# **Salomon Brothers**

Antti Ilimanen (212) 783-5833

# A Framework for Analyzing Yield Curve Trades

**Understanding the Yield Curve: Part 6** 

Reprinted with permission from *Understanding the Yield Curve*, United States Fixed-Income Research Portfolio Strategies, November 1995. Copyright 1995, Salomon Brothers, New York, NY. All rights reserved.

Introduction	1
Forward Rates and Their Determinants	1
<ul> <li>How Do the Main Determinants Influence the Yield Curve Shape?</li> </ul>	1
<ul> <li>Decomposing Forward Rates Into Their Main Determinants</li> </ul>	3
Decomposing Expected Returns of Bond Positions	8
Five Alternative Expected Return Measures	8
Link to Scenario Analysis	14
Appendix A Decomposing the Forward Rate Structure Into Its Main Determinants	18
Appendix B. Relating Various Statements About Forward Rates to Each Other	21

## FIGURES

1.	Splitting a Zero-Coupon Bond's One-Period Yield Change Into Two Parts	4
2.	Decomposing Forward Rates Into Their Components, Using Historical Average Risk Premia and Volatilities	6
3.	Decomposing Forward Rates Into Their Components, Using Survey Rate Expectations and Implied Volatilities	6
4.	Forward-Implied Yield Changes versus Survey-Expected Yield Changes in the Treasury Bill Market 1993-95	8
5.	Three-Month Expected Return Measures and Their Components, as of 26 Sep 95	11
6.	Expected Returns of a Three-Month Bill a Three-Year Bond, and a Ten-Year Bond, as of 26 Sen 95	12
7.	Scenario Analysis and Expected Bond Returns	15
8.	Various Yield Curve Scenarios	10
9.	Decomposing the Total Expected Return Into Four Components	

#### INTRODUCTION

In Part 1 of this series, *Overview of Forward Rate Analysis*, we argued that the shape of the yield curve depends on three factors: the market's rate expectations; the required bond risk premia; and the convexity bias. After examining these determinants in detail in Parts 2-5, we now return to the "big picture" to show how we can decompose the forward rate curve into these three determinants. Even though we cannot directly observe these determinants, the decomposition can clarify our thinking about the yield curve.

Our analysis also produces direct applications — it provides a systematic framework for relative value analysis of noncallable government bonds. Analogous to the decomposition of forward rates, the total expected return of any government bond position can be viewed as the sum of a few simple building blocks: (1) the yield income; (2) the rolldown return; (3) the value of convexity; and (4) the duration impact of the rate view. A fifth term, the financing advantage, should be added for bonds that trade "special" in the repo market.

The following observations motivate this decomposition. A bond's near-term expected return is a sum of its horizon return given an unchanged yield curve and its expected return from expected changes in the yield curve. The first item, the horizon return, is also called the rolling yield because it is a sum of the bond's yield income and the rolldown return (the capital gain that the bond earns because its yield declines as its maturity shortens and it "rolls down" an upward-sloping yield curve). The second item, the expected return from expected changes in the yield curve, can be approximated by duration and convexity effects. The duration impact is zero if the yield curve is expected to remain unchanged, but it may be the main source of expected return if the rate predictions are based on a subjective market view or on a quantitative forecasting model. The value of convexity is always positive and depends on the bond's convexity and on the perceived level of yield volatility.

We argue that both prospective and historical relative value analysis should focus on near-term expected return differentials across bond positions instead of on yield spreads. The former measures take into account all sources of expected return. Moreover, they provide a consistent framework for evaluating all types of government bond positions. We also show, with practical examples, how various expected return measures are computed and how our framework for relative value analysis is related to the better-known scenario analysis.

## FORWARD RATES AND THEIR DETERMINANTS

How Do the Main Determinants Influence the Yield Curve Shape?

We first describe how the market's rate expectations, the required bond risk premia<sup>1</sup>, and the convexity bias influence the term structure of spot and forward rates. The market's expectations regarding the future interest rate behavior are probably the most important influences on today's term structure. **Expectations for parallel increases in yields tend to make today's term structure linearly upward sloping, and expectations for** 

<sup>&</sup>lt;sup>1</sup> The bond risk premium is defined as a bond's expected (near-term) holding-period return in excess of the riskless short rate. Historical experience suggests that long-term bonds command some risk premium because of their greater perceived riskiness. However, our term "bond risk premium" also covers required return differentials across bonds that are caused by other factors than risk, such as liquidity differences, supply effects or market sentiment.

falling yields tend to make today's term structure inverted. Expectations for future curve flattening induce today's spot and forward rate curves to be concave (functions of maturity), and expectations for future curve steepening induce today's spot and forward rate curves to be convex.<sup>2</sup> These are the facts, but what is the intuition behind these relationships?

The traditional intuition is based on the *pure expectations hypothesis*. In the absence of risk premia and convexity bias, a long rate is a weighted average of the expected short rates over the life of the long bond. If the short rates are expected to rise, the expected average future short rate (that is, the long rate) is higher than the current short rate, making today's term structure upward sloping. A similar logic explains why expectations of falling rates make today's term structure inverted. However, this logic gives few insights about the relation between the market's expectations regarding future curve reshaping and the curvature of today's term structure.

Another perspective to the pure expectations hypothesis may provide a better intuition. The absence of risk premia means that all bonds, independent of maturity, have the same near-term expected return. Recall that a bond's holding-period return equals the sum of the initial yield and the capital gains/losses that yield changes cause. Therefore, if all bonds are to have the same expected return, initial yield differentials across bonds must offset any expected capital gains/losses. Similarly, each bond portfolio with expected capital gains must have a yield disadvantage relative to the riskless asset. If investors expect the long bonds to gain value because of a decline in interest rates, they accept a lower initial yield for long bonds than for short bonds, making today's spot and forward rate curves inverted. Conversely, if investors expect the long bonds to lose value because of an increase in interest rates, they demand a higher initial yield for long bonds than for short bonds, making today's spot and forward rate curves upward sloping. Similarly, if investors expect the curve-flattening positions to earn capital gains because of future curve flattening, they accept a lower initial yield for these positions. In such a case, barbells would have lower yields than duration-matched bullets (to equate their near-term expected returns), making today's spot and forward rate curves concave. A converse logic links the market's curve-steepening expectations to convex spot and forward rate curves.

The above analysis presumes that all bond positions have the same near-term expected returns. In reality, investors require higher returns for holding long bonds than short bonds. Many models that acknowledge bond risk premia assume that they increase linearly with duration (or with return volatility) and that they are constant over time. Parts 3 and 4 of this series showed that empirical evidence contradicts both assumptions. Historical average returns increase substantially with duration at the front end of the curve but only marginally after the two-year duration. Thus, the bond risk premia make the term structure upward-sloping and concave, on average. Moreover, it is possible to forecast when the required bond risk premia are abnormally high or low. Thus, the time-variation in the bond risk premia can cause significant variation in the shape of the term structure.

 $<sup>^2</sup>$  A concave (but upward-sloping) curve has a steeper slope at short maturities than at long maturities; thus, a line connecting two points on the curve is always below the curve. A convex (but upward-sloping) curve has a steeper slope at long maturities than at short maturities; thus, a line connecting two points on the curve is always above the curve.

Convexity bias refers to the impact that the nonlinearity of a bond's price-yield curve has on the shape of the term structure. This impact is very small at the front end but can be quite significant at very long durations. A positively convex price-yield curve has the property that a given yield decline raises the bond price more than a yield increase of equal magnitude reduces it. All else equal, this property makes a high-convexity bond more valuable than a low-convexity bond, especially if the volatility is high. It follows that investors tend to accept a lower initial yield for a more convex bond because they have the prospect of enhancing their returns as a result of convexity. Because a long bond exhibits much greater convexity than a short bond, it can have a lower yield and yet, offer the same near-term expected return. Thus, in the absence of bond risk premia, the convexity bias would make the term structure inverted. In the presence of positive bond risk premia, the convexity bias tends to make the term structure humped — because the negative effect of convexity bias overtakes the positive effect of bond risk premia only at long durations. An increase in the interest rate volatility makes the bias stronger and, thus, tends to make the term structure more humped.

The three determinants influence the shape of the term structure simultaneously, making it difficult to distinguish their individual effects. One central theme in this series has been that the shape of the term structure does not only reflect the market's rate expectations. Forward rates are good measures of the market's rate expectations only if the bond risk premia and the convexity bias can be ignored. This is hardly the case, even though a large portion of the short-term variation in the shape of the curve probably reflects the market's changing expectations about the future level and shape of the curve. The steepness of the curve on a given day depends mainly on the market's view regarding the rate direction, but in the long run, the impact of positive and negative rate expectations largely washes out. Therefore, the average upward slope of the yield curve is mainly attributable to positive bond risk premia. The curvature of the term structure may reflect all three components. On a given day, the spot rate curve is especially concave (humped) if market participants have strong expectations of future curve flattening or of high future volatility. In the long run, the reshaping expectations should wash out, and the average concave shape of the term structure reflects the concavity of the risk premium curve and the convexity bias.

**Decomposing Forward Rates Into Their Main Determinants Conceptually, each one-period forward rate can be decomposed to three parts: the impact of rate expectations; the bond risk premium; and the convexity bias.** So far, this statement is just an assertion. In this subsection, we show intuitively why this relationship holds between the forward rates and their three determinants. We provide a more formal derivation in Appendix A (where we take into account the fact that the analysis is not instantaneous but yield changes occur over a discrete horizon, during which invested capital grows). In Appendix B, we tie some loose strings together by summarizing various statements about the forward rates and by clarifying the relations between these statements.

Figure 1 shows how the yield change of an n-year zero-coupon bond over one period (dashed arrow) can be split to the rolldown yield change and the one-period change in an n-1 year constant-maturity spot rate  $s_{n-1}$  ( $\Delta s_{n-1} = s^*_{n-1} - s_{n-1}$ ) (two solid arrows).<sup>3</sup> A zero-coupon bond's price

<sup>&</sup>lt;sup>3</sup> All rates and returns in this report are expressed in percentage terms (200 basis points =  $2^{i}(i)$ , whereas in the equations in Parts 1 and 2 of this series they were expressed in decimal terms (200 basis points = 0.02).

can be split in a similar way (see Appendix A). Thus, an n-year zero's holding-period return over the next period  $h_n$  is:

$h_{\rm p}$ = return if the curve is unchanged	+ [return from the curve changes]	(1)
= rolling yield	+ [percentage price change (at horizon)]	
$\approx$ (one-period) forward rate	+ [-duration*( $\Delta s_{n-1}$ ) + (0.5*convexity*( $\Delta s_{n-1}$ ) <sup>2</sup> ].	

Equation (1) is based on the following relations. First, a bond's one-period horizon return given an unchanged yield curve is called the rolling yield. A zero-coupon bond's rolling yield equals the one-period forward rate  $(f_{n-1,n})$ . For example, if the four-year (five-year) constant-maturity rate remains unchanged at 9.5% (10%) over the next year, a five-year zero bought today at 10% can be sold next year at 9.5%, as a four-year zero; then the bond's horizon return is  $1.10^5/1.095^4 - 1 = 0.1202 = 12.02\%$ , which is the one-year forward rate between four- and five-year maturities (see Equation (12) in Appendix B). The second source of a zero's holding-period return, the price change caused by the yield curve shift, is approximated very well by duration and convexity effects for all but extremely large yield curve shifts.



It is more interesting to relate the forward rates to expected returns and expected rate changes than to the realized ones. We take expectations of both sides of Equation (1), split the bond's expected holding-period return into the short rate and the bond risk premium, and recall that  $E(\Delta s_{n-1})^2 \approx (Vol(\Delta s_{n-1}))^2$ . Then we can rearrange the equation to express the one-period forward rate as a sum of the other terms:

## Forward rate $\approx$ short rate + duration\* $E(\Delta s_{n-1})$ + bond risk premium + convexity bias, (2)

where bond risk premium =  $E(h_n - s_1)$  and convexity bias  $\approx -0.5$ \*convexity\*( $Vol(\Delta s_{n-1})$ )<sup>2</sup>.

If we move the short rate to the left-hand side of the equation, we decompose the "forward-spot premium"  $(f_{n-1,n} - s_1)$  into a rate expectation term, a risk premium term and a convexity term (see Equation (11) in Appendix A). We interpret the expectations in Equation (2) as the market's rate and volatility expectations and as the expected risk premium that the market requires for holding long-term bonds. The market's expectations are weighted averages of individual market participants' expectations.

Some readers may wonder why our analysis deals with forward rates and not with the more familiar par and spot rates. The reason is the simplicity of the one-period forward rates. A one-period forward rate is the most basic unit in term structure analysis, the discount rate of one cash flow over one period. A spot rate is the average discount rate of one cash flow over many periods, whereas a par rate is the average discount rate of many cash flows — those of a par bond — over many periods. All the averaging makes the decomposition messier for the spot rates and the par rates than it is for the one-period forward rate in Equation (2). However, because the spot and the par rates are complex averages of the one-period forward rates, they too can be conceptually decomposed into the three main determinants.

Because the approximate decomposition in Equation (2) is derived mathematically without making specific economic assumptions, it is true in general. In reality, however, it is hard to make this decomposition because the components are not observable and because they vary over time. Further assumptions or proxies are needed for such a decomposition. In Figure 2, we use historical average returns to compute the bond risk premia and historical volatilities to compute the convexity bias — together with the observable market forward rates (as of September 26, 1995) — and back out the only unknown term in Equation (2): the expected spot rate change times duration. We also could divide this term by duration to infer the market's rate expectations. The rate expectations that we back out in Figure 2 suggest that the market expects small declines in short rates and small increases in long rates.

If bond risk premia vary over time, the use of historical average risk premia may be misleading. As an alternative, we can use survey data or rate predictions based on a quantitative forecasting model to proxy for the market's rate expectations. In Figure 3, we use the consensus interest rate forecasts from the *Blue Chip Financial Forecast* to proxy for the market's rate expectations. In addition, we use implied volatilities from option prices to compute the convexity bias. These components can be used together with the one-year forward rates to back out estimates of the unobservable bond risk premia.

Figure 2. Decomposing Forward Rates Into Their Components, Using Historical Average Risk Premia and Volatilities



Note: The one-year forward rates are based on the Salomon Brothers Treasury Model curve on September 26, 1995. The bond risk premia are based on the historical arithmetic average returns of various maturity-subsector portfolios between 1970-94, expressed in excess of the riskless one-year return. The convexity bias is based on the historical volatilities of various on-the-run bonds' weekly yield changes between 1990 and 1995. The rate expectation term for each duration is then backed out as the difference — one-year forward rate - one-year spot rate - bond risk premium - convexity bias.

Figure 3. Decomposing Forward Rates Into Their Components, Using Survey Rate Expectations and Implied Volatilities



Note: The one-year forward rates are based on the Salomon Brothers Treasury Model curve on September 26, 1995. The market's rate expectations are proxied by the consensus interest rate forecasts from the *Blue Chip Financial Forecast* (October 1, 1995 issue, a survey conducted among industry economists and analysts on September 26-27. 1995). The convexity bias is based on the implied volatilities from the prices of Salomon Brothers' OTC options for various on-the-run bonds on September 26. The bond risk premium for each duration is then backed out as the difference — one-year forward rate - one-year spot rate - expected rate change \* duration - convexity bias.

A comparison of Figures 2 and 3 shows that the two decompositions look similar up to the seven-year duration, but quite different beyond that point. The similarity of the convexity bias components in these two figures suggests that the use of historical or implied volatilities makes little difference, at least in this case. It is also clear that the Blue Chip survey predicted small declines in the short rates and small increases in the long rates, just as the inferred rate expectations in Figure 2. However, the predicted increases in long rates were smaller in this survey (where the largest increase was four basis points) than the inferred forecasts of Figure 2 (where the largest increase was eight basis points). Because the forward rate curve is the same in both figures, the smaller predicted rate increases lead to higher bond risk premia in Figure 3 than in Figure 2.<sup>4</sup>

Figures 2 and 3 are snapshots of the forward rates and their components on one date. A comparison of similar decompositions over time would provide insights into the relative variability of each component. In Figure 4, we try to illustrate the impact of changing rate expectations and risk premia on the steepness of the U.S. Treasury bill curve. The figure shows that on four recent dates the forwards always implied larger increases in the three-month rate than the market expected, based on surveys of bond market analysts. The difference is proportional to the required bond risk premium of longer bills over shorter bills (because bills exhibit negligible convexity, its impact can be ignored). Not surprisingly, this difference is always positive; moreover, it varies over time.

The time-variation in the bond risk premium in Figure 4 appears economically reasonable. In December 1993, after a long bull market, market participants were complacent, neither expecting much higher rates nor demanding much compensation for duration extension. After the Fed began to tighten the monetary policy, the market expected further rate increases. In addition, the increase in volatility and in risk aversion levels (perhaps caused by own losses as well as the well-publicized losses of other investors) increased the required bond risk premia. By the end of 1994, the market was extremely bearish, expecting almost 100 basis points higher three-month rates in six months. However, the forwards were implying almost 200 basis points higher rates — the difference reflects an abnormally large risk premium. In 1995, the bond market rallied strongly as the market's expectations for further Fed tightening receded and turned into expectations of easing policy. However, a large part of this rally was caused by the collapse in the required bond risk premium, perhaps reflecting lower inflation uncertainty and higher wealth that reduced the market's risk perceptions and risk aversion. In general, the time-variation in required returns appears to have contributed as much to the changing slope of the yield curve as has the time-variation in the

<sup>&</sup>lt;sup>4</sup> We hasten to point out that these calculations are quite imprecise, especially at long durations. Even an error of a couple of basis points in our proxy for the market's rate expectation will have a large impact on any long bond's expected return (and thus on the estimated bond risk premium), because the expected yield change is scaled up by duration. Such sensitivity reduces the usefulness of this decomposition at long durations.

**market's interest rate expectations**.<sup>5</sup> Finally, we note that the time-variation in the estimated bond risk premia has been very market-directional over the past two years; this may not always be the case.





Note: Forward-implied yield change is the difference between the implied three-month rate six months forward  $(f_{0.5, 0.75})$  and the current three-month rate  $(s_{0.25})$  (see the data below). Survey-expected yield change is the difference between the expected three-month rate six months ahead (E(s0.25)) and the current three-month rate, where the market's rate expectation is proxied by the mean in the *Wall Street Journal*'s semiannual survey. Part 2 of this series discusses the survey data and their shortcomings. It also shows that the difference between the forward-implied yield change is proportional to the bond risk premium.

	<sup>\$</sup> 0.25	<sup>f</sup> 0.5, 0.75	E(s <sub>0.25</sub> )
Dec 93	3.09%	3.80%	3.40%
Jun 94	4.22	5.49	4.62
Dec 94	5.70	7.51	6.60
Jun 95	5.63	5.58	5.40

## DECOMPOSING EXPECTED RETURNS OF BOND POSITIONS

## Five Alternative Expected Return Measures

Our framework for decomposing the yield curve also provides a framework for systematic relative value analysis of government bonds with known cash flows. We can evaluate all bond positions' expected returns comprehensively, yet with simple and intuitive building blocks. We emphasize that relative value analysis should be based on near-term expected return differentials, not on yield spreads, which are only one part

<sup>5</sup> In the earlier parts of this series, we provide further evidence of the importance of time-varying risk premia. Why do so many market participants and analysts think that the rate expectations are much more important determinants of the yield curve steepness than are the bond risk premia, in spite of the contradicting empirical evidence? We offer one possible explanation: Individual market participants have their individual rate views and individual required risk premia. (Few investors think explicitly in terms of such premia, but they will extend duration only if they expect longer bonds to outperform shorter bonds by a margin sufficient to offset their greater risks.) However, what matters for the yield curve is the market's aggregate rate view and aggregate required risk premia. Both vary over time as individual rate views and risk perceptions and preferences change (and as the composition of market participants changes). We stress that the changing individual rate views may have smaller aggregate effects than the changing risk perceptions and preferences even if individual rate views are more volatile than individual risk perceptions and preferences. This effect can occur if the changes in risk perceptions and preferences are much more highly correlated across individuals than are changes in rate views. For example, when volatility is high, most market participants are likely to demand abnormally high bond risk premia even if they have widely different views about the future rate direction. Perhaps most market participants focus on the fact that (their) individual rate views vary more over time than do their risk perceptions and preferences, ignoring the fact that market rates are driven by the aggregate effects, for which correlations across individuals matter a lot

of them. That is, **total return investors should care more about expected returns than about yields.** Thus, our approach brings fixed-income investors closer to mean-variance analysis in which various positions are evaluated based on the trade-off between their expected return and return volatility.

Equation (1) shows that a zero's holding-period return is a sum of its return given an unchanged yield curve and its return caused by the changes in the yield curve. The return given an unchanged yield curve is called the rolling yield because it is a sum of the zero's yield and the rolldown return. The return caused by changes in the yield curve can be approximated well by duration and convexity effects. Taking expectations of Equation (1) and splitting the rolling yield into yield income and rolldown return, the near-term expected return of a zero is:

#### Expected return ~ Yield income (3) + Rolldown return + Value of convexity + Expected capital gain/loss from the rate "view"

For details, see Equation (9) in Appendix A or the notes below Figure 5. A similar relation holds approximately for coupon bonds, and we will describe the three-month expected return of some on-the-run Treasury bonds as the sum of the four components in the right-hand side of Equation  $(3).^{6}$ 

This framework is especially useful when evaluating positions of two or more government bonds, such as duration-neutral barbells versus bullets. We first compute expected return separately for each component and then compute the portfolio's expected return by taking a market-value weighted average of all the components' expected returns.

## It may be helpful to show step by step how the expected return measures are improved, starting from simple yields and moving toward more comprehensive measures:

• A bond's **yield income** includes coupon income, accrued interest and the accretion/amortization of price toward par value. Yield to maturity is the correct return measure if all interim cash flows can be reinvested at the yield and the bond can be sold at its purchasing yield.<sup>7</sup> Yield ignores the rolldown return the bond earns if the yield curve stays unchanged.

 $<sup>^{6}</sup>$  However, certain modifications are needed when Equation (3) is used to describe the expected returns of coupon bonds rather than those of zeros — and the approximation will be somewhat worse. We use each bond's rolling yield to measure the horizon return given an unchanged yield curve: this measure no longer equals the one-period forward rate. We also use the end-of-horizon duration and convexity as well as the change in the constant-maturity rate of a constant-coupon curve at horizon, and we adjust the duration and convexity effects for the fact that the bond's value increases to (1+ rolling yield/100) by the end of the horizon. Besides the approximation error of ignoring higher-order terms than duration and convexity effects, another source of error exists for coupon bonds: The reinvestment rate assumptions vary across bonds. Recall that the calculation of the yield to maturity implicitly assumes that all cash flows are reinvested at the bond's yield to maturity. This fact may lead to exaggerated estimates of yield income for long-term bonds if the yield curve is upward sloping, a problem common to all expected return measures that use the consistency of the reinvestment rate assumptions across bonds, any inconsistencies should have a relatively small impact on the overall level of bonds' expected returns.

 $<sup>^{7}</sup>$  The yield to maturity of a single cash flow is unambiguous, whereas the yield of a portfolio of multiple cash flows is a more controversial measure. The duration-times-market-value weighted yield is a good proxy for a portfolio's true yield to maturity (internal rate of return). However, a portfolio's market-value weighted yield may be a better estimate of the portfolio's *likely yield income over a short horizon* (its near-term expected return) than is its yield to maturity. The yield to maturity weighs longer cash flows more heavily and is more influenced by the built-in reinvestment rate assumptions. We will return to this topic in a future report.

• **Rolling yield** is a better expected return proxy if an unchanged curve is a reasonable base case. Yet, it ignores the value of convexity and, thus, implicitly assumes no rate uncertainty. Thus, the rolling yield measures expected return if no curve change and no volatility is expected.

• Combining the rolling yield with the value of convexity improves the expected return measure further. In Part 5 of this series, we showed that a bond's **convexity-adjusted expected return** equals the sum of the rolling yield and the value of convexity. This measure recognizes the impact of rate uncertainty but implies that no change is expected in the yield curve. Empirical evidence described in Part 2 suggests that an unchanged yield curve is often a reasonable base "view."

• If investors want, they can replace the prediction of an unchanged curve with some other rate (or spread) "view." One possibility is to use survey-based information of the market's current rate forecasts; such an approach may be useful for backing out the market's required return for each bond. Alternatively, investors may ignore the market view and input either their own rate views or an economist's subjective rate forecasts or rate predictions from some quantitative model. For example, the predictors identified in Part 4 of this series can be used to forecast the changes in the long rates. The impact of any rate view is approximated by the expected yield change scaled by duration (see Equation (10) in Appendix A), which may be added to the convexity-adjusted expected return. The sum gives us the "expected return with a view" — the four-term expected return measure in Equation (3). However, this equation is a perfect description of expected returns only for bonds that lie on the fitted curve. Thus, the relative value measures above ignore "local" or bond-specific richness or cheapness relative to the curve.

· Many technical factors can make a specific bond "locally" rich or cheap (relative to adjacent-maturity bonds), or they can make a whole maturity sector rich or cheap relative to the fitted curve. Such factors include supply effects (temporary price pressure on a sector caused by new issuance), demand effects (maturity limitations or preferences of important market participants --- for example, the richness of quarter-end bills), liquidity effects (lower transaction costs for on-the-runs versus off-the-runs, for 30-year bonds versus 25-year bonds, for Treasury bills versus duration-matched coupon bonds, etc.), coupon effects (motivated by tax benefits, accounting rules, etc.), and above all, the financing effects (the "special" repo income that is common for on-the-runs).8 Fortunately, it is easy to add to the four-term expected return measures the financing advantage and two local cheapness measures — the spread off the fitted curve and the expected cheapening toward the fitted curve. The five-term expected return measures are comprehensive measures of total expected returns — ignoring small approximation errors, they incorporate all sources of expected return for noncallable government bonds.9

<sup>8</sup> Whether such local cheapness effects appear as deviations from a fitted yield curve or as "wiggles" or "kinks" in the fitted curve depends on the curve-estimation technique. Recall that all curve-estimation techniques try to fit bond prices well while keeping the curve reasonably shaped. If the goodness of fit is heavily weighted, all bonds have small or no deviations from the fitted curve. However, a close fit may lead to "unreasonably" jagged forward rate curves. Based on Equation (2), the forward rate curve should be smooth, rather than jagged, because maturity-specific expectations of rate or volatility behavior are hard to justify and because arbitrageurs presumably are quick to exploit any abnormally large expected return differentials between adjacent-maturity bonds.

<sup>&</sup>lt;sup>9</sup> In our analysis, we include the local effects into the expected bond returns separately as a fifth term. As an alternative, we could include the financing advantage (repo income) and the spread off the curve in the yield income, and we could include the expected cheapening in the rolldown return. "Rich" bonds, such as the on-the-runs, are unlikely to roll down the fitted curve if the overall curve remains unchanged. More likely, they will lose their relative richness eventually. It may be reasonable to assume that an on-the-run bond's yield advantage and *expected* cheapening exactly offset its *expected* financing advantage.

As a numerical illustration, Figure 5 shows the various expected return measures for three bonds (the three-month Treasury bill and the three-year and ten-year on-the-run Treasury notes) and for the barbell combination of the three-month bill and the ten-year bond. In this example, we use as much market-based data as possible: for example, implied volatilities, not historical, to estimate the value of convexity, and the "view" (rate predictions) based on survey evidence of the market's rate expectations, not on a quantitative forecasting model. All the numbers are based on the market prices as of September 26, 1995.

Maturity	0.25	3	10	Barbell
Yield Income	1.349%	1.474%	1 568%	1 425%
+ Rolldown Return	0.000	0.065	0.108	0.038
= Rolling Yield	1.349	1.537	1.676	1.463
+ Value of Convexity	0.000	0.014	0.082	0.029
= Convexity-Adj. Expected Return	1.349	1.551	1.758	1.492
+ Duration Impact of the "View"	0.000	-0.056	-0.284	-0.099
= Expected Return with a View	1.349	1.495	1.474	1.393
+ Total Local Rich/Cheap Effect	-0.015	-0.039	0.025	-0.002
= Total Expected Return	1.334	1.456	1.499	1.391
Background Information				
Par Yield	5.507	6.011	6.408	NA
Rolldown Yield Change	NA	-0.026	-0.015	NA
Duration now	0.245	2.708	7.300	2,706
Duration at Horizon	0.000	2.500	7.168	2.500
Convexity now	0.002	0.090	0.669	0.235
Convexity at Horizon	0.000	0.077	0.643	0.224
Yield Volatility	NA	0.598	0.502	NA
Yield Change "View"	-0.046	0.022	0.039	NA
On-the-Run Yield	5.446	5.978	6.261	NA
Financing Advantage	0.000	0.038	0.463	NA
Spread to the Par Curve	-0.015	-0.008	-0.037	NA
Expected Cheapening Return	0.000	-0.068	-0.401	NA

## Figure 5. Three-Month Expected Return Measures and Their Components, as of 26 Sep 95

NA Not available.

Note: Barbell is a combination of 0.53 units of the ten-year par bond and 0.47 units of the three-month bill; these weights duration-match the barbell with the three-year par bond bullet. Yield income is the return that a par bond earns over three months if it can be sold at its yield and if any cash flows are reinvested at the yield. The yields are semiannually compounded and based on the Salomon Brothers Treasury Model's par yield curve. Rolldown return is the capital gain that a bond earns from the rolldown yield change. Rolling yield is a bond's horizon return given an unchanged yield curve. Value of convexity is approximated by  $0.5^{\circ}$  convexity at horizon '(yield volatility)<sup>2</sup> (1 + rolling yield/100), where yield volatility is the basis-point yield volatility over a three-month horizon. The latter is computed by multiplying the on-the-run bond's relative yield volatility — implied volatility based on the price of a three-month OTC option written on this bond - by its yield level and dividing by two (for deannualization). For the three-year bond, we interpolate between the implied volatilities of on-the-run twos and fives. Duration impact of the "view" is (-duration at horizon)\*(expected change in a constant-maturity rate over the next three months)\*(1 + rolling yield/100). In this example, the "view" reflects the market's yield curve expectations, as measured by the Blue Chip consensus forecasts (that are based on a survey of professional economists and market analysts conducted on September 26-27, 1995). The "expected return with a view" measures the expected return for a hypothetical par bond that lies exactly on the Model curve, ignoring any local cheapness or financing advantage of actual bonds. We can add to this four-term measure a fifth component called the total local rich/cheap effect. It is the sum of three additional sources of return for specific bonds: (1) the financing advantage (the difference between the three-month term repo rate for general collateral and the three-month special term repo rate for the on-the-run bond, divided by four for deannualization); (2) the spread between the on-the-run bond yield and the Model par yield, divided by four for deannualization; and (3) the bond's expected cheapening as it loses the richness associated with the on-the-run status (estimated by the Salomon Brothers Government Bond Strategy Group).

The top panel of Figure 5 shows how nicely the different components of expected returns can be added to each other. Moreover, the barbell's expected return measures are simply the market-value weighted averages of its components' expected returns. In this case, the yield income, the rolldown return and the value of convexity are all higher for the longer bonds. In contrast, the duration impact of the market's rate view is negative, because the Blue Chip survey suggested that the market expected small increases in the three-year and the ten-year rates over the next quarter. The local rich/cheap effect is negative for the shorter instruments but positive for the ten-year note; the reason is that the negative yield spread and the expected cheapening of the ten-year note are

not sufficient to offset the ten-year note's high repo market advantage. According to all five expected return measures, the barbell has a lower expected return than the duration-matched bullet.

Figure 6 shows the five different expected return curves plotted on the three bonds' durations. In this case, the simplest expected return measure (yield income) and the most comprehensive measure (total expected return) happen to be quite similar. In general, the relative importance of the five components may be dramatically different from that in Figure 5 where the yield income dominates. The longer the asset's duration and the shorter the investment horizon, the greater is the relative importance of the duration impact. It is worth noting that realized returns can be decomposed in the same way as the expected returns and that the duration impact typically dominates the realized returns even more.<sup>10</sup>

Figure 6. Expected Returns of a Three-Month Bill, a Three-Year Bond and a Ten-Year Bond, as of 26 Sep 95



The total expected returns, if estimated carefully, should produce the most useful signals for relative value analysis because they include all sources of expected returns. Yield spreads may be useful signals, but they are only a part of the picture. Therefore, we advocate the monitoring of broader expected return measures relative to their history as cheapness indicators — just as yield spreads are often monitored relative to their history.

The components of expected returns discussed above are not new. However, few investors have combined these components into an integrated framework and based their historical analysis on broad expected return measures. An additional useful feature of this framework is that all types of government bond trades can be evaluated consistently within it: the portfolio duration decision (market-directional view); the maturity-sector positioning and barbell-bullet decision (curve-reshaping view); and the individual issue selection (local cheapness view). With small modifications, the framework can be extended to include the

<sup>10</sup> Realized returns can be split into an expected part and an unexpected part, and both parts can be decomposed further. Equation (3) describes the decomposition of the expected part, while the unexpected part can be split into duration and convexity effects. This type of return attribution can have a useful role in risk management and performance evaluation, but these two activities are not our focus in this report.

cross-country analysis of currency-hedged government bond positions. Other possible future extensions include the analysis of foreign exchange exposure and the analysis of spread positions between government bonds and other fixed-income assets.

We note some reservations. Even if two investors use the same general framework and the same type of expected return measure, they may come up with different numbers because of different data sources and different estimation techniques.

• The whole analysis can be made with any raw material; we emphasize the importance of good-quality inputs. Various candidates for the raw material include on-the-run and off -the-run government bonds, STRIPS, Eurodeposits, swaps, and Eurodeposit futures. [This multitude of course opens the possibility of trading between these curves if we can assess how various characteristics (say, convexity) are priced in each curve.] The most common approach is first to estimate the spot curve (or discount function) using a broad universe of coupon government bonds as the raw material, and then to compute the forward rates and other relevant numbers. In European bond markets, the liquid swap curve (using cash Eurodeposits and swaps as the raw material) has gained more of a benchmark status. Of course, some credit and tax-related spread may exist between the swap curve and the government bond yield curve. Recently, yet another approach has become popular: Eurodeposit futures prices are used as the raw material. In this case, the forward rates are computed by adjusting for the convexity difference between a futures contract and a forward contract, and only then are spot rates computed from the forwards.

• Some components of expected returns are easier to measure — and less debatable — than others. The yield income is relatively unambiguous. The rolldown return and the local rich/cheap effects depend on the curve-fitting technique. The value of convexity depends on the volatility input and, thus, on the volatility estimation technique. The rate "view," the fourth term, can be based on various approaches, such as quantitative modeling or subjective forecasting that relies on fundamental or technical analysis. Even the quantitative approach is not purely objective because infinitely many alternative forecasting models and estimation techniques exist. Forecasting rate changes is of course the most difficult task as well as the one with greatest potential rewards and risks. Forecasting changes in yield spreads may be almost as difficult. The short-term returns of most bond positions depend primarily on the duration impact (rate changes or spread changes). However, even if investors cannot predict rate changes, they may earn superior returns in the long run — and with less volatility — by systematically exploiting the more stable sources of expected return differentials across bonds: yields; rolldown returns; value of convexity; and local rich/cheap effects. More generally, while the total expected return differentials are, in theory, better relative value indicators than the yield spreads, in practice, measurement errors conceivably can make them so noisy that they give worse signals. Therefore, it is important to check with historical data that any supposedly superior relative value tools would have enhanced the investment performance at least in the past.

#### Link to Scenario Analysis

Many active investors base their investment decisions on subjective yield curve views, often with the help of scenario analysis. Our framework for relative value analysis is closely related to scenario analysis. It may be worthwhile to explore the linkages further. An investor can perform the scenario analysis of government bonds in two steps. First, the investor specifies a few yield curve scenarios for a given horizon and computes the total return of his bond portfolio — or perhaps just a particular trade — under each scenario. Second, the investor assigns subjective probabilities to the different scenarios and computes the probability-weighted expected return for his portfolio. Sometimes the second step is not completed and investors only examine qualitatively the portfolio performance under each scenario. However, we advocate performing this step because investors can gain valuable insights from it. Specifically, the probability-weighted expected return is the "bottom-line" number a total return manager should care about. By assigning probabilities to scenarios, investors also can explicitly back out their implied views about the yield curve reshaping and about yield volatilities and correlations.

In scenario analysis, investors define the mean yield curve view and the volatility view implicitly by choosing a set of scenarios and by assigning them probabilities. In contrast, our framework for relative value analysis involves explicitly specifying one yield curve view (which corresponds to the probability-weighted mean yield curve scenario) and a volatility view (which corresponds to the dispersion of the yield curve scenarios). Either way, the yield curve view determines the duration impact and the volatility view determines the value of convexity (and these views together approximately define the expected yield distribution).

Figure 7 presents a portfolio that consists of five equally weighted zero-coupon bonds with maturities of one to five years and (annually compounded) yields between 6% and 7%. The portfolio's maturity — and its Macaulay duration — initially is three years. Over a one-year horizon, each zero's maturity shortens by one year. We specify five alternative yield curve scenarios over the horizon: parallel shifts of +100 basis points and -100 basis points; no change; a yield increase combined with a curve flattening; and a yield decline combined with a curve steepening (see Figure 8). We compute the one-year holding-period returns for each asset and for the portfolio under each scenario. In particular, the neutral scenario shows the rolling yield that each zero earns if the yield curve remains unchanged. We can evaluate each scenario separately. However, such analysis gives us limited insight — for example, the last column in Figure 7 shows just that bearish scenarios produce lower portfolio returns than bullish scenarios.

In contrast, if we assign probabilities to the scenarios, we can back out many numbers of potential interest. We begin with a simple example in which we only use the two first scenarios, parallel shifts of 100 basis points up or down. If we assign these scenarios equal probabilities (0.5), the expected return of the portfolio is 7.04% (= 0.5\*5.02 + 0.5\*9.06). On average, these scenarios have no view about curve changes; yet, this expected return is four basis points higher than the expected portfolio return given no change in the curve (that is, the 7% rolling yield computed in the neutral scenario). This difference reflects the value of convexity. If we only use one scenario, we implicitly assume zero volatility, which leads to downward-biased expected return estimates for positively convex bond positions. If we use the two first scenarios (bear and bull), we implicitly assume a 100-basis-point yield volatility; this assumption may or may not be reasonable, but it certainly is more reasonable than an assumption of no volatility. This example highlights the importance of using multiple scenarios to recognize the value of convexity. (The value is small here, however, because we focus on short-duration assets that have little convexity.)

## Figure 7. Scenario Analysis and Expected Bond Returns

	Bond					Portfolio
Initial Maturity	1	2	3	4	5	3
Horizon Maturity	0	1	2	3	4	2
Initial Yield	6.00%	6.25%	6.50%	6.75%	7.00%	_
Yield Change Scenarios						
(Of 1- to 5-Year Constant-Maturity Rates)						
Bear	1.00%	1.00%	1.00%	1.00%	1.00%	
Bull	-1.00	-1.00	-1.00	-1.00	-1.00	
Neutral	0.00	0.00	0.00	0.00	0.00	
Bear-Flattener	1.00	0.875	0.75	0.625	0.50	
Bull-Steepener	-0.50	-0.375	-0.25	-0.125	0.00	
One-Year Returns in Each Scenario						
Bear	6.00%	5.51%	5.02%	4 53%	4.05%	5.02%
Bull	6.00	7.51	9.04	10.59	12 15	9.02.0
Neutral	6.00	6.50	7 00	7 50	8.01	7 00
Bear-Flattener	6.00	5.51	5.26	5.26	5.51	5.51
Bull-Steepener	6.00	7.01	7.76	8.26	8.51	7.51
Assign Equal Probability					•	·
(0.2) to Each Scenario and Back Out Various Statistics						
Mean Return	6.00%	6 41%	6 82%	7 23%	7 65%	6 82%
Vol. of Return	0.00	0.80	1.52	2.17	2.78	1.45
Mean Yield Change	0.10	0.10	0.10	0.10	0.10	
Vol. of Yield Change	0.80	0.76	0.72	0.69	0.66	



Now we return to the example with all five yield curve scenarios in Figure 8. As an illustration, we assign each scenario the same probability ( $p_i = 0.2$ ). Then, it is easy to compute the portfolio's probability-weighted expected return:

$$E(h_p) = \sum_{i=1}^{5} p_i * h_i = 0.2 * (5.02 + 9.06 + 7.00 + 5.51 + 7.51) = 6.82$$
(4)

Given these probabilities, we can compute the expected return for each asset, and **it is possible to back out the implied yield curve views**. The lower panel in Figure 7 shows that the mean yield change across scenarios

is +10 basis points for each rate (because the bear-flattener and the bull-steepener scenarios are not quite symmetric in magnitude in this example), **implying a mild bearish bias but no implied curve-steepness views**. In addition, we can back out the implied basis-point yield volatilities (or return volatilities) by measuring how much the yield change (or return) outcomes in each scenario deviate from the mean. These yield volatility levels are important determinants of the value of convexity. The last line in Figure 7 shows that the volatilities range from 80 to 66 basis points, **implying an inverted term structure of volatility**. Finally, we can compute implied correlations between various-maturity yield changes; the curve behavior across the five scenarios is so similar that all correlations are 0.92 or higher (not shown). Note that all correlations would equal 1.00 if only the first three scenarios were used; the imperfect correlations arise from the bear-flattener and the bull-steepener scenarios.

Whenever an investor uses scenario analysis, he should back out these implicit curve views, volatilities and correlations — and check that any biases are reasonable and consistent with his own views. Without assigning the probabilities to each scenario, this step cannot be completed; then, the investor may overlook hidden biases in his analysis, such as a biased curve view or a very high or low implicit volatility assumption which makes positive convexity positions appear too good or too bad. If investors use quantitative tools — such as scenario analysis, mean-variance optimization, or the approach outlined in this report — to evaluate expected returns, they should recognize the importance of their rate views in this process. Strong subjective views can make *any* particular position appear attractive. Therefore, investors should have the discipline and the ability to be fully aware of the views that are input into the quantitative tool.

In addition to the implied curve views, we can back out the four components of expected returns discussed above. In this example, we only analyze bonds that lie "on the curve" and thus can ignore the fifth component, the local rich/cheap effects. (1) We measure the yield income from the portfolio by a market-value weighted average yield of the five zeros, which is 6.50% (see footnote 7). (2) Each asset's rolldown return is the difference between the horizon return given an unchanged yield curve and the yield income. Figure 7 shows that the horizon return for the portfolio is 7% in the neutral scenario; thus, the portfolio's (market-value weighted average) rolldown return is 50 basis points (= 7% - 6.5%). Note that the rolldown return is larger for longer bonds, reflecting the fact that the same rolldown yield change (25 basis points) produces larger capital gains for longer bonds. (3) The value of convexity for each zero can be approximated by 0.5\*convexity at horizon\*(basis-point yield volatility)<sup>2\*</sup>(1 + rolling yield/100). Using the implicit yield volatilities in Figure 7, this value varies between 0.6 and 4.5 basis points across bonds. The portfolio's value of convexity is a market-value weighted average of the bond-specific values of convexity, or roughly two basis points. (4) The duration impact of the rate "view" for each bond equals (-duration at horizon)\*(expected yield change)\*(1 + rolling yield/100). The last term is needed because each invested dollar grows to (1 + rolling yield/100) by the end of horizon when the repricing occurs. The core of the duration impact is the product of duration and expected yield change. The expected yield change refers to the change (over the investment horizon) in a constant-maturity rate of the bond's horizon maturity. In Figure 7, all rates

are expected to increase by ten basis points, and the duration impact on specific bonds' returns varies between 0 and -40 basis points. The portfolio's duration impact is a market-value weighted average of bond-specific duration impacts, or about **-20 basis points**.<sup>11</sup>

Figure 9 shows that the four components add up to the total probability-weighted expected return of 6.82%. Decomposing expected returns into these components should help investors to better understand their own investment positions. For example, they can see what part of the expected return reflects static market conditions and what part reflects their subjective market view. Unless they are extremely confident about their market view, they can emphasize the part of expected return advantage that reflects static market conditions. In our example, the duration effect is small because the implied rate view is quite mild (ten basis points) and the one-year horizon is relatively long (the "slower" effects have time to accrue). With a shorter horizon and stronger rate views, the duration impact would easily dominate the other effects.



#### Figure 9. Decomposing the Total Expected Return into Four Components

<sup>&</sup>lt;sup>11</sup> It is possible to compute the value of convexity and the duration impact in another way and end up with almost exactly the same numbers. If we compute the bond returns based on yield curve scenarios that are, on average, unbiased or "viewless" (that is, we subtract the ten-basis-point mean yield change from each scenario), the probability-weighted expected portfolio return is 7.02%. We can estimate the duration impact of a view by comparing the portfolio's probability-weighted expected return given the five scenarios with the rate view (6.82%) to the expected return given the "viewless" set of scenarios. The difference is again -20 basis points. Similarly, we can estimate the value of convexity by comparing the probability-weighted expected return of the portfolio under a "viewless" set of scenarios (7.02%) to the expected return under the neutral scenario that has the same average rate view — an unchanged yield curve — but assumes no volatility (7.00%). The difference, two basis points, measures the pure effect of the implicit volatility forecast without a bias from a yield curve view.

### APPENDIX A. DECOMPOSING THE FORWARD RATE STRUCTURE INTO ITS MAIN DETERMINANTS

In this appendix, we show how the forward rate structure is related to the market's rate expectations, bond risk premia and convexity bias. In particular, the holding-period return of an n-year zero-coupon bond can be described as a sum of its horizon return given an unchanged yield curve and the end-of-horizon price change that is caused by a change in the n-1 year constant-maturity spot rate  $(\Delta s_{n-1})$ . The horizon return equals a one-year forward rate, and the end-of-horizon price change can be approximated by duration and convexity effects. These relations are used to decompose near-term expected bond returns and the one-period forward rates into simple building blocks. All rates and returns used in the following equations are compounded annually and expressed in percentage terms.

$$\frac{\mathbf{h}_{n}}{100} = \frac{\mathbf{P}_{n-1}^{*} - \mathbf{P}_{n}}{\mathbf{P}_{n}} = \frac{(\mathbf{P}_{n-1}^{*} - \mathbf{P}_{n-1}) + (\mathbf{P}_{n-1} - \mathbf{P}_{n})}{\mathbf{P}_{n}}$$

$$= \left(\frac{\Delta \mathbf{P}_{n-1}}{\mathbf{P}_{n-1}} * \frac{\mathbf{P}_{n-1}}{\mathbf{P}_{n}}\right) + \left(\frac{\mathbf{P}_{n-1}}{\mathbf{P}_{n}} - 1\right),$$
(5)

where  $h_n$  is the one-period holding-period return of an n-year bond,  $P_n$  is its price (today),  $P_{n-1}^*$  is its price in the next period (when its maturity is n-1), and  $\Delta P_{n-1} = P_{n-1}^* - P_{n-1}$ . The second term on the right-hand side of Equation (5) is the bond's rolling yield (horizon return). The first term on the right-hand side of Equation (5) is the instantaneous percentage price change of an n-1 year zero, multiplied by an adjustment term  $P_{n-1}/P_n$ .<sup>12</sup>

Equation (6) shows that the zero's rolling yield  $(P_{n-1}/P_n - 1)$  equals, by construction, the one-year forward rate between n-1 and n. Moreover, the adjustment term equals one plus the forward rate.

$$1 + \frac{f_{n-1,n}}{100} = \frac{\left(1 + \frac{s_n}{100}\right)^n}{\left(1 + \frac{s_{n-1}}{100}\right)^{n-1}} = \frac{P_{n-1}}{P_n} .$$
(6)

Equation (7) shows the well-known result that the percentage price change  $(\Delta P/P)$  is closely approximated by the first two terms of a Taylor series expansion, duration and convexity effects.

$$100 * \frac{\Delta P}{P} \approx - Dur * (\Delta s) + 0.5 * Cx * (\Delta s)^2,$$
 (7)

where Dur  $\equiv$  -  $\frac{dP}{ds} * \frac{100}{P}$  and  $Cx \equiv \frac{d^2P}{d^2s} * \frac{100}{P}$ .

<sup>&</sup>lt;sup>12</sup> The adjustment term is needed because the bond's instantaneous price change occurs at the end of horizon, not today. The value of the bond position grows from one to  $P_{n-1}/P_n$  at the end of horizon if the yield curve is unchanged. The end-of-horizon value  $(P_{n-1}/P_n)$  would be subject to the yield shift at horizon.

Plugging Equations (6) and (7) into (5), we get:

$$h_n \approx f_{n-1,n} + (1 + \frac{f_{n-1,n}}{100})^* [-Dur_{n-1}^*(\Delta s_{n-1}) + 0.5^* C x_{n-1}^*(s_{n-1})^2].$$
 (8)

Even if the yield curve shifts occur during the horizon, for performance calculation purposes the repricing takes place at the end of horizon. This disparity causes various differences between the percentage price changes in Equations (7) and (8). First, the amount of capital that experiences the price change grows to  $(1 + f_{n-1,n}/100)$  by the end of horizon. Second, the relevant yield change is the change in the n-1 year constant-maturity rate, not in the n-year zero's own yield (the difference is the rolldown yield change).<sup>13</sup> Third, the end-of-horizon (as opposed to the current) duration and convexity determine the price change.

The realized return can be split into an expected part and an unexpected part. Taking expectations of both sides of Equation (8) gives us the n-year zero's expected return over the next year:

$$E(h_n) \approx f_{n-1,n} + (1 + \frac{f_{n-1,n}}{100})^* [-Dur_{n-1}^* E(\Delta s_{n-1}) + 0.5^* C x_{n-1}^* E(s_{n-1})^2].$$
(9)

Recall from Equation (6) that the one-period forward rate equals a zero's rolling yield, which can be split to yield and rolldown return components. In addition, the expected yield change squared is approximately equal to the variance of the yield change or the squared volatility,  $E(\Delta s_{n-1})^2 \approx (Vol(\Delta s_{n-1}))^2$ . This relation is exact if the expected yield change is zero. Thus, the zero's near-term expected return can be written (approximately) as a sum of the yield income, the rolldown return, the value of convexity, and the expected capital gains from the rate "view" (see Equation (3)).

We can interpret the expectations in Equation (9) to refer to the market's rate expectations. Mechanically, the forward rate structure and the market's rate expectations on the right-hand side of Equation (9) determine the near-term expected returns on the left-hand side. These expected returns should equal the required returns that the market demands for various bonds if the market's expectations are internally consistent. These required returns, in turn, depend on factors such as each bond's riskiness and the market's risk aversion level. Thus, it is more appropriate to think that the market participants, in the aggregate, set the bond market prices to be such that given the forward rate structure and the consensus rate expectations, each bond is expected to earn its required return.<sup>14</sup>

 $<sup>^{13}</sup>$  If we used bonds' own yield changes in Equation (8), these yield changes would include the rolldown yield change. In this case, we should not use the forward rate (which includes the impact of the rolldown yield change on the return, in addition to the yield income) as the first term on the right-hand side of Equation (8). Instead, we would use the spot rate.

<sup>&</sup>lt;sup>14</sup> Individual investors also can use Equation (9) but the interpretation is slightly different because most of their investments are so small that they do not influence the market rates; thus, they are "price-takers." Any individual investor can plug his subjective rate expectations into Equation (9) and back out the expected return given these expectations and the market-determined forward rates. These expected returns may differ from the required returns that the market demands; this discrepancy may prompt the investor to trade on his view.

Subtracting the one-period riskless rate  $(s_1)$  from both sides of Equation (9), we get:

$$E(h_{n} - s_{1}) \approx (f_{n-1,n} - s_{1}) + (1 + \frac{f_{n-1,n}}{100})^{*} [-Dur_{n-1}^{*}E(\Delta s_{n-1}) + 0.5^{*}Cx_{n-1}^{*}(Vol(s_{n-1}))^{2}].$$
(10)

We define the bond risk premium as  $BRP_n \equiv E(h_n - s_1)$  and the forward-spot premium as  $FSP_n \equiv f_{n-1,n} - s_1$ . The forward-spot premium measures the steepness of the one-year forward rate curve (the difference between each point on the forward rate curve and the first point on that curve) and it is closely related to simpler measures of yield curve steepness. Rearranging Equation (10), we obtain:

$$FSP_{n} \approx BRP_{n} + (1 + \frac{f_{n-1,n}}{100}) * [-Dur_{n-1} * E(\Delta s_{n-1}) + 0.5 * Cx_{n-1} * (Vol(s_{n-1}))^{2}].$$
(11)

In other words, the forward-spot premium is approximately equal to a sum of the bond risk premium, the impact of rate expectations (expected capital gain/loss caused by the market's rate "view") and the convexity bias (expected capital gain caused by the rate uncertainty). Unfortunately, none of the three components is directly observable.

The analysis thus far has been very general, based on accounting identities and approximations, not on economic assumptions. Various term structure hypotheses and models differ in their assumptions. Certain simplifying assumptions lead to well-known hypotheses of the term structure behavior by making some terms in Equation (11) equal zero --- although fully specified term structure models require even more specific assumptions. First, if constant-maturity rates follow a random walk, the forward-spot premium mainly reflects the bond risk premium, but also the convexity bias  $[E(\Delta s_{n-1}) = 0 \Rightarrow FSP_n \approx BRP_n + CB_{n-1}]$ . Second, if the local expectations hypothesis holds (all bonds have the same near-term expected return), the forward-spot premium mainly reflects the market's rate expectations, but also the convexity bias  $[BRP_n = 0 \Rightarrow FSP_n \approx$  $\text{Dur}_{n-1} * E(\Delta s_{n-1}) + CB_{n-1}$ ]. Third, if the unbiased expectations hypothesis holds, the forward-spot premium only reflects the market's rate expectations [BRP  $_{n}$  + CB $_{n-1}$  = 0  $\Rightarrow$  FSP $_{n} \approx$  Dur $_{n-1} \approx E(\Delta s_{n-1})$ ]. The last two cases illustrate the distinction between two versions of the pure expectations hypothesis.

#### APPENDIX B. RELATING VARIOUS STATEMENTS ABOUT FORWARD Rates to each other

In the series *Understanding the Yield Curve*, we make several statements about forward rates — describing, interpreting and decomposing them in various ways. The multitude of these statements may be confusing; therefore, we now try to clarify the relationships between them.

We refer to the spot curve and the forward curves on a given date as if they were unambiguous. In reality, **different analysts can produce somewhat different estimates of the spot curve on a given date if they use different curve-fitting techniques or different underlying data** (asset universe or pricing source). We acknowledge the importance of these issues — having good raw material is important to any kind of yield curve analysis — but in our reports we ignore these differences. We take the estimated spot curve as given and focus on showing how to interpret and use the information in this curve.

In contrast, the relations between various depictions of the term structure of interest rates (par, spot and forward rate curves) are unambiguous. In particular, once a spot curve has been estimated, any forward rate can be mathematically computed by using Equation (12):

$$\left(1 + \frac{f_{m,n}}{100}\right)^{n-m} = \frac{\left(1 + \frac{s_n}{100}\right)^n}{\left(1 + \frac{s_m}{100}\right)^m},$$
(12)

where  $f_{m,n}$  is the annualized n-m year interest rate m years forward and  $s_n$  and  $s_m$  are the annualized n-year and m-year spot rates, expressed in percent. Thus, a **one-to-one mapping exists between forward rates and** *current* **spot rates.** The statement "the forwards imply rising rates" is equivalent to saying that "the spot curve is upward sloping," and the statement "the forwards imply curve flattening" is equivalent to saying that "the spot curve, an unambiguous mapping exists between various types of forward curves, such as the implied spot curve one year forward ( $f_{1,n}$ ) and the curve of constant-maturity one-year forward rates ( $f_{n-1,n}$ ).

The forward rate can be the agreed interest rate on an **explicitly** traded contract, a loan between two future dates. More often, the forward rate is **implicitly** defined from today's spot curve based on Equation (12). However, arbitrage forces ensure that even the explicitly traded forward rates would equal the implied forward rates and, thus, be consistent with Equation (12). For example, the implied one-year spot rate four years forward (also called the one-year forward rate four years ahead,  $f_{4,5}$ ) must be such that the equality  $(1 + s_5/100)^5 = (1 + s_4/100)^4 * (1 + f_{4,5}/100)$  holds. If  $f_{4,5}$  is higher than that, arbitrageurs can earn profits by short-selling the five-year zeros and buying the four-year zeros and the one-year forward contracts four years ahead, and vice versa. Such activity should make the equality hold within transaction costs.

Forward rates can be viewed in many ways: the arbitrage interpretation; the break-even interpretation; and the rolling yield interpretation. According to the arbitrage interpretation, implied forward rates are such rates that would ensure the absence of riskless arbitrage opportunities between spot contracts (zeros) and forward contracts if the latter were traded. According to the break-even interpretation of forward rates, implied forward rates are such *future* spot rates that would equate holding-period returns across bond positions. According to the rolling yield interpretation, the one-period forward rates show the one-period horizon returns that various zeros earn if the yield curve remains unchanged. Footnotes 15-17 show that each interpretation is useful for a certain purpose: active view-taking relative to the forwards (break-even); relative value analysis given no yield curve views (rolling yield); and valuation of derivatives (arbitrage).

All of these interpretations hold by construction (from Equation (12)). Thus, they are not inconsistent with each other. For example, the one-period forward rates can be interpreted and used in quite different ways. The implied one-year spot rate four years forward  $(f_{4,5})$  can be viewed as either the break-even one-year rate four years into the future or the rolling yield of a five-year zero over the next year. Both interpretations follow from the equality  $(1 + s_5/100)^5 = (1 + s_4/100)^4 *$  $(1 + f_{4.5}/100)$ . This equation shows that the forward rate is the break-even one-year reinvestment rate that would equate the returns between two strategies (holding the five-year zero to maturity versus buying the four-year zero and reinvesting in the one-year zero when the four-year zero matures) over a five-year horizon. [Rewriting the equality as  $(1 + s_4/100)^4$ =  $(1 + s_5/100)^5/(1 + f_{4.5}/100)$  gives a slightly different viewpoint; the forward rate also is the break-even selling rate that would equate the returns between two strategies (holding the four-year zero to maturity versus buying the five-year zero and selling it after four years as a one-year zero) over a four-year horizon.] Finally, rewriting the equality as  $1 + f_{4.5}/100 = (1 + s_5/100)^5/(1 + s_4/100)^4$  shows that the forward rate is the horizon return from buying a five-year zero at rate s5 and selling it one year later, as a four-year zero, at rate  $s_4$  (thus, the constant-maturity four-year rate is unchanged from today). In this series, we focus on the last (rolling yield) interpretation.

Interpreting the one-period forward rates as rolling yields enhances our understanding about the relation between the curve of one-year forward rates ( $f_{0,1}$ ,  $f_{1,2}$ ,  $f_{2,3}$ ,..., $f_{n-1,n}$ ) and the implied spot curve one year forward ( $f_{1,2}$ ,  $f_{1,3}$ ,  $f_{1,4}$ ,..., $f_{1,n}$ ). The latter "break-even" curve shows how much the spot curve needs to shift to cause capital gains/losses that exactly offset initial rolling yield differentials across zeros and, thereby, equalize the holding-period returns. Thus, a steeply upward-sloping curve of one-period forward rates requires, or "implies," a large offsetting increase in the spot curve over the horizon, while a flat curve of one-period forward rates only implies a small "break-even" shift in the spot curve.<sup>15</sup> A similar link exists for the rolling yield differential between a duration-neutral barbell versus bullet and the break-even yield spread change (curve flattening) that is needed to offset the bullet's rolling yield advantage. These examples provide insight as to why an upward-sloping spot curve implies rising rates and why a concave spot curve implies a flattening curve.

<sup>&</sup>lt;sup>15</sup> Part 1 of this series describes one common way to use the break-even forward rates. Investors can compare their subjective views about the yield curve at some future date (or about the path of some constant-maturity rate over time) to the forward rates and directly determine whether bullish or bearish strategies are appropriate. If the rate changes that the forwards imply are realized, all bonds earn the riskless return because  $(1 + s_n/100)^n(1 + f_{1,n}/100)^{n+1} = (1 + s_n/100)^n$ . If rates rise by more than that, long bonds underperform short bonds. If rates rise by less than that, long bonds outperform short bonds. If rates rise by less than that, long bonds outperform short bonds (because their capital losses do not quite offset their initial yield advantage).

Appendix A shows that forward rates can be conceptually decomposed into three main determinants (rate expectations, risk premia, convexity bias). One might hope that the arbitrage, break-even or rolling yield interpretations could help us in backing out the relative roles of rate expectations, risk premia and convexity bias in a given day's forward rate structure. However, such hope is in vain. The three interpretations hold quite generally because of their mathematical nature. Thus, they do not guide us in decomposing the forward rate structure.

Therefore, even when two analysts agree that today's forward rate structure is an approximate sum of three components, they may disagree about the relative roles of these components. We can try to address this question empirically. It is closely related to the question about the forward rates' ability to forecast future rate changes and future bond returns. Ignoring convexity bias, if the forwards primarily reflect rate expectations, they should be unbiased predictors of future spot rates (and they should tell little about future bond returns). However, if the forwards mainly reflect required bond risk premia, they should be unbiased predictors of future rate changes). In Part 2 of this series, *Market's Rate Expectations and Forward Rates*, we present some empirical evidence indicating that the forward rates are better predictors of future bond returns than of future rate changes.<sup>16,17</sup>

Finally, our analysis does not reveal the fundamental economic determinants of the required risk premia or the market's rate expectations — nor does it tell us to what extent the nominal rate expectations reflect expected inflation and expected real rates. Macroeconomic news about economic growth, inflation rates, budget deficits, and so on, can influence both the required risk premia and the market's rate expectations. More work is clearly needed to improve our understanding about the mechanisms of these influences.

<sup>&</sup>lt;sup>16</sup> This evidence also suggests that the current spot curve is a better predictor of the next-period spot curve than is the implied spot curve one period forward. These findings imply that the rolling yields are reasonable proxies for the near-term expected bond returns — although even rolling yields capture a very small part of the short-term realized bond returns. Note that the poorer the forwards are in predicting future rate changes, the better they are in predicting bond returns — because then the implied rate changes that would offset initial yield advantages tend to occur more rarely. Note also that some investors may not care whether the forwards' ability to predict bond returns reflects rational risk premia or the market's inability to forecast rate changes; they want to earn any predictable profit irrespective of its reason.

<sup>&</sup>lt;sup>17</sup> One common misconception is that the forward rates are used in the valuation of swaps, options and other derivative instruments *because* the forwards are good predictors of future spot rates. In fact, the forwards' ability to predict future spot rates has nothing to do with their usefulness in derivatives pricing. Unlike forecasting returns, the valuation of derivatives is based on arbitrage arguments. For example, traders can theoretically construct, by dynamic hedging, a riskless combination of a risky long-term bond and an option written on it. The price of the option should be such that the hedged position earns the riskless rate — otherwise a riskless arbitrage opportunity arises. The forward rates are central in this valuation because the traders, via their hedging activity, can lock in these rates for future periods. This arbitrage argument implies that the yield curve option pricing models should be calibrated to be consistent with the market forward rates *despite* the fact that the forwards are quite poor predictors of future spot rates.