $df'(0)$ can't be positive, and $a_0 \geq 0$ says that the curve should have positive curvature at $t = 0$. The constraints involving θ tell us that

$$\theta \geq \max_{i=1,\ldots,I} \left[\left| \sum_{l=1}^{L_i} df(t_{il}) p_{il} \right| \right],$$

and by minimizing θ, we minimize the max-error; this inequality is split into $2I$ constraints so that all constraints are linear.

We have a linear optimization problem with a finite number of variables ($N+3$), but with an infinite number of constraints ($\forall\, t \in [0, T]$). To solve this problem, one could consider using one of the techniques developed specifically for (linear) semi-infinite optimization problems. Because of the nature of the problem, however, one can safely replace the conditions involving 'for all $t \in [0, T]$' by for all $t \in \{1/M, 2/M, \ldots, T/M\}$ with M sufficiently large; in the EpiSolutions Inc. implementation M is usually chosen so that the mesh size $(1/M)$ is 1 month. After this time-discretization, the problem becomes a linear programming problem that can be solved using a variety of commercial packages.

One important component of the *EpiCurves* solution is that the zero-curves are defined at every time t, there is never any need to resort to interpolations to fill in missing time-gaps. This, of course, gives us great flexibility in choosing the right approximations when building pricing mechanisms.

8 More examples

There remains only to 'look' and analyze some examples.

Bond-Portfolio. The first one is a Bond-portfolio. This data-set includes U.S. Treasury Bill and U.S. Treasury Bond data from August 3, 2001. There are 7 instruments in all, including 3 U.S. Treasury Bills with term-to-maturities ranging from 3 to 7 months and 4 U.S. Treasury Bonds with term-to-maturities ranging from 2 to 30 years. This data was obtained from the Bloomberg U.S. Treasuries web page; details can be found in Table 4 in the Appendix. As a point of comparison, we use the results of the BootStrapping technique supplied by Financial Provider no. 2; Financial Provider no. 1 BootStrapping functionality can't deal with a Bond-[portfolio. The results are graphed in Figure 7.

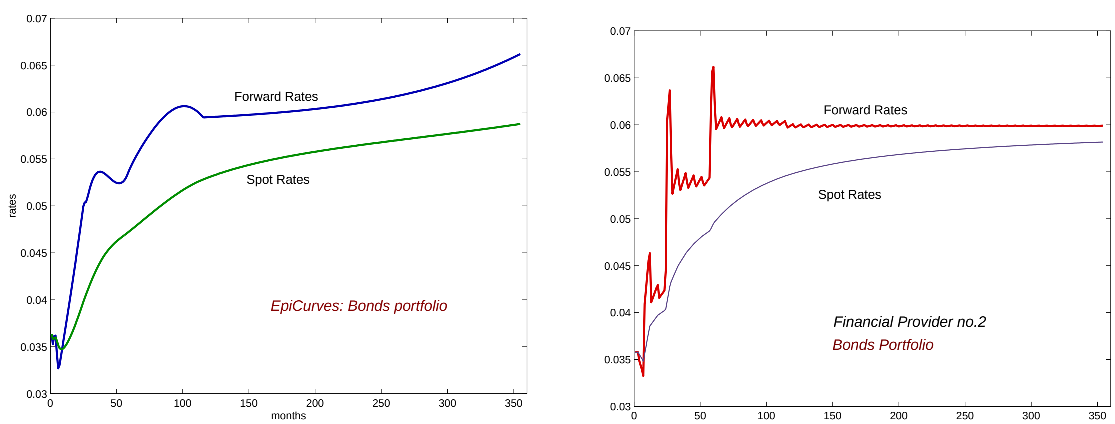

Figure 7: Spot and Forward Rates associated with a Bonds portfolio.

DFS2-Portfolio. This next example is a relatively challenging one. The portfolio includes 51 instruments: Deposits, Futures and Swaps from August 3, 2001 with quite a bit of overlap of maturity-dates. A short description of the composition of this portfolio follows here; details can be found in Tables 5–6. in the Appendix. There are 51 instruments in all, broken down as follows: 3 Eurodollar Deposits with term-to-maturities ranging from 1 to 6 months, 40 Eurodollar Futures with 90-day deposit maturities ranging from 4 months to 10 years, and 8 Swaps with term-to-maturities ranging from 1 to 10 years. This data was obtained from the Federal Reserve (Statistical Release H.15) and the Chicago Mercantile Exchange (CME).

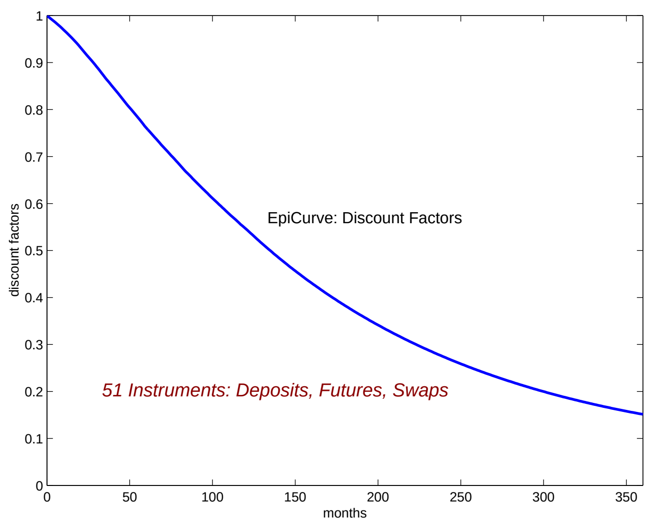

Figure 8: Discount Factors curve for the 51-instruments portfolio

In the EpiSolutions Inc. implementation of the *EpiCurves* methodology, there is an option that allows the user to fine tune the level of accuracy that will be acceptable; accuracy being defined in terms of the max-error, i.e. in terms of the objective of the optimization problem. Asking for a higher level of accuracy will usually result in a more jagged curve since one must accommodate/adjust more rapidly to even small changes in the cash flow. This is effectively illustrated by curves graphed in Figure 9. In the first one the tolerance is 5 basis points, in the second one just 1 basis point.

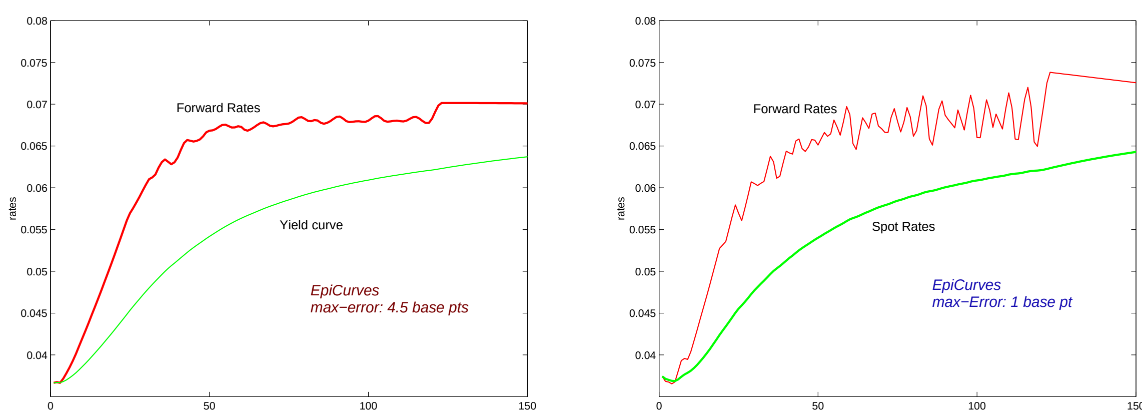

Figure 9: Variations of the zero-curves under max-error tolerance

Notwithstanding this fine tuning, *EpiCurves* is really the only methodology that provides 'serious' zero-curves associated with such a portfolio. The only BootStrapping approach that one could rely on, if one can use the word 'rely,' is the one implemented by EpiSolutions Inc. But the results are less than satisfactory. Both the Spot and the Forward Rate curves were derived for this portfolio with the Forward Rate curve generated by BootStrapping spiking up to 30% at one point and then immediately thereafter going negative! Both the Spot and the Forward Rate curves were derived for this portfolio with the Forward Rate curve generated by BootStrapping spiking up to 30% at one point and then immediately thereafter going negative! Of course, introducing a *convexity adjustment* to the futures substantially improves the BootStrapping results, although the Forward Rate curve still comes with some abrupt rate changes; note that this convexity adjustment has only a minor effect on the *EpiCurves* results.

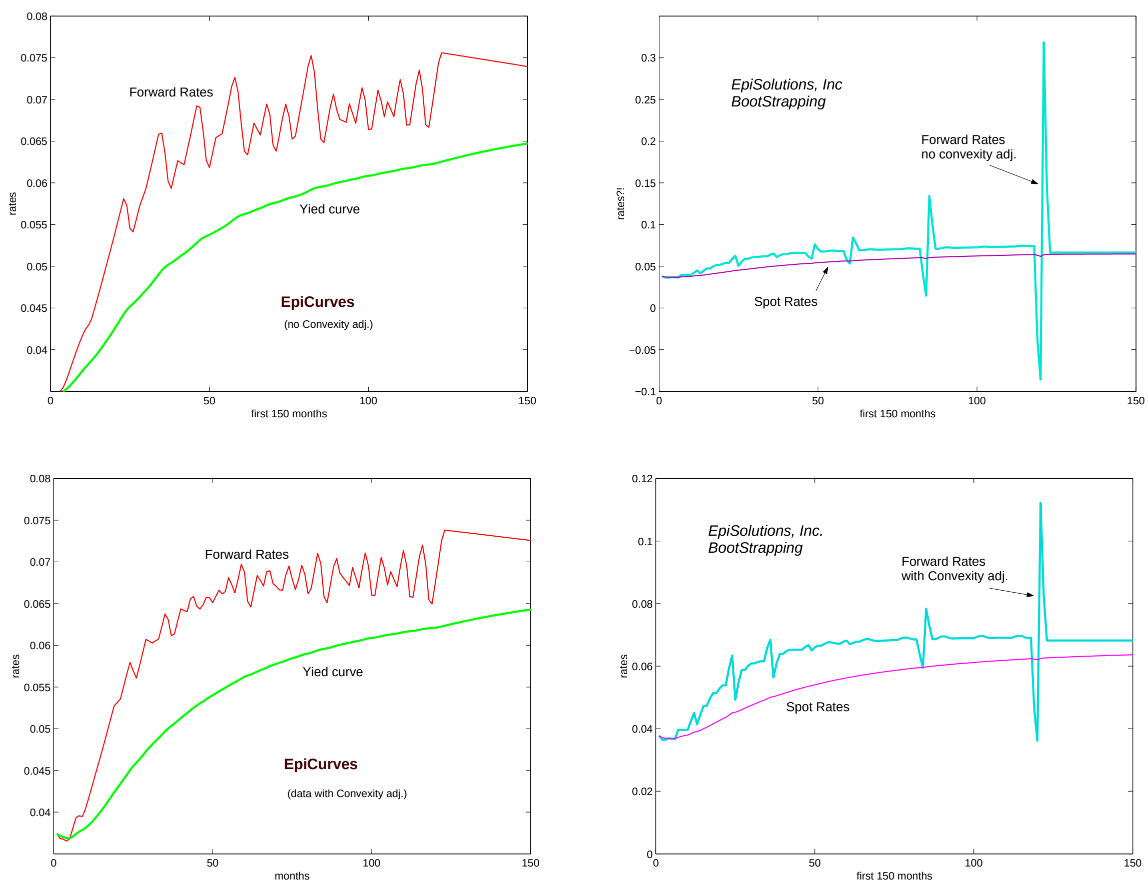

Figure 10: Spot and Forward Rates with and without Convexity adjustments

9 Summary

The major objective in developing the *EpiCurves* methodology was to overcome the inconsistent assumptions and limitations of the standard BootStrapping technique and its Maximum Smoothness variant. This is accomplished by allowing for inclusion of the *complete portfolio* of term structure instruments, while at the same time providing the *smoothness* so crucial to practitioners as the solid foundation on which to build believable valuations, forecasts, and other financial analytics.

10 Appendix

Settle	*Maturity*	*Rate*
06/10/97	06/11/97	0.054688
06/11/97	07/11/97	0.056250
06/11/97	08/11/97	0.057500
06/11/97	09/11/97	0.057344
06/11/97	12/11/97	0.058125
06/11/97	06/11/98	0.061875

Table 1: Eurodollar Deposits

Settle	*Maturity*	*Rate*
06/11/97	06/11/99	0.062010
06/11/97	06/11/00	0.064320
06/11/97	06/11/01	0.065295
06/11/97	06/11/02	0.066100
06/11/97	06/11/04	0.067860
06/11/97	06/11/07	0.068575

Table 2: Swaps .

Delivery	*Maturity*	*Price*	*Delivery*	*Maturity*	*Price*
06/16/97	09/16/97	94.195	06/19/00	09/19/00	93.220
09/15/97	12/15/97	94.040	09/18/00	12/18/00	93.190
12/15/97	03/16/98	93.820	12/18/00	03/19/01	93.190
03/16/98	06/16/98	93.725	03/19/01	06/19/01	93.120
06/15/98	09/15/98	93.610	06/18/01	09/18/01	93.080
09/14/98	12/14/98	93.510	09/17/01	12/17/01	93.050
12/14/98	03/15/99	93.410	12/17/01	03/18/02	92.980
03/15/99	06/15/99	93.390	03/18/02	06/18/02	92.980
06/14/99	09/14/99	93.360	06/17/02	09/17/02	92.940
09/13/99	12/13/99	93.330	09/16/02	12/16/02	92.900
12/13/99	03/13/00	93.260	12/16/02	03/17/03	92.830
03/13/00	06/13/00	93.250	03/17/03	06/17/03	92.830

Table 3: Eurodollar Futures

Settle	*Maturity*	*Price*
08/03/01	11/01/01	3.44
08/03/01	01/31/02	3.36
08/03/01	02/28/02	3.33

Settle	*Maturity*	*Coupon*	*Price*
08/03/01	07/31/03	0.03875	$99 + 30/32$
08/03/01	05/15/06	0.04625	$99 + 26/32$
08/03/01	02/15/11	0.05000	$98 + 25/32$
08/03/01	02/15/31	0.05375	$99 + 00/32$

Table 4: Bond portfolio — U.S Treasury Bills and Bonds

Delivery	*Maturity*	*Pric*	*Convex.*	*Delivery*	*Maturity*	*Price*	*Convex.*
09/17/01	12/17/01	96.430	0.023	09/18/06	12/18/06	93.340	13.999
12/17/01	03/18/02	96.295	0.117	12/18/06	03/19/07	93.250	15.364
03/18/02	06/18/02	96.150	0.274	03/19/07	06/19/07	93.290	16.794
06/17/02	09/17/02	95.805	0.494	06/18/07	09/18/07	93.260	18.283
09/16/02	12/16/02	95.420	0.775	09/17/07	12/17/07	93.230	19.826
12/16/02	03/17/03	95.020	1.121	12/17/07	03/17/08	93.140	21.444
03/17/03	06/17/03	94.780	1.533	03/17/08	06/17/08	93.180	23.128
06/16/03	09/16/03	94.514	2.005	06/16/08	09/16/08	93.150	24.870
09/15/03	12/15/03	94.315	2.539	09/15/08	12/15/08	93.125	26.665
12/15/03	03/15/04	94.100	3.138	12/15/08	03/16/09	93.035	28.537
03/15/04	06/15/04	94.030	3.803	03/16/09	06/16/09	93.075	30.474
06/14/04	09/14/04	93.905	4.529	06/15/09	09/15/09	93.050	32.468
09/13/04	12/13/04	93.800	5.314	09/14/09	12/14/09	93.030	34.514
12/13/04	03/14/05	93.660	6.166	12/14/09	03/15/10	92.940	36.641
03/14/05	06/14/05	93.650	7.085	03/15/10	06/15/10	92.980	38.831
06/13/05	09/13/05	93.575	8.064	06/14/10	09/14/10	92.950	41.079
09/19/05	12/19/05	93.505	9.183	09/13/10	12/13/10	92.930	43.376
12/19/05	03/20/06	93.395	10.294	12/13/10	03/14/11	92.845	45.756
03/13/06	06/13/06	93.420	11.379	03/14/11	06/14/11	92.885	48.199
06/19/06	09/19/06	93.375	12.707	06/13/11	09/13/11	92.855	50.700

Table 5: 51-Instruments portfolio — Eurodollar Futures

Settle	*Maturity*	*Rate*
08/03/01	09/03/01	0.0366
08/03/01	11/03/01	0.0359
08/03/01	02/03/02	0.0360

Settle	*Maturity*	*Rate*
08/03/01	08/03/02	0.0385
08/03/01	08/03/03	0.0444
08/03/01	08/03/04	0.0491
08/03/01	08/03/05	0.0523
08/03/01	08/03/06	0.0547
08/03/01	08/03/08	0.0576
08/03/01	08/03/11	0.0598
08/03/01	08/03/31	0.0632

Table 6: 51-Instruments portfolio — Eurodollar Deposits & Swaps

References

[1] K.J. Adams and D.R. Van Deventer. Fitting yield curves and forward rate curves with maximum smoothness. *Journal of Fixed Income*, ?:52–62, 1994.

[2] M. Buono, R.B. Gregoru-Allen, and U. Yaari. The efficacy of term structure estimation techniques: A Monte Carlo study. *Journal of Fixed Income*, 1:52–59, 1992.

[3] F.B. Hildebrand. *Introduction to Numerical Analysis*. Dover Publications Inc., 1987.

[4] G.S. Shea. Term structure estimation with exponential splines. *Journal of Finance*, 40:319–325, 1985.

[5] O.A. Vasicek and H.G. Fong. Term structure modeling using exponential splines. *Journal of Finance*, 37:339–356, 1977.