THE XRM YIELD CURVE METHODOLOGY

James A. Girola

U.S. Department of the Treasury

November 4, 2024

I'd like to thank Treasury's Office of Economic Policy for supporting the Yield Curve Project for many years. I'm especially grateful to Eric Van Nostrand for encouraging and supporting the writing of this monograph, and for reading multiple drafts and providing comments and discussions that greatly improved the monograph's content. I'd also like to thank Benjamin Harris who initially supported the monograph work. And I received very helpful comments from Tara Sinclair and Hanjo (Terry) Kim.

This Office of Economic Policy Working Paper presents original work by the staff of the Office of Economic Policy. It is intended to generate discussion and critical comment while informing and improving the quality of the analysis conducted by the Office. The paper is a work in progress and subject to revision. Comments are welcome, as are suggestions for improvements, and should be directed to the authors. This Working Paper may be quoted without additional permission.

Contents

1. Introduction

This monograph describes the Extended Regressions on Maturity Ranges (XRM) Yield Curve Methodology for constructing yield curves developed by the U.S. Department of the Treasury Office of Financial Analysis within the Office of Economic Policy. XRM is a general methodology for computing yield curves for a wide variety of fixed income market sectors.

The first section of this chapter sets out the distinguishing features of XRM and the second section summarizes yield curves done by XRM that are currently published. The third section outlines the subsequent chapters.

Much of the discussion in this monograph draws from and reproduces or quotes from parts of Girola (2007, 2010, 2011, 2014, 2015, 2016, 2019, 2022), as well as previous work on this methodology in U.S. Department of the Treasury (2005a, 2005b, 2006, and 2007). The monograph provides more information about this material including detailed exposition of foundations and mathematical development.

XRM Features

The XRM methodology is a general methodology that can be used to develop different types of yield curves for various sectors of fixed income markets. The methodology was invented to meet the need for a practical yield curve approach that generates stable and accurate estimates and that is grounded in a conceptual foundation derived from established bond market characteristics. In the XRM methodology, the conceptual foundation is based on maturity ranges, and the resulting estimates are smooth and robust over time regardless of the state of the bond market. Furthermore, the XRM methodology is also designed to fulfill additional important requirements for yield curve estimation that are often ignored by other approaches.

The first of these additional requirements is that the methodology must provide yield curves that realistically represent the various sectors of the U.S. corporate bond market as well as the market for U.S. Treasury securities. Most other yield curve approaches are aimed at Treasury securities. Corporate bonds, however, are much more heterogeneous than Treasuries, many have special features that must be taken into account, and there is often a large amount of random noise in corporate bond data that may obscure market behavior. Therefore, the methodology must be powerful enough to handle the complexities of corporate bonds as well as Treasuries.[1]

Another important requirement is that the yield curve methodology must be able to combine corporate bonds of different qualities from disparate markets into a single yield curve. The HQM yield curve that is described in this monograph, originally developed for pension funding as mandated by the Pension Protection Act of 2006 (PPA), Pub. L. 109-280, blends together high quality corporate bonds, that is, bonds rated AAA, AA, or A, into a single well-defined yield curve. As explained later, the HQM yield curve uses special regression variables to do the blending.

The requirement for the inclusion of regression variables in the HQM yield curve to combine qualities is an illustration of the general requirement that the yield curve methodology must include regression variables as needed. These regression variables are in addition to the discount function in the methodology that is used to compute present values of future cash flows of bonds.

Regression variables are essential for capturing the special attributes of individual bonds. The conventional yield curve approach does not provide for regression variables and cannot be applied to bonds with special attributes. In addition to variables for combining markets in the HQM yield curve, the yield curves presented in this monograph include a hump regression variable described later for the bulge in yields that is often seen around 20 years maturity.

In addition to these requirements, there is the critical requirement that the yield curve methodology must be able to project the yield curve results for indefinitely long maturities beyond the farthest maturity of the bonds used to construct the yield curve, and the projections must be consistent with actual yields in bond markets. The projections show what yields would be for such long-dated bonds even if these bonds don't exist at the time in markets.

In the case of the HQM yield curve, for example, pension liabilities that need to be discounted by the yield curve can extend well beyond the 30 years maturity that is the maximum maturity of bonds used to estimate HQM. Consequently, the HQM projection extends out for another 70 years maturity thereby providing yields for discounting up through a total of 100 years maturity. The projected yields at such long maturities must be consistent with yields before 30 years maturity and must reflect market beliefs about long-term investment returns.

The requirement that the yield curve methodology generate good projections is important for all yield curves, not just HQM, and is important even if the projections are not actually used. This is because the projection of a yield curve extends the pattern of yields in the bond set used to construct the yield curve to higher maturities beyond those in the bond set. Therefore, an effective projection ensures that the pattern of yields in the bond set has been accurately estimated.

[1] An extended treatment of corporate bonds is contained in Fabozzi, ed. (2005).

In addition to these requirements, there is the practical requirement that a yield curve methodology used for actual applications must produce reliable market indicators. This implies that yield curve estimates have to be robust and stable with respect to random transitory perturbations in the market. Therefore, the estimates cannot be unduly buffeted by market disturbances, and at the same time the estimates must capture all significant market movements.

In particular, in order to produce reliable results, it is important that the yield curve methodology is not subject to the criticism that it is arbitrary, in the sense that there are equally valid methodologies or equally valid ways of implementing a given methodology that produce different results from the same data. Clearly, one cannot have confidence in an estimate of a yield curve when there is another legitimate estimate of the yield curve that gives different results.

Yield curves can be arbitrary when they are derived from overly complicated models. This is because bond data may not show clear enough patterns to distinguish among competing complicated models, with the result that choices among such models cannot be justified empirically despite the fact that such choices can significantly affect results.

And especially, yield curves are arbitrary when there are free parameters in the methodology without a rationale for setting them. In that case, yield curve results depend upon the preferences of the analyst computing the yield curve.

The XRM methodology is not arbitrary because it is grounded in analysis of bond markets that breaks bond trading into maturity ranges and derives mathematical equations that represent each range. The choice and modeling of the maturity ranges implies the use of a cubic spline to capture all the ranges, and the specifications of the spline are given by the maturity ranges. Therefore, the spline is not arbitrarily chosen in advance without justification: rather, it arises naturally from the use of maturity ranges.

Furthermore, the cubic spline is constructed so that it includes mathematical constraints that ensure that the end of the spline at the farthest maturity 30 years of the bonds in the market can be carried forward to provide a projection that is consistent with bonds of less than 30 years maturity. Specifically, the calculated forward rate at 30 years maturity is held fixed indefinitely beyond 30 years maturity for the projection.

All of this is set out in detail in the following chapters. Maturity ranges ensure that XRM yield curve results are stable and robust, are not overly complicated or arbitrary, don't have free parameters, and accurately reflect markets.

In sum, it can be seen from this discussion that the XRM methodology contains features that are not present in conventional yield curves. XRM methodology makes use of maturity ranges to produce stable and accurate yield curves. The ranges are derived from analysis of markets.

In addition, the XRM methodology constructs the yield curve so that it can be extended to higher maturities in order to project yields beyond the maturities of bonds trading in the market, and the projections are consistent with actual bonds in the market. Yield curves done by XRM currently project yields beyond 30 years out through 100 years maturity, and the projected yields are consistent with yields up through 30 years maturity and with long-term returns provided in the market.

Also, the XRM methodology includes regression variables in addition to the discount function to account for special characteristics of individual bonds. The regression variables make it possible to include different kinds of bonds in the same yield curve. And the hump variable, which is a regression variable, picks up movement in yields at longer maturities that are not part of the discount function.

Therefore, as its name indicates, the Extended Regressions on Maturity Ranges XRM methodology focuses on the maturity ranges and derives the yield curve using nonlinear regressions that are extended both to project yields beyond 30 years maturity and to take account of additional bond characteristics by using special regression terms.

The XRM methodology can be compared to other approaches. One well-known approach is presented in Nelson and Siegel (1987) and Svensson (1994).[2] Well-known spline approaches include the generalized cross validation (GCV) approach with smoothing splines of Fisher, Nychka, and Zervos (1995).[3]

XRM Yield Curves

The purpose of the XRM methodology is to produce yield curves. Yield curves have many applications. The primary types of yield curves currently done by XRM are par yield curves and spot yield curves, and the latter are critical in funding applications for pensions and other programs that pay cash flows in the future.

A yield curve provides information about a particular bond market sector at a point in time, including information about yields in the sector at different maturities. Yield curves can also provide additional information and statistics about the sector, and the XRM methodology produces an extensive set of such information.

Currently there are three sectors for which yield curves are done by XRM. These yield curves are set out in the following two chapters.

The first sector for which XRM produces a yield curve is the corporate bond high quality market. This sector comprises U.S. corporate bonds that are in the top three credit qualities A, AA, and AAA. The yield curve for this sector is termed the High Quality Market (HQM) Yield Curve, and it is the original yield curve that motivated the development of the XRM methodology.

In addition to the HQM yield curve, XRM produces the Treasury Nominal Coupon-Issue (TNC) Yield Curve that pertains to the fixed income sector comprising U.S. Treasury nominal coupon issues, both notes and bonds. Nominal Treasury coupon issues provide principal and interest cash flows in nominal dollars. Nominal Treasuries are sometimes called regular Treasuries.

And third, there is the sector comprising U.S. Treasury Inflation-Protected Securities or TIPS. The XRM yield curve for this sector is called the Treasury Real Coupon-Issue (TRC) Yield Curve. TIPS notes and

[2] An application of the Svensson approach to the U.S. Treasury market is presented in Gürkaynak, Sack, and Wright (2007a and 2007b).
[3] See also Fisher and Zervos (1996).

bonds provide principal and interest payments in real terms, and payments are converted to dollars using the not seasonally adjusted Consumer Price Index for All Urban Consumers (CPI-U).[4]

The first two of these yield curves, HQM and TNC, are called nominal yield curves because they provide results in nominal or dollar terms. The TRC yield curve is called a real yield curve because its results are in real terms.

In addition, a fourth curve is also done, and this is called the Treasury Breakeven Inflation (TBI) Curve. This curve combines the TNC and TRC yield curves to calculate breakeven inflation.

The XRM methodology computes yield curves for these sectors by developing price equations for the sectors. For each sector, the price equation shows the price of a bond in that sector as a function of the bond's future payments and other information.

These price equations are estimated from sets of bonds made up of bond data for Treasury notes and bonds and bond data for corporate bonds. The bond data are mostly comprised of bonds whose characteristics are similar to Treasuries, including payment of coupons. In this discussion, such bonds are termed standard bonds, and they are described in detail in Chapter 4.

Using the price equations computed from the sets of bonds, the XRM methodology can generate an array of bond statistics, including prices of bonds with different features as functions of the bond cash flows, and, inversely, different streams of cash flows that are consistent with specified prices.

In particular, with the price equations, XRM can generate various types of yield curves. This monograph focuses on two types of yield curves: the par yield curve and the spot yield curve. Both yield curves are described in detail in succeeding chapters. Because the price equations are estimated from standard bonds, the resulting yield curves are consistent with standard bonds.

The par yield curve shows yields for standard bonds trading at par. These bonds can be called standard par bonds, and they comprise standard bonds whose flat price (excluding accrued interest) is equal to principal. Par yield curves have many uses including pricing of new bonds and as indicators of the state of financial markets.

The spot yield curve shows yields for nonstandard bonds that provide single payments at maturity, that is, zero coupon bonds with no coupon payments. The yields are called spot rates. Even though the zero coupon bonds are nonstandard, the spot rates are computed from the par yield curve and therefore the spot rates are calculated so as to be consistent with the par yield curve for standard bonds.

In addition, each price equation generates a long-term forward rate, which is the projected average forward rate for bond maturities above 30 years which is the usual maximum maturity for bonds that exist in markets. The long-term forward rate can then be used to project the par and spot yield curves out beyond 30 years maturity and the projections are consistent with market results up through 30 years maturity. The long-term forward rate is defined and described in detail in succeeding chapters.

---

[4] In earlier work Girola (2005 and 2006), the XRM methodology was used to construct yield curves for TIPS which were then used to project the real return on Treasury securities for Social Security.

It is the average forward rate in the farthest maturity range as defined later, and is held fixed for all maturities above 30 years.

One of the principal uses of spot rates is to discount future cash flows. This was the original purpose in developing the HQM yield curve for the Pension Protection Act. Each cash flow is discounted by the spot rate whose maturity is the same as the time when the cash flow is paid, and this gives the present value of the cash flow. Summing up the present values of the cash flows gives the amount of money that must be put aside at present to fund the future cash flows. For a pension plan, this is the present funding of future pension payments.

Spot Yield Curves for April 2024

As a preview of the following chapters, the next figure shows the monthly average spot yield curves for April 2024 for the three yield curves HQM, TNC, and TRC. Each spot yield curve is the average of the 22 business days in April, and each is plotted for the 200 maturities ½ year up through 100 years. The values above 30 years maturity are the projection range.

Figure 1.1

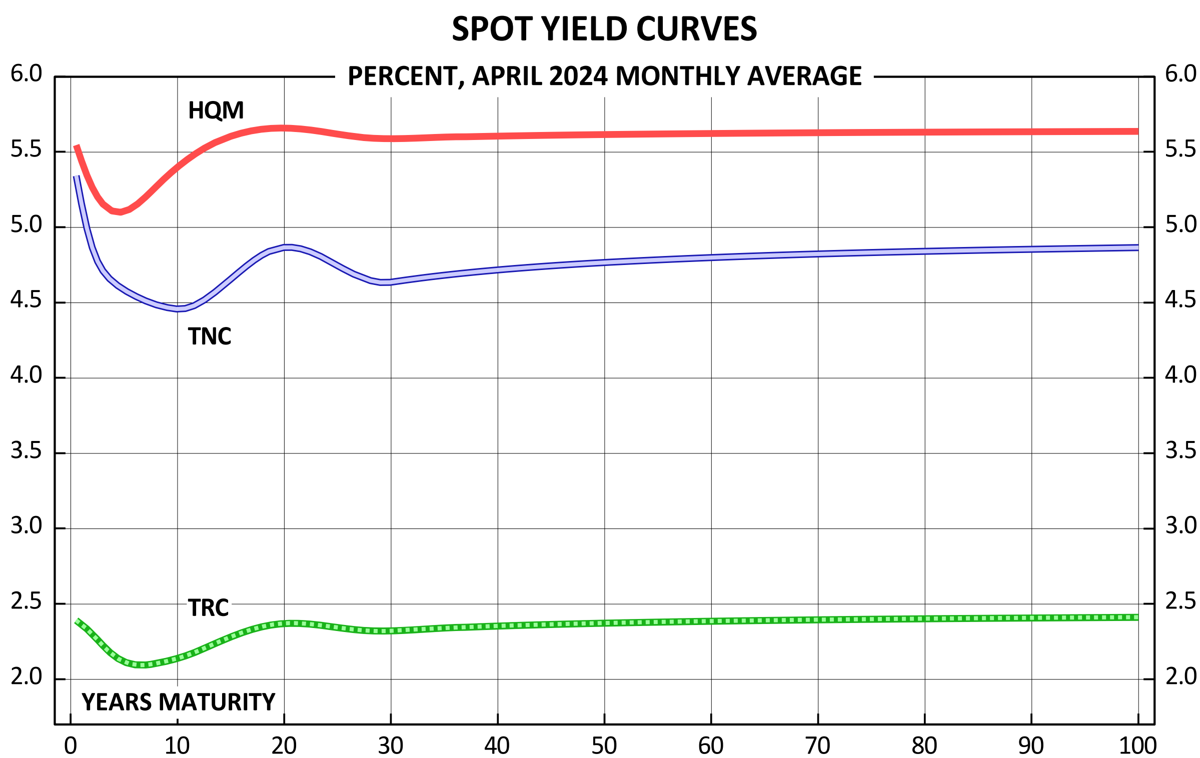

Each of the three yield curves starts high at ½-year maturity and declines, reflecting the federal funds rate in April at 5¼ percent or above. The yield curves have humps: the HQM hump is at 20 years maturity, and the TNC and TRC humps are both at 20½ years maturity.

The yield curves gently rise in the projection range. The HQM yield curve starts at 5.59 percent at 30 years maturity and rises to 5.64 percent at 100 years maturity. The TNC yield curve starts at 4.64 percent at 30 years maturity and rises to 4.87 percent at 100 years maturity. And the TRC yield curve in real terms starts at 2.32 percent and rises to 2.41 percent at 100 years maturity.

The TRC yield curve is positive in April. This does not always happen and TRC can be negative. In contrast, HQM and TNC, being nominal yield curves, are expected to be positive.

Outline of the Chapters

This section outlines the rest of the chapters.

The next two chapters discuss the development of the XRM methodology and the yield curves to which it is currently applied. Chapter 4 describes in detail the bond sets and the standard bonds that are used for XRM estimation.

Chapter 5 describes the discount function and companion forward rate, which are building blocks for constructing bond price equations. And Chapter 6 uses the discount function to write down the conceptual form of the price equation.

Chapters 7-9 develop the mathematical form of the price equation that makes specific the conceptual price equation so that it can be used for estimation. Chapter 7 discusses the fundamental concept of maturity ranges that is the foundation of the XRM methodology. Chapter 8 summarizes basic features of B-splines that are applied to the maturity ranges in order to set up the spline function for estimating the price equation. And Chapter 9 extends the B-splines to enable projection of yields out through 100 years maturity.

Chapters 10-12 describe the regression variables that are currently being using in XRM. Chapter 10 presents the hump variable that is used in all the XRM yield curves to account for the hump in yields of standard bonds often seen around 20 years maturity. Chapter 11 describes the HQM credit variables that blend together corporate bonds of A, AA, and AAA qualities into a single yield curve. And Chapter 12 discusses the variables for on-the-run Treasuries and variables for historical Treasuries.

Chapter 13 outlines estimation techniques. Chapter 14 describes the computation of the par and spot yield curves from the price equations out through 100 years maturity, including forward spot rates. And Chapter 15 describes the computation of breakeven inflation from the TNC and TRC yield curves.

Finally, Chapters 16-18 present selected results from the yield curve computations including figures. These chapters provide illustrations of the concepts discussed in previous chapters and demonstrate practical results that show how the concepts are used.

2. Development of the HQM Yield Curve

This chapter shows how the required applications of the yield curves led to the special features in the XRM methodology that are applied to the HQM yield curve.

For the three sectors listed in Chapter 1, the first yield curve done was the HQM yield curve for high quality corporate bonds. The motivation for HQM was the Pension Protection Act of 2006 that mandated a yield curve for funding of single-employer pension plans.

The Pension Protection Act

The Pension Protection Act of 2006 (PPA), Pub. L. 109-280, signed into law on August 17, 2006, mandated that spot rates from a high quality corporate bond yield curve must be used to discount actuarial estimates of future pension payments from single-employer pension plans for the purpose of determining required funding of these plans.

This mandate was a significant innovation in pension funding: previously, discounting had been done by a single fixed interest rate. At the time of the PPA, much of the discussion focused on the fact that use of a yield curve would allow pension plans to adjust pension funding based on the age of their workforce. Of course, use of a yield curve allows funding to reflect all kinds of conditions in addition to age.

Although the proposal to use a yield curve was widely applauded at the time, there was still the problem of which yield curve to use. This was especially true because, unfortunately, yield curve technology at the time was unable to produce a suitable yield curve. The technology wasn't grounded in market characteristics such as maturity ranges with the result that available yield curve approaches at that time gave different results with no way to choose among them: the approaches were arbitrary.

Also, the available approaches frequently produced unacceptable negative nominal interest rates or interest rates with spurious movements, especially at long maturities, and such results were inconsistent with markets. Another problem was that the approaches were often numerically ill-conditioned such that they gave multiple results or didn't converge at all giving no result.

In view of these problems, there were concerns that a workable yield curve for pension discounting could not be developed. And to make matters more difficult, the PPA included additional challenges in its requirements for the yield curve.

First, the yield curve mandated by the PPA was for corporate bonds rather than Treasury securities. Most yield curve work at that time had been centered on Treasuries. As discussed in the previous chapter, corporate bonds have many special complexities that Treasuries don't have, greatly increasing the difficulty in formulating a successful yield curve.

Moreover, the PPA mandated the use of a single yield curve blending together corporate bonds rated A, AA, and AAA in an average. There existed no approach to calculate a single yield curve like this. Furthermore, the PPA mandated the use of zero coupon spot rates rather than yields on bonds with coupons, and therefore a complete set of spot rates had to be estimated that accurately reflect market conditions despite the fact that there does not exist a fully developed zero coupon corporate bond market. There was no way to do this either.

Finally, it was necessary to project spot rates out through 100 years maturity because many pension liabilities extended well beyond the 30 years maturity that was the usual maximum for actual bonds in markets. Again there was no way to do this. At the time, it was considered normal simply to use the yield at 30 years maturity for yields beyond 30 years even though the 30-year yield was usually biased down.

HQM Features for the PPA

The purpose of the HQM yield curve was to solve all these problems, both general problems of yield curve construction and problems arising from the specific requirements of the PPA, and to provide a new yield curve methodology that would be accurate and acceptable to users in the business and pension communities. The features of the HQM yield curve using the XRM methodology accomplish this purpose.

The starting point for HQM is the use of bond market analysis to derive the maturity ranges. The initial reason why other yield curve approaches are inaccurate is that they aren't grounded in market analysis.

Consequently, the HQM yield curve is justified by market analysis and is not arbitrary. Furthermore, the maturity ranges generate robust yield curve estimates that are well-conditioned, converge rapidly, and are not affected by market instabilities. The estimation technique for HQM is a custom-written algorithm that is stable regardless of starting point.

Moreover, the requirement that the HQM yield curve be able to produce reliable projections out through 100 years maturity constrains the yield curve results to be economically reasonable and eliminates the problem of negative nominal interest rates at long maturities.

And to satisfy the challenges imposed by PPA requirements, the XRM methodology is sufficiently robust that it can be applied effectively to corporate bonds using bond sets comprising thousands of securities with significant random movements and noise. This is currently being done in the HQM yield curve. The HQM credit regression variables described later enable HQM to blend together into a single yield curve the corporate bonds rated A, AA, and AAA. And the XRM yield curve technology enables spot rates to be lifted out of HQM estimates derived from standard bonds with coupons, and these spot rates are consistent with bond markets even though a fully developed corporate zero coupon bond market doesn't exist.

As a result, the HQM yield curve solves the general problems of yield curve construction that prevented accurate yield curve results 20 years ago, and also meets the requirements of the PPA.

The HQM yield curve was first published in Internal Revenue Service (IRS) Notice 2007-81 containing data for August 2007. As of this writing, the HQM yield curve is published in U.S. Department of the Treasury (2024a) and also published on the Internal Revenue Service website. Results are published every month and include monthly average and end of month spot rates and selected par yields. All results are projected out through 100 years maturity.

When first introduced in connection with the PPA, the HQM yield curve went back through October 2003. The HQM yield curve was carried back farther through January 1984 as mandated by the Moving Ahead for Progress in the 21st Century Act of 2012 (MAP-21). Although at present published results are monthly, the yield curve is actually estimated daily for each business day back through 1984.

As mentioned, the HQM yield curve contains the two credit variables that enable the blending of high quality bonds of different credit qualities. In February 2024 the hump variable was added to HQM. Also in February 2024, standard bonds with end calls, described in Chapter 4, were also added, and this addition greatly increased the number of bonds used for estimation.

The next chapter describes the TNC and TRC yield curves, which are the other two XRM yield curves discussed here.

3. The Suite of Treasury Yield Curves

Building on the success of the HQM yield curve, the Office of Economic Policy was asked to extend the HQM yield curve to Treasury securities because more Treasury yield curve data were needed for required applications. The result was the TNC yield curve for nominal Treasuries, followed by the TRC yield curve for TPS and the TBI curve that uses TNC and TRC in combination to derive breakeven inflation.

The reasons why new Treasury yield curves were needed were similar to the reasons for creating the HQM yield curve:

First, more accurate Treasury yield curves were needed that better capture market behavior than other approaches. The use of maturity ranges significantly increases accuracy.

In addition, applications such as government pensions need projections out through 100 years maturity. The use of the XRM projection methodology provides such projections as it does for the HQM yield curve.

Also, there is the requirement for regression variables for a variety of reasons. Both the TNC and TRC yield curves use the hump regression variable for better estimates.

In sum, the use of the XRM methodology for Treasury yield curves generates accurate par yields from which can be derived spot rates. Spot rates are a critical requirement in applying Treasury yield curves to the discounting of future government liabilities.

The TNC Yield Curve

The TNC yield curve pertains to nominal Treasury coupon issues, both notes and bonds, so it is a nominal yield curve. The TNC yield curve originally went back through 2003 and was extended back through 1976 to give a half century of data. The TNC yield curve is published on the Main Treasury website at U.S. Department of the Treasury (2024b). Similar to the HQM yield curve, the website publishes monthly average and end of month spot rates for TNC projected out through 100 years maturity and selected par yields. The TNC yield curve is estimated daily the same as HQM.

An important feature of the TNC yield curve is that it uses all nominal Treasury coupon issues that have been issued, no issues are omitted. At the present time, when-issued Treasuries are not included.

The use of all Treasury coupon issues is in contrast to conventional yield curve approaches that omit on-the-run Treasuries, and omit callable Treasuries and flower bonds from historical data. Because the TNC yield curve uses all coupon issues, it is an extensive dataset of Treasury yields available over the last half century.

As discussed in Chapter 12, special regression variables are included in historical yield curves to account for callable Treasuries and flower bonds. Conventional yield curves that omit callables leave out a substantial amount of historical data especially at long maturities for the early period 1976-1984. Callables and flower bonds no longer exist.

As discussed in later chapters, the TNC yield curve uses the same maturity ranges as HQM, the same constrained B-splines, and the same hump variable. The TNC yield curve doesn't have the two HQM credit variables because Treasuries are default risk free.

The TNC yield curve comprises all issued coupon issues including on-the-run. Conventional yield curves for Treasury securities omit the on-the-run and first off-the-run because the prices of these securities are considered to be subject to special liquidity and other influences and cannot be mixed in with off-the-run securities. In contrast, the TNC yield curve includes all these securities and includes dummy variables for on-the-run and first off-the-run: Chapter 12 discusses the dummies.

The presence of the dummies in effect removes on-the-run securities from the yield curve estimation, with the result that the TNC yield curve is off-the-run, that is, it pertains to the market for off-the-run Treasury securities. Nevertheless, the dummies still produce estimates of on-the-run effects on price that generate on-the-run yields to be contrasted with off-the-run. The webpage with TNC data cited above contains selected on-the-run estimated yields.

In addition, it should be noted that another important application of TNC spot rates is that these spot rates are indicators of nominal social rates of time preference. This is because these spot rates are default risk free, and they are off-the-run without on-the-run influences. Therefore, the aggregate preference of society is indifferent between a future cash flow and its present value discounted by TNC. One application of social time preference is to help make decisions about public construction and infrastructure.

The TRC Yield Curve

The TRC yield curve pertains to real Treasury coupon issues or TIPS, both notes and bonds, so it is a real yield curve. The TRC yield curve goes back through 2003 when the TIPS market became sufficiently established to provide stable results, so the TRC yield curve provides almost a quarter century of data. The TRC yield curve is published on the Main Treasury website at U.S. Department of the Treasury (2024b). Similar to the TNC yield curve, the website publishes monthly average and end of month spot rates projected out through 100 years maturity and selected par yields for TRC all in real terms. As with HQM and TNC, the TRC yield curve is estimated daily.

Analogous to TNC, the TRC yield curve uses all real Treasury coupon issues that have been issued. In contrast to nominal Treasuries, there are no special features attached to on-the-run real issues, so no dummy variables are used in TRC.

As shown in later chapters, the TRC yield curve uses the same maturity ranges as HQM and TNC, the same constrained B-splines, and the same hump variable. As will become apparent, the mathematical structure of the constrained B-splines plus hump variable is a standardized structure that fits well all fixed income securities of high quality credit over the last half century.

The TRC yield curve doesn't have the two HQM credit variables because Treasuries are default risk free. So for regression variables, the TRC yield curve has only the hump variable.

Similar to TNC, TRC spot rates are indicators of social rates of time preference, where in the case of TRC, the time preference is real rather than nominal. Analogous to TNC, TRC spot rates have no on-the-run influences and are risk free. Therefore, the aggregate preference of society is indifferent between a future cash flow in real terms and its present value discounted by TRC.

The TBI Curve

The TBI curve combines the TNC and TRC yield curves to derive breakeven inflation. This curve goes back through 2003 same as TRC. TBI data are published on the Main Treasury website at U.S. Department of the Treasury (2024b). Breakeven inflation is discussed in Chapter 15.

The TBI curve is more accurate than conventional calculations of breakeven inflation for several reasons. First, conventional breakeven calculations use on-the-run nominal Treasuries while TBI rates use the TNC curve which is off-the-run. Therefore, TBI is not distorted by on-the-run influences. Moreover, the algebraic formula for computing breakeven inflation is correctly applied in TBI, whereas conventional calculations typically simply subtract real Treasuries from nominal Treasuries.

However, the most important reason why TBI rates are more accurate is that TBI uses spot rates for nominal and real Treasuries while conventional calculations use yields with coupons. Spot rates are the correct interest rate measure for breakeven inflation because breakeven inflation is intended to represent the inflation rate over a time span that provides equal returns for nominal and real Treasuries, and calculation of equal returns from one point in time to another requires spot rates.

Consequently, it appears that a large part of the apparent inaccuracy in conventional breakeven inflation calculations is eliminated when the calculation is done correctly using spot rates rather than yields.

The next chapter discusses bond sets to be used by XRM to estimate price equations and yield curves.

4. Bond Structure

The starting point for the application of the XRM methodology is the choice of the set of bonds for which the yield curve will be done. Based on the characteristics of this bond set, XRM sets up a price equation and estimates the price equation from the bond set. Yield curves and other statistics will be derived from the estimated price equation. For the three yield curves discussed, HQM, TNC, and TRC, there will be three respective bond sets for three price equations. This chapter describes the bond sets.

For each of the three yield curves, the bonds in each bond set are based on what are here termed standard bonds. Standard bonds, defined in detail below, are bonds that look like Treasury coupon issues: they have a fixed maturity, pay a fixed coupon semiannually, and have no additional characteristics such as embedded options. Standard bonds are the foundation class of bonds trading in U.S. bond markets, and other classes of bonds are analyzed and traded relative to standard bonds. By focusing on standard bonds, the resulting yield curve estimates are consistent with fundamental yields and returns that underlie U.S. bond markets. The concept of standard bonds has been developed in the Office of Financial Analysis.

The bond sets for the three yield curves also include nonstandard bonds as discussed below, but at the present time the nonstandard bonds are chosen so as to be closely related to standard bonds. The bond sets avoid nonstandard bonds because nonstandard bonds are frequently priced differently from standard bonds and may require additions to the price equation such as special regression variables. Furthermore, nonstandard bonds often include other features that affect their prices, again requiring additional regression variables. Fortunately, there are more than enough standard bonds and closely related bonds that the bond sets for the yield curves can produce accurate estimates of fundamental bond market behavior.

The first three sections discuss standard bonds and the following sections discuss bonds included in the yield curves. Also discussed are concepts of yield that will apply to the yield curve formulas in Chapter 14.

Standard Bonds

This section sets out the characteristics that define standard bonds. The definition is given in detail because bonds can be complicated, especially corporate bonds, and the bonds used for the yield curves need to be clearly described. In particular, it's necessary to derive from the bond set the essential features of each bond that are needed for the estimation process. The discussion in the rest of this chapter is meant to set out these features conceptually. The discussion applies to all standard bonds.

In considering bond features, it must be stressed that each bond has attached to it an end of month convention that determines whether or not the bond pays cash flows at end of month, and a day count convention that determines how to calculate days between dates. Therefore, in calculating the concepts described below such as length of half-years and accrued interest for a particular bond, the conventions for that bond should be used.The concepts set out here provide the guide to what information is needed from each bond for yield curve estimation.

The two typical day count conventions are actual/actual for Treasuries using actual days and 30/360 or some variant for corporate bonds using a 360-day year and 30-day months. These conventions can give different results in actual computations. It would be too much of a digression to describe these conventions in detail; the references provide details on these conventions as well as more information in general about the material in this chapter.[5]

Standard bonds are straightforward. A standard bond has a fixed accrual date when interest starts to accrue and a fixed maturity date when the principal is repaid and the bond expires. For analysis, the size of the bond is scaled so that the principal equals 100. The timespan from accrual date to maturity date is divided into $\breve{n}$ coupon periods, with the requirement $\breve{n} \geq 3$ to make sure the market trades it as a bond rather than as cash.

Each of the $\breve{n}$ coupon periods has a length in half-years denoted as $\breve{h}_{\breve{\imath}}$, $\breve{\imath} = 1, \ldots, \breve{n}$, where a half-year is the time period of half a year. Regular coupon periods have a length of one half-year, and because all the coupon periods are required to be regular except the first and last, all the intermediate coupon periods have lengths of one half-year: $\breve{h}_{\breve{\imath}} = 1,\ \ \breve{\imath} = 2, \ldots, \breve{n} - 1$.

The first coupon period is said to be regular if its length in half-years $\breve{h}_1 = 1$ too. If it has a different length, it's called odd, and it's short if its length $\breve{h}_1 < 1$ and long if $\breve{h}_1 > 1$. Analogously, if the last coupon period is regular, $\breve{h}_{\breve{n}} = 1$; otherwise the last coupon period is odd, in which case $\breve{h}_{\breve{n}} < 1$ if the last coupon period is short and $\breve{h}_{\breve{n}} > 1$ if it's long.

At one time Treasury securities could have odd first coupon periods, but no longer: all the coupon periods of Treasuries currently being traded are regular. In contrast, corporate bonds can have

[5] References for information about bonds and yields include Garbade (1996), Krgin (2002), and Stigum and Robinson (1996).

odd first or last coupon periods. Also, with rare exceptions the term of a Treasury security is given by $\frac{\breve{n}}{2}$, so a 10-year Treasury note has $\breve{n} = 20$.

Payments from a standard bond are determined by a fixed annual coupon rate κ that is paid per year on each 100 of principal. The coupon rate is strictly positive: $\kappa > 0$. The coupon rate can't be zero because standard bonds pay a coupon, they are not zero coupon. And even though theoretical coupon rates can be negative as discussed in Chapter 14, actual bonds in the market always have a positive coupon.

Based on the coupon rate κ, a standard bond provides a payment $\breve{c}_{\breve{\imath}}$ on the last day of each coupon period, which day is called the coupon payment date or the coupon anniversary date. The amount of the payment is calculated as half the coupon rate for each half-year of the coupon period. In addition, the principal of 100 is paid back on the last day of the last coupon period which is the same as the maturity date. The payment scheme is as follows, taking into account odd coupon periods. The amount of the coupon rate paid for a regular coupon period is called a regular coupon, and the amount paid for an odd coupon period is called an odd coupon which can be a short or long coupon:

$$\breve{c}_{\breve{\imath}} = \breve{h}_{\breve{\imath}} \frac{\kappa}{2}, \breve{\imath} = 1, \dots, \breve{n} - 1 \qquad (4.1a)$$

$$\breve{c}_{\breve{n}} = \breve{h}_{\breve{n}} \frac{\kappa}{2} + 100 \qquad (4.1b)$$

A standard bond has no other features, so this completes the definition of a standard bond. In particular, a standard bond has no embedded options, floating coupons, or anything else.

Price and Cash Flows

Each price equation is estimated using a bond set that is traded on a particular day and has that day's date attached to the price equation. The resulting yield curve is referred to by the same date. However, even though the bond set is traded on that day, the actual payment for and transfer of ownership of the bonds occurs a bit later on the settlement date. As of this writing, the settlement date is one business day after the date of trade for the bond sets discussed here.

Also, for each bond in this discussion, the settlement date must be the same date as or after the issue date of the bond. At the present time, the yield curves including the Treasury yield curves do not include when-issued bonds.

And finally, the issue date of the bond must be the same as or after the accrual date of the bond. This means that interest must start accruing on the bond before or at the same time as the settlement date.

The next step is to examine trading of a standard bond on a particular day with settlement a business day later. The first thing is to lay out the remaining coupon periods of this bond at the time of settlement. If the settlement date is the same as the accrual date, the remaining coupon periods are the same as the original periods. However, if the bond is settled after the accrual date, there could be fewer coupon periods left before the bond's maturity and the first coupon period remaining may be shorter than the original coupon period.

Let the lengths in half-years of the remaining coupon periods be designated as $\tilde{h}_{\tilde{\imath}}, \tilde{\imath} = 1, \dots, \tilde{n}$, with $\tilde{n} \leq \check{n}$.The last $\tilde{n} - 1$ coupon periods are the same as the original, but the length of the first remaining coupon period $\tilde{h}_1$ may be less than the length $\bar{h}$ of the original coupon period that contains it.

The price of this bond excluding accrued interest, called the flat price or clean price, is denoted as p_{FLAT}, with $p_{\text{FLAT}} > 0$. The amount of accrued interest on the bond is $(\bar{h} - \tilde{h}_1)\frac{\kappa}{2}$. Therefore, the full or dirty price of the bond, designated p, is given as:

$$p = p_{\text{FLAT}} + \left(\bar{h} - \tilde{h}_1\right)\frac{\kappa}{2} \tag{4.2}$$

For payments, the traded bond pays n cash flows $c_i, i = 1, \dots, n$, at the end of the coupon periods, with $n = \tilde{n}$. The number of years from the settlement date to the payment of each cash flow is designated by $h_i, i = 1, \dots, n$, given as:

$$h_i = \sum_{i'=1}^{i} \frac{\tilde{h}_{i'}}{2} \tag{4.3}$$

So the payments c_i are assumed to be done on the last days of their respective coupon periods. However, this assumption ignores weekends and holidays, because payments are actually done on the next business day following a weekend or holiday[6]. So to be more precise about payment dates, τ_i can be defined as the number of years to actual payment date given as the number of days from the settlement date to the actual payment date divided by 365.25 representing the average number of days in a year including leap years. The value 365.25. is approximate but close enough for practical applications.

In sum, a traded bond for inclusion in a yield curve is represented by $3n + 1$ values including its price p, the n cash flows c_i, the n years to ends of coupon periods h_i, the n years to payment τ_i, and any additional information needed to define the regression variables. Additional information could include ratings for corporate bonds. Armed with these concepts, the next section computes yields on traded standard bonds.

[6] Because trading of issued bonds is done in private markets, Good Friday is here included as a holiday. See Dershowitz and Reingold (1997) for the mechanics of calculating the date of Good Friday.

Yields

This section sets out bond yield formulas that will be used later in computing yield curves. Note that different presentations use different terminologies for these formulas; the scheme here is meant to be consistent with notation for yield curves later. Following market practice, all these formulas will use semiannual compounding.

The first yield formulation is for y_{CURR} often call current yield. This is included for comparison because current yield doesn't have much use in yield curves:

$$y_{\text{CURR}} = 100\frac{\kappa}{p} \tag{4.4}$$

A more useful concept is yield to maturity, or frequently simply referred to as yield. The following formulas for yield to maturity follow the usual market convention for yield calculations that ignores weekends and holidays and uses half-years. There are several cases requiring different formulas.

To start, consider the case of a zero coupon bond with price p that pays a single cash flow c at the end of h years. This could be a zero coupon bond originally or a standard bond with one payment left. If $h < \frac{1}{2}$, the yield y as a percent is given by the following formula calculated as simple interest:

$$p = \frac{c}{\left(1+2h\frac{y}{200}\right)} \tag{4.5}$$

If $h \geq \frac{1}{2}$, the yield y for this bond is the value that solves:

$$p = \frac{c}{\left(1+\frac{y}{200}\right)^{2h}} \tag{4.6}$$

Note that if $h = \frac{1}{2}$, the two formulas give the same result. Equation (4.6) is the formula for the spot rate from a zero coupon bond that is used for the spot yield curve in Chapter 14.

Equation (4.5) also applies to a standard bond that has only one cash flow left to be paid. Turning to a standard bond with multiple cash flows: each of these cash flows can be viewed as a zero coupon bond and the yield for each cash flow can be set such that the prices of the cash flows sum to p. If the yields for the cash flows are constrained to be the same, the yield becomes in some sense the average yield for the whole bond. This is the yield to maturity, and it is the value of y that solves the following:

$$p = \sum_{i=1}^{n} \frac{c_i}{\left(1+\frac{y}{200}\right)^{2h_i}} \tag{4.7}$$

Equation (4.7) is typically applied to corporate bonds that use a 30/360 day count, and this formula is often called the street convention for calculating yield. Note again that if $n = 1$ and $h_1 = \frac{1}{2}$, Equations (4.5) and (4.7) give the same result. Equation (4.7) is used to compute the par yield curve in Chapter 14.

A different convention for calculating yield is the Treasury convention. As its name implies, it is used for Treasury securities both nominal and TIPS. For securities with one payment and at most a half-year to maturity, the same formula Equation (4.5) is used in the Treasury convention. But for other securities, Equation (4.7) is replaced with the following equation that uses simple interest up to the first half-year. For this equation, define $h_0 = h_1$ if $h_1 \leq \frac{1}{2}$ and $h_0 = h_1 - \frac{1}{2}$ otherwise. The case where $h_1 > \frac{1}{2}$ is for long first coupons, and these no longer exist in Treasury securities currently being traded although they did exist historically.

$$p = \sum_{i=1}^{n} \frac{c_i}{\left(1+2h_0\frac{y}{200}\right)\left(1+\frac{y}{200}\right)^{2(h_i-h_0)}} \tag{4.8}$$

Note that Equation (4.7) and Equation (4.8) are the same if $h_1 = \frac{1}{2}$, which is the assumption for the par and spot yield curves derived in Chapter 14.

The formulas so far have used half-year coupon periods and ignored weekends and holidays. There is another concept of yield called true yield y_{TRUE} that uses the actual years to cash flow payments given in the τ_i defined above. The formulas are entirely analogous to Equations (4.5) and (4.7). Here is the formula equivalent to Equation (4.5) for bond with maturity τ:

$$p = \frac{c}{\left(1+2\tau\frac{y_{\text{TRUE}}}{200}\right)} \tag{4.9}$$

And here is the equivalent to Equation (4.7):

$$p = \sum_{i=1}^{n} \frac{c_i}{\left(1+\frac{y_{\text{TRUE}}}{200}\right)^{2\tau_i}} \tag{4.10}$$

For all the formulas set out above, it is clear that the cash flows in the numerator on the right-hand side are positive because $\kappa > 0$. Moreover, the yield can be anything above -100 percent, so a solution exists for each formula. For nominal bonds, it should always be true that $p < \sum_{i=1}^{n} c_i$ with the result that the yield has to be positive. However, for TIPS this doesn't have to be true and the yield can be nonpositive.

These formulas show that there are two ways to compute yield: conventional yield using half-years ignoring weekends and holidays, and true yield using actual years to payments. The price equation estimation uses the second approach: true yield and actual years to payments. In contrast, once having estimated the price equation using true yield, the computations of the spot and par yield curves use the market convention of half-years.

Bond Set for TNC and TRC

Having defined standard bonds, it's now possible to list which bonds are included in the bond set for estimating the TNC and TRC Treasury yield curves. As much as possible, all Treasury coupon issues trading in the market are included in the computation of these yield curves. Prices for estimation are bid prices.

However, coupon issues that have only one cash flow left or for whom the maturity $\tau_n \leq \frac{1}{2}$ are excluded because they are so short-term that they are part of the cash market and are no longer traded as bonds. And Treasury bills and floating rate notes are excluded because they are not coupon issues.

All other Treasury coupon issues are included of whatever maturity. In particular, for TNC, on-the-run and first off-the-run coupon issues are included with special dummy variables. Also, in historical data, Treasury callable bonds and flower bonds are included with special regression variables even though they're not standard bonds. These securities are discussed in Chapter 12.

For TIPS, the concept of on-the-run doesn't apply nor are there any nonstandard bonds. So all TIPS coupon issues are included except the very short-term issues as mentioned above.

Bond Set for HQM

The HQM bond set also includes as many standard bonds as possible. Same as Treasuries, prices for estimation are bid prices.

All bonds for HQM must be nominal corporate bonds issued by U.S. corporations and denominated in dollars. The bonds must be rated A, AA, or AAA by nationally recognized statistical rating organizations.

Analogous to Treasuries, HQM bonds that have only one cash flow left or for whom the maturity $\tau_n \leq \frac{1}{2}$ are excluded. Also excluded are bonds with maturity $\tau_n > 30$; there are only a few bonds with maturity greater than 30 years, not enough to extend reliably the maturity ranges beyond 30 years. Maturities below 1 year are currently filled in with Federal Reserve AA financial and AA nonfinancial commercial paper rates. To ensure sufficient liquidity, each of the corporate bonds in HQM must have a minimum size in terms of par amount outstanding: the current minimum is $250 million.

At the present time, callable bonds with a call schedule are excluded, while make whole calls are included. Putable bonds and bonds with sinking funds are also excluded. All these bonds are nonstandard bonds.

However, there is an exception in that bonds with a single call date within the final year before maturity are included in HQM. These bonds are here called end calls, bonds with an end-call date. End calls have become very important in bond markets starting around 2016, and they need to be included to capture market behavior.

In principle, end calls are nonstandard bonds because the presence of the call date could affect the bond's price. This might mean that another regression variable is required to pick up the end-call effects on price. However, so far there has been no indication that the end call is actually affecting price.

As a result, the presence of the end call can be ignored in yield curve estimation and doesn't need to be accounted for, and the end-calls are considered to be standard bonds that can be put in the price equation without adjustment for the call. End-calls were introduced into the HQM yield curve in February 2024.

Other nonstandard bonds are excluded from HQM: A bond that does not have a fixed coupon is excluded; in particular, bonds with floating coupons are excluded because their interest payments are unpredictable and they therefore provide little if any information about rates of return over future time periods. Bonds without a fixed maturity date when the principal is to be returned are also excluded. And convertible bonds are excluded because their price depends on the companion equity and they do not have an unambiguous price for their cash flows. Bonds that are capital securities or hybrid preferred stock are excluded, as are bonds issued by a government-sponsored enterprises. Asset-backed bonds are also excluded.

The discussion in this chapter has described the sets of bonds for XRM yield curve estimation. Starting with the next chapter, the price equation is constructed that will use these bond sets.

5. The Discount Function and Forward Rate

This chapter describes the discount function and the companion forward rate. The discount function is the foundation of the price equation, and using the discount function as described in this chapter, the price equation can be written down conceptually in the next chapter. The price equation will be estimated using the set of bonds described in the previous chapter, and the estimation results will be used to compute yield curves and related statistics.

The discount function is a general concept in finance, and this discussion stresses the particular implementation of the discount function for the XRM methodology.

The Discount Function for Nominal Bonds

In this section, the discount function for nominal bonds is described including the HQM and TNC yield curves whose payments are in dollars. The next section will discuss the discount function for TIPS that are used in the TRC yield curve and whose payments are in real terms.

Nominal bonds have a nominal discount function that measures economic time preference. Time preference means that payments received earlier in time are more valuable. The reason why an earlier payment is better is because the earlier payment can earn interest or be used for other purposes, its real value will not decline as much if there's inflation, and in general a shorter period is less risky. The valuation of cash payments at different future times is basic to bond pricing.

The discount function $\delta(\tau)$ pertains to the settlement date for which the yield curve is being calculated and maturities are computed forward from that point. The discount function $\delta(\tau)$ is defined as follows: for each future maturity in years $\tau \geq 0$, $\delta(\tau)$ gives the present market price of $1 that will be received τ years in the future. Because of time preference, the market price is less than $1, and furthermore the price declines as τ rises because payments further out in time are less valuable. Written conversely, the discount function implies that the payment to be received in τ years for a present investment of \$1 is $\frac{1}{\delta(\tau)}$.

Each sector of the bond market has its own discount function reflecting the characteristics of cash flows in that sector. Specifically, high quality bonds in the HQM yield curve have a different discount function than nominal Treasury securities in the TNC yield curve. In particular, the HQM discount function is usually below the TNC discount function because future cash flows from high quality corporate bonds with default risk are worth less than Treasury cash flows without default risk, although transient factors such as liquidity can cause HQM to be above TNC at short maturities around one year.

The discount function is a theoretical concept because typically there isn't any market for future cash flows based on time preference. So the discount function can't be traded directly but is embedded in bond prices. While it's true that some bond sectors such as Treasuries trade spot or zero coupon bonds, these bonds don't represent the discount function directly because their prices may be influenced by factors other than time preference. The regression variables discussed later pick up these factors.

Based on this discussion, the features of the discount function are specified as follows. First, the discount function starts at unity at the present time $\tau = 0$ because \$1 paid at present gives back \$1. The discount function must be positive throughout. And the discount function declines because payments received later are worth less:

$$\delta(0) = 1;\ \delta(\tau) > 0 \tag{5.1}$$

$$\frac{\mathrm{d}\delta(\tau)}{\mathrm{d}\tau} < 0 \tag{5.2}$$

The Discount Function for the TRC Yield Curve

TIPS securities in the TRC yield curve are real bonds, and they have a real discount function that gives the present market price in real terms to receive \$1 in real terms in the future. This discount function has the same features as Equation (5.1).

However, in contrast to nominal discount functions, the TRC real discount function doesn't necessarily decline and can even be greater than unity for $\tau > 0$. The reason for this is that the TRC discount function must be consistent with the TNC discount function. To illustrate, let $\delta_{\mathrm{TNC}}(\tau)$ and $\delta_{\mathrm{TRC}}(\tau)$ be the TNC and TRC discount functions and let p_τ and p_0 be the expected price level at maturity τ and the actual price level at present. Then the following relationship holds approximately for the two discount functions:

$$\delta_{\mathrm{TNC}}(\tau)\frac{p_\tau}{p_0} \approx \delta_{\mathrm{TRC}}(\tau) \tag{5.3}$$

The left-hand side of this equation is the price to receive \$1 in real terms at maturity τ formulated using the TNC discount function. This must approximate the equivalent price from the TRC discount function, with differences reflecting such things as different market conditions between real and nominal bonds and inflation uncertainty. This equation shows that as expected inflation and $\frac{p_\tau}{p_0}$ rises, increasing amounts of dollars must be invested at present to get the same real \$1 in the future, with the result that $\delta_{\text{TRC}}(\tau)$ must rise too maybe even above unity if $\delta_{\text{TNC}}(\tau)$ does not decline enough. This is the situation when nominal interest rates don't rise as fast as inflation causing real interest rates to be negative.

The Forward Rate

This section defines the forward rate associated with the discount function. The definition will show that the discount function and the forward rate are in a sense the inverse of each other and each one implies the other.

To define the forward rate, assume that there is a loan of \$1 which using the discount function promises to pay $\frac{1}{\delta(\tau_1)}$ at maturity τ_1. If this loan is extended for a short time to $\tau_2 > \tau_1$, the interest rate on this extension would be:

$$\frac{\frac{1}{\delta(\tau_2)}-\frac{1}{\delta(\tau_1)}}{\frac{1}{\delta(\tau_1)}(\tau_2-\tau_1)} = -\frac{\delta(\tau_2)-\delta(\tau_1)}{\delta(\tau_2)(\tau_2-\tau_1)}$$

As τ_2 approaches τ_1, this interest rate approaches the forward rate, with the result that the forward rate $\phi(\tau)$ is defined as follows:

$$\phi(\tau) = -\frac{\mathrm{d}\delta(\tau)}{\mathrm{d}\tau}\frac{1}{\delta(\tau)} = -\frac{\mathrm{dlog}(\delta(\tau))}{\mathrm{d}\tau} \qquad (5.4)$$

Note that the forward rate is instantaneous, that is, the period of time for which the loan is extended is infinitesimal. Moreover, as given in this definition, the forward rate is a simple interest rate expressed in decimals: for analysis and comparison with other interest rates, it can be converted to a percentage by multiplying by 100.

The definition shows that for nominal bonds where the discount function is declining, the forward rate is always positive. However, for TIPS, the forward rate is negative when the discount function is rising, and it is zero when the discount function is unchanged.

The definition gives the forward rate as the relative curvature of the discount function. Consequently, the forward rate is higher when the discount function is falling more rapidly indicating a faster increase in time preference. This implies that the forward rate is high at a particular maturity when markets see less opportunity and more risk at that maturity. Therefore, at each maturity, the forward rate summarizes in a single number the market views about risks and rewards at that maturity, and as a result, forward rates can be compared both over time and across different maturities. For this reason, it's usually easier to estimate the price equation using the forward rate and derive the discount function from it rather than estimate the discount function directly. This is the approach followed in XRM.

By integrating each side of Equation (5.4) and using the fact that $\log(\delta(0)) = 0$, the following inverse formula is obtained for the discount function in terms of the forward rate:

$$\delta(\tau) = \exp\left(-\int_{\alpha=0}^{\tau} \phi(\alpha)\mathrm{d}\alpha\right) \tag{5.5}$$

When the price equation is estimated using the forward rate, this is the formula that will be used to derive the discount function.

Constant Forward Rate

When the forward rate is constant and positive for all maturities beyond a certain maturity, the discount function becomes an exponential decline. This fact is useful because as discussed in Chapter 9 the forward rate will be held constant at the long term forward rate ϕ^* in the projection range for $\tau \geq \tau^*$, where τ^* is 30 years maturity.

If the forward rate is a constant $\phi^* > 0$ starting at τ^*, Equation (5.5) takes the following form for $\tau \geq \tau^*$. In this equation, the discount function has an asymptote at zero. If $\tau^* = 0$ so that the forward rate is constant for all maturities, the discount function is an exponential decline throughout.

$$\begin{aligned} \delta(\tau) &= \exp\left(-\int_{\alpha=0}^{\tau^*} \phi(\alpha)\mathrm{d}\alpha\right) \times \exp\left(-\int_{\alpha=\tau^*}^{\tau} \phi^*\mathrm{d}\alpha\right) \\ &= \exp\left(-\int_{\alpha=0}^{\tau^*} \phi(\alpha)\mathrm{d}\alpha\right) \times \exp(-\phi^*(\tau - \tau^*)) \end{aligned} \tag{5.6}$$

Normally ϕ^* is positive for nominal yield curves. It's usually positive for the TRC yield curve too. However, on about three dozen days over the last 20 years it has been negative, all of which were in the midst of the COVID-19 pandemic. For $\tau \geq \tau^*$, if it's negative with $\phi^* < 0$, the discount function rises exponentially indefinitely. This situation reflects temporary market anomalies because the discount rate reflecting time preference should eventually decline.

If $\phi^* = 0$, the discount function remains constant throughout the projection range. Such a situation is also abnormal in markets because it implies the absence of time preference. For all days starting in 2003 forward for which the TRC yield curve is done, there's hasn't been one day in which $\phi^* = 0$, although there have been times when it was near zero.

The Discount Spot Rate

The discount function can be converted into an interest rate called the discount spot rate that gives the return from investing in the discount function: that is, for $1 invested now the discount function provides $\$\frac{1}{\delta(\tau)}$ at maturity τ. The discount spot rate is distinct from the spot rate in the spot yield curve discussed in Chapter 14, because the latter spot rate includes the effects of regression variables. The discount spot rate concept was developed in the Office of Financial Analysis.

Normally it's not possible in actual markets to invest in the discount function. But the concept of the discount spot rate makes it easier to visualize the discount function in comparison to other interest rates, and is useful in analyzing the asymptotic properties of the actual spot rate implied by the market in spot yield curves.

The discount spot rate $r_{\mathrm{D}}(\tau)$ as a function of τ is given by the following formula for the discount function and the forward rate. In the formula, $r_{\mathrm{D}}(\tau)$ is expressed as a percentage, and following standard market practice, the compounding is semiannual as indicated by the number 2 in the exponent and the denominator.

$$\left(1+\frac{r_{\mathrm{D}}(\tau)}{200}\right)^{2\tau}=\frac{1}{\delta(\tau)} \tag{5.7a}$$

$$\Rightarrow r_{\mathrm{D}}(\tau)=200\times\left(\left(\frac{1}{\delta(\tau)}\right)^{\frac{1}{2\tau}}-1\right) \tag{5.7b}$$

$$=200\times\left(\exp\left(\frac{\int_{\alpha=0}^{\tau}\phi(\alpha)\mathrm{d}\alpha}{2\tau}\right)-1\right)$$

From this formula, the derivative of the discount spot rate is given as:

$$\frac{\mathrm{d}r_{\mathrm{D}}(\tau)}{\mathrm{d}\tau}=200\times\exp\left(\frac{\int_{\alpha=0}^{\tau}\phi(\alpha)\mathrm{d}\alpha}{2\tau}\right)\times\left(\frac{\phi(\tau)}{2\tau}-\frac{\int_{\alpha=0}^{\tau}\phi(\alpha)\mathrm{d}\alpha}{2\tau^2}\right) \tag{5.8}$$

The sign of this derivative is determined by the expression:

$$\phi(\tau)-\frac{\int_{\alpha=0}^{\tau}\phi(\alpha)\mathrm{d}\alpha}{\tau}$$

Therefore, the discount spot rate at any maturity is rising when the forward rate at that maturity is greater than the forward rate average up to that maturity. One implication is that the curve of forward rates can have a hump, but this does not necessarily imply that the discount spot rate must also have a hump.

As discussed later, the forward rate is assumed to settle down to the long-term forward rate ϕ^* for maturities greater than τ^* which equals 30 years. Inserting this assumption into Equation (5.7b) for $\tau>\tau^*$ gives:

$$r_{\mathrm{D}}(\tau)=200\times\left(\exp\left(\frac{\int_{\alpha=0}^{\tau^*}\phi(\alpha)\mathrm{d}\alpha+\int_{\alpha=\tau^*}^{\tau}\phi^*\mathrm{d}\alpha}{2\tau}\right)-1\right) \tag{5.9}$$

This equation implies that as τ goes to infinity, the discount spot rate converges to the long-term discount spot rate r_{D}^*:

$$r_{\mathrm{D}}^* = 200 \times \left(\exp\left(\frac{\phi^*}{2}\right) - 1\right) \tag{5.10a}$$

This equation is the same as:

$$\left(1 + \frac{r_{\mathrm{D}}^*}{200}\right)^2 = \exp(\phi^*) \tag{5.10b}$$

Therefore, r_{D}^* and ϕ^* are the same interest rate except that r_{D}^* is semiannually compounded following market convention and expressed as a percent while ϕ^* is continuously compounded because it's based on the instantaneous forward rate and is expressed as a decimal. So the discount spot rate converges to the long-term forward rate.

However, a caveat here is that the convergence may be of limited practical use because even though the formulas show that convergence eventually occurs, it may not be achieved or come close enough even within the long timeframe of 100 years maturity of the projection range. The speed of convergence depends on the bond data at a particular time. Therefore, the convergence may be more of an indicator of general tendency rather than a number that can be used for decisions in markets.

6. The Price Equation

This chapter uses the discount function and companion forward rate from the previous chapter to write down the conceptual form of the price equation. Subsequent chapters will specify the mathematical form of the discount function and the price equation for estimation using the set of bonds set out in Chapter 4.

Chapter 4 shows that each bond used in estimation is described by the bond price p including accrued interest at the settlement date, n cash flows c_i plus years from settlement τ_i when the cash flows are received, $i = 1, \dots, n$, and any additional information required to compute the regression variables. The price equation models the bond price as a function of the future cash flows and the additional information.

The first section of this chapter states the purpose of the price equation. The next section sets out the conventional price equation for yield curve analysis using the discount function alone. The third section discusses the extended price equation used by XRM.

Purpose of the Price Equation

The purpose of the price equation is to model the prices of the types of bonds discussed in Chapter 4 as functions of cash flows and other variables. The resulting price equation when estimated provides information about yield curves and other statistics characterizing the bond market. There are separate price equations for HQM, TNC, and TRC yield curves.

The price equation is broken into two parts. The first part shows the effects of time preference on the cash flows from the bonds in determining the bond price. The second part shows the effects of other factors on the price separate from the effects of time preference. These other factors are measured by regression variables in the price equation.

It's important to realize that the price equation must be constructed so as to represent the standard bond set in Chapter 4 even if the equation is used to estimate results for other types of bonds. This is because the standard bond set is used for estimation, so to get accurate estimates, the price equation must pertain to standard bonds. Calculations for nonstandard bonds, such as spot rates for zero coupon bonds, will then be consistent with the standard bonds.

For example, the HQM price equation pertains to standard high quality corporate bonds with coupons as described in Chapter 4. However, the HQM yield curve is typically used to estimate spot or zero coupon rates for pension discounting and other purposes even though a fully developed corporate bond zero coupon market does not exist. By calculating the spot rates from the HQM price equation, the spot rates are consistent with standard corporate bonds with coupons and therefore with the bond market as it currently exists, and the spot rates show what a corporate bond zero coupon market would look like if it did exist along with standard corporate bonds.

Conventional Price Equation

Conventional yield curve analysis views a bond as made up of n separate zero coupon bonds paying respectively c_i at τ_i. Given the discount function, the price of each of these zero coupon bonds based on time preference can be computed by applying the discount function to the cash flows, and all these prices can be summed to get the total price of the bond:

$$p = \sum_{i=1}^{n} \exp\left(- \int_{\alpha=0}^{\tau_i} \phi(\alpha) \mathrm{d}\alpha\right) c_i \qquad (6.1)$$

This formulation is the conventional price equation typically used by yield curve approaches. This equation only includes effects of time preference. The bulk of the research effort in yield curve analysis over the last several decades has been devoted to devising a workable functional form for the discount function.

The point to be emphasized is that by including only the discount function, the conventional price equation mixes in other effects on price separate from time preference with the time preference effects. For example, effects that frequently generate a hump in yields around 20 years maturity that are examined in Chapter 10 are forced to be in the discount function although these effects are separate from time preference.

By mixing these effects together, the resulting estimates from the conventional price equation are imprecise. Consequently, it's necessary to include regression variables in addition to the discount function. Regression variables are included in the XRM methodology and are shown in the extended bond price equation set out in the next section.

Furthermore, all the bonds used for estimation in the conventional price equation have to be homogeneous: there's no provision for bonds that have different features because there's no way to account for them. This is another reason why regression variables are needed: for example, the HQM yield curve includes corporate bonds rated A, AA, and AAA, and the differences in price among these different ratings are picked up by regression variables.

Another limiting aspect of the conventional price equation is that it doesn't have any constraints on the forward rate. Constraints are required to ensure stability of the forward rate and to enable projection of the forward rate beyond 30 years maturity.

Extended Price Equation

The following is the extended price equation including regression variables used by XRM:

$$p = \sum_{i=1}^{n} \exp\left(-\int_{\alpha=0}^{\tau_i} \phi(\alpha)\mathrm{d}\alpha\right) c_i + \sum_{j=1}^{m} \theta_j x_j \qquad (6.2)$$

This equation has the same discount function as the conventional price equation but adds m regression variables x_j with coefficients θ_j, $j = 1, \dots, m$. Therefore, this equation has the two parts discussed above: the discount function for time preference and regression variables for effects on price separate from time preference. And it is assumed that the discount function has constraints that allow projection beyond 30 years maturity.

Note that in this conceptual formulation of the extended price equation, the characteristics of the regression variables aren't specified. This is because in general the characteristics can be anything, the variables aren't limited. In particular, the value of a regression variable for a bond is often a function of the time to maturity of the bond, as is the case with the hump variable, but this is not made explicit in the formulation.

It also should be noted that the regression variables are linear and added linearly to the discount function. In principle the regression variables could be nonlinear, and it might be useful to develop nonlinear variables in some circumstances. Here the linearity is an approximation, as it is for most regressions, although the results below will show that it works well.

Another point regarding linearity is that linearity assumes that the regression variables don't vary too much from a mean so that any nonlinear affects can be adequately approximated by linearity. This seems to be true for the variables included in the yield curves in this discussion.

Also, it should be noted that there's often a lot of noise in bond data so the signal to noise ratio can be low. While it may be possible to develop an elaborate nonlinear formulation that picks up subtle bond effects, it may be impossible to estimate such a formulation given all the randomness in the data. In fact, such a formulation may end up generating spurious random estimates. So it may be better to stick with a straightforward linear formulation which is the best that can be done with the data. In general, given bond data limitations, it's often better to stick with simple regression variables.

The next step for estimating the price equation is to determine the mathematical form of the discount function, and after that to specify the regression variables. The mathematical form of the discount function based on a spline over maturity ranges is described in Chapters 7-9. The regression variables currently used in the yield curves are described in Chapters 10-12.

7. Maturity Ranges

This chapter describes the maturity ranges that are the basis of the functional form of the discount function and forward rate for estimating the price equation. Using the maturity ranges, the following two chapters then formulate the spline function that represents the discount function and forward rate.

The Forward Rate Spline

The functional form must model the forward rate over all the maturities zero from 30 years that cover the bond set used in estimation. The XRM methodology develops this functional form by dividing these maturities into ranges.

The idea of maturity ranges is that at any time bond trading can be divided into maturity ranges such that trading in each range reflects market views of the rewards and risks of the bonds in that range. Therefore, trades in each range are related, because the trades reflect a common average market opinion at that time. As a result, values of the forward rate within each range, which are indicators of risk and reward, are related.

Once the maturity ranges are chosen, the forward rate in each range could be approximated by a constant that equals the forward rate average in that range. However, there can be significant movement in the forward rate within a range that needs to be captured so that the forward rate is accurately represented. This is especially true insofar as the last range is big. Moreover, fixed averages do not connect across ranges, and in order to get a well-behaved representation of the forward rate, the values of the forward rate in the ranges must join together smoothly.

A smooth forward rate can be obtained without departing too far from averages, and without introducing excessive volatility, by using a cubic polynomial to represent the forward rate values over the maturities in each range. The degree of the polynomial is set at cubic, because cubic polynomials exhibit sufficient flexibility to capture market movements of the forward rate in a range while having a small number of parameters to be estimated.

These cubic polynomials, one for each range, are strung together in a smooth fashion such that they are continuous with continuous first and second derivatives. The result is a smooth piecewise cubic polynomial that represents the forward rate over all maturities and that can be estimated from the bond set. But in fact this is the definition of a cubic spline, with the spline knot points being the maturities that delineate the maturity ranges.

Therefore, in XRM, the use of the cubic spline arises naturally from the smooth joining of the forward rate across the maturity ranges, and the choice of knots for the spline comes from the maturity ranges themselves. The choice of a cubic spline is not imposed upon the forward rate but is justified by the maturity ranges. Consequently, in contrast to conventional yield curve approaches, rather than choosing a spline initially, XRM focuses on maturity ranges and the choice of the spline arises from the ranges.

In addition, conventional yield curves that use splines and that don't have maturity ranges have difficulty picking the knots for the splines. Usually the choice of knots in such approaches is based on some rule of thumb that may depend on characteristics of the bond set itself. Consequently, the knots are arbitrary, and different knots can give different results. Moreover, choosing knots based on the set of bonds itself can be unstable because it's circular: rather than fitting a model to a set of bonds, the model itself depends on the bonds.

In contrast, the choice of knots in XRM is done before estimation and comes from the maturity ranges separate from the set of bonds used in the estimation. Therefore, the knots aren't arbitrary. And fixing the knots in advance of estimation provides exceptionally stable numerical results over time regardless of market conditions. Moreover, in the XRM methodology with fixed knots, results from different bond sets can be directly compared using the same model, where the bond sets can be from the same day or from different days.

A spline is a flexible mathematical formulation that is often used to model statistical functions. The next chapter will write out in detail the B-spline polynomials that are used to set up the forward rate spline and for estimation. The spline is chosen to be third-degree cubic because cubic polynomials provide plenty of flexibility without excessive complication or spurious movements often seen with higher degrees. Because maturities of the set of bonds used in estimation range from ½ year to 30 years, the span of maturities over which the spline is calculated is taken to be from zero to 30 years maturity.

The Maturity Ranges

The previous section set out the concept of maturity ranges. This section defines the maturity ranges used in the yield curves. The three yield curves presented here, HQM, TNC, and TRC, use the same maturity ranges.

The method of choosing the maturity ranges focuses on the set of central maturities that market observers watch as indicators of the state of the bond market. Bond traders develop the pricing and trading strategy for each bond by considering the central maturity to which it is closest. So it is reasonable to define the maturity ranges in which the forward rates and bond trades are related as ranges around the central maturities.

The central maturities are 1 year, 2 years, 5 years, 10 years, and 30 years, so each one of these maturities gives rise to a maturity range. In addition, there are subsidiary central maturities of 3 years, 7 years, and 20 years that help determine the endpoints of the ranges.

The reason why these maturities are central is that bond market analysis and commentary consistently use these maturities to describe and discuss bond market characteristics and have been doing so for many decades. This shows the importance of these maturities in the minds of bond market participants. Moreover, the fact that these maturities are central is supported by the fact that Treasury notes and bonds are issued at these maturities, including the Treasury bill at 1 year.

Therefore, there is a total of five maturity ranges, each one based on a central maturity that anchors the bond trading in that range and consists of a span of maturities around the central maturity. Here are the maturity ranges:

- The first maturity range comprises the shortest bonds. This range is anchored by the short-run maturity of 1 year and runs from zero through 1½ years maturity. However, as noted in Chapter 4, bonds with maturities less than ½ year are not used in the yield curves, so this range really starts at ½ year. And it ends at 1½ years which is halfway between the central maturity of 1 year and the next central maturity of 2 years.
- The second maturity range also comprises bonds that are still considered short-term. This range is anchored by the central maturity of 2 years and runs from 1½ years through the central maturity of 3 years. The 3-year maturity is considered by markets as the borderline between short-term bonds and the medium-term central maturity of 5 years.
- The third maturity range is medium-term. This range is anchored by the central maturity of 5 years and runs from 3 years through the central maturity of 7 years for a span that's 2 years on either side of the 5-year point. Analogous to the second maturity range, the 7-year maturity is considered to be the borderline between medium-term bonds and the central maturity of 10 years. The third range includes the 4-year maturity which used to be a central maturity in the Treasury market but is no longer issued.
- The fourth maturity range covers the core maturities of bonds, that is, bonds that are most important in market trading and that are neither short-term nor long-term. This range is anchored by the central maturity of 10 years, which is the most important maturity point that reflects bond market conditions, and runs from 7 years through 15 years. The 15-year maturity is halfway between the 10-year maturity and the long-term central maturity of 20 years, so the 15-year maturity can be considered the borderline for long-term bonds.

- Finally, the last or farthest maturity range encompasses long-term maturities from 15 years maturity up through the central maturity of 30 years and includes the central maturity of 20 years. Because the farthest maturity is 30 years for the set of bonds used in estimation, this last range plus the first four ranges include the entire bond set. The last range is a big range, and additional movement of the forward rate within this range will be picked up by the hump variable discussed later. Special note: for Treasury securities, this last range goes up a bit higher to 30.51 years maturity to encompass all recent Treasury securities. This will still be referred to as 30 years maturity in the discussion.
- In addition to these five maturity ranges at 30 years maturity and below, there is also a projection range of 30 years maturity through 100 years maturity. Because there are insufficient numbers of bonds beyond 30 years maturity to estimate the forward rate in the projection range, a fixed long-term forward rate estimated through 30 years maturity is used to provide yield data out through 100 years maturity. The projection range is discussed in Chapter 9.

From this discussion it follows that the maturities delineating these ranges are the six points $[0,1½,3,7,15,30]$. Therefore, these will be the knots for the cubic spline that represents the forward rate $\phi(\tau)$. The next chapter will show how to use B-splines to set up this spline for estimation.

Finally, there is the question whether these maturity ranges need to change over time. The XRM methodology allows for changes if needed. However, as a matter of fact, these maturity ranges have worked well for both Treasury and high quality and investment grade corporate bonds for over half a century. This has been true largely because the central maturities used to construct these ranges have been chosen to be the same as the maturities of Treasury issues, and because Treasury maturities have largely defined central maturities in the corporate bond market and, of course, in the Treasury market. And it should be noted that one consequence of having the same maturity ranges across Treasury and corporate bond markets and over time is that the estimated spline coefficients form a time series that can be compared over the decades.

8. B-Splines

The previous chapter showed that knots $[0,1½,3,7,15,30]$ derived from the maturity ranges will be used to generate the cubic spline that represents the forward rate. This chapter describes the mechanics for setting up this spline for estimation.

The spline discussion presented here is somewhat different from spline textbooks, so the methodology is written out in some detail. It's important to understand the features of the spline methodology, because in the next chapter the spline will be modified by constraints that aren't part of typical spline applications. The references contain a fuller exposition of spline mathematics; this chapter stresses aspects of splines that are relevant for the yield curves and is not a complete exposition of splines.[7]

The XRM methodology is based on B-splines or basis splines for constructing the spline for the forward rate. The B-splines are computed from the maturity range knots $[0,1½,3,7,15,30]$ and the forward rate spline is expressed as a linear combination of the resulting B-splines. The coefficients of the linear combination are estimated using the set of bonds.

The next section defines cubic B-splines and sets out the cubic polynomial equations that comprise them. The section after that presents the set of B-splines derived from the maturity ranges. The expositions in this chapter and the next will be in terms of general sets of knots, and after each stage the results will be applied to the specific knots from the maturity ranges that are actually used for the yield curves.

---

[7] Treatment of B-splines can be found in de Boor (1978) and Piegl and Tiller (1997).

B-Spline Definition

The usual approach for describing B-splines is to use a general recursion formula that applies to splines of any degree. However, here the degree is already chosen to be cubic, so it's possible to write out directly the algebra of the cubic spline polynomials. Furthermore, although the algebra is extensive, it's easier than general recursions. So that's the approach used here.

To build the definition of B-splines, the first step is to recognize that a B-spline is completely determined by five knots u_a, u_b, u_c, u_d, u_e.The five knots are sorted in ascending order, and they can be distinct or some of the knots can have duplicate values. In principle, any of the knots can be duplicated, but for the yield curves, it's enough to assume that only the first or the last is duplicated and at least two of the five knots are distinct. As examples, the set of five knots [3,7,15,30,30] has the last knot duplicated once and [0,0,0,0,1½] has the first knot duplicated three times.

The next step in defining B-splines is to construct the following cubic polynomials from the five knots. The polynomials are functions of τ which designates maturity. If some of the knots are duplicated, the respective polynomial doesn't exist because of division by zero, but that has no effect on the resulting B-spline because that polynomial is ignored in subsequent calculations as shown in the definition.

$$B^{\mathrm{I}}(\tau; u_a, u_b, u_c, u_d, u_e) = \frac{(\tau-u_a)^3}{(u_d-u_a)(u_c-u_a)(u_b-u_a)} \tag{8.1a}$$

$$B^{\mathrm{II}}(\tau; u_a, u_b, u_c, u_d, u_e) = \frac{(\tau-u_a)^2(u_c-\tau)}{(u_d-u_a)(u_c-u_a)(u_c-u_b)} + \frac{(\tau-u_a)(u_d-\tau)(\tau-u_b)}{(u_d-u_a)(u_d-u_b)(u_c-u_b)} + \frac{(u_e-\tau)(\tau-u_b)^2}{(u_e-u_b)(u_d-u_b)(u_c-u_b)} \tag{8.1b}$$

$$B^{\mathrm{III}}(\tau; u_a, u_b, u_c, u_d, u_e) = \frac{(\tau-u_a)(u_d-\tau)^2}{(u_d-u_a)(u_d-u_b)(u_d-u_c)} + \frac{(u_e-\tau)(\tau-u_b)(u_d-\tau)}{(u_e-u_b)(u_d-u_b)(u_d-u_c)} + \frac{(u_e-\tau)^2(\tau-u_c)}{(u_e-u_b)(u_e-u_c)(u_d-u_c)} \tag{8.1c}$$

$$B^{\mathrm{IV}}(\tau; u_a, u_b, u_c, u_d, u_e) = \frac{(u_e-\tau)^3}{(u_e-u_b)(u_e-u_c)(u_e-u_d)} \tag{8.1d}$$

Using these polynomials, the definition of a B-spline $B(\tau; u_a, u_b, u_c, u_d, u_e)$ is as follows:

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = 0;\ \tau < u_a \tag{8.2a}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = B^{\mathrm{I}}(\tau; u_a, u_b, u_c, u_d, u_e);\ u_a \le \tau < u_b \tag{8.2b}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = B^{\mathrm{II}}(\tau; u_a, u_b, u_c, u_d, u_e);\ u_b \le \tau < u_c \tag{8.2c}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = B^{\mathrm{III}}(\tau; u_a, u_b, u_c, u_d, u_e);\ u_c \le \tau < u_d \tag{8.2d}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = B^{\mathrm{IV}}(\tau; u_a, u_b, u_c, u_d, u_e);\ u_d \le \tau < u_e \tag{8.2e}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = 0;\ \tau \ge u_e;\ u_b \ne u_e \tag{8.2f}$$

$$B(\tau; u_a, u_b, u_c, u_d, u_e) = 1;\ \tau \ge u_e;\ u_b = u_e \tag{8.2g}$$

This definition is set up so that if two knots are duplicated, the respective polynomial is omitted. For example, if $u_a = u_b$, Line (8.2b) can't be true so it's omitted.

Line (8.2g) takes care of the special case in which the last four knots are the same. In that case the B-spline is given by B^{I} alone with the value of unity at the last knot u_e and at all maturities greater than u_e. This, in fact, is the configuration of the final B-spline from the maturity ranges, as shown in the next section. Setting the value of the final B-spline to unity beyond u_e is different from the usual definition of B-splines in which all B-splines are set to zero beyond the last knot. Setting to unity is important for the long-term forward rate as will be clear in the next chapter.

B-Splines for the Forward Rate

In this section, the definition of B-splines is first applied to knots in yield curves in general and then to the knots derived from the maturity ranges.

To set out the B-splines in general terms not tied to specific maturity ranges, let there be $q + 1$ distinct knots designated as $[u_0, u_1, \dots, u_q]$. It is assumed that $q \geq 5$ because at least six knots will be needed for a yield curve. The knots derived from the maturity ranges in the previous chapter are an example of these $q + 1$ knots: $[0,1½,3,7,15,30]$ with $q = 5$.

From these $q + 1$ knots it's possible to derive $q + 3$ B-splines by choosing sequential sets of five knots allowing for three duplicates at beginning and end and applying the B-spline definition in the previous section. The first B-spline, designated $B_1(\tau)$, is the B-spline given by the 5 knots $[u_0, u_0, u_0, u_0, u_1]$. The second B-spline, designated $B_2(\tau)$, is the B-spline given by the 5 knots $[u_0, u_0, u_0, u_1, u_2]$, and the third B-spline is given by $[u_0, u_0, u_1, u_2, u_3]$. The last B-spline $B_{q+3}(\tau)$ is given by the knots $[u_{q-1}, u_q, u_q, u_q, u_q]$.

For the specific case of the six knots from the maturity ranges, there are eight B-splines derived from the following eight sequences of five knots each:

$[0,0,0,0,1½]$

$[0,0,0,1½,3]$

$[0,0,1½,3,7]$

$[0,1½,3,7,15]$

$[1½,3,7,15,30]$

$[3,7,15,30,30]$

$[7,15,30,30,30]$

$[15,30,30,30,30]$

Given these sequences of knots, define the coefficients $\breve{\beta}_k, k = 1, \dots, q + 3$ to be estimated from the bond set. The forward rate is represented as a linear combination of the $q + 3$ B-splines as follows:

$$\phi(\tau) = \sum_{k=1}^{q+3} \breve{\beta}_k B_k(\tau) \tag{8.3}$$

Note that because of Equation 8.2g above, this equation gives a fixed $\phi(\tau) = \check{\beta}_{q+3}$for all $\tau \geq u_q$ beyond the last knot.

Summation of the B-Splines

An important feature of the B-splines that will be used in the next chapter is that they sum to unity at any maturity τ even beyond the furthest knot. This can be verified with algebra on the B-spline polynomials:

$$\sum_{k=1}^{q+3} B_k(\tau) = 1, \tau \geq u_0 \tag{8.4}$$

Price Equation with B-Splines

Using this spline representation of the forward rate, the price equation can be written as follows:

$$p = \sum_{i=1}^{n} \exp\left(-\int_0^{\tau_i} \left(\sum_{k=1}^{q+3} \check{\beta}_k B_k(\alpha)\right) \mathrm{d}\alpha\right) c_i \tag{8.5}$$

This is a conventional price equation because it omits the regression variables, and also because the forward rate spline has no constraints.

Constraints on the forward rate spline are the subject of the next chapter. Without constraints, the spline can exhibit spurious movements. Especially important is the constraint at maturity 30 years, that not only ensures that the forward rate attains a value at the last maturity that is consistent with markets, but also enables bond yields to be projected out beyond 30 years maturity.

The B-Splines from the Maturity Ranges

The following chart depicts the eight B-splines given by the knots from the maturity ranges. The B-splines for the knot sequences from the maturity ranges are left to right in the figure as shown by their humps:

Figure 8.1

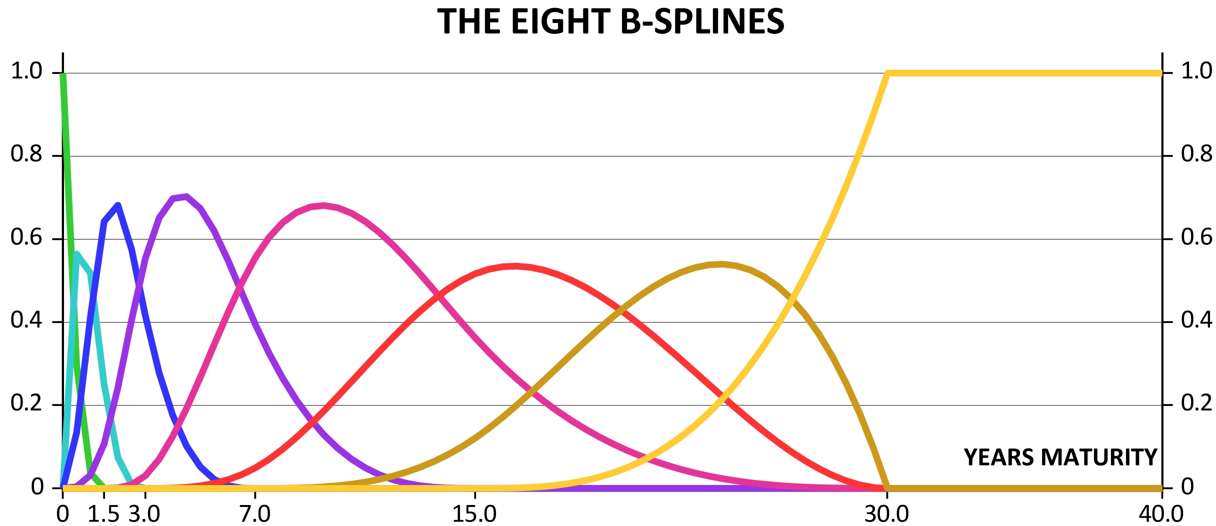

The B-splines in the figure are projected out through 40 years maturity for illustration. For actual application to the projection range, they would be run out through 100 years maturity.

The figure shows that all eight B-splines are nonnegative. The first and last B-splines equal unity at zero and 30, respectively. Furthermore, the last B-spline remains fixed at unity for maturities above 30 years while the other B-splines are zero. Therefore, the coefficient on the last B-spline will equal the fixed long-term forward rate as described in the next chapter. At each maturity, the vertical sum of all the B-splines is unity. The forward rate is given as a linear combination of the eight B-splines in this figure. However, before estimating the coefficients on this linear combination, the splines have to be constrained.

9. The Projection Range

This chapter describes the constraints on the B-splines that are needed to ensure smooth behavior of the forward rate at the nearest and furthest maturities, and that enable projection of the yield curve results in the projection range out through 100 years maturity. These constraints were developed in the Office of Financial Analysis.

There are three constraints. All three are linear and are implemented by replacing selected individual B-splines with linear combinations of the B-splines. Each constraint reduces the number of spline coefficients by one. Constraints need to be imposed because even though splines are flexible and fit bond data well, they can exhibit spurious behavior at the beginning maturity zero and at the furthest maturity 30 years.

The first constraint is near-term and requires that the forward rate have a zero second derivative at maturity zero. This effectively linearizes the forward rate at zero.

The second two constraints are long-term and are located at the last maturity of 30 years. One of the major problems with conventional yield curve approaches is that they don't impose any requirements at the 30-year point. As a result, the forward rate can vary significantly from market yields at 30 years. An example of this behavior, and a big problem in previous yield curve work, has been the fact that the forward rate for nominal yield curves frequently turned negative near 30 years maturity.

The second two constraints first ensure that the forward rate settles down at 30 years maturity by requiring that the derivative be zero. Second, the constraints cause the forward rate at 30 years maturity to equal the long-term forward rate, which is taken to be the average forward rate in the last maturity range of 15 to 30 years. These two constraints together ensure that the forward rate at 30 years maturity is consistent with market rates before 30 years maturity and can be projected forward in the projection range through 100 years maturity. In sum, the forward rate settles down to the long-term forward rate at 30 years maturity and remains at that rate with zero derivative for all higher maturities.

Constraint at Zero Maturity

As discussed in Chapter 3, the minimum maturity for the set of bonds for estimation is ½ year. Because of this minimum, there's at least a half-year gap between the zero maturity point and the lowest maturity bond in the bond set. Without constraint, it's possible that the forward rate spline could exhibit spurious behavior in this gap.

To reduce this possibility, the following linear constraint is imposed at zero maturity:

$$\frac{\mathrm{d}^2\phi(0)}{\mathrm{d}\tau^2} = 0 \tag{9.1}$$

This constraint linearizes the forward rate at zero thereby reducing spurious movements.

This constraint is imposed on the B-splines as follows. The following discussion will use notation from the last chapter and pertain to knots in general rather than the knots from the maturity ranges.

First, the only B-splines with nonzero second derivatives at the first knot u_0 are the first three, so these are the only three that need to be considered. The formulas for the derivatives of these three B-splines can be derived from the cubic polynomials in the previous chapter:

$$\begin{aligned} &\check{\beta}_1 \frac{\mathrm{d}^2 B_1(u_0)}{\mathrm{d}\tau^2} + \check{\beta}_2 \frac{\mathrm{d}^2 B_2(u_0)}{\mathrm{d}\tau^2} + \check{\beta}_3 \frac{\mathrm{d}^2 B_3(u_0)}{\mathrm{d}\tau^2} \\ &= \frac{6}{(u_1-{u_0}^2)}\check{\beta}_1 + \left(\frac{-6}{(u_2-u_0)(u_1-u_0)} + \frac{-6}{(u_1-u_0)^2}\right)\check{\beta}_2 + \frac{6}{(u_2-u_0)(u_1-u_0)}\check{\beta}_3 \\ &= 0 \end{aligned} \tag{9.2}$$

$$\Rightarrow \check{\beta}_2 = \frac{(u_2-u_0)}{(u_1-u_0)+(u_2-u_0)}\check{\beta}_1 + \frac{(u_1-u_0)}{(u_1-u_0)+(u_2-u_0)}\check{\beta}_3$$

Therefore, the constraint is implemented by replacing the first three B-splines $[B_1(\tau), B_2(\tau), B_3(\tau)]$ by $\left[B_1(\tau) + \frac{(u_2-u_0)}{(u_1-u_0)+(u_2-u_0)}B_2(\tau), B_3(\tau) + \frac{(u_1-u_0)}{(u_1-u_0)+(u_2-u_0)}B_2(\tau)\right]$ and removing the second coefficient $\check{\beta}_2$, thereby reducing the number of coefficients $\check{\beta}_k$.by one. Note that the two replacement constrained B-splines are still nonnegative and sum to the same result as the three original B-splines.

For the knots given by the maturity ranges, the first three B-splines are replaced by

$\left[B_1(\tau) + \frac{3}{1.5+3}B_2(\tau), B_3(\tau) + \frac{1.5}{1.5+3}B_2(\tau)\right]$.

The Long-Term Forward Rate

Analogous to the zero maturity point, the value of the forward rate spline can also veer off course at 30 years maturity. In particular, sometimes there can be a hump in the forward rate at earlier maturities that causes the forward rate to plunge downward as it approaches 30 years maturity. This can result in the forward rate falling so rapidly that by the time it reaches 30 years maturity, it is significantly different from market yields. Sometimes the forward rate can even turn negative for nominal yield curves if the spline coefficients are not constrained to be positive, a result that's inconsistent with markets. A negative forward rate was a significant problem in conventional yield curves in the past and reduced the usefulness of conventional results.

Therefore, it's necessary to constrain the spline coefficients so that the forward rate at maturity 30 years is consistent with market yields. In addition, the forward rate has to be constrained for maturities greater than 30 years so that yields can be projected.

The constraints at 30 years maturity are done in two stages. First, the derivative of the forward rate spline is constrained to zero at maturity 30 to make the forward rate flatten out at that point:

$$\frac{\mathrm{d}\phi(30)}{\mathrm{d}\tau} = 0 \tag{9.3}$$

Second, the forward rate beyond 30 years maturity is assumed to be a constant long-term forward rate ϕ^*. This is because there are insufficient bonds in the bond set to estimate the forward rate beyond 30 years maturity, so the closest that can be done is to fix the long-term forward rate at a constant representing the average forward rate beyond 30 years. For smoothness, the forward rate is made to attain the long-term forward rate at 30 years maturity and stay at that rate for all higher maturities. Moreover, because the derivative of the forward rate at 30 years maturity is constrained to zero, the forward rate flattens out smoothly to the long-term forward rate at 30 years maturity. And the derivative of the long-term forward rate for all maturities above 30 years continues to be zero.

So the question is how to estimate the long-term forward rate. In the yield curves done here, the long-term forward rate is taken to be the average forward rate in the last maturity range 15 years to 30 years, and this is constrained to equal the forward rate at maturity 30 years. Note that this constraint is imposed simultaneous with the estimation of the price equation:

$$\phi^* = \frac{\int_{\tau=15}^{30} \phi(\alpha)\mathrm{d}\alpha}{30-15} = \phi(30) \tag{9.4}$$

The reason for choosing this average for the long-term forward rate is that the maturities of bonds in the last maturity range are sufficiently distant in time that market assessments of risk and reward for such bonds as indicated by the forward rate are similar to assessments for bonds above 30 years maturity.

Furthermore, as required by Equation 8.2g in the B-spline definition, the final B-spline in the forward rate is unity at maturity 30 years and at all higher maturities, while the other B-splines are zero at these maturities. So once this constraint is imposed, the forward rate will remain at the long-term forward rate at 30 years maturity and beyond out through 100 years maturity.

In normal market conditions, the long-term forward rate will be positive. In the case of nominal bonds, the forward rate normally will be positive from zero maturity up through 30 years maturity, ensuring that the average forward rate in the last maturity range will be positive too.

In the case of TRC and TIPS, the forward rate will usually be positive on average in the last maturity range even if the forward rate is negative in earlier ranges. However, sometimes the TRC forward rate will be negative in the last maturity range as noted in Chapter 5, with the consequence that the long-term forward rate can be negative at 30 years maturity and remain negative out through 100 years maturity.

The Long-Term Constraints

This section lays out the implementation of the two long-term constraints.

The first constraint is the zero derivative at the last knot. Analogous to the linear constraint at the first knot, this discussion will apply to knots in general and will not be specific to the knots from the maturity ranges. Derivatives for all B-splines except the last two are zero and don't need to be dealt with. As in the case of the constraint at the first knot, formulas for the derivatives at the last knot can be derived from the cubic polynomials in the previous chapter:

$$\begin{aligned}
&\breve{\beta}_{q+2}\frac{\mathrm{d}B_{q+2}(u_q)}{\mathrm{d}\tau}+\breve{\beta}_{q+3}\frac{\mathrm{d}B_{q+3}(u_q)}{\mathrm{d}\tau} \\
&\quad=\frac{-3}{(u_q-u_{q-1})}\breve{\beta}_{q+2}+\frac{3}{(u_q-u_{q-1})}\breve{\beta}_{q+3} \\
&\quad=0 \\
&\Rightarrow \breve{\beta}_{q+3}=\breve{\beta}_{q+2}
\end{aligned} \tag{9.5}$$

Therefore, the constraint of zero derivative at the last knot u_q implies that the last two spline coefficients must be equal. So this constraint is imposed by replacing the last two B-splines $B_{q+2}(\tau)$ and $B_{q+3}(\tau)$ by their sum $B_{q+2}(\tau)+B_{q+3}(\tau)$ and removing the last coefficient $\breve{\beta}_{q+3}$.

To impose the additional constraint for computing the long-term forward rate, it's necessary to calculate the average forward rate in the last maturity range 15 years to 30 years maturity. Because the only B-splines with nonzero values in the last range are the last four $B_k(\tau), k=q,q+1,q+2,q+3$, the other B-splines can be ignored.

To simplify notation, the $A_k$can be defined:

$$A_k=\frac{\int_{\alpha=15}^{30}B_k(\tau)\mathrm{d}\alpha}{30-15}, k=q,q+1,q+2,q+3 \tag{9.6}$$

Using Equation 8.4 on summing B-splines, this implies:

$$\sum_{k=q}^{q+3}A_k=\frac{\int_{\alpha=15}^{30}\sum_{k=q}^{q+3}B_k(\tau)\mathrm{d}\alpha}{30-15}=\frac{\int_{\alpha=15}^{30}1\mathrm{d}\alpha}{30-15}=1 \tag{9.7}$$

The constraint requires the following, using Equation 8.2g, Equation 9.5, Equation 9.7, and the fact that $\phi(u_q) = \check{\beta}_{q+3}$:

$$\check{\beta}_q A_q + \check{\beta}_{q+1} A_{q+1} + \beta_{q+2}(A_{q+2} + A_{q+3}) = \check{\beta}_{q+2} \tag{9.8}$$

$$\Rightarrow \beta_{q+2} = \frac{A_q}{A_q + A_{q+1}} \check{\beta}_q + \frac{A_{q+1}}{A_q + A_{q+1}} \check{\beta}_{q+1}$$

Therefore, both long-term constraints are implemented simultaneously by replacing the last four B-splines $[B_q(\tau), B_{q+1}(\tau), B_{q+2}(\tau), B_{q+3}(\tau),]$ by the two constrained B-splines $\left[B_q(\tau) + \frac{A_q}{A_q + A_{q+1}}\left(B_{q+2}(\tau) + B_{q+3}(\tau)\right), B_{q+1}(\tau) + \frac{A_{q+1}}{A_q + A_{q+1}}\left(B_{q+2}(\tau) + B_{q+3}(\tau)\right)\right]$ and removing the last two coefficients $\check{\beta}_{q+2}$ and $\check{\beta}_{q+3}$.

Putting the three constraints together, including the one constraint at the first knot of zero and the two constraints at the last knot of 30, the number of constrained B-splines after replacing the original B-splines with the constrained B-splines is q. The set of constrained B-splines is given the notation $B_k^{\mathrm{C}}(\tau)$ with associated coefficients β_k, $k = 1, \dots, q$.

Note that all the constrained B-splines are nonnegative, the same as the original unconstrained B-splines. This implies that if the coefficients β_k are all positive, as is expected with nominal bonds, the forward rate at all maturities is positive and the discount function is declining in accord with Equation 5.5.

Application to the Knots from the Maturity Ranges

These three constraints are applied as follows to the forward rate from the maturity ranges. The six knots from the maturity ranges $[0, 1½, 3, 7, 15, 30]$ are converted to five constrained B-splines $B_k^{\mathrm{C}}(\tau), k = 1, \dots, 5$, as summarized in the following schematic in the form of matrix multiplication. In order to create this matrix, $\frac{A_q}{A_q + A_{q+1}}$ and $\frac{A_{q+1}}{A_q + A_{q+1}}$ have to be calculated: the respective results to two decimal places are 0.24 and 0.76, both when the last knot is 30 and when it is 30.51 as used for Treasury securities.

$$\begin{bmatrix} B_1^{\mathrm{C}}(\tau) \\ B_2^{\mathrm{C}}(\tau) \\ B_3^{\mathrm{C}}(\tau) \\ B_4^{\mathrm{C}}(\tau) \\ B_5^{\mathrm{C}}(\tau) \end{bmatrix} = \begin{pmatrix} 1 & 0.67 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.33 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0.24 & 0.24 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0.76 & 0.76 \end{pmatrix} * \begin{bmatrix} B_1(\tau) \\ B_2(\tau) \\ B_3(\tau) \\ B_4(\tau) \\ B_5(\tau) \\ B_6(\tau) \\ B_7(\tau) \\ B_8(\tau) \end{bmatrix}$$

The next step is to write the price Equation (6.2) using the constrained B-splines. Because the price equation uses the integral of the forward rate, it is necessary to define a set of five integrated constrained B-splines $B_k^{\mathrm{IC}}(\tau)$ as follows:

$$B_k^{\mathrm{IC}}(\tau) = \int_{\alpha=0}^{\tau} B_k^{\mathrm{C}}(\alpha)\mathrm{d}\alpha, k = 1, \dots, 5,\ \tau \geq 0 \tag{9.9}$$

Because the B-splines are cubic polynomials, integration is straightforward: indefinite integrals of the cubic polynomials are computed as fourth-degree polynomials and the constant terms in the new polynomials are adjusted so that the integrals are continuous over the maturity ranges and the projection range.

Then, with the five integrated constrained B-splines, the extended price equation can be written as:

$$p = \sum_{i=1}^{n} \exp\left(-\sum_{k=1}^{5} \beta_k B_k^{\mathrm{IC}}(\tau_i)\right) c_i + \sum_{j=1}^{m} \theta_j x_j \tag{9.10}$$

This is the form of the price equation to be estimated. It has five spline coefficients β_k and m regression coefficients θ_j. The spline coefficients have been described, and the next three chapters describe the regression variables.

The Constrained B-Splines

The following Figure 9.1 depicts the five constrained B-splines:

Figure 9.1

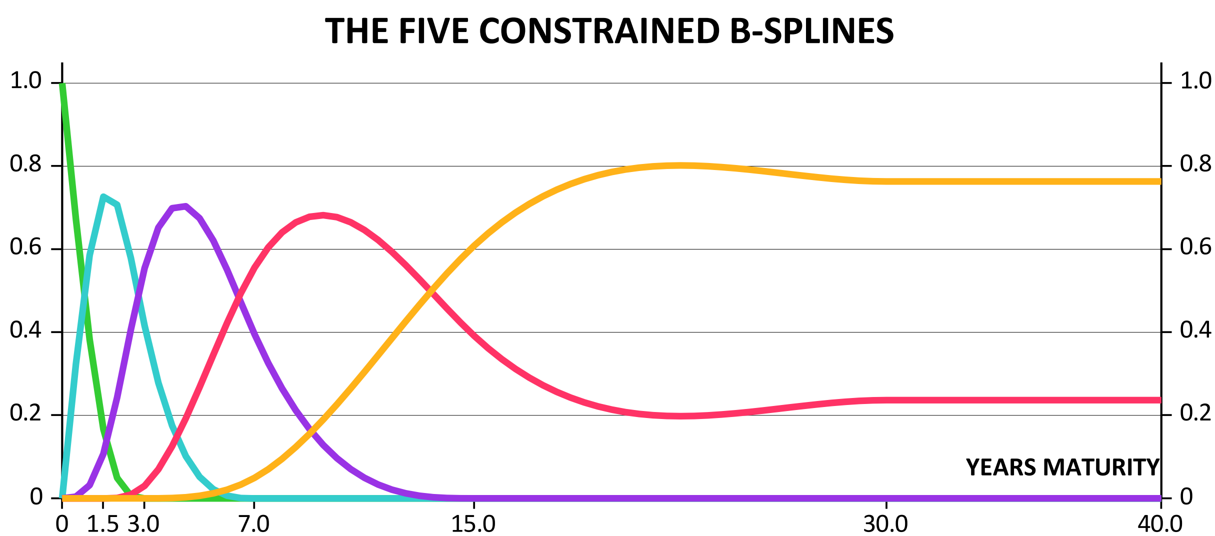

Because the five constrained B-splines are positive linear combinations of the eight original B-splines, the five constrained B-splines are nonnegative. Because of the constraints, the fourth and fifth constrained B-splines at 30 years maturity are nonzero, which means that the long-term forward rate will be given by a combination of the estimated coefficients for the forward rate on these constrained B-splines.

Therefore, for nominal bonds, if the estimated forward rate coefficients on the constrained B-splines are positive, the resulting forward rates will be positive and the discount function will decline, all in accord with the discussion in Chapter 5. And as shown in Chapter 13 on estimation, for the nominal yield curves HQM and TNC, the forward rate coefficients are constrained to be positive.

The following Figure 9.2 shows the five integrated constrained B-splines:

Figure 9.2

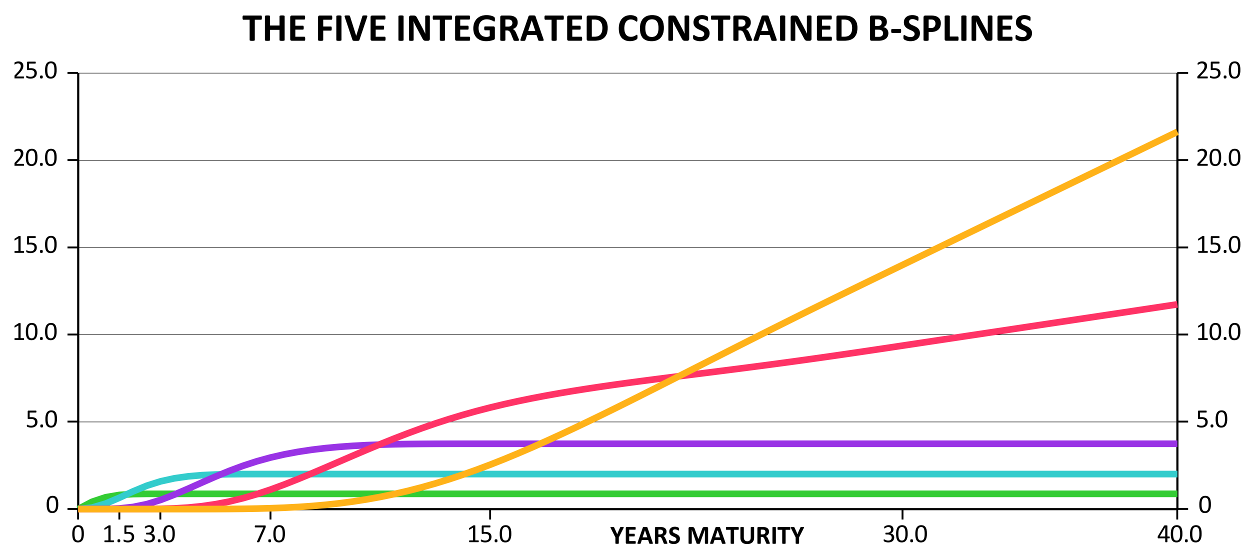

The first three integrated constrained B-splines reach a constant at 30 years maturity and remain there for higher maturities, while the last two rise continuously beyond 30 years

The mathematical form for estimation of the discount function and forward rate has now been described. The next three chapters define the regression variables.

10. The Hump Variable

This chapter describes the hump variable, which is the regression variable that captures the hump or higher level of yields often seen in standard bonds around 20 years maturity. The hump variable has been included in the TNC and TRC yield curves from the beginning going back through 1976 for TNC. The hump variable was added to the HQM yield curve in February 2024. The hump variable was developed in the Office of Financial Analysis.

The next sections discuss the need for the hump variable and describe its construction.

Purpose of the Hump Variable

Frequently, yields on bonds with maturities around 20 years are high relative to yields on bonds around 10 years or 30 years maturity, that is, 20-year yields exhibit a hump. This doesn't occur all the time, but it occurs often enough that it is considered one of the characteristics of bond markets.

Various explanations have been suggested for the existence of the hump. The fact that it does not always occur would seem to indicate that it's not inherently part of bond structure, but rather stems from factors separate from the bonds themselves. The best explanation appears to be that bonds around 20 years maturity are less liquid than bonds around 10 years or 30 years maturity, and therefore such bonds are less desirable with the result that they sell at lower prices and must offer higher yields.

One way that this difference in liquidity could come about is that the bond market is complex, and the complexity and size of the bond market make it efficient for traders to divide the market into parts so that trading activity and market analysis can focus on a limited set of bonds rather than on all bonds at the same time. This is a decision made by bond institutions and doesn't affect the value of individual cash flows from bonds of different maturities. In particular, traders have chosen to center attention on 10-year and 30-year bonds as opposed to 20-year bonds, and this makes the 20-year bonds less liquid.

Regardless of the reason for the hump, it's clear that the hump effect is not part of time preference or the discount function. In particular, cash flows from 20-year bonds are no different from cash flows from other bonds, so time preference does not affect them differently. Rather, it's the fact that these cash flows are packaged into a 20-year coupon issue that is causing the hump. Therefore, the hump reflects other factors than time preference and needs its own regression variable separate from the discount function.

It should also be noted that the source of the hump is in the set of standard bonds used in estimation as set out in Chapter 4. A different set of bonds, such as bonds from a fully developed zero coupon market, might not have a hump. Nevertheless, because the hump can exist in the bond set that's being used to estimate, the hump must be taken into account with the regression variable.

The hump in bond yields does not always exist. Frequently there's no hump at all, such as when liquidity concerns at 20 years maturity are small.

Or another possibility is that there is an inverse hump in which bond yields dip around 20 years maturity rather than rise. The inverse hump might stem from a market decision to issue bonds or sell existing bonds, and the bonds chosen for sale are at 10 or 30 years maturity because it's more convenient and more liquid to sell at those maturities rather than 20 years. As demonstrated below, the hump variable coefficient is negative when there is a hump but turns positive when there is an inverse hump.

Conventional yield curves that don't have regression variables fold the hump effect into the discount function. This flattens out the hump and spreads the hump effect into long-term yields. If the hump is small or nonexistent, omitting the hump variable makes the resulting yield curve simpler and produces smoother spot rates.

However, if the hump is large, leaving out the hump variable can significantly distort the discount function, and in particular the omission can bias up yields in the projection range that should depend on time preference alone. The hump variable is needed to eliminate this bias.

Another way to look at the hump variable is that it picks up additional market movements in the last maturity range. Because the last maturity range of 15 to 30 years is large, there are market movements in this range that need another variable to be captured. The hump effect is the most important feature of this range separate from time preference itself. Therefore, the hump variable works in tandem with the discount function to capture market movements in this big range.

This implies that the discount function and forward rate as given by the five constrained B-splines together with the hump variable can be viewed as an integrated mathematical structure for yield curve estimation. This structure, which comprises six coefficients and therefore has six degrees of freedom, has worked very well over the last half century in capturing the behavior of U.S. bonds of high credit quality. Without modification, for five decades this structure has tracked high quality corporate bonds as well as both nominal and real Treasury securities.

Construction of the Hump Variable

Analogous to the maturity ranges, the construction of the hump variable assumes that the market focuses on central maturities and bond pricing at other maturities is done relative to the central maturities. For the hump, the central maturities are 10, 20, and 30 years maturity, and the hump variable assumes that at various times there may be factors that affect the prices of bonds at 20 years maturity relative to the prices of bonds at 10 or 30 years maturity.

To start, the hump variable is set to unity at 20 years maturity and set to zero for maturities up through 10 years and for maturities 30 years and greater. This reflects the fact that the effect on bond price because of the hump occurs at the central maturity of 20 years relative to 10 years and 30 years. Then the value of the hump variable at other maturities greater than 10 years or less than 30 years is made to depend upon how close these maturities are to 20 years.

Therefore, to obtain a smooth hump function, the value of the hump variable for any bond is given by the cubic B-spline $2 \times B(\tau_n; 10,10,20,30,30)$, where, following previous notation, τ_n is the maturity of the bond. The B-spline is multiplied by 2 to set it to unity at its maximum value at 20 years maturity. The hump variable starts at zero at 10 years maturity and below, rises to unity at 20 years, and then declines back to zero at 30 years and beyond. The hump coefficient to be estimated in the price equation is denoted by θ_{HUMP} and the value of the hump variable for a particular bond as a function of the bond's maturity is denoted by $x_{\mathrm{HUMP}}(\tau_n)$, which gives:

$$\theta_{\mathrm{HUMP}}\, x_{\mathrm{HUMP}}(\tau_n) = \theta_{\mathrm{HUMP}} \times 2 \times B(\tau_n; 10,10,20,30,30) \qquad (10.1)$$

If there is a hump in the price equation at a particular time, it's clear that the coefficient on the hump variable θ_{HUMP} must be negative because the existence of the hump reduces the price and increases yield around 20 years maturity. Conversely, if the hump is inverse, the coefficient will be positive because in that case the presence of the hump raises price and reduces yield. And, of course, when there is no hump, the coefficient will be near zero. Therefore, the hump variable coefficient provides a test as to the existence of the hump, and the t-ratio on the coefficient shows its significance.

It should be noted that the actual position of the hump in a par or spot yield curve is determined by the discount function and hump in combination. In particular, even though the 20-year maturity point gives the maximum value unity of the hump variable, the position of the hump in the yield curve may not necessarily be at 20. This will be shown in results later.

Picture of the Hump Variable

The following chart depicts the hump variable starting at zero maturity and running out through 40 years maturity.

Figure 10.1

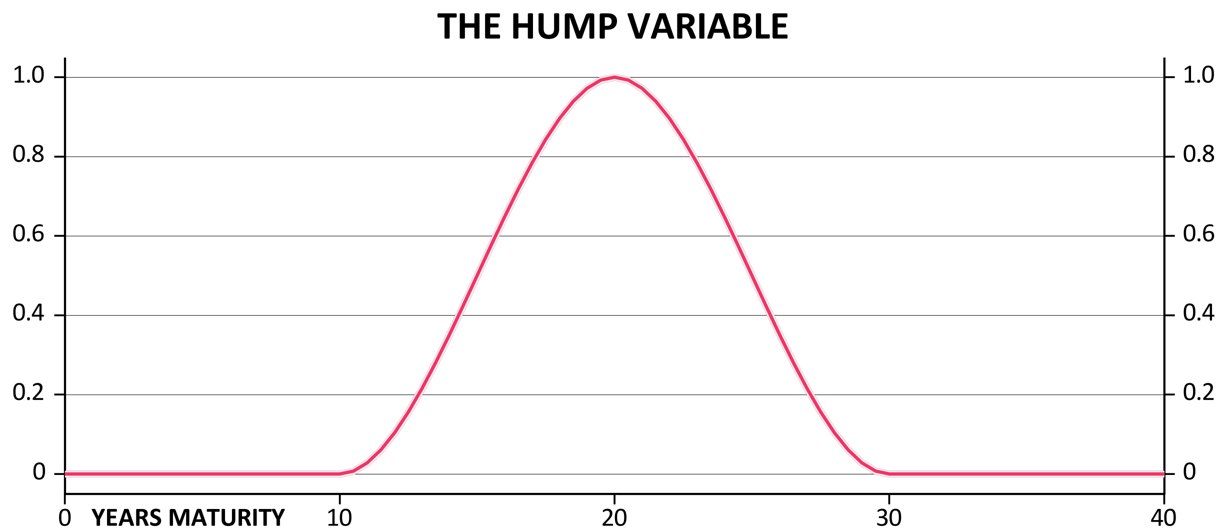

The hump variable is continuous and smooth for the span 10 years to 30 years, starting at zero and rising to unity at maturity 20 years, then falling back down to zero. It's symmetric around 20 years maturity. The maximum effect of the hump is at maturity 20 years.

11. HQM Credit Variables

This chapter describes the two credit variables in the HQM yield curve. The credit variables are regression variables in HQM. The credit variables were developed in the Office of Financial Analysis.

The purpose of the credit variables is set out and the construction of the variables is shown. The credit variables are in addition to the hump variable. Therefore, the HQM yield curve has a total of three regression variables which comprise the hump variable plus the two credit variables, with the result that the HQM yield curve has eight parameters to be estimated which comprise the five coefficients on the splines plus the three coefficients on the regression variables.[8]

Purpose of the Credit Variables

In the Pension Protection Act (PPA) discussed in Chapter 3, the idea of the HQM yield curve was to have a single yield curve that represents the average yields weighted by market size of high quality bonds rated A, AA, and AAA. Conventional yield curves would typically have approached this problem by calculating three separate yield curves, one for each quality AAA, AA, and A, and computing the weighted average.

However, the problem with separate yield curves is that the number of bonds is small for AAA and not so large for AA either, with the result that individual curves can be unstable. Nevertheless, even

[8] In some of the previous versions of the HQM yield curve, a variable was included to capture the effects of call options in bonds with a call schedule: see U.S. Department of the Treasury (2005a, 2005b, and 2006). The results from this variable demonstrated that effects of calls can be modeled with regression variables and without using option adjusted spreads. However, the final version of the HQM yield curve omits these types of callable bonds and this variable, except that end-calls are included in the HQM yield curve without a regression variable as discussed in Chapter 4.

if unstable, the shapes of the three curves are similar because the markets for the three qualities are similar. So it's efficient to make use of the similarity by putting bonds from all three ratings into a single price equation producing one discount function and hump variable that applies to bonds with market-weighted average rating. Moreover, the resulting blended yield curve eliminates the instability of individual yield curves and so it is robust in capturing the overall structure of the entire high quality bond market.

When putting three qualities together, it's still necessary to account for differences in the qualities. This can be done with two credit variables added to the discount function/hump variable in the price equation. This is an example of regression variables in which the bond set for estimation does not have to be homogeneous, and characteristics of different bonds, in this case quality characteristics, can be picked up by the regression variables. The next sections specify the credit variables.

Therefore, in the HQM price equation, the discount function/hump variable is assumed to apply to bonds with market-weighted average quality across the three qualities AAA, AA, and A. None of the bonds for estimation has such market-weighted quality. So the two credit variables adjust the price of each bond used in estimation to what it would be if the bond were of weighted average quality, and the resulting adjusted bond is then used to estimate the weighted average discount function/hump variable that is then used to construct yield curves that comply with the Pension Protection Act.

Taking advantage of bond similarities across qualities by using all bonds simultaneously is analogous to the econometric technique of seemingly unrelated regression. In seemingly unrelated regression, similarities in different data sets are exploited to provide a more robust result by fitting the model to the data sets all at the same time.

Credit Shares

As discussed, two credit variables are needed when merging the three bond qualities into a single price equation. The credit variables are based on the relative par amounts outstanding of the three qualities of the bonds, that is, the credit shares. So the credit shares need to be discussed first.

To begin, let Γ_{AAA}, Γ_{AA}, and Γ_{A}, be the sum of the par amounts outstanding for the bonds in the bond set used in estimation at the respective qualities AAA, AA, and A. Then define the two credit shares:

$$\omega_1 = \frac{\Gamma_{\mathrm{AA}}}{\Gamma_{\mathrm{AAA}}+\Gamma_{\mathrm{AA}}} \tag{11.1a}$$

$$\omega_2 = \frac{\Gamma_{\mathrm{A}}}{\Gamma_{\mathrm{AAA}}+\Gamma_{\mathrm{AA}}+\Gamma_{\mathrm{A}}} \tag{11.1b}$$

The first share ω_1 is the total par amount outstanding of AA bonds relative to AAA plus AA bonds, while the second share ω_2 is the total par amount outstanding of A bonds relative to all the high quality bonds AAA, AA, and A.

Recently, the first credit share ω_1 has been a bit below 90 percent and the second credit share ω_2 has been around 75 percent.

Constructing the Credit Variables

Using the two credit shares ω_1 and ω_2, the two credit variables can be constructed. The respective coefficients on the credit variables are θ_{CR1} and θ_{CR2} which are estimated in the price equation.

θ_{CR1} is the difference between the price of an AAA-rated bond and the price of an AA-rated bond per year of maturity of the bond. So for a bond with 10 years maturity, the difference in price between AAA and AA is $\theta_{\mathrm{CR1}}\tau_n = 10 \times \theta_{\mathrm{CR1}}$, where as before τ_n is the maturity of the bond. The price difference is analogous to insurance in that the amount of protection that a higher rating provides to a bondholder is greater as the length of time of the bond, that is maturity, is greater. Consequently, the price of the higher rating rises with maturity.

Analogously, θ_{CR2} is the difference per year of maturity between the price of a bond of market-weighted average AAA and AA quality ignoring A bonds and the price of an A-rated bond.

Usually both θ_{CR1} and θ_{CR2} are positive. However, it's possible for these coefficients to be negative. For example, in the crisis of 2008, θ_{CR1} was negative on various dates in October and November 2008 and in the early months of 2009. Negative values for θ_{CR1} mean that AA bonds have a higher price than AAA bonds, and could indicate that bond qualities as perceived by markets are different from ratings assigned to the bonds. The following discussion and the construction of the credit variables is still valid if one or both of these coefficients is negative.

Using these data items, the price p of an AAA bond with maturity τ_n can be adjusted to what it would be if the bond had market-weighted average quality.

First, the AAA price is adjusted to be market-weighted average AAA-AA quality. If ω_1 is near zero, there aren't many AA-rated bonds, so the weighted average price is already near AAA and little adjustment is needed. In contrast, if ω_1 is near unity, a large adjustment in the price is needed because the average is near AA but the bond price is AAA. So the amount of adjustment depends on ω_1, and consequently, the adjustment of p to provide a market-weighted average AAA-AA price is to subtract $\omega_1\theta_{\mathrm{CR1}}\tau_n$ from p. Note that the adjustment is linear in ω_1.

Finally, to complete the adjustment of price on an AAA bond, the market-weighted average AAA-AA price $p - \omega_1\theta_{\mathrm{CR1}}\tau_n$ is adjusted to be the price for market-weighted average quality of all HQM bonds. The reasoning is analogous to the previous paragraph using a linear adjustment in ω_2, and so the market-weighted average price is $p - \omega_1\theta_{\mathrm{CR1}}\tau_n - \omega_2\theta_{\mathrm{CR2}}\tau_n$.

To continue for the price on an AA-rated bond: following analogous reasoning, for an AA-rated bond the market-weighted average price is $p + (1 - \omega_1)\theta_{\mathrm{CR1}}\tau_n - \omega_2\theta_{\mathrm{CR2}}\tau_n$. The price is first adjusted to equal the market-weighted average AAA-AA price, and then adjusted to get the market-weighted average price for all HQM bonds.

Finally, for an A-rated bond, the market-weighted average price is $p + (1 - \omega_2)\theta_{\mathrm{CR2}}\tau_n$. In this case, the closer ω_2 is to zero, the more the price must be adjusted to get the market-weighted average HQM price. In contrast, the closer ω_2 is to unity, the more the average is already A quality, so little adjustment of the price is necessary.

To summarize, the two credit variables are $x_{\mathrm{CR1}}\tau_n$ with coefficient θ_{CR1} and $x_{\mathrm{CR2}}\tau_n$ with coefficient θ_{CR2}, where:

$$
\begin{aligned}
x_{\mathrm{CR1}} &= \omega_1 \text{ for an AAA-rated bond;} \\
&= (\omega_1 - 1) \text{ for an AA-rated bond;} \\
&= 0 \text{ for an A-rated bond.} \\
x_{\mathrm{CR2}} &= \omega_2 \text{ for an AAA- or AA-rated bond;} \\
&= (\omega_2 - 1) \text{ for an A-rated bond.}
\end{aligned}
$$

These two variables $x_{\mathrm{CR1}}\tau_n$ and $x_{\mathrm{CR2}}\tau_n$ are added to the HQM price equation along with the hump variable, and all coefficients including θ_{CR1} and θ_{CR2} are estimated. Note that the credit variables are derived by straightforward linear scaling: the variables and associated coefficients are linear, as is the scaling by maturity. More complicated variables could be devised, but given the low signal to noise ratio of bond data especially corporate bond data, additional complications can lead to spurious results. A general rule of yield curve regression variables is that simpler variables work best in bringing out patterns and characteristics of bond markets.

12. TNC On-the-Run Variables

This chapter describes the regression variables used in the TNC nominal Treasury yield curve. The TNC regression variables were developed in the Office of Financial Analysis.

There are seven on-the-run dummies in the TNC yield curve plus seven first off-the run dummies. The TNC yield curve also includes the hump variable, so the total number of variables is 15. Together with the five spline coefficients, the total number of coefficients to be estimated in TNC is 20.

In contrast to nominal Treasury coupon issues, there are no distinctions between on-the-run and off-the-run issues for TIPS. Therefore, the TRC yield curve doesn't have any dummy variables and has a total of six coefficients to be estimated, comprising the five spline coefficients and one hump coefficient.

This chapter also describes regression variables used for historical TNC data going back through 1976. The features of the historical data that required these variables no longer exist in Treasury coupon issues at the present time, so these variables are no longer used. The goal of the TNC and TRC yield curves is to include in the price equation at any time all Treasury coupon issues both historical and current that are already issued and available for trading at that time regardless of when they were issued. This includes both on-the-run and off-the-run securities as well as callable bonds and flower bonds, securities with odd first coupons, and securities of every term ranging from 1 year through 40 years. The regression variables make it possible to include all issues by sorting out differences in features among the issues.

The hump variable is also included for historical Treasury data. However, the hump variable is left out before December 1985 because before that date there aren't enough Treasury securities with long maturities and without call schedules to estimate the hump.

The historical results show that the XRM methodology provides stable results that capture market movements over the last half century. In particular, the ill-conditioned numerical results and multiple convergence points often seen in other yield curve approaches are not found in XRM.

On-the-Run Variables

The most recently issued Treasury nominal coupon issues of each term are called on-the-run coupon issues. At present, the on-the-run securities are the most recently issued notes of maturities 2 years, 3 years, 5 years, 7 years, and 10 years, and the most recently issued bonds of maturities 20 years and 30 years. Treasury securities that aren't on-the-run are called off-the-run.

In the past, Treasury also issued an on-the-run note at 4 years maturity, but this note has not been issued after December 1990. Other odd maturities were also issued in the past, but they were not considered on-the-run.

Note that at any point in time all the on-the-run securities have actually been issued, that is, the point in time is on or after the issue date of each security. There is another set of recent Treasuries at any time that are traded before issue, and they are called when-issued securities. At the present time, when-issued securities are not included in the TNC or TRC yield curves and not discussed here.

The seven current maturities of the on-the-run securities are included in the central maturities used to define the maturity ranges in Chapter 7. The choice of these seven maturities for on-the-run Treasury securities indicates the importance of these maturities in markets and shows why these maturities are central maturities in market trading and in delineating the maturity ranges.

The reason for the interest in on-the-run securities is that they are generally thought to be priced differently from off-the-run. The reasons for the difference in pricing are analogous to the hump variable: on-the-run securities are thought to be more liquid than other Treasuries and are traded for special purposes. Market institutions are set up for special trading of on-the-run thereby causing special effects on the on-the-run prices. So even though the individual cash flows of on-the-run securities are no different from off-the-run and subject to the same time preference, the package of cash flows into an on-the-run Treasury security causes the security to be priced differently.

And because on-the-run securities are priced differently, they cannot be mixed in with off-the-run in the price equation. Therefore, analogous to the hump variable, each on-the-run security is given its own dummy variable. The inclusion of the dummy has the effect of removing the on-the-run security from the price equation estimation, and at the same time provides an estimate of the on-the-run price relative to off-the-run that can be used for calculations such as comparison of on-the-run yield to off-the-run or examination of the significance of the on-the-run coefficient. Conventional yield curve approaches that don't have regression variables remove the on-the-run securities completely from the price equation, thereby providing no estimates of their prices.

The removal of on-the-run securities from the yield curve estimation by the dummies implies that the resulting TNC yield curve is off-the-run, that is, it provides yield estimates that pertain to the off-the-run Treasury market. In calculating the TNC yield curves as described in Chapter 14, these dummy variables are set to zero.

Coefficients on the dummies do find particular use in calculating on-the-run yields at the seven on-the-run maturities, and selected on-the-run TNC estimates are published on the website. It should be noted that these on-the-run estimates pertain to the seven on-the-run maturities exactly. The actual on-the-run security in the market for a any term has a maturity that is a bit less than the original term because it is already issued. For example, the 10-year on-the-run note in the market has a maturity a bit less than 10 years. However, the 10-year on-the-run TNC estimate pertains to a note of exactly 10 years maturity. This is important because often there is a desire to estimate what the yield would be on a newly issued on-the-run security of the full term, and the dummy variable provides this information.

In addition to the seven on-the-run securities, markets believe that the seven first off-the-run securities for the same seven maturities are also priced differently. First off-the-run means the second most recent issue of each term. Therefore, the seven first off-the-run securities are also given their own dummy variables, which also takes them out of the price equation estimation and provides separate price estimates for them.

There's disagreement as to the time when on-the-run securities were distinguished by market participants from off-the-run. It appears that this happened when Treasury markets became sufficiently regularized in the early 1980s. Historical TNC yield curves include the on-the-run and first off-the-run dummies back through the beginning of TNC data in 1976. However, published on-the-run TNC estimates using the coefficients of the dummy variables start in 1986.

In summary, in addition to the hump variable, the TNC yield curve has seven on-the-run dummies that are set to unity for the respective security and zero otherwise

$$x_{\mathrm{ON2}}, x_{\mathrm{ON3}}, x_{\mathrm{ON5}}, x_{\mathrm{ON7}}, x_{\mathrm{ON10}}, x_{\mathrm{ON20}}, x_{\mathrm{ON30}}$$

with respective coefficients

$$\theta_{\mathrm{ON2}}, \theta_{\mathrm{ON3}}, \theta_{\mathrm{ON5}}, \theta_{\mathrm{ON7}}, \theta_{\mathrm{ON10}}, \theta_{\mathrm{ON20}}, \theta_{\mathrm{ON30}}$$

and first-off-the fun dummies

$$x_{\mathrm{OFF2}}, x_{\mathrm{OFF3}}, x_{\mathrm{OFF5}}, x_{\mathrm{OFF7}}, x_{\mathrm{OFF10}}, x_{\mathrm{OFF20}}, x_{\mathrm{OFF30}}$$

with respective coefficients

$$\theta_{\mathrm{OFF2}}, \theta_{\mathrm{OFF3}}, \theta_{\mathrm{OFF5}}, \theta_{\mathrm{OFF7}}, \theta_{\mathrm{OFF10}}, \theta_{\mathrm{OFF20}}, \theta_{\mathrm{OFF30}}$$

Callable Treasury Bonds

These Treasury bonds have a call option that enables them to be redeemed by Treasury before maturity. There were 28 such bonds available for trade in markets at some time during the period 1976-2002, with a minimum of 11 such bonds at any time. The terms of the callable bonds could be 20, 25, 30, and 31 years. Each of the callable bonds could be called at any time during the last 5 years before maturity, except for one old 25-year bond that matured in May 1985 and could be called during the last 10 years before maturity. Of the 28 callable bonds, 24 were actually called and 4 were allowed to mature without being called.

The callable bonds can't be mixed in with noncallable bonds because the presence of the calls distorts their prices. Therefore, the conventional yield curve approach is to remove all of the callables from the estimation of the price equation. However, this is a big distortion because the callable bonds constitute a major source of information about historical bond markets. For example, in the period 1976 through 1984, all traded bonds with more than 20 years until maturity (except possibly for flower bonds, see below) were callable, so to omit them is to omit the entire long end of the yield curve. Even as late as 1995, the entire section of the Treasury yield curve from about 10 years maturity to 20 years maturity contains only callable bonds.

Therefore, to achieve the goal of the TNC yield curve to include all Treasury coupon issues, the callable bonds are included with a regression variable that captures the effects on price of the presence of the call option. This variable is defined as the number of years remaining in the life of the bond when the bond might be called multiplied by the annual coupon rate of the bond. If the bond is sufficiently old that it could be called at any time after the settlement date of the yield curve, then this variable is defined as the number of years to maturity of the bond times coupon rate. This variable is a measure of risk: it measures the amount of coupon payments that the bondholder risks in buying the bond if the bond is called before maturity.

This variable captures the reduction in bond price caused by the existence of the call, especially from December 1985 forward when there are sufficient numbers of noncallable bonds so that this variable can be easily identified. This is a simple variable, but as seen in the other regression variables, simple variables work best in yield curves because there is a lot of noise in bond data, and complicated regression variables, while they may seem good in theory, can often pick up spurious effects. A simple variable picks up strong effects that come through the noise of the data and dominate the movements of yield curves.

Estate Tax Anticipation Bonds

This is another set of historical Treasury bonds that conventional yield curves leave out. These bonds if held by the estate of a deceased bondholder enable the estate to redeem them at par to pay federal estate taxes. These bonds are typically called flower bonds.

There were 11 flower bonds available for trade in the period 1976-2002, and the last flower bond matured in November, 1998. The terms of these bonds ranged from 20 years to 40 years. Five of the flower bonds were callable, so they overlapped the 28 callable bonds discussed in the previous section.

These bonds are also important for bond analysis and should be included in the price equation. However, similar to the case of callable bonds, the estate tax effects in these bonds distort the prices of these bonds such that they can't be mixed in with other Treasury securities.

So analogous to callable bonds, flower bonds are included with a regression variable to pick up estate tax effects on prices. This variable is defined as years to maturity, or if the bond is callable, the average of years to maturity and years to the first possible call date. The idea is that the longer is the time to maturity of the bond, the longer the bond provides insurance for estate taxes against a fall in bond prices.

Again this is a simple regression variable, but it captures well the greater demand for flower bonds because of their estate tax effects and the consequent higher prices that bondholders were willing to pay for them.

The HQM, TNC, and TRC price equations are now specified, including the spline for the discount function and forward rate and the regression variables and their coefficients. The next chapter will show how to estimate these coefficients from the bond set in Chapter 4. And the two chapters after that will derive yield curves using the estimates.

13. Estimation

This chapter describes the estimation of the price equation from the bond set. The price equation was set out as Equation (9.10) and is reproduced below as Equation (13.1):

$$p = \sum_{i=1}^{n} \exp\left(-\sum_{k=1}^{5} \beta_k B_k^{\mathrm{IC}}(\tau_i)\right) c_i + \sum_{j=1}^{m} \theta_j x_j \qquad (13.1)$$

As previously discussed, each of the three yield curves HQM, TNC, and TRC contains the five constrained B-splines described in Chapter 9 comprising five coefficients plus the hump variable in Chapter 10 for an additional coefficient. In addition, the HQM yield curve also contains the two credit variables for a total of eight coefficients to be estimated, and the TNC yield curve also contains the 14 dummies for a total of 20 coefficients to be estimated. The TRC yield curve doesn't have additional regression variables and contains the B-splines plus the hump variable for a total of six coefficients to be estimated.

The method of estimation of the coefficients is nonlinear least squares. This method chooses the estimates of the coefficients so as to minimize the sum of squared residuals between the prices of the bonds in the bond set and the fitted prices given by the coefficients. Use of nonlinear least squares enables standard statistical tests to be performed; in particular, the usual covariance matrix can be calculated, and t-ratios can be calculated for coefficients such as the hump variable and the credit variables to test for significance at a particular time and to compare movements over time.[9]

In addition, the fact that each of the three yield curves uses the same five constrained B-splines implies that the coefficients on the B-splines can be directly compared across yield curves both at a point in time and over time. This provides additional market information.

[9] For discussion of estimation techniques, see Gill, Murray, and Wright (1981) and Judge, Griffiths, Hill, Luetkepohl, and Lee (1985). The basic technique used for estimation is Gauss-Newton with a line search. With the mathematical form chosen here for the yield curves, this technique converges rapidly on the bond sets of data.

For the HQM yield curve, it should be noted that the yield curve is fitted directly to the high quality bond set without the use of a preliminary Treasury yield curve. Sometimes other yield curve work has estimated a corporate yield curve from the Treasury yield curve. However, such an approach substantially increases the complexity of the curve fitting process and reduces accuracy.

Before calculating the yield curves, the bond set data are weighted, as described in the next section.

Weighting of the Bond Data before Estimation

The weighting of the bond data for the HQM yield curve is applied in two stages:

In the first stage of weighting, equal weights are assigned to commercial paper rates, and the par amounts outstanding of the bonds are rescaled so that their sum equals the sum of the weights on the commercial paper rates. Then the data are multiplied by the square root of their weights, that is, the commercial paper weights and the rescaled par amounts. This produces the result that the squared residuals of the bond prices in the least squares fit of the bond price equation are weighted by par amounts outstanding.

The purpose of this first stage is to give greater weight to larger bonds because they are more important in the market and generally more liquid. And commercial paper rates are assigned a big weight in order to anchor the yield curve at the short end of maturities.

In the second stage of weighting for the HQM yield curve, for bonds in the bond set with Macaulay duration greater than unity, the weighted bond data from the first stage are divided by the square root of duration. The purpose of the second stage is to correct for heteroscedasticity: bonds with higher duration are more volatile, with the result that the error terms of such bonds in the price equation have higher variance.

In contrast to the HQM yield curve, the two Treasury yield curves TNC and TRC do not use the first stage of weighting. First, there are no short-term data analogous to commercial paper rates that can be included in the Treasury yield curves. Also, the Treasury market is so large and liquid that it's not necessary or useful to weight individual issues by size.

However, the second stage of weighting is done for Treasuries: specifically, before estimation, the Treasury bond data are weighted by the square root of the inverse of Macaulay duration.

Estimation

After weighting, the estimation of the yield curves is done by nonlinear least squares. In addition, for the nominal yield curves HQM and TNC, the spline coefficients are constrained to be positive by putting a low floor equal to a tenth of a basis point below the acceptable values of the coefficients.

The reason for this constraint is that estimation of nominal yield curves assumes that the spline coefficients are positive in order to avoid the unacceptable anomaly of negative forward rates. Without this constraint, there can be times when short-term interest rates are low and the B-splines can have small spurious movements near zero maturity that can cause a spline coefficient and a forward rate to be slightly negative. This constraint eliminates such an anomaly. The constraint is not applied to TIPS in the TRC yield curve because the forward rate for TIPS can be negative.

The algorithm used for computation of the nonlinear least squares estimates is Gauss-Newton with a line search. The XRM methodology is very stable, and the algorithm typically converges with only about five iterations at most and no line search. Frequently, other yield curve approaches do not converge so robustly, and some don't converge at all. The use of the maturity ranges helps for stability.

The next two chapters will set out yield curves computed from the price equation using the coefficient estimates described in this chapter.

14. Spot and Par Yield Curves

This chapter shows how to calculate yield curves from the price equation using the five spline coefficients and the coefficients for the hump variable and other regression variables. The values of the coefficients can be given by the least squares estimates in the previous chapter or by other estimates. Yield curve calculations using regression variables were developed in the Office of Financial Analysis.

As already defined, a yield curve shows the yield at each maturity for a specific type of bond in a chosen sector of the bond market. In this discussion, the sectors are high quality corporate, nominal Treasury, and TIPS. The two types of bonds are the standard bond as described in Chapter 4 selling at par for which par yields are calculated and the zero coupon bond for which spot rates are calculated.

The next section sets out preliminary information on the yield curves, and the following sections present par and spot yield curves and forward rates.

Yield Curve Structure

The starting point for a yield curve is the price equation for the chosen sector of the bond market and its associated coefficients.

The price equation pertains to the settlement date for which the yield curve is done. The yield curve provides yields for future maturities relative to the settlement date, and so the maturities must be specified. In this discussion, the maturities start with the settlement date and extend out for 100 years, with the first 30 years of maturities pertaining to bonds existing on the settlement date and the maturities from 30 years through 100 years making up the projection range. In addition, in the discussions of long-term convergence at the end of the chapter, the maturities are assumed theoretically to go out to infinity.

Analogous to the discussion in Chapter 4, the time period of the maturities from settlement date to 100 years is broken into 200 half-years starting at ½ year and increasing by half-years out through 100 years. Therefore, the maturities in years for which the yields in the yield curve are calculated take on the 200 values $\frac{1}{2}, 1, 1\frac{1}{2}, \ldots, 100$. For each yield in the yield curve, the maturity of the yield in terms of years is designated by τ, so τ takes on the values ½ year through 100 years. Because the first maturity is exactly $\frac{1}{2}$, accrued interest is zero.

In principle, yields in the yield curve could be produced for any maturities, and it will be apparent how the formulas can be extended to apply to every maturity. But for most applications the set of 200 half-year maturities is sufficient.

In previous chapters, the cash flows c_i for individual bonds and associated years from settlement to payments τ_i take account of weekends and holidays. In contrast, the maturities for yield curves are here set at half-year intervals, so weekends and holidays are ignored. This follows market practice which ignores weekends and holidays when calculating yields as discussed in Chapter 4.

The yield curve formulas will be written for a generic price equation which could be HQM, TNC, or TRC, although differences among these sectors will be noted. The combined value of the regression variables multiplied by their estimated coefficients at maturity τ will be written as $V_{\mathrm{R}}(\tau)$. In the general case, price equations can include any regression terms, and the formulas will be similarly general. However, for HQM, TNC, and TRC as presently constructed, the only variable included in $V_{\mathrm{R}}(\tau)$ is the hump coefficient times the hump variable value. For HQM, the credit variables are all set to zero in the yield curve calculation because the discount function/hump variable by construction pertains to market-weighted average high quality. For TNC, the dummy variables are set to zero because the TNC yield curve is off-the-run. For TRC, there are no other regression variables than the hump.

And to reemphasize, the method of yield calculation for the yield curves is semiannual compounding. This is the standard in bond markets, and by employing this method, the data produced by the yield curves can be directly used in financial market applications. Some bond analysis uses continuous compounding or some other form of compounding, but such results aren't consistent with actual bond markets. The only exception to the use of semiannual compounding is the case of breakeven inflation rates in the next chapter, which use annual compounding to be consistent with the market standard for inflation data.

The Par Yield Curve

The type of bond for the par yield curve is the standard bond described in Chapter 4 selling at par with the following assumptions. At each maturity τ, the principal of this bond is 100, and its coupon rate as a percentage is $\kappa(\tau)$. This implies that the bond pays $\frac{\kappa(\tau)}{2}$ at the end of each half-year up through τ years maturity plus the principal of 100 at maturity. The first payment date is one half-year from settlement and there is no accrued interest.

In addition, the type of bond for the par yield curve is a bond selling at par. A bond selling at par can be defined in three ways: flat or clean price excluding accrued interest equals 100, full or dirty price including accrued interest equals 100, or coupon rate equals yield. Because it is assumed that there is no accrued interest for the bond being using here for the par yield curve, the first two definitions are the same, and as shown below, the third also turns out to be the same. So the type of bond for the par yield curve can be described simply as a standard bond whose price is 100.

With these definitions, the following price equation uses the discount function $\delta(\tau)$ to show the coupon rate $\kappa(\tau)$ for a standard bond selling at par:

$$100 = \frac{\kappa(\tau)}{2} \sum_{\iota=1}^{2\tau} \delta\left(\frac{\iota}{2}\right) + 100\delta(\tau) + V_{\mathrm{R}}(\tau) \qquad (14.1a)$$

Rearranging this equation gives the coupon rate as follows:

$$\kappa(\tau) = 2 \times \frac{100(1-\delta(\tau)) - V_{\mathrm{R}}(\tau)}{\sum_{\iota=1}^{2\tau} \delta\left(\frac{\iota}{2}\right)} \qquad (14.1b)$$

This equation shows that the coupon rate for the nominal bonds in HQM and TNC is normally positive. The denominator of this equation is positive because the discount function is positive. And apart from regression variables, the numerator is also positive for nominal bonds because the discount function is declining for such bonds such that $1 - \delta(\tau)$ is positive.

However, with the inclusion of regression variables, the numerator even for nominal bonds can in principle be negative. Nevertheless, the only regression variable included for either the HQM or TNC yield curve is the hump variable. This is because, as noted, the credit variables are set to zero when calculating the HQM yield curve and the on-the-run dummies are set to zero when calculating the TNC yield curve.

And if there is a hump, the term $V_{\mathrm{R}}(\tau)$ will be negative ensuring a positive numerator. This makes sense because the presence of a hump in the yield curve implies a hump in coupon rates. On the other hand, if the hump is inverse, in all normal circumstances the hump will be small. Consequently, the numerator is assumed to be positive for nominal bonds, and this implies that the coupon rate for a nominal bond selling at par is also positive.

Even though the on-the-run dummies are set to zero in calculating the TNC yield curve, the values of the dummies may be used in calculating the coupon rate for selected on-the-run maturities. For example, to calculate the 10-year on-the-run coupon rate and compare it to the off-the-run coupon rate, the coefficient on the 10-year dummy is used along with $\tau = 10$. However, in every normal circumstance, the on-the-run effects are small relative to the yield curve as a whole, so again the numerator is assumed to be positive.

In contrast to the HQM and TNC, the TRC real coupon rate can be nonpositive. This is because the TRC discount function can be above unity, as shown in Chapter 5, and even apart from regression variables, this can make the numerator negative.

As a general caution, for all yield curves it's important to analyze whether the regression variables are distorting the coupon rates, especially if additional regression variables are included beyond what's used here. The fact that XRM allows for regression terms means that in principle any regression variables can be inserted. If the resulting yield curves come out unreasonable, such as with unexpected negative coupon rates, the solution can usually be found by reexamining the bond data and the specification of the variables.

Using the coupon rate given by Equation (14.1b), it's possible to calculate the par yield $y(\tau)$, that is, the yield on a bond selling at par, as follows. The par yield is the value of $y(\tau)$ that solves the following equation. This equation is the same as Equation (4.7) in Chapter 4:

$$100 = \frac{\kappa(\tau)}{2} \sum_{\iota=1}^{2\tau} \frac{1}{\left(1+\frac{y(\tau)}{200}\right)^{\iota}} + \frac{100}{\left(1+\frac{y(\tau)}{200}\right)^{2\tau}} \tag{14.2}$$

Using the fact that

$$\sum_{\iota=1}^{2\tau} \frac{1}{\left(1+\frac{y(\tau)}{200}\right)^{\iota}} \left(\frac{1}{\left(1+\frac{y(\tau)}{200}\right)} - 1 \right) = \frac{1}{\left(1+\frac{y(\tau)}{200}\right)^{2\tau+1}} - \frac{1}{\left(1+\frac{y(\tau)}{200}\right)}$$

the solution to Equation (14.2) is

$$y(\tau) = \kappa(\tau) \tag{14.3}$$

The par yield equals the coupon rate. So all that is necessary to derive the par yield curve is to compute the coupon rate at all maturities. This result confirms the definition given above that for standard bonds and for the specification of the par yield curve being used here, when the clean price is 100 the coupon rate equals the yield.

It should be noted that even though Equation (14.1b) gives the coupon rate required for a bond to sell at par, it's possible that at a particular time no bond in the market actually has this coupon rate. For example, in the case of TIPS, the coupon rate of actual securities is always positive, so if a negative coupon rate is required for a par bond, a par bond cannot exist.

Nevertheless, the main purpose of the par yield curve is to be an indicator of bond yields at different maturities at a particular time for all bonds regardless of coupon rates, and as such to show market conditions. The par yield curve works well for this purpose even if no actual par bonds exist, as demonstrated in later chapters showing results with market data. Some type of bond must be chosen to do the yield curve, and the par yield curve is the standard market concept.

Moreover, the par yield curve provides the framework for constructing spot rates that are consistent with market yields. It is through the par yield curve that the spot rates are computed. This is done in the next section.

The Spot Yield Curve

The type of bond for the spot yield curve is a zero coupon bond, that is, a bond with a single payment of the principal 100 at maturity and no coupon payments. The spot yield curve gives the spot rate for such bonds at each of the 200 maturities from ½ year through 100 years. Because there are no coupons, there is, of course, no accrued interest. The spot rate is calculated following market convention with semiannual compounding using the following formula, where the price of the bond is p, the maturity of the bond is τ, and the spot rate as a percent is $r(\tau)$. This is the same as Equation (4.6) in Chapter 4:

$$p = \frac{100}{\left(1+\frac{r(\tau)}{200}\right)^{2\tau}} \tag{14.4}$$

The price equation from which the par yield curve was derived was constructed for standard bonds. This same price equation will be used for zero coupon bonds, and the spot rate will be calculated so as to be consistent with this price equation and therefore with par yields. Consistency means that the spot rate can be used to price standard bonds and the same result will be obtained as with the price equation. Therefore, consistency means that the spot rate will give the same results as a market for zero coupon bonds even if such a market does not exist. For example, there's no developed market for zero coupon corporate bonds, but spot rates from the HQM yield curve show results if such a market did exist. Moreover, spot rates that are consistent with standard bonds can be approximated in practice with combinations of standard bonds.

Specifically, the spot rate is consistent with the par yield curve if the individual payments from a standard par bond can be discounted by the spot rates and summed up and the sum equals 100. The following formula sets out this consistency requirement for the spot rate $r(\tau)$:

$$100 = \frac{\kappa(\tau)}{2}\sum_{\iota=1}^{2\tau}\frac{1}{\left(1+\frac{r\left(\frac{\iota}{2}\right)}{200}\right)^{\iota}} + \frac{100}{\left(1+\frac{r(\tau)}{200}\right)^{2\tau}} \qquad (14.5)$$

Going back to Equation (14.1a), the first thing to note about this equation is that if the regression variables are zero at every maturity τ, that is, $V_{\mathrm{R}}(\tau) = 0$ for all τ, the spot rate $r(\tau)$ at each maturity τ that is consistent with the price equation and discount function is the discount spot rate $r_{\mathrm{D}}(\tau)$ from Equation (5.7a). This is because the discount spot rate is derived directly from the inverse of the discount function and ignores the regression variables.

However, if regression variables are not zero, they must be taken into account in the calculation of the spot rate because they affect the coupon rate from Equation (14.1b). The first step in doing this is to recognize from Equation (14.5) that the first spot rate at maturity ½ year is the same as the first coupon rate. After that, the rest of the spot rates are calculated sequentially for successive maturities one at a time by using the spot rates at lower maturities. The procedure is shown below:

$$r(\tau) = \kappa(\tau),\ \ \tau = 0.5 \qquad (14.6a)$$

$$r(\tau) = 200 \times \left(\left(\frac{\frac{\kappa(\tau)}{2}+100}{100-\frac{\kappa(\tau)}{2}\sum_{\iota=1}^{2\tau-1}\frac{1}{\left(1+\frac{r\left(\frac{\iota}{2}\right)}{200}\right)^{\iota}}}\right)^{\frac{1}{2\tau}} - 1\right),\ \ \tau > 0.5 \qquad (14.6b)$$

A general way to determine the sign of the spot rate in this procedure is to recognize from Equation (14.5) that the sign of the spot rate is the same as the sign of the coupon rate at each maturity. This can be seen by recognizing that if $r(\tau) = 0$, the right side of Equation (14.5) will be greater than 100 if $\kappa(\tau) > 0$ and it will be less than 100 if $\kappa(\tau) < 0$. So to make both sides of the equation equal, $r(\tau)$ must be raised from zero when the coupon rate is positive and lowered from zero when it is negative. This logic works for nominal yield curves and TIPS.

The Forward Spot Rate

The forward spot rate for future bonds can be calculated from the spot rate. The approach to do this is to create a synthetic future zero coupon bond by buying and selling zero coupon bonds at present. Of course, if a fully developed spot market doesn't exist, synthetic future bonds can't be fully realized and can only be approximated by bonds that are actually available. Nevertheless, the forward spot rate shows what such future bonds would be expected to look like.

One use of the forward spot rate is to measure market expectations of future interest rates. However, this measure may be biased up by a term premium that reflects future uncertainty about interest rates, which leads lenders to demand higher interest rates to compensate.

The forward spot rate must be distinguished from the forward rate $\phi(\tau)$ presented earlier in connection with the discount function. The main distinction is that the forward spot rate incorporates the effects of regression variables whereas $\phi(\tau)$ comes from the discount function alone. If there were no regression variables, the forward rate $\phi(\tau)$ would conceptually be the same as the forward spot rate, although $\phi(\tau)$ is instantaneous while the forward spot rate pertains to periods of time.

In order to calculate the forward spot rate, the spot rate $r(\tau)$ can be expanded to include a second time variable that indicates number of years in the future when the zero coupon bond begins. So the spot rate can be denoted as $r(\tau_1, \tau_2)$, which is the spot rate with a maturity of τ_1 years beginning τ_2 years in the future; that is, the forward spot rate for τ_1 years τ_2 years hence. In this notation, the spot rate beginning in the present is denoted as $r(\tau_1, 0)$.

For the forward spot rate, this notation implies:

$$\left(1 + \frac{r(\tau_1+\tau_2,0)}{200}\right)^{2(\tau_1+\tau_2)} = \left(1 + \frac{r(\tau_2,0)}{200}\right)^{2\tau_2} \left(1 + \frac{r(\tau_1,\tau_2)}{200}\right)^{2\tau_1} \qquad (14.7a)$$

$$\Rightarrow\ r(\tau_1, \tau_2) = 200 \times \left(\left(\frac{\left(1+\frac{r(\tau_1+\tau_2,0)}{200}\right)^{\tau_1+\tau_2}}{\left(1+\frac{r(\tau_2,0)}{200}\right)^{\tau_2}}\right)^{\frac{1}{\tau_1}} - 1\right) \qquad (14.7b)$$

Convergence of the Par Yield

This section derives the convergence of the par yield to the long-term par yield $\kappa^*$when maturity goes to infinity. The derivation in this section assumes that the values of the regression variables $V_{\mathrm{R}}(\tau)$ are zero for $\tau > \tau^*$, where $\tau^* = 30$ years maturity. The derivation uses Equation (14.1b) above and the long-term forward rate ϕ^*. The analysis in this section pertains to $\phi^* > 0$. As mentioned in Chapter 5, the cases for TIPS in which $\phi^* < 0$ or $\phi^* = 0$ are not normal in markets and are not analyzed. Any convergence analysis in such cases tends to be unstable and inaccurate because these cases arise from temporary market anomalies.

For $\phi^* > 0$, the first step is to write down the discount function for $\tau \geq \tau^*$ following Equation (5.6):

$$\delta(\tau) = \delta(\tau^*)\exp\left(-\phi^*(\tau - \tau^*)\right) \tag{14.8}$$

For convergence, the numerator and denominator of Equation (14.1b) can be analyzed separately. Using Equation (14.8), the numerator converges to 200 as τ goes to infinity.

For $\tau > \tau^*$, again using Equation (14.8), the denominator can be written as

$$\sum_{\iota=1}^{2\tau^*}\delta\left(\frac{\iota}{2}\right) + \delta(\tau^*)\sum_{\zeta=1}^{2(\tau-\tau^*)}\exp\left(-\frac{\phi^*\zeta}{2}\right)$$

Using the fact that the last part of the denominator is a geometric series, the denominator converges to:

$$\sum_{\iota=1}^{2\tau^*}\delta\left(\frac{\iota}{2}\right) + \delta(\tau^*)\sum_{\zeta=1}^{\infty}\exp\left(-\frac{\phi^*\zeta}{2}\right) = \sum_{\iota=1}^{2\tau^*}\delta\left(\frac{\iota}{2}\right) + \delta(\tau^*)\frac{1}{\exp\left(\frac{\phi^*}{2}\right)-1} \tag{14.9}$$

Putting numerator and denominator together, this gives:

$$\kappa^* = \frac{200}{\sum_{\iota=1}^{2\tau^*}\delta\left(\frac{\iota}{2}\right)+\delta(\tau^*)\frac{1}{\exp\left(\frac{\phi^*}{2}\right)-1}} \tag{14.10a}$$

Therefore, in contrast to the long-term discount spot rate r_{D}^*, the long-term par yield κ^* does not necessarily equal ϕ^*.

However, if $\tau^* = 0$ so that ϕ^* is constant at all maturities, Equation (14.10a) becomes:

$$\kappa^* = 200 \times \left(\exp\left(\frac{\phi^*}{2}\right) - 1\right) \tag{14.10b}$$

which is the same as Equation (5.10a). So in the special case of a constant long-term forward rate for all maturities, the par yield converges to this constant.

Convergence of the Spot Rate

The discussion in Chapter 5 has already shown that the discount spot rate converges to the long-term forward rate as maturity goes to infinity. This section shows that the spot rate including effects of the regression variables also converges to the long-term forward rate. As in the previous section, the derivation in this section assumes that the regression variables are zero for maturities above 30 years and assumes $\phi^* > 0$.

The first step is to put Equations (14.1a) and (14.5) together because they both equal 100, with the discount function in Equation (14.1a) written in terms of the discount spot rate:

$$\frac{\kappa(\tau)}{2}\sum_{\iota=1}^{2\tau}\frac{1}{\left(1+\frac{r\left(\frac{\iota}{2}\right)}{200}\right)^{\iota}}+\frac{100}{\left(1+\frac{r(\tau)}{200}\right)^{2\tau}}=\frac{\kappa(\tau)}{2}\sum_{\iota=1}^{2\tau}\frac{1}{\left(1+\frac{r_D\left(\frac{\iota}{2}\right)}{200}\right)^{\iota}}+\frac{100}{\left(1+\frac{r_D(\tau)}{200}\right)^{2\tau}}+V_{\mathrm{R}}(\tau) \tag{14.11}$$

The following analysis uses the notation:

$$\Delta_{\mathrm{S}}(\tau)=\frac{1}{\left(1+\frac{r(\tau)}{200}\right)^{2\tau}}-\frac{1}{\left(1+\frac{r_{\mathrm{D}}(\tau)}{200}\right)^{2\tau}} \tag{14.12}$$

$\Delta_{\mathrm{S}}(\tau)$ is the difference between the discount function computed from the spot rate with the regression variables and the regular discount function $\delta(\tau)$. Convergence of the spot rate means that as maturity goes to infinity, $\Delta_{\mathrm{S}}(\tau)$ goes to zero.

Inserting Equation (14.12) into Equation (14.11) and rearranging gives:

$$\left(\frac{\kappa(\tau)}{2}+100)\right)\Delta_{\mathrm{S}}(\tau)=-\frac{\kappa(\tau)}{2}\sum_{\iota=1}^{2\tau-1}\Delta_{\mathrm{S}}\left(\frac{\iota}{2}\right)+V_{\mathrm{R}}(\tau) \tag{14.13}$$

The summation term in this equation can be simplified by writing Equation (14.13) for the previous half-year $\tau-\frac{1}{2}$ and rearranging:

$$\frac{\kappa\left(\tau-\frac{1}{2}\right)}{2}\sum_{\iota=1}^{2\tau-1}\Delta_{\mathrm{s}}\left(\frac{\iota}{2}\right)=-100\Delta_{\mathrm{S}}\left(\tau-\frac{1}{2}\right)+V_{\mathrm{R}}\left(\tau-\frac{1}{2}\right) \tag{14.14}$$

Inserting Equation (14.14) into Equation (14.13) gives:

$$\left(\frac{\kappa(\tau)}{2}+100)\right)\Delta_{\mathrm{S}}(\tau)=-\frac{\kappa(\tau)}{2}\frac{-100\Delta_{\mathrm{S}}\left(\tau-\frac{1}{2}\right)+V_{\mathrm{R}}\left(\tau-\frac{1}{2}\right)}{\frac{\kappa\left(\tau-\frac{1}{2}\right)}{2}}+V_{\mathrm{R}}(\tau) \tag{14.15}$$

As τ goes to infinity, $\kappa(\tau)$ goes to κ^* which is the long-term par yield defined in the previous section, and $V_{\mathrm{R}}(\tau)$ is zero. So for large τ:

$$\Delta_{\mathrm{S}}(\tau)\approx-\frac{100}{\frac{\kappa^*}{2}+100}\Delta_{\mathrm{S}}\left(\tau-\frac{1}{2}\right) \tag{14.16}$$

This is a difference equation with coefficient less than unity in absolute value, so it converges to zero.

Consequently, the long-term spot rate r^* to which the spot rate $r(\tau)$ converges equals the long-term discount spot rate r_{D}^* to which the discount spot rate $r_{\mathrm{D}}(\tau)$ converges, and they both equal the long-term forward rate ϕ^*.

15. The TBI Curve

This chapter continues the examination of spot rates to derive the Treasury Breakeven Inflation (TBI) Curve. As described in Chapter 3, the TBI curve combines the TNC and TRC yield curves to generate the inflation rate that equates returns on nominal Treasury securities and TIPS. Markets call this inflation rate the breakeven inflation rate.

The first section in this chapter defines the breakeven inflation rate as it is derived in the TBI curve. Sections after that examine features of the TBI rate and compute the forward breakeven rate.

The TBI Curve

This is the definition of breakeven inflation: breakeven inflation over a period is the inflation rate that if realized, would equate over that period the real return for nominal Treasury securities and the real return for TIPS.

Using the spot rate derivations from the previous chapter, let $r_{\mathrm{N}}(\tau)$ be the spot rate for nominal Treasury securities with maturity τ and $r_{\mathrm{R}}(\tau)$ be the spot rate for TIPS. Let $r_{\mathrm{P}}(\tau)$ be the inflation rate from the present through maturity τ. As done previously, the two spot rates are semiannually compounded, but, in contrast, the inflation rate is annually compounded in accord with market practice. The inflation measure for TIPS and $r_{\mathrm{P}}(\tau)$ is the (not seasonally adjusted) Consumer Price Index for All Urban Consumers (CPI-U).

By this definition, the breakeven inflation rate equates the real return received at maturity τ by investing \$1 in nominal Treasuries and deflating by the inflation rate and the real return received at τ by investing \$1 in TIPS. Consequently, if the breakeven inflation rate is realized, the real returns from both markets are equal.

The following equations calculate the breakeven inflation rate following this definition. The left-hand side of the first equation is the return on nominal Treasuries deflated by inflation and the right-hand side is the return on TIPS.

$$\frac{\left(1+\frac{r_{\mathrm{N}}(\tau)}{200}\right)^{2\tau}}{\left(1+\frac{r_{\mathrm{P}}(\tau)}{100}\right)^{\tau}} = \left(1+\frac{r_{\mathrm{R}}(\tau)}{200}\right)^{2\tau} \tag{15.1a}$$

$$\Rightarrow\ r_{\mathrm{P}}(\tau) = 100 \times \left(\left(\frac{\left(1+\frac{r_{\mathrm{N}}(\tau)}{200}\right)}{\left(1+\frac{r_{\mathrm{R}}(\tau)}{200}\right)}\right)^{2} - 1\right) \tag{15.1b}$$

The definition of breakeven inflation implies that the breakeven inflation rate approximates market expectations of inflation. This is because returns in nominal and real Treasury markets would be different if expectations diverged from breakeven inflation, with the consequence that trading of Treasuries would cause them to be made equal. The primary use of breakeven inflation is to measure inflationary expectations in the market.

Even so, there are factors that can cause breakeven inflation to differ from expectations. The most important factor is that inflation uncertainty increases the risk of the return on nominal Treasuries. Returns on TIPS don't have this risk because TIPS returns don't depend on inflationary expectations. Because of this risk, nominal Treasuries contain a term premium that pushes up returns relative to TIPS, with the result that the breakeven inflation rate is biased up relative to what markets actually expect. This bias must be taken into account when using breakeven inflation to monitor expectations. Nevertheless, when calculated using XRM methodology, it appears that this bias is lower than often thought.

In addition, there are special factors that can cause returns to differ regardless of inflationary expectations. For example, liquidity differences between nominal and real Treasury markets are sometimes more important than relative returns. This can be especially true at short maturities.

Additional Issues for Breakeven Inflation

Users of breakeven inflation often compute it simply as the difference between the yield on a nominal Treasury coupon issue and the yield on a TIPS coupon issue of the same maturity. There are several problems with this.

The biggest problem is that this computation uses yields and not spot rates. If the purpose of breakeven inflation is to equate returns over a period for nominal and real Treasuries, the previous section showed that the appropriate interest rate measure is the respective spot rates. Spot rates give amounts received at the end of a period, whereas yields give streams of coupon payments received throughout the period. The use of yields biases down the measure of breakeven inflation.

Moreover, the duration of the nominal Treasury yield is shorter than the TIPS yield for the same maturity. This is because the TIPS yield provides larger nominal cash flows as maturity rises. So in this sense, nominal and real Treasury yields aren't comparable. Spot rates, in contrast, at a given maturity for real and nominal have the same duration and it is equal to the maturity.

Another problem is that the nominal Treasury coupon issue used for this calculation is typically on-the-run, which may be more liquid than the TIPS coupon issue and may also bias down the breakeven calculation. In contrast, the TBI breakeven rate uses off-the-run securities in TNC that do not have special on-the-run characteristics and are similar to TRC securities. Therefore, this downward bias is eliminated.

And, of course, a simple subtraction of two yields is incorrect as demonstrated in the last two equations. And also breakeven inflation should be annually compounded as opposed to semiannually compounded as is typical for yields.

The Forward Breakeven Rate

Analogous to the forward spot rate in the previous chapter, the forward TBI breakeven inflation rate can be calculated. The forward TBI rate can be used as an indicator of expected inflation in future periods.

As in the case of the forward spot rate, the forward TBI rate may be biased up by a term premium, and this term premium may increase for future periods further out that are considered riskier.

Analogous to the forward spot rate, the TBI rate $r_{\mathrm{P}}(\tau)$ can be expanded to include a second time variable that indicates number of years in the future when the inflation begins. So the TBI rate can be denoted as $r_{\mathrm{P}}(\tau_1, \tau_2)$, which is the TBI rate over a time of τ_1 years beginning τ_2 years in the future; that is, the forward TBI rate for τ_1 years τ_2 years hence. In this notation, the TBI rate beginning in the present has $\tau_2 = 0$.

For the forward TBI rate, using annual compounding, this notation implies:

$$\left(1 + \frac{r_{\mathrm{P}}(\tau_1+\tau_2,0)}{100}\right)^{\tau_1+\tau_2} = \left(1 + \frac{r_{\mathrm{P}}(\tau_2,0)}{100}\right)^{\tau_2} \left(1 + \frac{r_{\mathrm{P}}(\tau_1,\tau_2)}{100}\right)^{\tau_1} \quad (15.2a)$$

$$\Rightarrow r_{\mathrm{P}}(\tau_1, \tau_2) = 100 \times \left(\left(\frac{\left(1+\frac{r_{\mathrm{P}}(\tau_1+\tau_2,0)}{100}\right)^{\tau_1+\tau_2}}{\left(1+\frac{r_{\mathrm{P}}(\tau_2,0)}{100}\right)^{\tau_2}} \right)^{\frac{1}{\tau_1}} - 1 \right) \quad (15.2b)$$

16. Selected HQM and TNC Results

This chapter and the next two chapters present selected results from calculated yield curves that illustrate the concepts in earlier chapters. This chapter concentrates on the nominal HQM and TNC yield curves and presents results for Friday August 30, 2024, which was the last business day in August and is representative of recent market trading behavior.

The date of the yield curves is 08/30/2024 because that is the date of the bond trades that are used for estimation. For both yield curves, the settlement date is the next business day, which is September 3, 2024 after allowing for Labor Day.

The first section in this chapter describes the HQM yield curve on 08/30/2024 and the second section describes the TNC yield curve. The sections following the first two present charts.

The HQM Yield Curve on 08/30/2024

The HQM yield curve on this date was computed using 3,738 securities including 3,731 bonds and 7 AA commercial paper rates from the Federal Reserve. Convergence in the computation was achieved with 4 iterations, showing the stability of the XRM methodology.

The 3,731 bonds include 73 AAA rated, 652 AA rated, and 3006 A rated. The two credit shares ω_1 and ω_2 are 89.2 percent for AA bonds relative to AAA+AA and 75.4 percent for A rated bonds relative to all high quality bonds.

For the maturity ranges, there are 389 bonds with maturities in the first maturity range from zero through 1½ years maturity, 530 bonds in the second maturity range, 937 bonds in the third, 650 bonds in the fourth, and 1,225 bonds in the farthest maturity range from 15 through 30 years maturity.

These figures show that the HQM yield curve covers the high quality bond market with plenty of bonds of all three ratings and in each maturity range.

Of the 3,731 bonds, 1,300 do not have end calls but almost twice as many 2,431 do have end calls. This shows the importance of end calls in current bond markets and confirms that they must be included to get accurate yield curves. As previously discussed, because the presence of the end call doesn't affect bond price, the end-call bonds are treated as standard bonds and the end-call feature is ignored in estimation.

The five spline coefficients on the constrained B-splines expressed as percentages are 5.07, 3.75, 4.32, 5.81, and 5.46. The coefficients on the two credit variables are 14 basis points and 15 basis points. And the coefficient on the hump variable is -50 basis points. The hump coefficient shows a small hump effect, although the charts below show that there is no actual hump on this date and the hump effect blends with the discount function to provide a more accurate shape to the yield curve.

The TNC Yield Curve on 08/30/2024

The TNC yield curve on 08/30/2024 was computed using 316 Treasury securities which include every Treasury note or bond coupon issue available for trade on that day with at least two payments still to be made and maturity greater than a half year. Convergence in the computation was achieved just as smoothly as in the HQM yield curve with 4 iterations.

For the maturity ranges, there are 54 securities with maturities in the first maturity range, 69 securities in the second maturity range, 95 securities in the third, 20 securities in the fourth, and 78 securities in the farthest maturity range. Each maturity range has enough securities for a good estimate of the yield curve.

The five spline coefficients on the constrained B-splines for the TNC yield curve expressed as percentages are 4.95, 2.96, 3.98, 3.65, and 5.03. Each spline coefficient is below the equivalent HQM spline coefficient. The coefficient on the hump variable is -2.93, indicating a strong hump effect that will be seen in the charts below.

Discount Functions

Figure 16.1

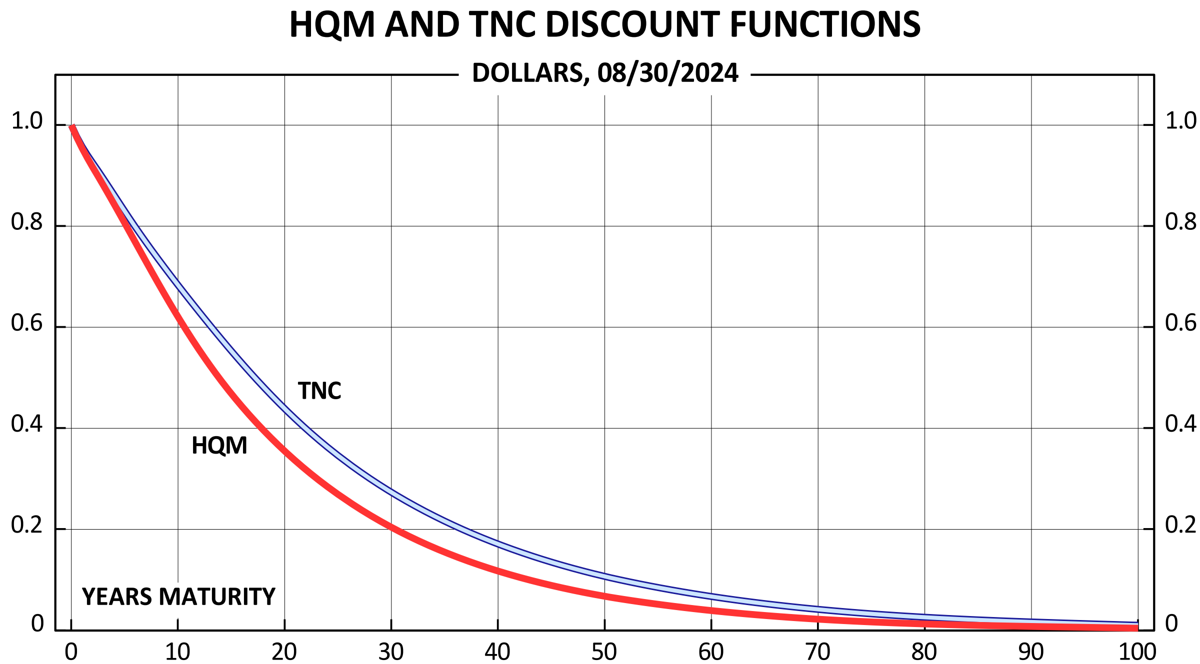

This chart shows the discount functions from the HQM and TNC yield curves for this date. Both discount functions start at unity at maturity zero, and because the spline coefficients are positive for both yield curves, the discount functions decline continuously throughout their ranges. In the projection range at 30 years maturity and above, the decline is exponential.

The TNC discount function is above the HQM discount function, indicating that the market is willing to pay more at each maturity for a contract to receive $1 at that maturity with Treasury characteristics than with high quality corporate bond characteristics. The reason is that the Treasury contract is less risky because Treasuries have no default risk, in contrast to corporate bonds.

Forward Rates

Figure 16.2

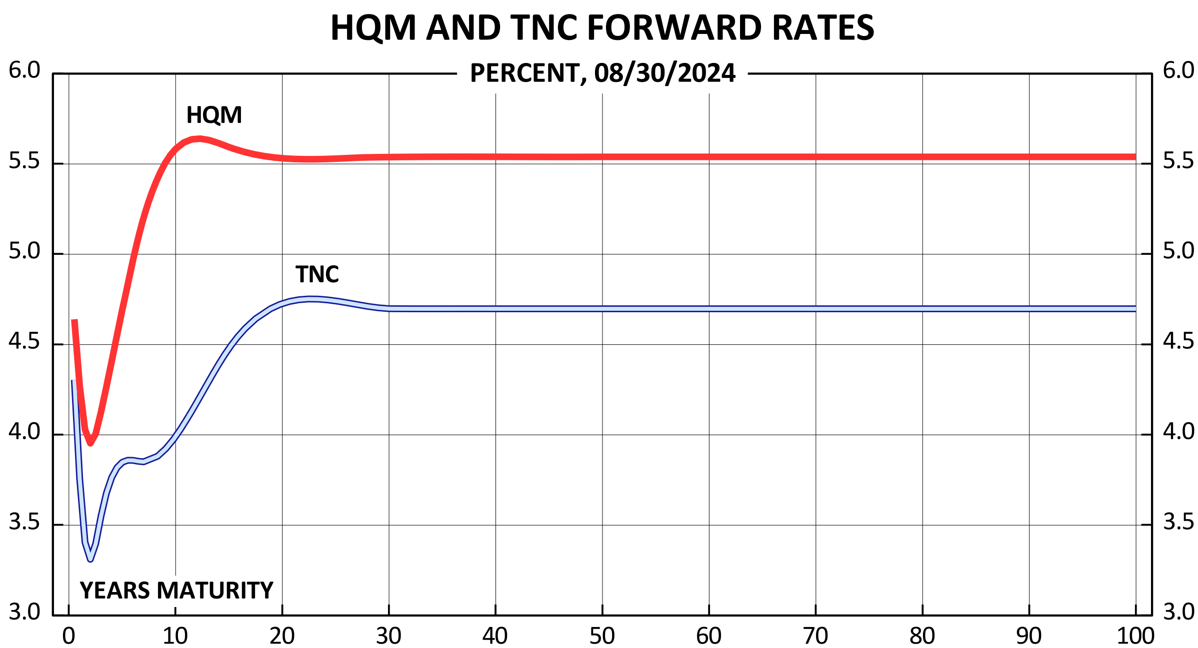

This chart shows the forward rates derived from the discount functions in the previous chart.

The three spline constraints are visible in the chart: The second-derivative constraint at maturity zero causes the forward rate to become linear at zero. The first-derivative constraint at maturity 30 years causes the forward rate to flatten out and move smoothly into the long-term forward rate.

And the long-term forward rate constraint causes the forward rate at 30 years maturity to equal the average forward rate in the farthest maturity range from 15 to 30 years and to remain constant at that average throughout the projection range.

In this chart, the long-term forward rate for the HQM yield curve is 5.54 percent and the long-term forward rate for TNC is 4.70 percent.

The chart shows that the forward rates have a hump around 10 to 20 years maturity. A hump like this is often seen in the forward rate, and it doesn't necessarily imply that there will be a hump in the par or spot yield curves.

HQM Par Yield Curve

Figure 16.3

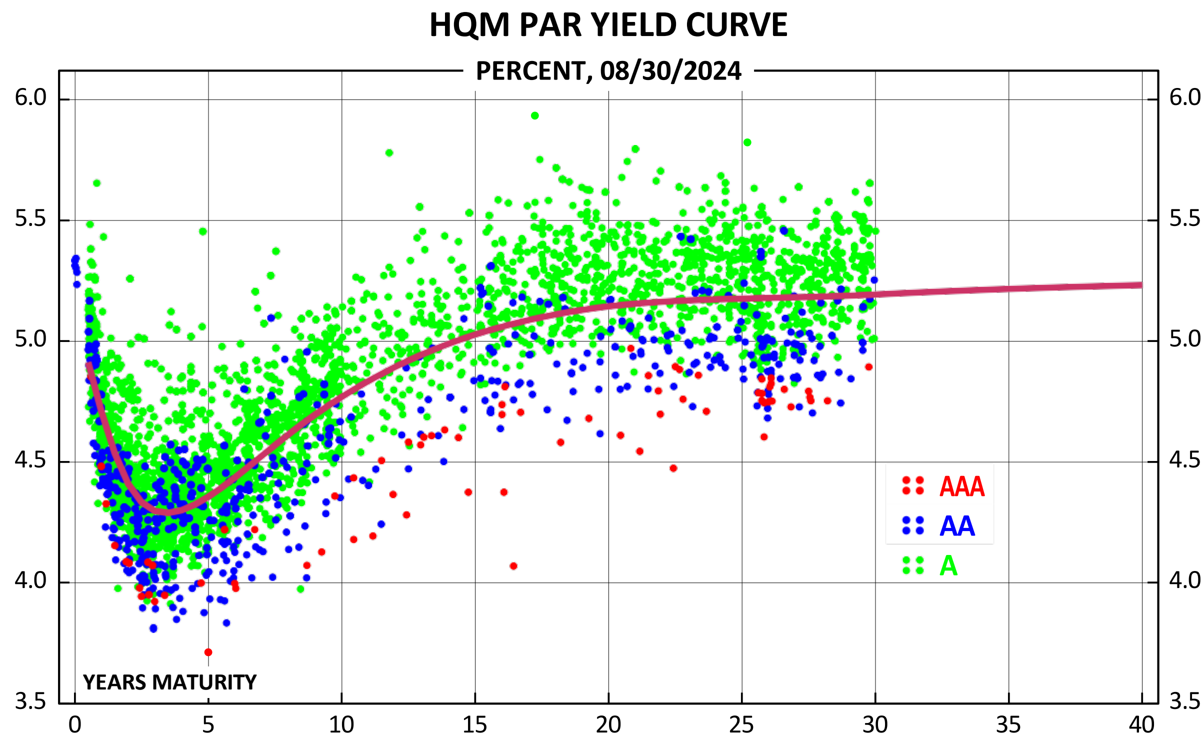

This chart shows the HQM par yield curve for this date together with a scatter diagram of the yields and maturities of the 3,738 securities that were used to estimate the curve. Yields for AAA securities are red dots, blue for AA, and green for A. As expected, AAA yields are the lowest on average and A yields are highest. The yields in the scatter diagram were calculated using the street convention formula Equation (4.7) in Chapter 4.

The scatter diagram illustrates that there's a lot of noise in corporate bond data. However, the yield curve manages to pick up the pattern and run through the dots.

This yield curve doesn't have a hump. Nevertheless, the hump variable picks up the flattening in the dots right after 20 years maturity and helps produce a more accurate yield curve than the splines alone. Therefore, even when there isn't a hump, the hump variable can contribute to accuracy.

This yield curve is shown projected out 10 years through 40 years maturity. Full projections go out through 100 years maturity.

TNC Par Yield Curve

Figure 16.4

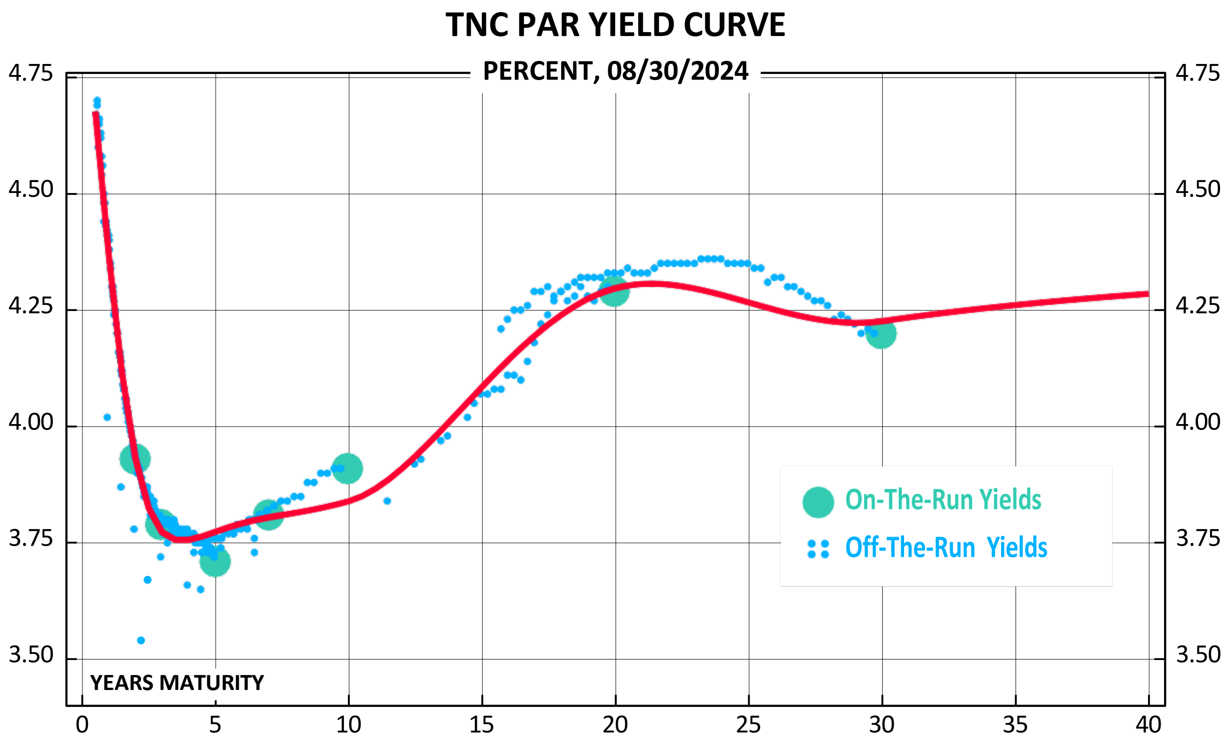

Analogous to the previous chart, this chart shows the TNC par yield curve for the end of August together with a scatter diagram of the yields and maturities of the 316 securities used to estimate the curve. The yields in the scatter diagram were calculated using the Treasury convention formula Equation (4.8) in Chapter 4.

The yield curve is not expected to run through the yields. That is because this yield curve is for bonds selling at par, whereas many of the individual yields pertain to bonds that may be way off par. One reason why bonds can be selling far from par is that they're old with coupon rates very different from current rates. Nevertheless, the par yield curve should track the shape of the market, and the TNC and TRC yield curves are accurate in doing this.

In addition, this yield curve has a significant hump derived from the hump variable, as was mentioned earlier in the chapter. This hump helps capture the hump in the scatter diagram that is apparent in the chart.

Par and Spot Yield Curves

Figure 16.5

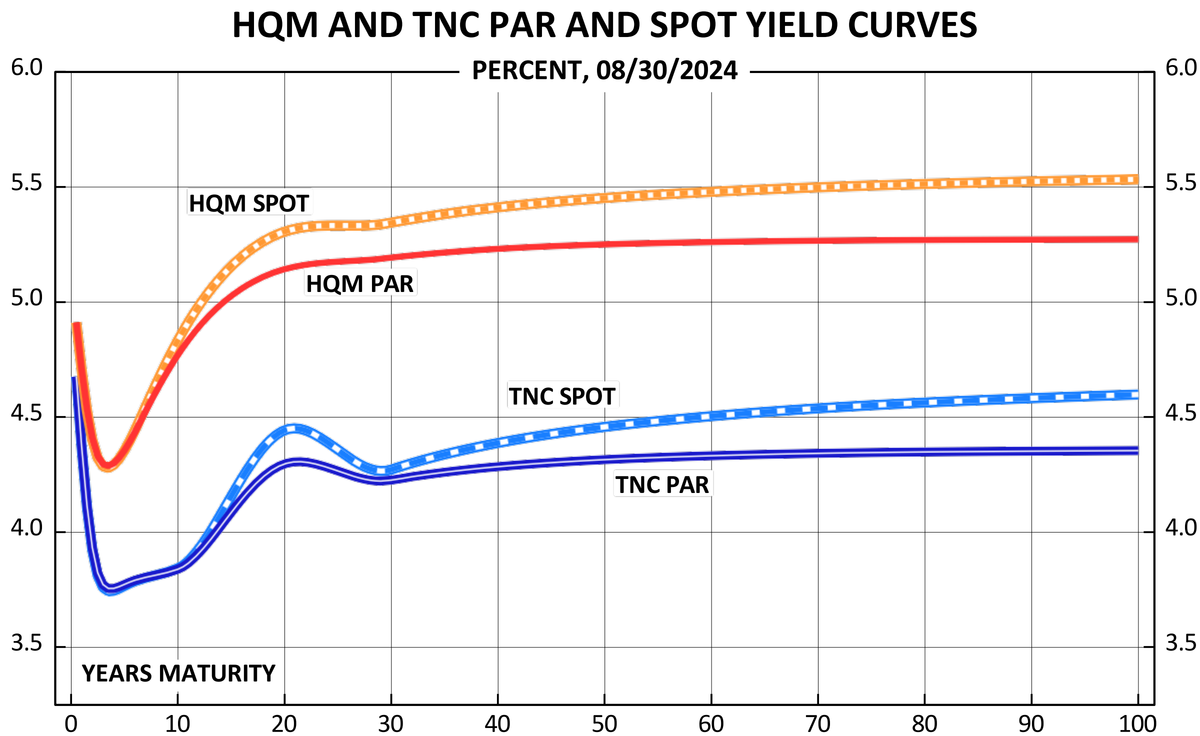

This chart brings the together the par yield curves in the previous two charts and includes the companion spot yield curves, with everything projected out through 100 years maturity. The spot yield curves include the effects of the regression variables, which here is the hump variable.

The spot yield curves are above the par yield curves at higher maturities. As expected, the HQM yield curves are above TNC. The projections beyond 30 years maturity rise gently.

HQM Spot Yield Curves and the Forward Rate

Figure 16.6

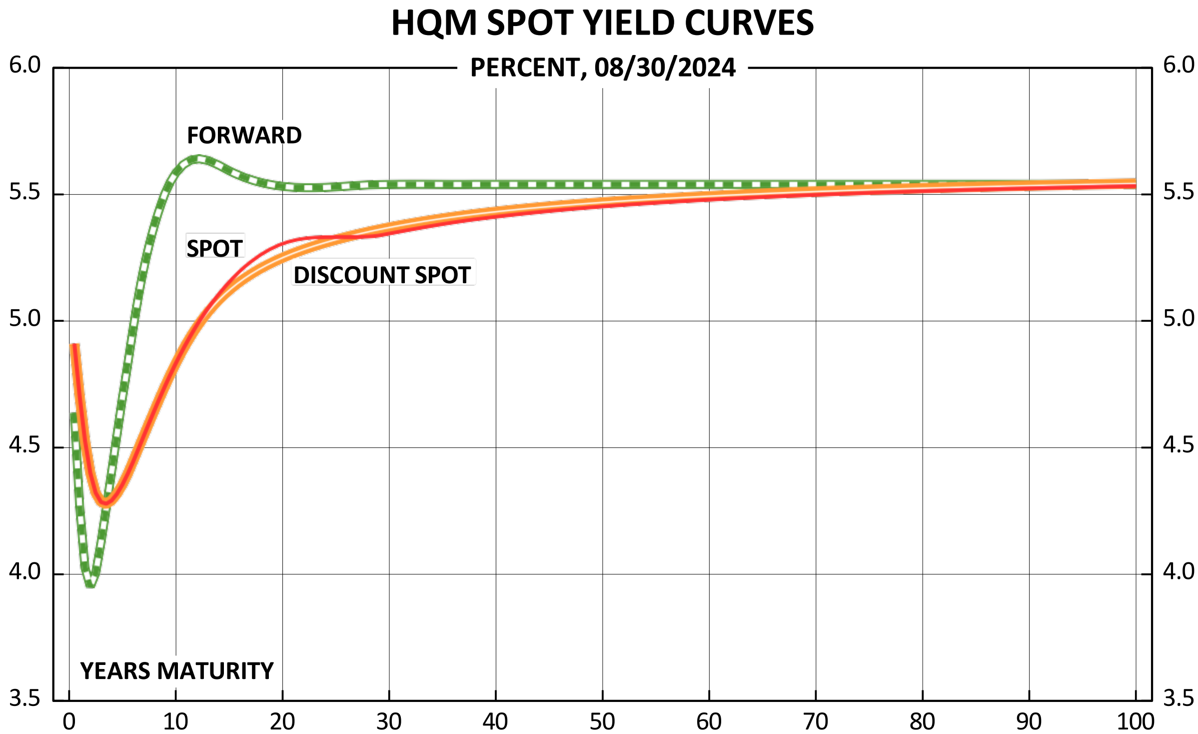

This chart compares HQM spot rates with the forward rate. The forward rate is the same as in Figure 16.2, and the spot rate is the same as in Figure 16.5. The discount spot rate is defined in Chapter 5, and it is the spot rate computed solely from the discount function ignoring the hump variable.

The two spot rates show that the hump variable creates a hump around the spot rate. Nevertheless, both spot rates converge as maturity rises, as was shown in Chapter 14.

The spot rates start out below the forward rate at 30 years maturity, then rise to converge to the fixed long-term forward rate. This was shown in Chapters 5 and 14.

Note that Chapter 5 also shows that the spot rates will eventually end up higher than the long-term forward rate as maturity is extended out because the rates use different formulas for measurement even though they are actually the same: the spot rates are semiannually compounded and the forward rate is continuously compounded.

17. Additional HQM and TNC Results

This chapter displays additional charts for the nominal yield curves that supplement the charts in the previous chapter. These charts show differing effects of the hump variable. Also included are three historical charts.

HQM Par Yield Curve with Hump

Figure 17.1

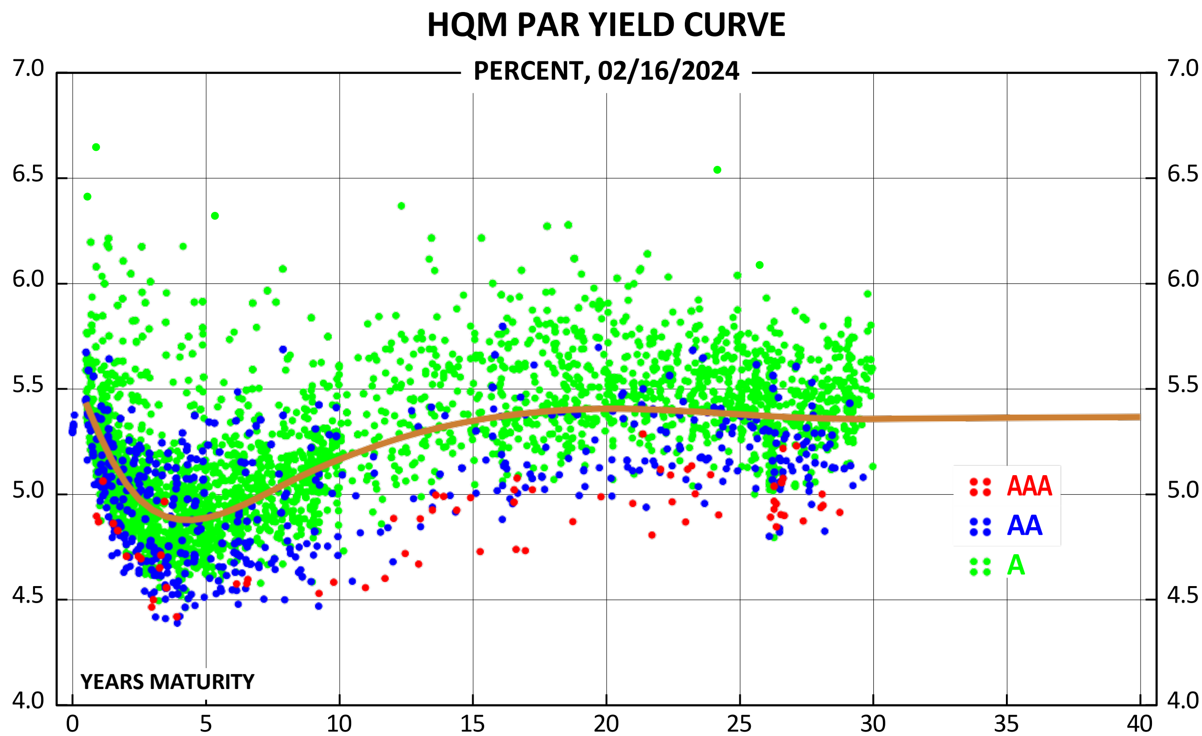

This chart shows another recent HQM par yield curve plus scatter diagram for 02/16/2024, and can be compared to the HQM par yield curve in the previous chapter for 08/30/2024. The number of bonds in this chart is 3,465 plus 6 commercial paper rates.

In contrast to the previous chapter's HQM yield curve, this HQM yield curve has a mild hump around 20 years maturity. The coefficient on the hump variable for this chart is -80 basis points, which is a big higher than the value -50 basis points already seen for the yield curve on 08/30/2024. The scatter diagram shows a hump in yields that is picked up by the yield curve.

TNC Yield Curve without Hump

Figure 17.2

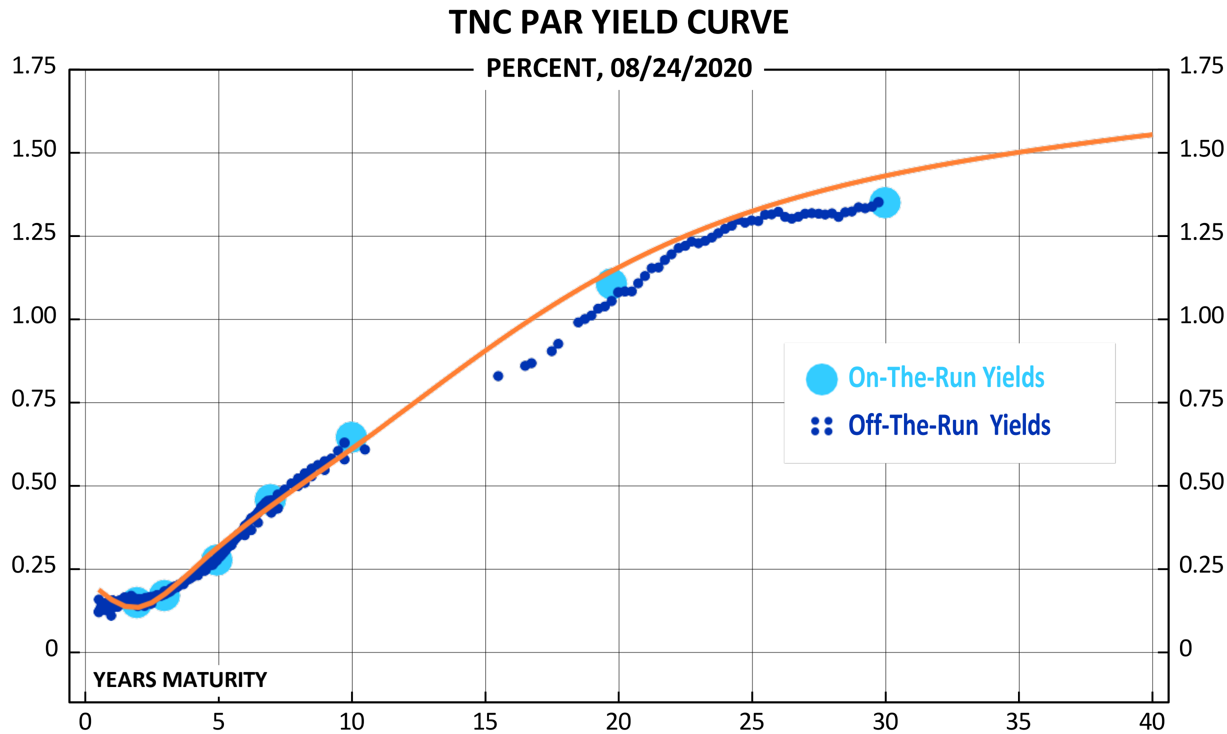

This chart shows the TNC par yield curve plus scatter diagram for 08/24/2020, which can be compared to the analogous yield curve in the previous chapter for 08/30/2024. The number of Treasury securities in this chart is 286.

This yield curve doesn't have a hump, in contrast to the previous chapter's TNC yield curve. The value of the coefficient on the hump variable is near zero in this yield curve, in contrast to the strong -2.92 value for 08/30/2024. So this yield curve shows that when there is no hump in the data, the XRM methodology assigns a zero to the hump coefficient. Nevertheless, the scatter diagram in the chart shows that this yield curve tracks the pattern of the market.

In this chart, the second spline coefficient is also zero, which effectively removes the second constrained B-spline from the yield curve. So it's possible at times that the spline that fits the data best has one of the spline coefficients set to zero.

HQM History

Figure 17.3

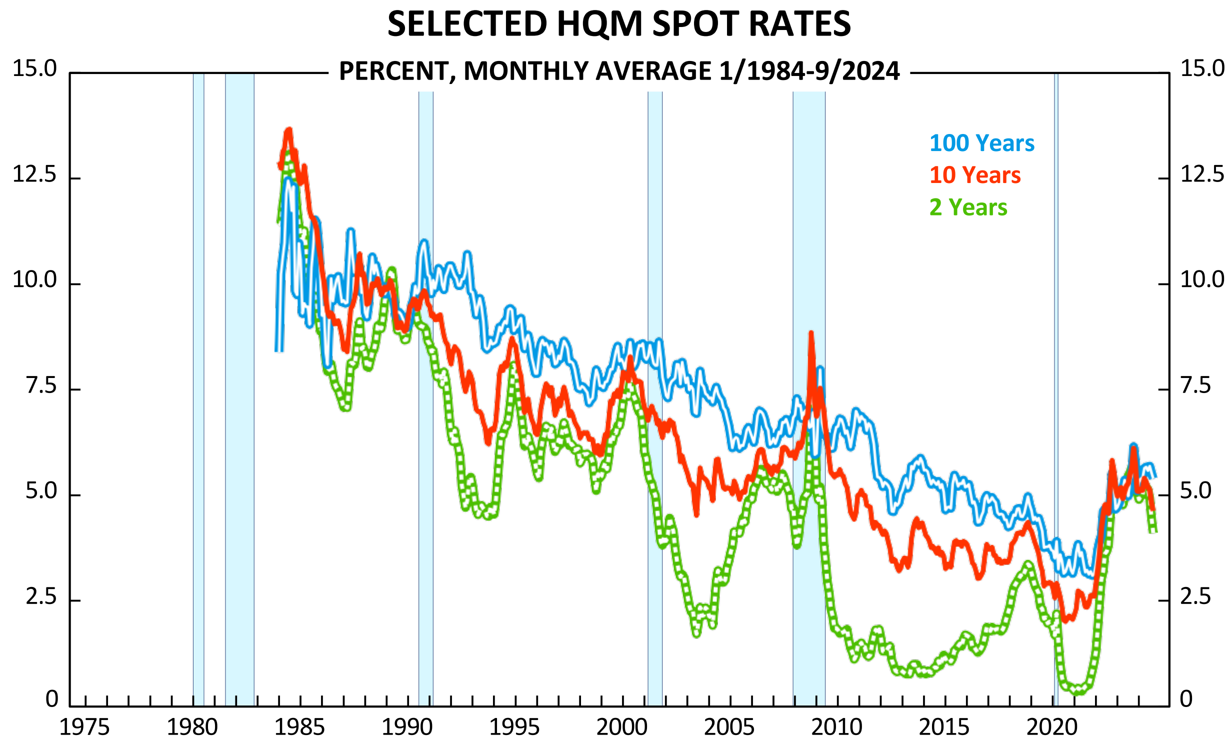

This chart shows the history of the HQM yield curve spot rates for 10, 100, and 2 years maturity from the beginning of the HQM yield curve in 1984 to present. The blue shaded areas indicate periods of recession.

The chart suggests that the rates at the different maturities tend to bunch up near recessions.

TNC History

Figure 17.4

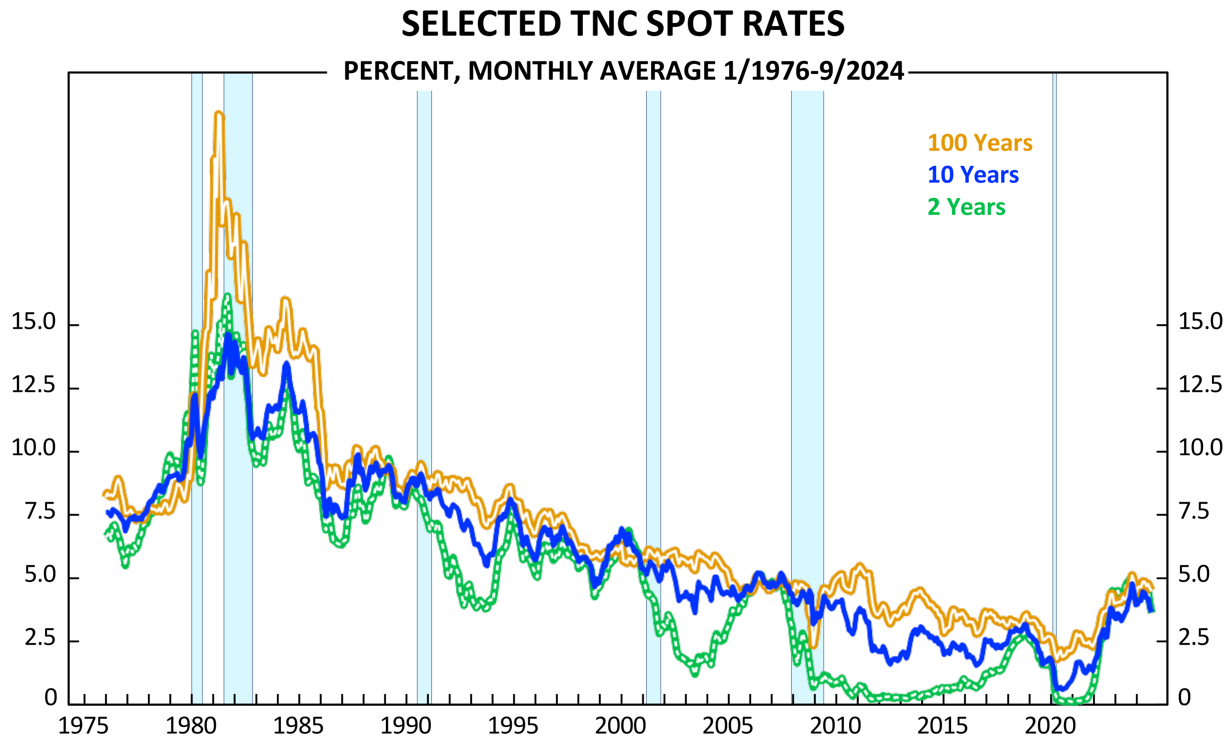

Analogous to the previous chart, this chart shows history for the TNC yield curve. Again the chart indicates that the spot rates bunch around recessions.

This chart indicates that the spot rate projected at 100 years maturity was very high around 1980. The period around 1980 saw very high inflation rates, leading some observers at the time to worry that inflation was out of control. The high 100-year spot rate shows that the market was very concerned and that high expectations for long-term interest rates were embedded in the market. Therefore, the XRM projection methodology picked up the inflation concerns at that time.

10-Year Spot Rate History

Figure 17.5

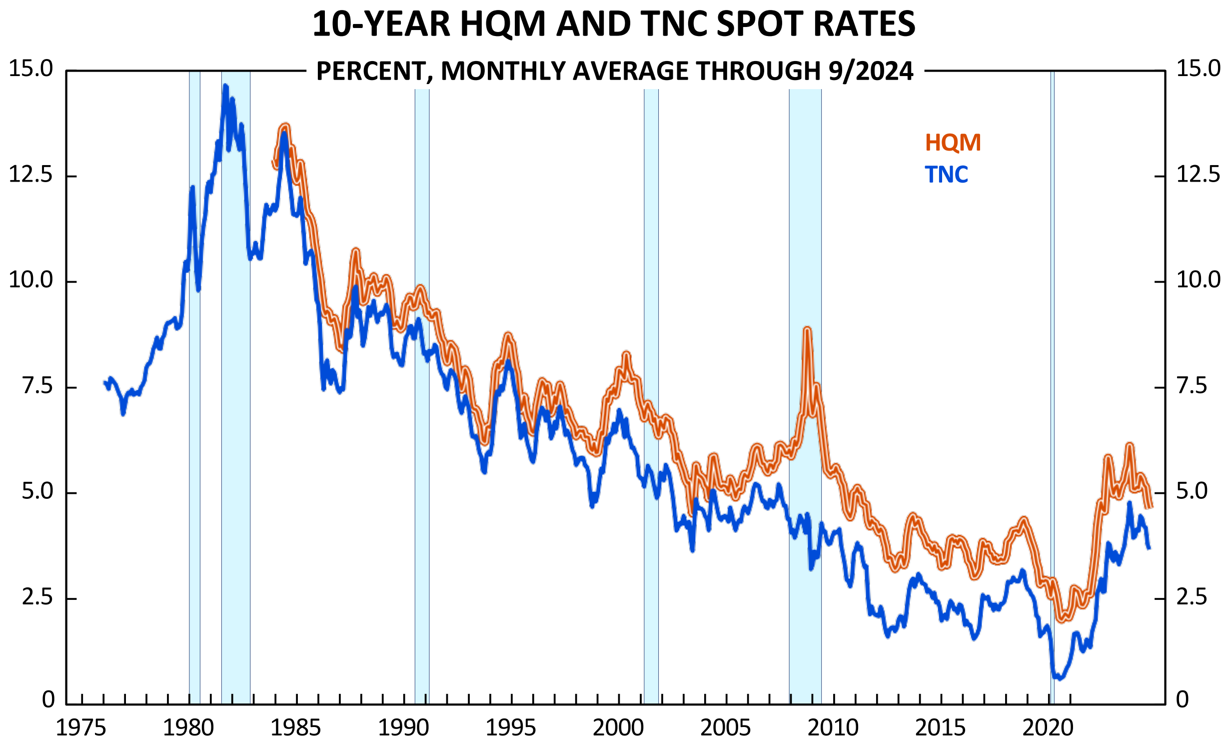

This chart compares the HQM and TNC spot rate at 10 years maturity over time.

The chart shows that the two rates generally tracked each other. However, there were times, such as in the 2008 recession, when the spread between the two changed temporarily.

18. Selected TRC Results

Analogous to Chapter 16, this chapter presents selected results for the TRC real yield curve based on TIPS. Initially, the results are for 08/30/2024 so they can be compared to Chapter 16. But another day 08/31/2020 is also presented, which is four years earlier and shows a very different market pattern for TIPS illustrating how different TIPS trading can be at different times.

The TRC Yield Curve on 08/30/2024

The TRC yield curve on 08/30/2024 was computed using 50 TIPS which include every Treasury real note or bond in the market on that day with at least two payments remaining and maturity greater than one half year. Convergence in the computation was achieved with 4 iterations.

For the maturity ranges, there are 5 bonds with maturities in the first maturity range, 7 bonds in the second maturity range, and 16, 7, and 15 bonds in the third, fourth, and fifth maturity ranges.

The five spline coefficients on the constrained B-splines for this TRC yield curve on this date expressed as percentages are 3.75, 0.74, 1.56, 2.02, and 2.29. The hump variable coefficient is -1.35, indicating the presence of a hump.

The TRC Yield Curve on 08/31/2020

The TRC yield curve on 08/31/2020 was chosen to be different from the yield curve on 08/30/2024 so as to illustrate potential differences in the TIPS market at different times.

Same as for 08/30/2024, the yield curve on 08/31/2020 was computed using 44 TIPS which include all the Treasury real notes and bonds in the market on that day with at least two payments remaining and maturity greater than one half year. Convergence in the computation was achieved with 5 iterations.

For the maturity ranges, there are 3 bonds with maturities in the first maturity range, 5 bonds in the second maturity range, and 14, 11, and 11 bonds in the third, fourth, and fifth maturity ranges.

However, the five spline coefficients for the 08/31/2020 TRC yield curve are different from the coefficients for the 08/30/2024 yield curve. For the former, the coefficients expressed as percentages are -1.25, -1.66, -1.41, -0.31, and 0.29. The first four are negative, and charts below will show that the resulting yields are largely negative. The hump variable coefficient is -2.47, indicating the presence of hump effects.

Discount Functions

Figure 18.1

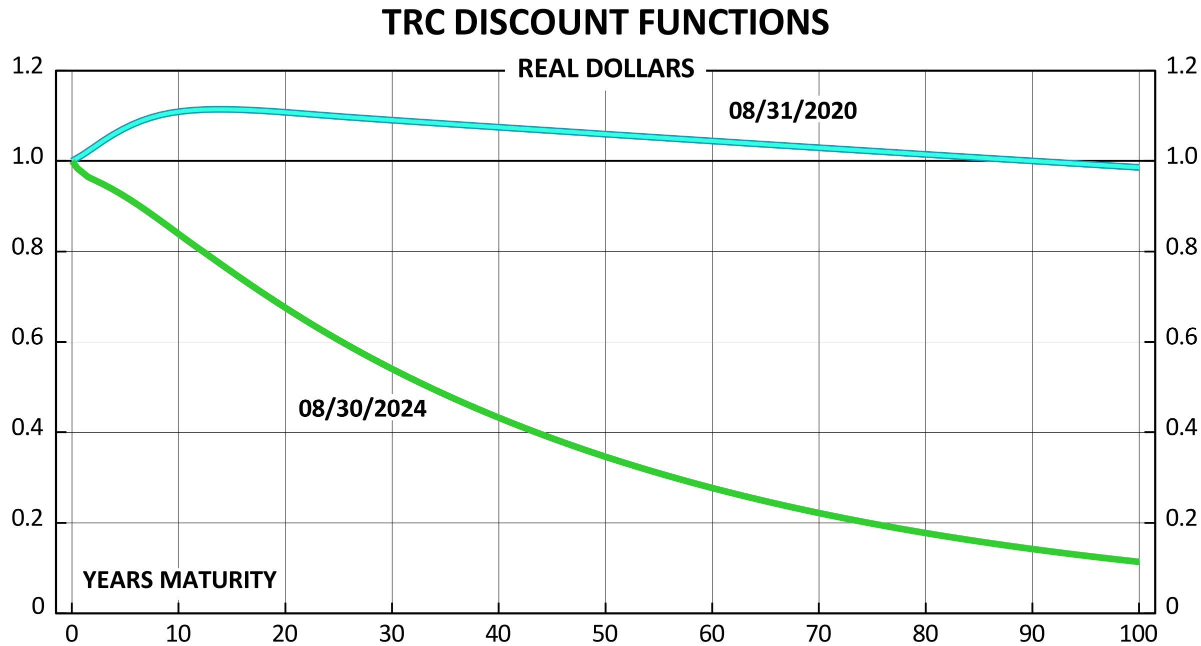

This chart shows the discount functions for the two TRC TIPS yield curves. Both functions start at unity and are positive throughout.

Because the spline coefficients are positive, the 08/30/2024 discount function declines throughout its range. It declines exponentially in the projection range and has an asymptote at zero.

In contrast, the negative spline coefficients for the 08/31/2020 yield curve cause the discount function to be above unity way into the projection range. So the market is willing to pay more than $1 in real terms to get $1 in the future. As shown in Chapter 5, this implies that the discount spot rate is negative. And, in addition, this discount function rises at the beginning, which implies a negative forward rate.

Forward Rates

Figure 18.2

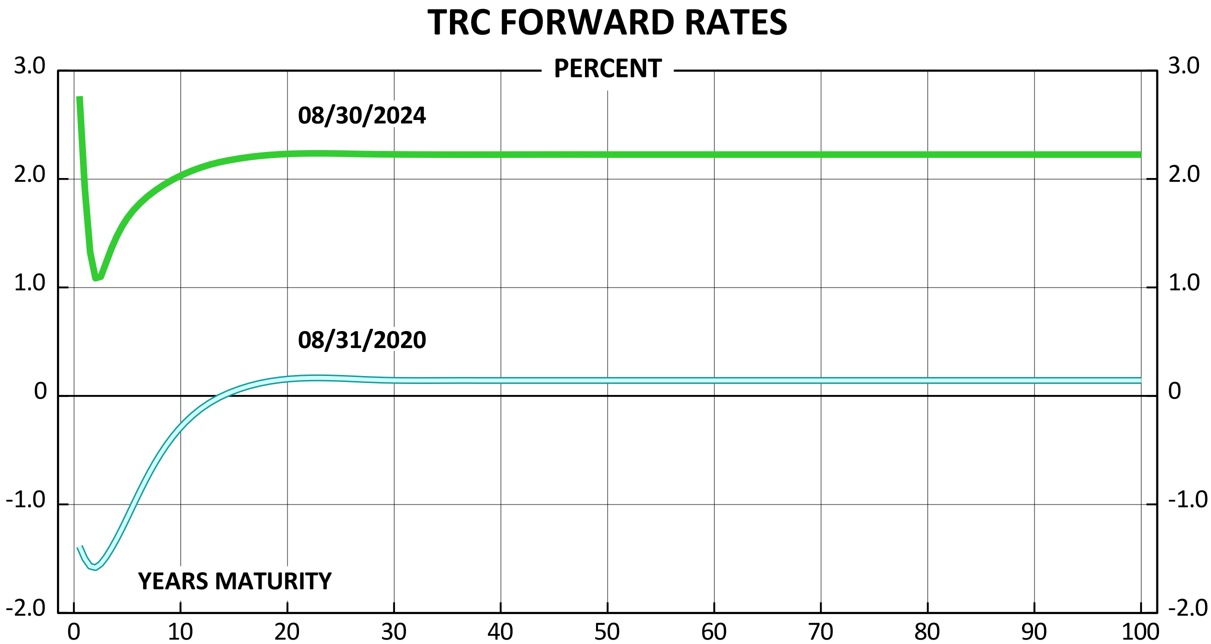

This chart shows the forward rates from the discount functions in the previous chart.

As expected, the forward rate for the 08/30/2024 yield curve is positive throughout because the spline coefficients are positive.

In contrast, the 08/31/2020 forward rate is negative until around 15 years maturity corresponding to the period when the discount rate is rising. Nevertheless, even when the discount rate remains above unity, it still starts to decline at around 15 years maturity, which causes the forward rate to become positive. Note that when the forward rate initially turns positive, the discount spot rate is still negative.

The long-term forward rate for the 08/30/2024 yield curve is 2.23 percent and the long-term forward rate for the 08/31/2020 yield curve is 0.14 percent. Note for the latter that the fact that the forward rate is positive at higher maturities results in a positive long-term forward rate despite negative forward rates at earlier maturities.

TRC Par Yield Curve for 08/30/2024

Figure 18.3

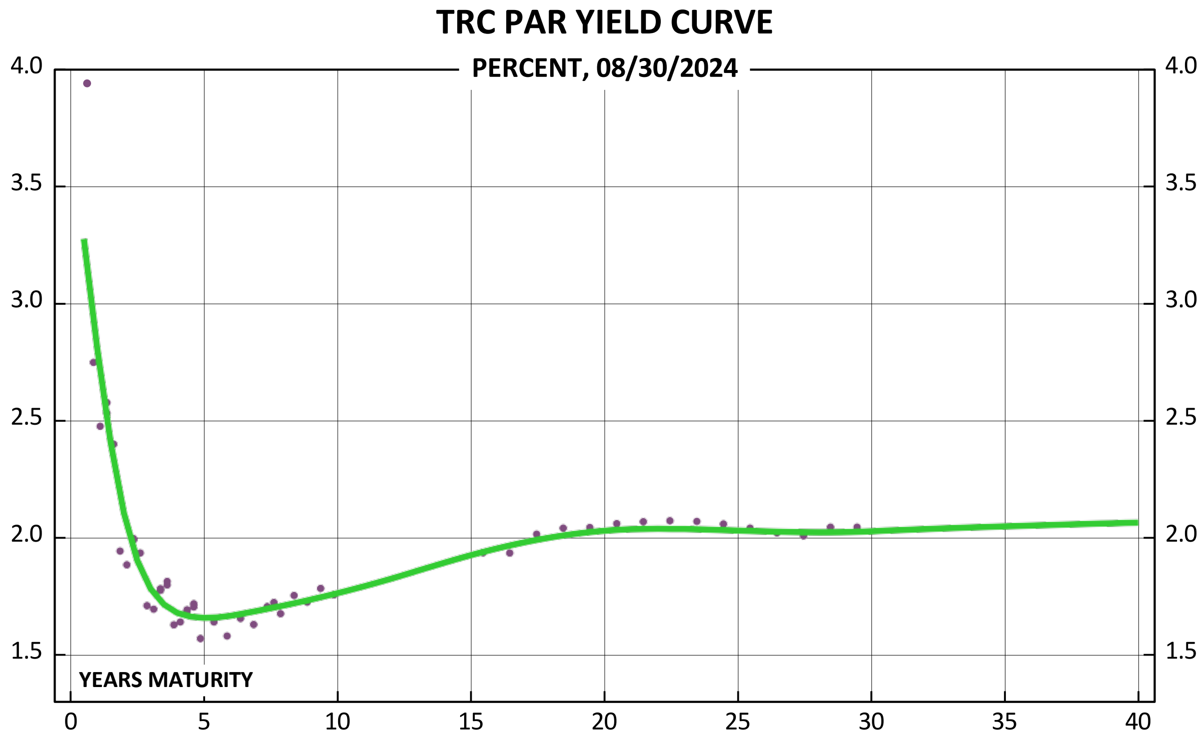

This chart shows the TRC par yield curve and associated scatter diagram for 08/30/2024.

As expected, because the spline coefficients are positive, the par yield curve is positive too. The yield curve has a hump as indicated by the hump coefficient. Similar to the TNC yield curves, this par yield curve tracks the pattern of bond yields in the market.

TRC Par Yield Curve for 08/31/2020

Figure 18.4

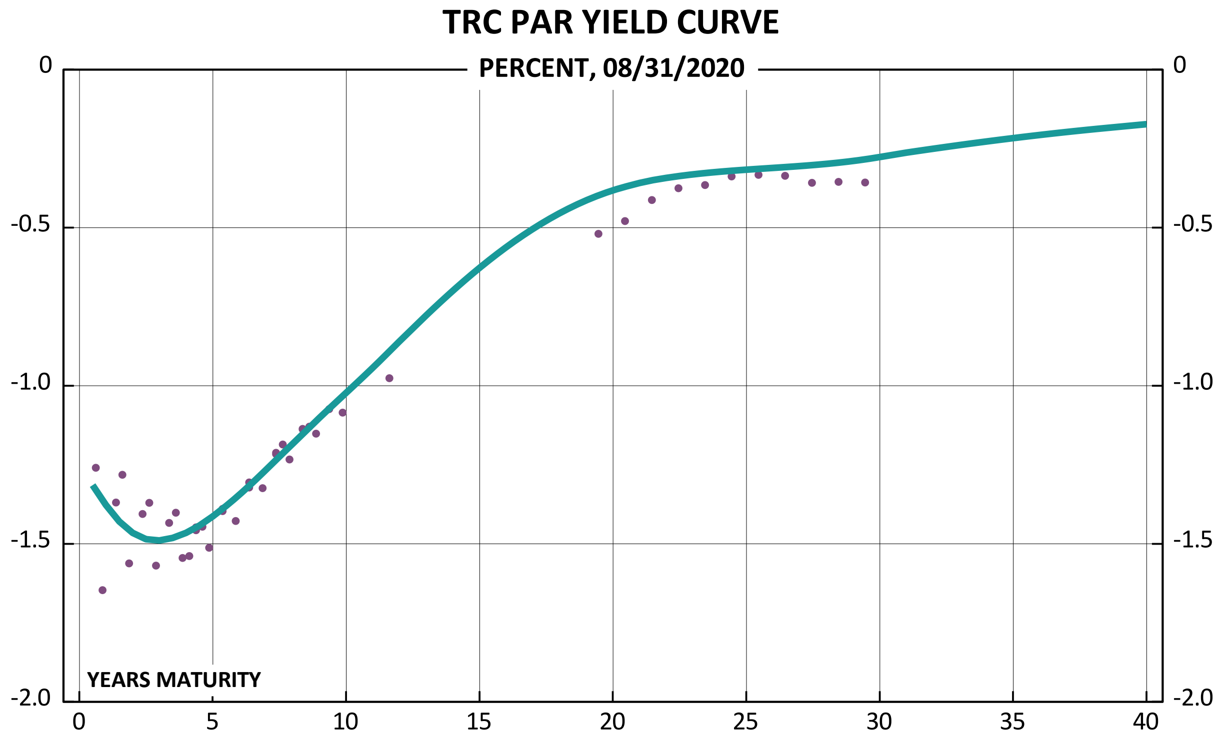

This chart shows the TRC par yield curve for 08/31/2020 and the associate scatter diagram of yields.

This par yield curve is negative throughout. However, such a par yield curve cannot actually exist in the market and there cannot actually be any bond trading at par because the negativity implies that the coupon rate on such a bond would have to be negative, and all coupon rates on TIPS are set to be positive. So this par yield curve should be used as an indicator of market conditions rather than an opportunity for trading. And, similar to the previous chart, even though the par yield curve is negative, it still captures the shape of the individual yields in the market and so it proves to be a useful indicator.

To note, the scatter diagram yields are computed using the Treasury convention yield formula applied to the actual bonds in the market. Yields on the scatter diagram are negative because these bonds are trading away from par.

The negative hump variable coefficient of -2.47 would suggest that there might be a hump. However, there is no hump. So in this chart, similar to the TNC yield curve for 08/30/2024, the hump variable flattens out the par yield curve after 20 years maturity and enables a more precise fit to the market without actually producing a hump.

TBI Curve

Figure 18.5

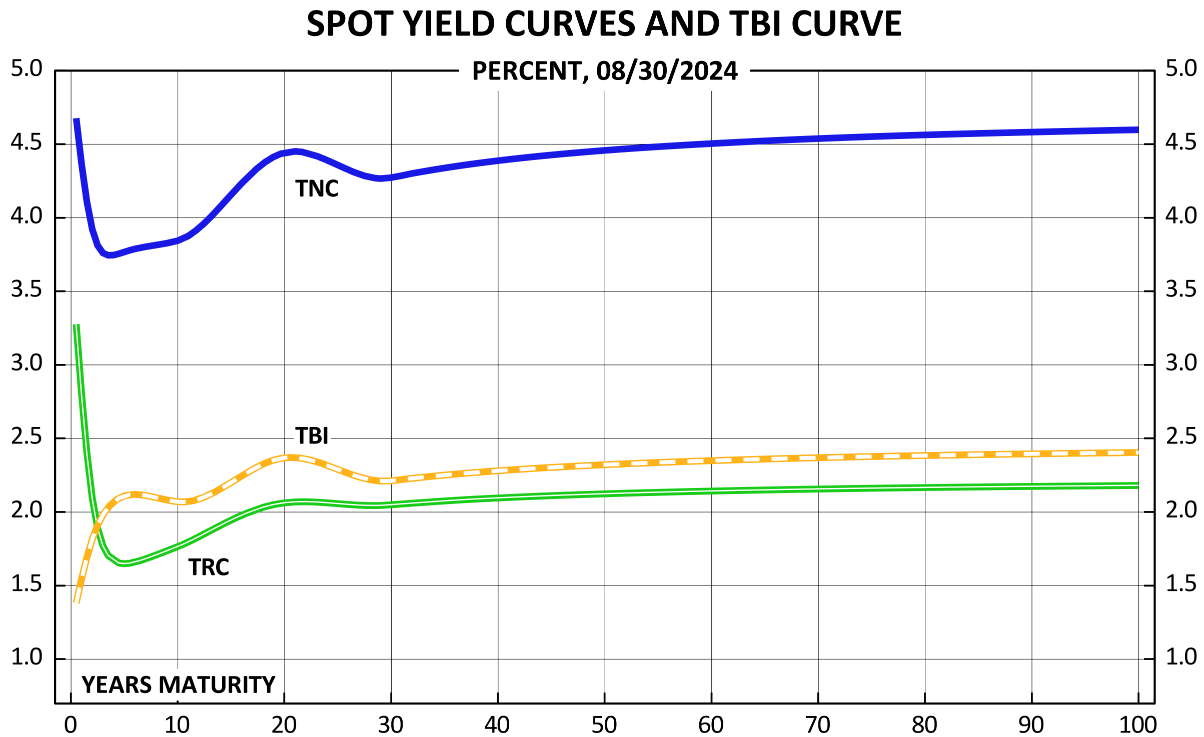

This chart shows the TBI curve for 08/30/2024 along with the TNC and TRC spot yield curves from which it is derived.

The TBI rate is around 2.0 percent through about 15 years maturity. After that the rate rises a bit but is still below 2.5 percent

19. Conclusion

This monograph has described the XRM methodology for constructing yield curves using extended regressions on maturity ranges, including techniques for representing the market behavior of forward rates in the maturity ranges by a cubic spline. The spline is constructed so that it calculates a fixed long-term forward rate that can generate yield curve projections at long maturities.

The XRM also has regression variables, and the use of regression variables was also explained. The hump variable that is used to account for yields around 20 years maturity was described. And credit variables were created for the HQM yield curve to combine bonds in the three top qualities AAA, AA, and A into a single yield curve.

The XRM methodology was applied to create three yield curves: the HQM yield curve for high quality corporate bonds as mandated by the Pension Protection Act of 2006, and the TNC and TRC yield curves for nominal Treasury coupon issues and TIPS.

The description of the XRM methodology and the results for the three yield curves show that XRM can be effectively applied to high quality corporate and Treasury fixed income securities to derive accurate estimates of yield curves and related information. Therefore, the XRM methodology is a general approach for fixed income markets.

References

de Boor, Carl, 1978, A Practical Guide to Splines, Springer-Verlag New York Inc.

Dershowitz, Nachum, and Edward M. Reingold, 1997, Calendrical Calculations, Cambridge University Press, Cambridge.

Fabozzi, Frank J., ed., 2005, The Handbook of Fixed Income Securities Seventh Edition, McGraw-Hill, New York.

Fisher, Mark, Douglas Nychka, and David Zervos, 1995, "Fitting the term structure of interest rates with smoothing splines," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, 95-1, January.

Fisher, Mark, and David Zervos, 1996, "YieldCurve," in Hal R. Varian, ed., Computational Economics and Finance Modeling and Analysis with *Mathematica*, Springer-Verlag New York, Inc.

Garbade, Kenneth D., 1996, Fixed Income Analytics, The MIT Press, Cambridge.

Gill, Philip E., Walter Murray, and Margaret H. Wright, 1981, Practical Optimization, Academic Press, San Diego.

Girola, James A., 2005, "The Long-Term Real Interest Rate for Social Security," U.S. Department of the Treasury Economic Policy Research Paper 0502, March 30, https://home.treasury.gov/system/files/226/long-term-rates-socialsecurity.pdf.

Girola, James A., 2006, "Implications of Returns on Treasury Inflation-Indexed Securities for Projections of the Long-Term Real Interest Rate," U.S. Department of the Treasury Economic Policy Research Paper 0601, March, https://home.treasury.gov/system/files/226/silr-0306-330x.pdf.

Girola, James A., 2007, The Corporate Bond Yield Curve for the Pension Protection Act, U.S. Department of the Treasury, October 11, https://home.treasury.gov/system/files/226/hqm_pres.pdf.

Girola, James A., 2010, Introduction to the HQM Yield Curve, U.S. Department of the Treasury, April 12, https://home.treasury.gov/system/files/226/ycp_apr2010.pdf.

Girola, James A., 2011, The HQM Yield Curve: Basic Concepts, U.S. Department of the Treasury, October 17, https://home.treasury.gov/system/files/226/ycp_oct2011.pdf.

Girola, James A., 2014, The Yield Curve for Treasury Nominal Coupon Issues, U.S. Department of the Treasury, May 16, https://home.treasury.gov/system/files/226/trp_apr2014.pdf.

Girola, James A., 2015, The Treasury Real Yield Curve and Breakeven Inflation, U.S. Department of the Treasury, July 21, https://home.treasury.gov/system/files/226/tii_may2015.pdf.

Girola, James A., 2016, Treasury Yield Curves and Discount Rates, U.S. Department of the Treasury, February 27, https://home.treasury.gov/system/files/226/ycds_feb2016.pdf.

Girola, James A., 2019, The Treasury Breakeven Inflation Curve, U.S. Department of the Treasury, March 26, https://home.treasury.gov/system/files/226/bri_mar2019_acc.pdf.

Girola, James A., 2022, Origins of the Yield Curves, U.S. Department of the Treasury, February 24, https://home.treasury.gov/system/files/226/ycrems.pdf.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007a, "The U.S. Treasury yield curve: 1961 to the present," Journal of Monetary Economics, 54, 2291-2304.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007b, "The TIPS Yield Curve and Inflation Compensation," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, 2008-05, November.

Judge, George G., W.E. Griffiths, R. Carter Hill, Helmut Luetkepohl, and Tsoung-Chao Lee, 1985, The Theory and Practice of Econometrics Second Edition, John Wiley and Sons, New York.

Krgin, Dragomir, 2002, Handbook of Global Fixed Income Calculations, John Wiley & Sons, Inc., New York.

Nelson, Charles R., and Andrew F. Siegel, 1987, "Parsimonious Modeling of Yield Curves," The Journal of Business, 60, October, 473-489.

Piegl, Les, and Wayne Tiller, 1997, The *NURBS* Book, Second Edition, Springer-Verlag, New York.

Stigum, Marcia, and Franklin L. Robinson, 1996, Money Market and Bond Calculations, Richard D. Irwin, Chicago.

Svensson, Lars E. O., 1994, "Estimating and Interpreting Forward Interest Rates: Sweden 1992-1994," National Bureau of Economic Research Working Paper No. 4871, September.

U.S. Department of the Treasury, 2005a, Creating a Corporate Bond Spot Yield Curve for Pension Discounting, February 7, Washington, D.C.

U.S. Department of the Treasury, 2005b, *Supplement to* Creating a Corporate Bond Spot Yield Curve for Pension Discounting, February 7, Washington, D.C.

U.S. Department of the Treasury, 2006, The Corporate Bond High Quality Centered Yield Curve for Pension Discounting, January 24, Washington, D.C.

U.S. Department of the Treasury, 2007, The Treasury High Quality Market Corporate Bond Yield Curve for the Pension Protection Act of 2006, October 9, Washington, D.C.

U.S. Department of the Treasury, 2024a, Corporate Bond Yield Curve Papers and Data, August 28, Washington, D.C., https://home.treasury.gov/data/treasury-coupon-issues-and-corporate-bond-yield-curve/corporate-bond-yield-curve.

U.S. Department of the Treasury, 2024b, Treasury Coupon Issues, August 28, Washington, D.C., https://home.treasury.gov/data/treasury-coupon-issues-and-corporate-bond-yield-curves/treasury-coupon-issues.