

Investments

Lecture 3: Yield curves

Philip H. Dybvig
Washington University in Saint Louis

- *The* term structure of interest rates
- Relations among yield curves: intuition
- Conventions and complications in practice
- Continuous compounding
- Duration and effective duration

The term structure of interest rates

Generically, we refer to the *term structure* of interest rates (*term structure* for short) the pattern of interest rates for different maturities implicit in quoted bond prices on a single date. Even in riskless securities (still our main focus in this lecture), there are different ways of representing the term structure using *the forward rate curve*, or *the yield curve*. Indeed, there are many different yield curves, for discount bonds and for coupon bonds of different moneyness.

Using the basic no-arbitrage relationships and algebra, we can understand the connections among the various sorts of yield curves. The forward, zero-coupon, and par-coupon yield curves all start at the same place at short time-to-maturity, but the forward rate curve is steepest, the zero-coupon next-steepest, and the par-coupon the least steep. This is because the zero-coupon rate at some maturity is a sort of average of forward rates at that and earlier maturities, and the par-coupon rate at some maturity is a similar average that puts more weight on earlier maturities.

Connecting various yield curves: intuition


The various interest rates and yields correspond to different time-patterns of investing. The forward rate corresponds to investing during the future year. The zero-coupon rate corresponds to investing from now until some future date, and is approximately the average of all the forward rates along the way. The par-coupon rate combines investing to maturity with investing for shorter periods until the various coupons are paid. Like the zero-coupon rate, the par coupon rate is an average of futures rates from now until maturity, only with more weight placed on nearby forward rates than for the zero-coupon rate.

The diagram on the following slide illustrates the intuition. For forward investing (corresponding to $f(0, 10)$), the position is undertaken 9 years from now and liquidated 10 years from now. For an investment in a zero-coupon bond (corresponding to $z(0, 10)$), the position is undertaken now and liquidated 10 years from now. The value is constant in time 0-value terms (although it is falling in dollar terms). Not surprisingly, the appropriate yield is an average of forward rates across the 10 years. For a par coupon bond ($c(0, 10)$, assumed to be 5%), the value invested is falling over time as coupons are paid. For a coupon bond, the

yield is an average of forward rates that puts some more weight on early rates than does the zero-coupon rate. The par-coupon rate is also a weighted average of zero-coupon rates. Finally, for a self-amortizing bond (coupons but no principal, like a mortgage), the present value declines more quickly as coupons are paid, so in this case the yield puts even more weight on early forward rates.

Rates intuition diagram


Time pattern of investing, forward lending




Time pattern of investing, zero-coupon



Time pattern of investing, par-coupon



Time pattern, self-amortizing



Connecting the different term structures: some algebra

Forward rates and zero-coupon rates are connected by the expression

$$z(0, T) = \left(\prod_{s=1}^T (1 + f(0, s)) \right)^{1/T} - 1 \approx \frac{1}{T} \sum_{s=1}^T f(0, s)$$

This can be proven using the simple expressions for the discount factor in terms of the forward rates and in terms of the zero-coupon rates. The approximation is very good provided the interest rates are not too large.

Zero-coupon and par-coupon rates are connected by the expression

$$c(0, T) = \sum_{s=1}^T w(0, s) f(0, s)$$

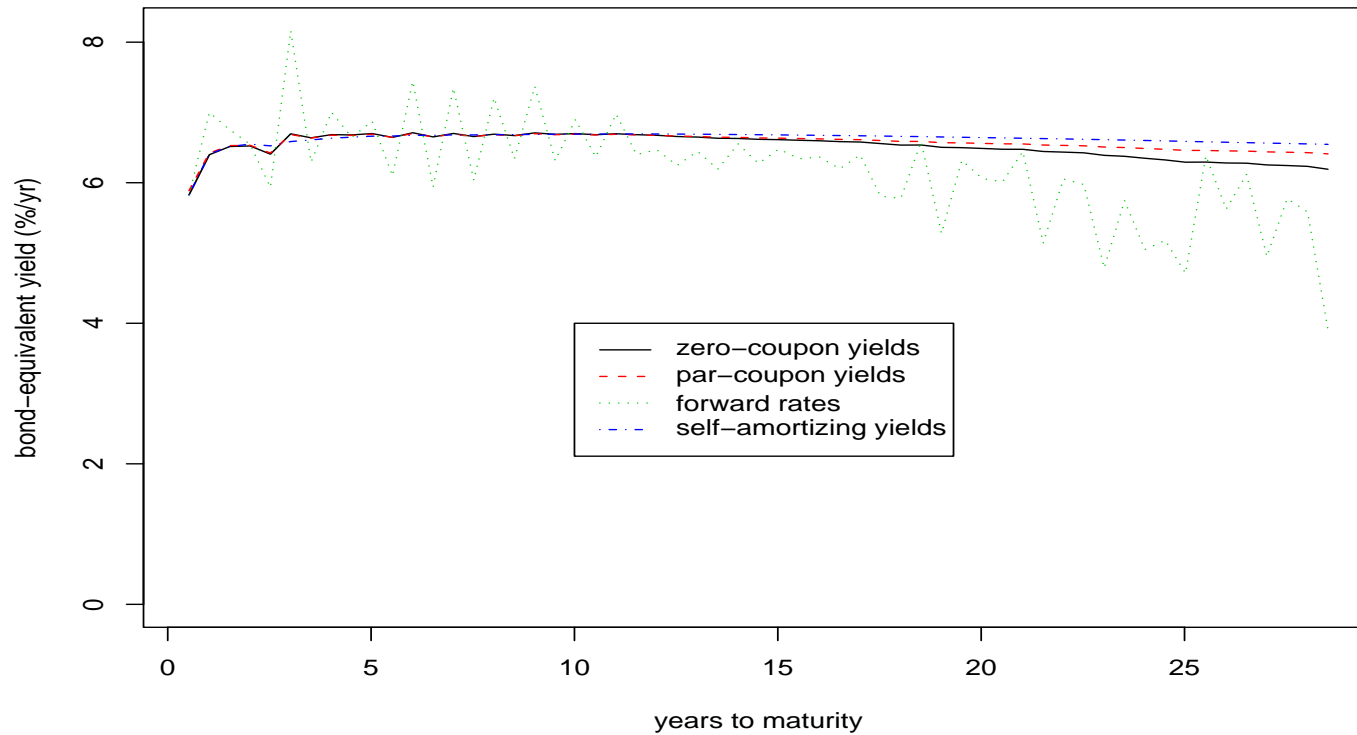
where the positive weights $w(0, s)$ are defined by

$$w(0, s) = \frac{D(0, s)}{\sum_{t=1}^T D(0, t)}$$

sum to one and are positive and decreasing in s .

A recent yield curve

Various yields (from strip coupon interest), May 4, 2000



STRIP price source: Wall Street Journal, May 5, 2000

Computing the term structure from Treasury STRIPs

The previous plot was computed using Treasury STRIPs that are claims to individual principal or coupon payments from Treasury Bonds or Notes. Originally, claims to individual cash flows from Treasury issues were created by investment banks as claims to funds that held the Treasury issues in trust. Now, these claims are created as a service by the Treasury, and the Treasury will strip or reconstitute securities you hold for a modest fee. One curious feature of the program is that cash flows from principal repayment (at the end) and coupon interest (through the life of the bond) are not interchangeable. In order to reconstitute a bond, you need the right principal (corpus) but you can use interest stripped off other bonds. This means the market for coupon interest strips is, at least in principle, more liquid than the market for principal strips. However, the two types of STRIP usually trade within a few basis points of each other.

Constructing the yield curve: issues

Some details need to be decided upon when constructing a yield curve, even once we have decided to use Treasury STRIP prices. For example,

- Use bid price, ask price, spread midpoint?
- Frequency of data to use (monthly? semi-annual?)
- Smooth data or use quotes directly?
- Round to an integer number of periods?
- Use continuously-compounded yield or bond-equivalent yield?
- How to extrapolate beyond available data or interpolate between available data?

Often these choices do not matter much, especially if you are consistent. It depends in part on how you will use the data and the precision required.

STRIP quotes and zero-coupon yield

Strips are quoted in dollars and 32nds per \$100 of face value. For example, the strip maturing in May 2010 used in constructing the yield curve in the previous picture is quoted (the ask price) at 51 : 21 or $51 + 21/32 = \$51.65625$ per \$100 of payment 10 years from now. The bond equivalent zero-coupon yield y used in the figure can be computed approximately from the formula

$$D_{0,10} = .5165625 = (1 + y/2)^{-10*2}$$

(Recall that by definition, $a^{-b} = 1/a^b$.) Solving for y , we have that

This gives the semi-annually compounded interest rate quoted as an annual rate. In the computation used in the figure, there is an adjustment for the time between the settlement of a trade today and the middle of the month when a coupon would be received. This adjustment is approximate (I assume settlement the next day and a whole number of months from the 15th for each stripped coupon) but for more precise work you would look at the actual number of days until each stripped coupon.

Forward rates

The forward rates are computed using pairs of STRIP prices. (Recall the arb from the previous lecture: forward lending is replicated by buying a longer STRIP using proceeds from selling short just enough of a shorter STRIP.) For example, we can use the previously noted May, 2010 strip price of 51:21 (= 51 21/32 or 51.65625) and the Nov 2009 strip price of 53:14 (= 53 14/32 or 54.4375) to compute the forward rate for lending from 9 1/2 years out until 10 years out as


$$y = 2 \left(\frac{53.4375}{51.65625} - 1 \right) \\ \approx 6.89\%$$

One surprising feature of the yield curve plot is the irregular appearance of the forward rate. This is actually spurious detail. The forward rates come from differencing STRIP prices, and differencing magnifies relative errors. (For example, think about $100 - 99$, when 100 and 99 are both accurate plus or minus 2%.) Given the problems with the quotes and the spread, we cannot rule out the shape in the plot, but we also cannot rule out more reasonable-looking shapes.

Smoothing to remove spurious forward rate irregularities


The following two figures illustrate the effect of smoothing (using a technique to be described shortly) on the discount STRIP prices and on the forward rates. It is interesting to see how smoothing that cannot be seen easily in the plot of zero-coupon rates can make the forward rate curve much smoother. This underlines the fact that the detail in the forward rate curve is not significant. Using the smoothed curve is good psychologically because it eliminates distracting and irrelevant features. It also can be used as an input to simulation or other analysis (as we will do later) in which we want to be sure the results are not driven by spurious features.

Zero-coupon bond prices, with and without smoothing



Yield curves from smoothed zero-coupon bond prices

Various yields (from smoothed strip coupon interest), May 4, 2000



STRIP price source: Wall Street Journal, May 5, 2000

Smoothing the yield curve

We have seen that an almost imperceptible adjustment to the original STRIP data makes the forward rate curve much smoother. (The only part that is really unclear is the very short end. The short end is important in practice; to have a closer look we would use many more STRIPs and T-Bills to nail it down.) This smoothing is important for looking at the yield curve and for communication without the spurious detail.

I used linear regression to smooth the yield curve. The functions I chose as covariates (or independent variables) in the regression were picked to capture the overall features of the yield curve without too many sudden changes. The specific regression I fit to the yield curve was:

$$z(0, t) = a + b \exp(-t) + c \exp(-t/3) + d \exp(-t/9) + e \exp(-t/27) + \varepsilon$$

I used ordinary least squares to estimate the parameters a , b , c , d , and e , and I replaced the zero-coupon rates by the fitted values (the same equation without the error term). Then I derived the discount factors and various rates from the fitted values.

Smoothing the yield curve: tips

- Using too many functions *overfits* and does not smooth
- Using too few functions *underfits* and will not match the STRIPS well
- Fit the yield curve, not the raw bond prices
- Use functions that can mimic features you believe to be important
- Avoid functions that have unwanted sharp changes

Self-amortizing loan yield curve

A self-amortizing loan pays the same amount, say c per period, in every period until the loan is paid off. The yield is computed in the usual way, so the price (or amount borrowed initially) should be the present value of the cash flows given the yield. If y is the bond-equivalent yield of a self-amortizing loan with semi-annual coupons maturing $n/2$ years from now, we have that

$$\begin{aligned} P &= \sum_{s=1}^n cD_{0,s/2} \\ &= \sum_{s=1}^n \frac{c}{(1 + y/2)^s} \\ &= \frac{c}{y/2} \left(1 - \frac{1}{(1 + y/2)^n} \right) \end{aligned}$$

where the last expression is the annuity formula (What is the underlying arb?). Therefore, the self-amortizing loan yield $s(0, t)$ is the value of y that solves the equation:

$$\frac{2}{y} (1 - (1 + y/2)^{-2t}) - \sum_{s=1}^n D_{0,s/2} = 0$$

Since the left-hand side is a decreasing function of y , it is easy to solve this numerically. Note that we can also compute the payments (given the amount borrowed) from these expressions.

Continuously-compounded yields

Compounding k periods per year at a fixed annual rate r grows our money in T years by a factor

$$(1 + r/k)^{kT}$$

As k increases, this factor gets larger due to *interest on interest* or the magic of compounding. It is an interesting mathematical fact that as k increases without bound, this factor tends to the limit

$$e^{rT} = \exp(rT) \tag{1}$$

where $e \approx 2.71828\dots$ is a transcendental number called the base of the natural logarithm. Naturally, the growth factor $\exp(rT)$ is called continuous compounding, and r is called the continuously-compounded interest rate. There is corresponding discounting with a factor $1/\exp(rT) = \exp(-rT)$. We have been working with bond-equivalent yields, which assume compounding twice a year; continuous yields correspond to continuous discounting. In some cases (for example, in computing the zero-coupon rate) it is useful to use the (natural) logarithm,

which is the inverse of the exponential function: $\exp(\log(x)) = \log(\exp(x)) = x$. This is the logarithm base e . Usually, beyond high school we usually use the natural logarithm (or log base 2 for some applications in information science), while in high school and before we usually use the log base 10 (the inverse of 10^x : $10^{\log_{10}(x)} = \log_{10}(10^x) = x$).