

A VaR-based Model for the Yield Curve*

Ruben D. Cohen, Citigroup, London E14 5LB, UK.

Abstract

An intuitive model for the yield curve, based on the notion of value-at-risk, is presented. It leads to interest rates that hedge against potential losses incurred from holding an underlying risky security until maturity. This result is also shown to tie in directly with the Capital Asset Pricing Model via the Sharpe Ratio. The conclusion here is that the normal yield curve can be characterised by a constant Sharpe Ratio, non-dimensionalised with respect to \sqrt{T} , where T is the bond maturity.

Among other features of the model are that it is able to explain, qualitatively if not quantitatively, the existence of (1) a normal yield curve at times of “normal economic growth”, (2) an inverted curve during periods of “high uncertainty”, “high interest rates” or “low economic growth”,¹ (3) a flat yield curve in more certain times and (4) a liquidity trap when economic growth is expected to be negative.

1. Introduction

The yield curve, which is a graphical depiction of interest rates plotted against the maturities of the underlying fixed income instrument—most commonly treasury bonds—constitutes an important area of economics and finance, and its determination has proven to be one of the most challenging endeavours undertaken in both theoretical and practical research. Constructing original framework models to predict these curves lies at the forefront of activities in every major financial institution, as any successful trading strategy that originates from these is guaranteed to generate vast amounts of profits.

Although interest in the shape of yield curves began as early as 1913 with work on the business cycle (Mitchell, 1913), it was Kessel (1965) who first focused specifically on the behaviour of the term spreads. Since then, work, both theoretical and empirical, has advanced to cover the use of term spreads to forecast the economy, as it was always recognised that the yield curve is not only affected by economic policies and shocks (Evans & Marshall, 2001) but that it could also serve as a leading indicator for a variety of economic parameters, including changes in *GDP*, *GNP*, consumption, investment, etc. (Diebold *et al*, 2005; Ang *et al*, 2006; among others).

Bearing in mind the above, as well as other criteria [such as the real rate, inflation premium and interest rate risk premium] that go into forming the yield curve and many of which are widely discussed in the literature [see, for example, Esterella (2005) for a literature review], it is accepted that two independent schools of thought lie behind what shapes the curve. These are the (a) market segmentation hypothesis and (b) expectations hypothesis (Fabozzi, 1996 & 1999). Whereas in the former the

shape of the yield curve is determined by the supply and demand for securities within each maturity sector, assuming that neither investors nor borrowers are willing to shift from one maturity sector to another, in the latter, broadly speaking, it is presumed that the underlying forward rates signal the market’s expectations of future actual rates.

These two hypotheses together lead to the following three attributes, namely (i) interest rates for different maturities move together in time, (ii) yield curves tend to slope upward when short rates are low and downward when short rates are high and, finally, (iii) yield curves are typically upward sloping. However, since the expectations hypothesis explains only points i and ii above and segmentation only iii, a combined theory, being the liquidity premium, was later put in place to account for all three observations. In summary, this theory states that investors prefer to hold shorter bonds to long, owing to a liquidity premium that increases with maturity. This rather subjective characterisation leads to the third behaviour that yield curves are typically upward sloping.

While the preceding paragraphs collectively reflect the quasi-static or quasi-equilibrium nature of yield curves and ignore their time-dependent movements, a whole new quantitative area, emphasising specifically on the dynamics, has developed within the past 2 to 3 decades. This field, which has attracted great attention from theoreticians and practitioners alike, centres on the time-dependent stochastic nature of the yield curve and has found its niche in the pricing of fixed-income instruments in a dynamic setting. As this area is out of the scope of this work, we shall not delve into it any further and, instead, refer the interested reader to some background literature, including, among many others, Brace *et al* (1997), Benninga & Wiener (1998) and references therein.

I express these views as an individual, not as a representative of companies with which I am connected.

*I would like to thank Aaron Brown, Head of Credit Risk Architecture at Morgan Stanley, NY, and Simeon Tsonev, currently at the Department of Economics, Columbia University, NY, for their constructive comments and criticisms on various aspects of this work. Also, thanks to Simeon for providing the data.

Thus, given the various aspects of the yield curve, especially in the light of the quasi-equilibrium models, along with the different governing theories summarised above, we now proceed to develop the framework for an alternative approach. The model that results from this should, hopefully, provide us with not only an intuitive explanation for why a yield curve assumes a certain shape—be it normal, inverted or flat—it will also, among other things, open up a direct link with the Capital Asset Pricing Model (CAPM), placing the yield curve well within the confines of classical risk-return principles in portfolio management theory.

2. Methodology

The framework for the model is derived here by considering investing in a risky security, or portfolio of securities, taking into account three different scenarios. These securities are not pre-defined. Rather, they are “underlying”, similar in context to option-pricing methodologies, where characteristics such as “implied” return and volatility are extracted from market data. In addition, as with any other model, certain simplifying assumptions are made. As we construct the model and introduce these assumptions, every attempt is made to rationalise and justify them.

2A. Scenario I—Investing in a Risky Security with Expected Return Higher than the Interest Rate

Consider the more common scenario² of investing in a risky security with expected returns higher than the interest rate. For this, one needs to:

- first take a view on the total return, r , and volatility or standard deviation, σ , of the security going forward.
- Plan to hold the security for a certain length of time, T .
- Borrow sufficient funds to purchase the security, with intent to sell it back at time T and pay off the interest with the proceeds. Issuing a zero bond with maturity T and yield to maturity b_T could do this.
- The yield could be selected such that the cumulative interest paid at time T —i.e. $b_T T$ —is set at less than or equal to $rT - \sigma\sqrt{T}$,³ which is the security’s total return at time T at risk level -1σ . It is understood that, although at maturity the net expected return, $(r - b_T)T$, is positive owing to $r > b_T$, as per the definition for Scenario I, there is still some likelihood, due to the security’s risk, that the investment may under perform and produce a return of less than $b_T T$.

Note that Step 4 is analogous to a value-at-risk (VaR) approach, whereby the quantity $b_T T$ is set at 1σ below the mean of the distribution of the T -year return from the risky security, i.e. $rT - \sigma\sqrt{T}$. This, effectively, represents a level of risk aversion where the resulting interest rate, b_T , acts as a hedge against potential losses incurred from holding the underlying risky security until maturity.

The situation is exemplified in Figure 1, which illustrates the movement of the security’s return along $\pm 1\sigma$ in the total return vs. time plane. The diagram basically portrays something similar to, but not exactly, a parabola⁴ enveloping the time-varying mean, rt , the latter moving linearly with respect to time. In addition, two yield values, b_1 and b_2 , are included, maturing at T_1 and T_2 , respectively. These maturity times denote the points where lines $b_1 t$ and $b_2 t$ intersect the lower VaR level, being $rT - \sigma\sqrt{T}$. Why should

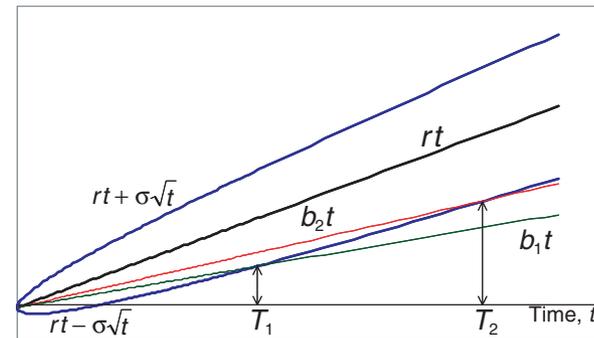


Figure 1: Movement of the risky security in time, with mean return r and standard deviation σ .

one limit the VaR levels to $\pm 1\sigma$ and not another value? Purely for convenience, as any other specification would produce a consistent outcome.

The bounds of $\pm 1\sigma$ are also shown, moving at $rt + \sigma\sqrt{t}$ and $rt - \sigma\sqrt{t}$, respectively, and yield values, b_1 and b_2 , intersecting the lower line [i.e. $rt - \sigma\sqrt{t}$] at times T_1 and T_2 , respectively.

In view of the above, therefore, we write

$$b_T T \leq rT - \sigma\sqrt{T} \quad (1)$$

which, when limited to the equality⁵, gives

$$b_T T = rT - \sigma\sqrt{T} \quad (2)$$

or

$$b_T = r - \frac{\sigma}{\sqrt{T}} \quad (3)$$

after dividing both sides by T . Equation 3, therefore, expresses the yield, b_T , as a function of maturity, T , with the expected r and σ remaining constant over the time horizon. This behaviour, as illustrated in Figure 2, portrays an upward sloping curve, typical of a normal yield curve.

With reference to earlier works, we should mention Siegel and Nelson (1988), who, through a totally different approach, derived a similar relationship, but with a maturity dependence of $1/T$ instead of $1/\sqrt{T}$. This was obtained by assuming that the yield decays as $1/T$ around an asymptotic value reached at $T \rightarrow \infty$.

Returning to Equation 3, an alternative way of plotting, which might be more useful for testing purposes, is to represent b_T as a function of $1/\sqrt{T}$. This leads to a straight line with a y-intercept of r and slope of $-\sigma$, as demonstrated in Figure 3. The advantage to this type of representation is that the expected return of the underlying security, along with its implied volatility, could be easily extracted.

2B. Scenario II—Investing in a Risky Security with Expected Return Lower than the Interest Rate

This scenario, as depicted in Figures 4a–b, differs from the previous in that the expected average return of the security is “low”. Here “low” means either low enough to cause the lower bound, $rT - \sigma\sqrt{T}$, of the quasi-parabola

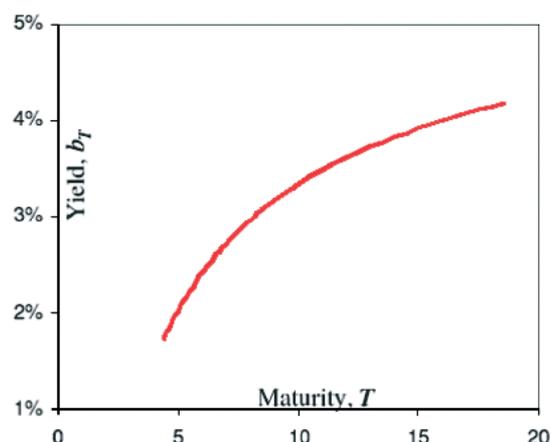


Figure 2: An illustration of the yield curve, obtained based on the methodology leading to Equation 3. Note that this is a normal yield curve.

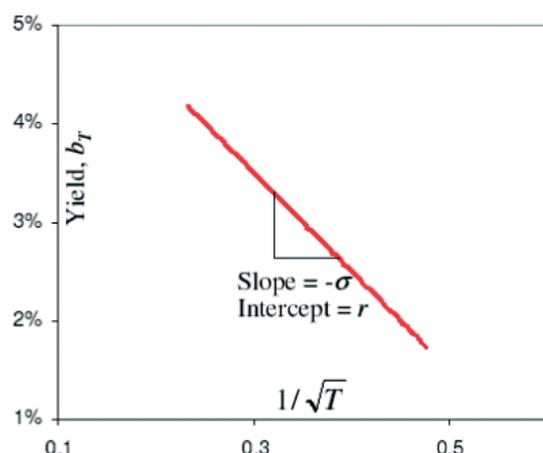
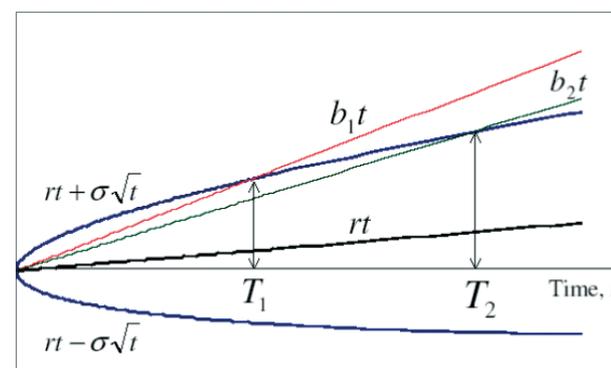


Figure 3: Equation 3 plotted as b_T vs. $1/\sqrt{T}$. The outcome is a straight line with slope $-\sigma$ and intercept r , equal to the implied [negative] volatility and mean return, respectively, of the underlying security.

to fall entirely in the negative territory, as illustrated in Figure 4a [i.e. this could be caused by a large σ], or that the interest rates are higher than the security's expected return, as shown in Figure 4b. The latter may happen even when σ is not sufficiently large to cause the bottom portion to fall in the negative region. It is, never the less, important to emphasise that the "low-growth" scenario does not necessarily mean expected security returns that are [or an economic growth rate that is] negative.⁶

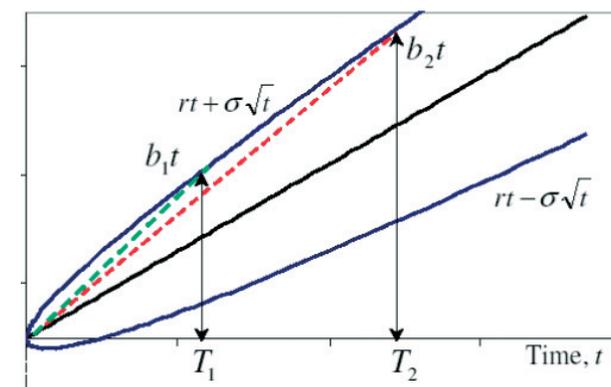
In either case above, although one must again first take a view on the total expected return, r , and standard deviation, σ , of the risky security going forward, as in Scenario I, the rules of borrowing the cash to invest in the security, as outlined there, do not apply. Instead, one must

- raise the cash by selling the underlying security, with intent to buy it back at time T , and
- use the cash to long a bond with yield and maturity b_T and T , respectively.



(a)

Figure 4a: Investing in a risky security with low r and high σ . This situation involves selling the security and longing the bond. Again, as in Figure 1, the bounds of $\pm 1\sigma$ are also shown, moving at $rt + \sigma\sqrt{t}$ and $rt - \sigma\sqrt{t}$, respectively, the latter falling completely in the negative territory. The yields, b_1 and b_2 , intersecting the upper line [i.e. $rt + \sigma\sqrt{t}$] at times T_1 and T_2 , respectively, are also shown.



(b)

Figure 4b: Investing in a risky security with returns lower than the interest rate. Again, the bounds of $\pm 1\sigma$ are also shown, moving at $rt + \sigma\sqrt{t}$ and $rt - \sigma\sqrt{t}$, respectively, although, in contrast to Figure 4a, the latter does not fall completely in the negative. The yields, b_1 and b_2 , intersecting the upper line [i.e. $rt + \sigma\sqrt{t}$] at times T_1 and T_2 , respectively, are also shown.

This approach is illustrated in both Figures 4a and 4b, where, in contrast to the previous scenario, the yield line, $b_T T$, intersects the *top* curve, $rt + \sigma\sqrt{T}$, at time T . Therefore, to profit from this investment, the following criterion must be met:

$$b_T T \geq rT + \sigma\sqrt{T} \quad (4)$$

where, again, the equality shall be used to give:

$$b_T T = rT + \sigma \sqrt{T} \tag{5}$$

Equation 4 depicts a net positive return of $(b_T - r)T$ to the portfolio, but with a VaR limited to $+1\sigma$.

Once more, dividing both sides of Equation 5 by T leads to

$$b_T = r + \frac{\sigma}{\sqrt{T}} \tag{6}$$

which represents the yield curve for Scenario II. This is illustrated in Figures 5 and 6, where the yield, b_T , is plotted against T and $1/\sqrt{T}$, respectively. The curve in Figure 5 is downward sloping, consistent with an inverted yield curve.

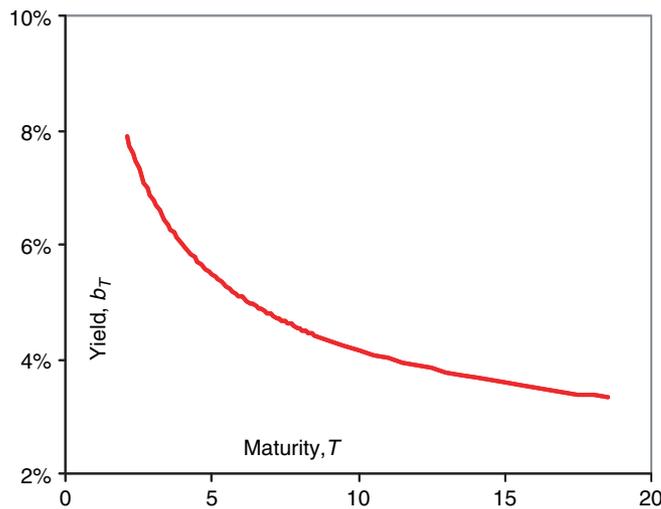


Figure 5: An illustration of the yield curve pertaining to Scenario II, represented by Equation 6 and plotted as b_T vs. T . Note that this is an inverted yield curve.

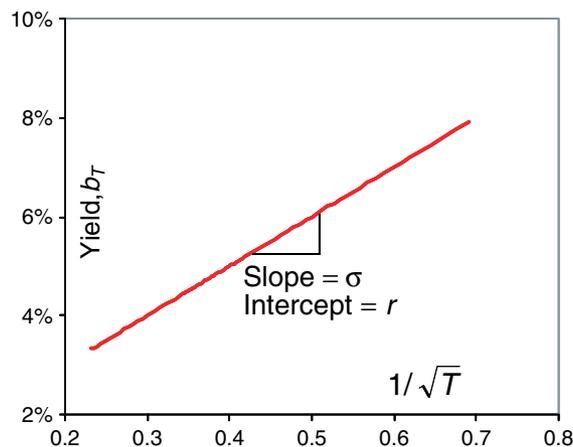


Figure 6: Equation 6 plotted as b_T vs. $1/\sqrt{T}$. The outcome is a straight line with slope σ and intercept r , equal to the implied volatility and mean expected return, respectively, of the underlying security.

With regards to inverted yield curves, it is accepted that they generally occur when a recession is anticipated, although it is still widely debated as to what exactly causes the inversion. For instance, while some say that an expectation of falling short rates leads to the inversion, others argue that growing uncertainties about the future is what causes it. We note here that both explanations fall within the scope of Scenario II. To illustrate, falling short rates are expected when interest rates are simply deemed to be too high. This situation, which reflects the behaviour of the treasury yield curves during the early part of the 1980s decade, is clearly portrayed in Figure 4b. In contrast, the case of growing uncertainties about the future is displayed in Figure 4a, where σ , which proxies the degree of uncertainty, is large enough to force an investor to follow the hedging strategy consistent with that of Scenario II.

2C. Scenario III—Investing in a Risky Security With Negative Expected Returns

This situation is illustrated in Figure 7, where the risky security is expected to have a negative mean return. Here, also, in anticipation of falling returns, the investor is likely to follow the strategy of selling the security and longing the bond, as outlined in Section 2B. Consequently, the driver of the yield curve is, once more, similar to that in Scenario II, namely the upper curve $rt + \sigma\sqrt{t}$. What is notably different, however, is that, owing to the security's return, r , being negative, the investor is able to maximise profits by selecting a specific maturity, \tilde{T} , corresponding to a yield of \tilde{b} , as illustrated in Figure 7. Ignoring the trivial details, one could demonstrate that, with r negative, $rt + \sigma\sqrt{t}$ is maximised at some maturity \tilde{T} , where

$$\tilde{T} = \frac{\sigma^2}{4r^2} \tag{7a}$$

which coincides with an interest rate \tilde{b} equal to the negative of the security's return, i.e.

$$\tilde{b} = -r \tag{7b}$$

What this suggests, therefore, is that, in anticipation of a declining market or economy, investors will concentrate their investments around a single, well-defined maturity, \tilde{T} , which could be relatively short, depending

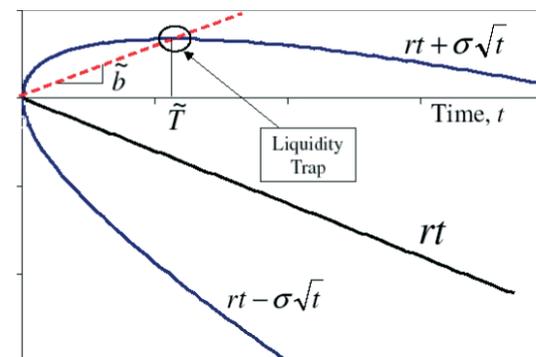


Figure 7: Investing in a risky security with negative expected returns, r , as described in Section 2C. The liquidity trap is shown to occur where the upper portion of the curve is at a maximum.

on the magnitudes of σ and r in Equation 7a—for example, the higher the certainty [i.e. lower σ] about falling markets, the shorter the maturity, \tilde{T} , of the bond investment. As a result, with a definite expectation of negative market returns [or economic growth] and, thus, a σ that approaches zero, the investment turns into a short-term saving, which, nonetheless, is consistent with the notion of the liquidity trap, where short-term savings are the norm, all in anticipation of improving conditions and, ultimately, increasing interest rates.

3. Limitations

A limitation that immediately comes to attention here concerns Scenario I. Returning to Figure 1, it is observed that a portion of the lower curve, $rT - \sigma\sqrt{T}$, becomes negative, with no positive yield identifiable at or below the $-\sigma$ VaR boundary. This region, which is defined by characteristic time scale, τ , can be shown to fall within

$$0 \leq \tau \leq \frac{\sigma^2}{r^2}$$

or

$$\frac{1}{\sqrt{\tau}} \geq \frac{r}{\sigma}$$

according to Equation 3. This implies that the maturities associated with the yield curves for Scenario I, as described in Section 2A, must satisfy the condition of $T > \tau$, thus allowing one to express the above as

$$T > \frac{\sigma^2}{r^2} \quad (8a)$$

or

$$0 \leq \frac{1}{\sqrt{T}} \leq \frac{r}{\sigma} \quad (8b)$$

If we let $\sigma^2/r^2 \sim 1$ represent a *conservative* order-of-magnitude estimate of the realistic cases detailed in Section 4A below, we conclude that what leads to the normal yield curves of Scenario I is limited strictly to debt maturities T greater than 1 year.

4. Observations

This part consists of certain observations, both empirical and theoretical, which deal with the model's validity, as well as some of its characteristics. The empirical observations, which comprise a summary comparison of the results with data, are included in Section 4A. The remaining sections, in contrast, focus on the theoretical aspects by (i) establishing a link with the CAPM, (ii) discussing how the underlying security's risk premium manifests itself here and why it could possibly be associated with the equity premium puzzle and, finally, (iii) addressing the implications of the no-arbitrage principle.

4A. Empirical

To examine the empirical validity of this model, it is helpful to plot the yield curves as b_T vs. $1/\sqrt{T}$, in accordance with Figures 3 and 6. Figures 8 and 9 are such plots, representing the US government [discount bond] yields at different dates. Each curve has a best-fit straight line passing

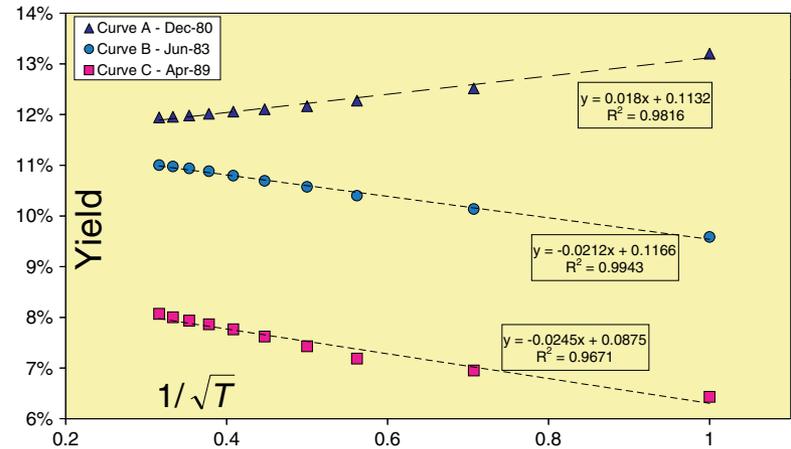


Figure 8: Samples of yield curves, illustrating time regimes of normal and inverted yield curves and their linearity when plotted against $1/\sqrt{T}$, where T is the maturity in years. These data are from McCulloch and Kwon (1993), spanning from Jan-52 to Feb-91 and extending to $T = 10$ years.

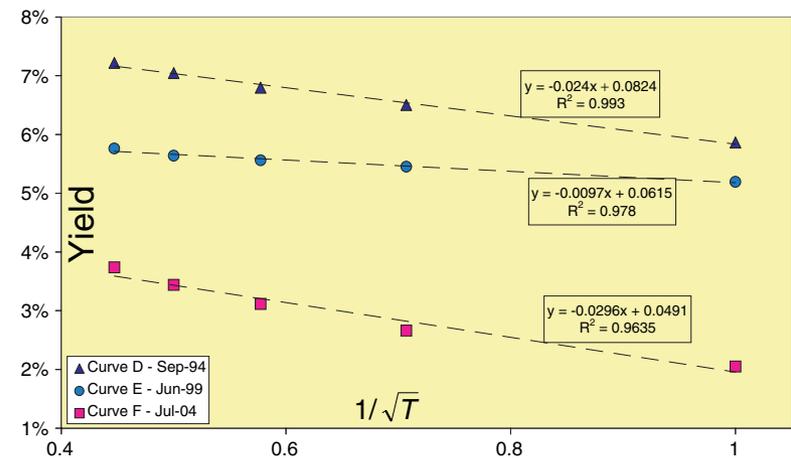


Figure 9: Samples of yield curves, illustrating time regimes of normal yield curves and their linearity when plotted against $1/\sqrt{T}$, where T is the maturity in years. These are the Fama—Bliss⁷ data, spanning from Jan-74 to Dec-04 and extending to $T = 5$ years.

through it, specifying the *implied* return, r , and volatility, σ , of the underlying security, as well as the value of the R^2 . These properties are summarised in Table 1.

For interest, we have provided a selection of different economic regimes—e.g. when interest rates were relatively high [Curves A], as in Scenario II, which led to inverted curves, as well as eras of relatively high market returns and/or low volatility [Curves B-F]⁸, as in Scenario I, which subsequently produced normal curves. It should be noted that while the above do not represent an exhaustive sample of the curves that behave according to the model, there are others that are inconsistent with the model's predictions. In addition, the limitations discussed in Section 3 above constrain the curves in Figures 8 and 9 to maturities longer than one year.

TABLE 1: CLASSIFICATION AND IMPLIED [ANNUAL] PROPERTIES [R, σ AND R²] OF THE YIELD CURVES DEPICTED IN FIGURES 8 AND 9. NOTE THAT ALL THE IMPLIED EXPECTED RETURNS ARE POSITIVE.

Curve	Time frame	Classification	Implied, expected mean annual return, r(%)	Implied expected annual volatility, σ (%)	R ²
A	Dec-80	Inverted	11.32	1.80	0.9816
B	Jun-83	Normal	11.66	2.12	0.9943
C	Apr-89	Normal	8.75	2.45	0.9671
D	Sep-94	Normal	8.24	2.40	0.9930
E	Jun-99	Flat to Normal	6.15	0.97	0.9780
F	Jul-04	Normal	4.91	2.96	0.9635

Furthermore, an order of magnitude comparison of the implied returns places the underlying security at somewhere between a typical US-based index [returning about 10% per year, as estimated from a risk-free rate of about 5% plus a risk premium of roughly 5%] and the annual, nominal GDP growth rate, which normally lies at around 4%. Lastly, none of the data consisted of a negative implied return, which could have potentially led to a liquidity trap.

4B. The Sharpe Ratio

Further to the above, we now demonstrate that the model ties in neatly with the CAPM. This link comes via the Sharpe Ratio, which appears here in a “dimensionless” form, as opposed to its classical, dimensional counterpart.

Generally speaking, the Sharpe Ratio, Sh , is expressible by

$$Sh \equiv \frac{r - b}{\sigma} \tag{9}$$

where r and σ are, as defined earlier, the expected return and standard deviation, respectively, of the underlying security and b is some risk-free rate. As to what exactly b should be has been widely debated since the inception of the CAPM. For instance, some insist that it must be the shortest available government treasury rate, while others argue that it must be the government rate, i.e. b_T , with a maturity that matches the investor’s holding period. All in all, these arguments tend to be subjective, to the point that the use of the risk-free rate in CAPM still happens to be a matter of personal preference.

It is suggested here that the root of this problem lies in how the Sharpe Ratio is defined. Acquiring a dimension of $[1/\sqrt{time}]$,⁹ as indicated by Equation 9, rather than being dimensionless, effectively means that the measurement time scale or frequency of r , σ and b —i.e. whether these parameters are measured monthly, quarterly, annually, etc.—can considerably affect the magnitude of Sh . A non-dimensional depiction of Sh should, never the less, be able to eliminate this problem.

To illustrate, we introduce a maturity-adjusted Sharpe Ratio, Sh_T , such that

$$Sh_T \equiv \frac{[r - b_T]T}{\sigma \sqrt{T}} \tag{10a}$$

$$Sh_T \equiv \frac{[r - b_T]\sqrt{T}}{\sigma} \tag{10b}$$

which, in effect, characterises the Sharpe Ratio as a function of the maturity of the investment. For instance, if r , b_T and σ were all annualised parameters and T the maturity of the investment given in number of years, then the products rT and $b_T T$, which appear in the numerator of Equation 10a, are, respectively, the cumulative or total expected return and interest paid at the end of T years and $\sigma \sqrt{T}$ [in the denominator] the expected T -year-based standard deviation of the returns. The expression in 10, therefore, extends the measurement time scale of the Sharpe ratio from its classical, annualised form given in Equation 9, which is 1 year, to T years.

Let us now return to Equation 3 and re-write it as

$$\frac{r - b_T}{\sigma} = \frac{1}{\sqrt{T}} \tag{11a}$$

or simply

$$\frac{[r - b_T]\sqrt{T}}{\sigma} = 1 \tag{11b}$$

Comparison of the above with Equation 10b leads to us to conclude that along the yield curve produced by the present model, the maturity-adjusted Sharpe Ratio remains constant, in this case being 1.¹⁰

The above, in effect, carries with it two important messages. First, any yield curve generated by the proposed model is characterisable by a constant [risk] preference, bearing in mind the relationship between investor’s preference and Sharpe Ratio. The constant preference here signifies a yield curve that is determined by moving along $rT - \sigma \sqrt{T}$, which is the lower portion of the quasi-parabola depicted in Figure 1. Second, the risk-free rate to be implemented when defining the Sharpe Ratio must be such that the maturity, T , of the borrowing should match the holding period of the investment. This makes sense because a borrowing rate with a maturity of T years is indeed risk free, owing to the fact that it remains unchanged throughout the time period T and provided that the security is also held over the same length of time.

4C. The Risk Premium of the Underlying Security

It is useful now to refer to the risk premium of the underlying security in the context of this work and see how it relates to the equity risk premium (ERP) and the associated puzzle. This is important because, as noted earlier (Cohen, 2002), even though there is little argument that the risk-free rate should be based on a government-issued security, questions still abound on what maturity it should take. We shall try to address this issue here.

Let us begin by expressing Equation 3 as

$$r - b_T = \frac{\sigma}{\sqrt{T}} \tag{12}$$

which is the yield-vs.-maturity relationship for a normal yield curve and recall that, if T were to be given in number of years, r and σ must be in



annual frequency¹¹. This, subsequently, implies that the risk premium of the underlying security with respect to the T -year risk-free rate is σ/\sqrt{T} , which means, for instance, that the risk premium relative to the 1-year risk-free rate is σ , relative to the 5-year risk-free rate is $\sigma/\sqrt{5}$ and so on.

How does all this relate to the *ERP* puzzle? This puzzle originates from the fact that the typical, empirical measure of roughly 5% for the *ERP*, based on the historical S&P index, is an order of magnitude or so higher than that which can be explained by any economic or financial theory. This value, it must be noted, is generally calculated with respect to short-term government treasuries—namely, 1 year or even less. The measured equity return, r , on the other hand, is typically averaged over dozens of years, which makes it a *long-run* parameter.

The common practice of measuring the *long-run* view of the security's return against a *short-term* interest rate could point to an inconsistency in how the *ERP* is obtained. To demonstrate, let us return to Equation 12 and note that with r representing an infinite-horizon view of the underlying security's average return, its difference relative to the 1-year risk-free rate, proxied here by the 1-year government yield, would be σ . An estimate of this based on the values of σ presented in Table 1, therefore, leads to a risk premium of about 2%, which is, by order of magnitude, consistent with the *ERP* puzzle. In light of this, as the tenor, T , of the risk-free rate increases, the security's risk premium declines with $1/\sqrt{T}$. Subsequently, the risk premium, and hence the puzzle, can be eliminated if the underlying security's return were to be compared against a risk-free rate with maturity $T \rightarrow \infty$.

Finally, how effective in practice is this approach in reducing the severity of the *ERP* puzzle? Perhaps quite effective, if we were dealing with values for σ on the order of those presented in Table 1 and recognising that the highest maturity for a US government bond is 30 years. The reason is that with a σ of about 2%, as per those in Table 1, if the risk premium were to be based on a 30-year yield, it would then be roughly $2\%/\sqrt{30} \approx 0.35\%$. Although this is still higher than what the *ERP* should be in theory¹², it does represent a significant improvement over the original 2% calculated otherwise on the basis of the 1-year risk-free rate.

4D. The No-arbitrage Principle

The no-arbitrage principle plays an important role in every security-pricing model. The literature on this subject is extensive, but for bonds and bond yields, in general, it can be divided into two areas (Fisher, 2001): (i) equilibrium and (ii) dynamic. Since the model proposed here is equilibrium, we shall avoid going into the details of the latter and, instead, refer the interested reader to the literature [e.g. Benninga & Wiener (1998) and Fisher (2001) and references therein].

Whether equilibrium or not, the no-arbitrage principle manifests itself in bond-pricing and yield-curve models through the expectations hypothesis, which has already been discussed in the Introduction section. The main point here is that the forward rates that result from the yield curve must be such that they satisfy the no-arbitrage condition at every point along the different tenors on the curve.

In reference to the present model, however, the theoretical issue that arises applies particularly to the normal curves, which are derivable in the way outlined in Section 2A. As portrayed in Figure 1 and re-iterated later in Section 3, the problem with this model is that shorter yields [i.e.

less than 1 year or so] cannot be extracted because they turn out to be negative.

Aside from the above-mentioned limitation and notwithstanding the use of the expectations hypothesis to extract the forward rates, which is a universal and well-established procedure, there is another source of arbitrage that crops up here. This relates to a flat yield curve in the manner discussed in Benninga & Wiener (1998). Briefly, the problem here is that if one were to construct a portfolio of bonds with different maturities based on a flat term structure at, say, time t , then the price of the same portfolio at a later time, say $t + \Delta t$, assuming the curve remains flat, always goes up regardless of whether the interest rate increases or decreases. This, subsequently, creates an unambiguous arbitrage opportunity, independent of which direction the rate moves.

How is this reflected in the current model? In reference to Equations 3 and 6, setting the expected volatility, σ , equal to zero leads to a flat yield curve. This makes sense because, by definition, a flat yield curve, in association with the expectations hypothesis, implies no uncertainty about the future and, thus, no anticipated rate changes going forward. But, never the less, the very fact that the model can generate a flat curve, albeit under the constraint of $\sigma = 0$, makes it vulnerable to arbitrage, provided that the curve is flat and expected to remain flat going forward.

One, however, could argue against this vulnerability in two ways. Firstly, how likely is it to see a flat yield curve in practice, let alone a flat curve that maintains its shape over time? Secondly, a flat curve under the expectations hypothesis implicitly rules out any changes in the forward rates. Therefore, the arbitrage scenario described in the preceding paragraph is self-contradictory and, thus, cannot be supported.

In view of the above, therefore, we conclude that the only area of concern with regards to not satisfying the no-arbitrage principle is that this model cannot provide shorter rates, owing to the limitation discussed in Section 3. Other than that, forward rates that satisfy no arbitrage could be easily derived for this model, given Equation 3 and the standard methods that are described in the literature [e.g. see Fisher (2001)].

5. Practical Applications

The model, as it stands now, is perhaps at best a framework in its most basic form, but, never the less, it can potentially lead to an approach that is different from, and more intuitive than, what is now available. To be more useful in practice, it will, of course, need some major calibration. One possible way to achieve this is to relax the risk-aversion constraint of moving strictly along the lower portion of the quasi-parabola—i.e. along the -1σ line for all the maturities—as one generates the normal yield curve from it. This calibration could, for example, be undertaken by introducing an error term at each tenor to account for the differences between theory and data. A comprehensive historical analysis of this error could form the basis of this calibration.

Another type of calibration could involve a historical mapping of the empirical values of r and σ , such as those in Table 1, onto, let's say, *GDP* [or some market index] growth rate and volatility, respectively. Such an analysis would then enable the user to work backwards—i.e. take the economic projections of *GDP* growth and uncertainty going forward, map them into r and σ and insert these parameters into the model to estimate

a projected yield curve. The outcome of this could possibly be used to trade futures in bonds, swap rates, etc., or for trading interest rates and swaps across currencies if forecasts of relevant economic data are available for both currencies.

A further potential application is sensitivity analysis. For instance, how is the term structure impacted as economic projections of, let's say, GDP and uncertainty are changed or revised?

And last, but not least, the question might arise on whether or not this model has any applicability to option pricing. One way, for instance, might be to substitute the risk-free rate that comes out of here directly into an option-pricing model, such as the Black-Scholes. This has the benefit of reducing the number of exogenous inputs into the equation. The answer to this question, however, is no, simply because, owing to the limitations discussed in Section 3 here, this model works for maturity time scales longer than one year, which is above the practical range of most typical vanilla options, such as calls and puts.

6. Conclusions

An intuitive model for the yield curve has been presented. Based on the notion of VaR, the approach leads to interest rates that hedge against potential losses incurred from holding a risky security until maturity.

Several simplifying assumptions are made, one of which is that the normal, symmetric diffusion process governs the stochastic behaviour of the underlying security—namely that the expected return grows linearly with time [i.e. constant mean] and a volatility that widens as the square root of time [i.e. constant variance]. Although severe, this assumption is common to virtually all practical stochastic models. Nonetheless, the restriction could be relaxed by utilising more realistic features of security returns distributions, which are asymmetric and with fat tails.

Another assumption is that the VaR is bounded by $\pm 1\sigma$. It is noted that changing these arbitrary limits of risk aversion would in no way alter the overall qualitative outcome of the model. The model would still produce both normal and inverted yield curves that progress linearly with $1/\sqrt{T}$, as suggested by Equations 3 and 6, respectively, the only difference being slopes that are different from $\pm 1\sigma$. This, of course, is quite limiting and could perhaps, along with the previously mentioned assumptions, explain the failure of the model in some instances.

In view of the above assumptions, several conclusions have been reached. One is that the model theoretically leads to three types of shapes for the yield curve, which, in all cases, relate directly to $1/\sqrt{T}$. The normal curve is predicted to occur under normal conditions when interest rates are lower than the expected return of the underlying security, which is generally higher owing to the security's risk premium. The inverted curve is found to show up in periods of high interest rates, as explained in Section 2B, or when uncertainties, as proxied by the standard deviation of the risky security, are relatively large. Finally, a flat yield curve is predicted during times of low expected uncertainty or as $\sigma \rightarrow 0$.

There are, in addition, other findings. One is that the model fits in tightly with the CAPM, linking the interest rates and their respective maturities to the risk and return of the underlying security via the Sharpe Ratio, Sh . Basically, after adjusting Sh for maturity, the yield curve would be describable by a constant Sh . In view of the relationship between Sh

and investor preference in classical CAPM, one could, subsequently, argue that a yield curve that follows this model belongs to a risk preference that is characterised by VaR limited to -1σ ; i.e. moving along the lower portion of the quasi-parabola shown in Figure 1. The second finding simply identifies the risk-free rate for use in the Sharpe Ratio and defines it as the yield to maturity, b_T . The logic here is that with the interest rate, b_T , remaining constant and, essentially, free of any rollover-related risks, it should then truly represent the risk-free rate for the borrowing, which has a maturity that matches the holding period of the security.

A further finding involves the well-known ERP puzzle. The conclusion here is that comparing an infinite-horizon view of the underlying security's [e.g. an equity index] return against a short-term risk-free rate may be a potential cause. The puzzle could, in theory, be explained and/or eliminated if one were to compute the risk premium based on a risk-free rate that has an infinite maturity. However, as this is not practical, one could, at least, reduce the severity of it by focusing more on the longer maturity government rates [e.g. 30 years] rather than the shorter ones [e.g. 1 year or less].

As for practical applications, it is suggested that the model undergo some calibration in order to be useful. This calibration could take the form of an error analysis, as well as a mapping of the model's outputs of the implied r and σ onto measured economic or market variables, such as GDP or market growth rates and volatility.

We finally conclude here by reiterating that this is only a model, which, like any other, has its advantages and disadvantages. One disadvantage is that it is not dynamic, although it has the potential for serving as a platform for a more sophisticated dynamic or stochastic setting. Moreover, although several examples of its validity have been provided that illustrate the predicted linearity in $1/\sqrt{T}$ [e.g. in Figures 8 and 9], other situations [although none has been provided here in the interest of space] can be shown not to follow suit. Among the advantages, notwithstanding, are that (i) it is intuitive, (ii) it explains the typical variations in yield curves—i.e. normal, inverted and flat, (iii) it predicts the existence of the liquidity trap and, lastly, (iv) it links the yield curve with CAPM—all these achieved within a single framework.

FOOTNOTES & REFERENCES

1. Terminology such as "high uncertainty", "high interest rates" and "low growth" will be defined in due course.
2. This scenario, which leads to the "normal" yield curve, is more common than the one discussed later in Scenarios II and III. The reason for the margin between the security's expected mean return and the risk-free rate is the security's risk premium, which is addressed in more detail in Section 4C.
3. Obviously, we are assuming here that the risky security follows the simple diffusion process. Although this is a rough assumption, it happens to be the most common and forms the foundation of almost every pricing methodology that revolves around stochastic behaviour.
4. The shape becomes a parabola as r approaches zero.
5. The inherent argument supporting the equality is that the lender demands the maximum he can get and the borrower [investor in the underlying security] is happy to provide

it as long as it fits within his preference criteria. Admittedly, this is debatable, but we shall proceed, leaving it outside the scope of this paper.

6. The negative expected security return is, of course, characteristic of a declining economic performance.

7. The Fama-Bliss data come from Section 4.3 of *CRSP Monthly US Treasury Database Guide* [Version CA90.200502.1], Center for Research in Security Prices, University of Chicago Graduate School of Business.

8. Curve D suggests a low implied volatility compared to the others and can, therefore, be considered nearly flat.

9. This is because r and b both have dimension [1/time], while σ acquires dimension $[1/\sqrt{\text{time}}]$.

10. The number of standard deviations determines the value of the constant that one chooses to limit the VaR .

11. Likewise, if T were to be given in number of quarters, for instance, then r and σ must be quarterly in frequency.

12. This gets worse for an equity index, such as the S&P, whose annual volatility is around 10% or so.

■ Ang, A., Piazzesi, M. and Wei, M. (2006) "What does the Yield Curve tell us about GDP growth?" *Journal of Econometrics* 131, pp. 359–403.

■ Benninga, S. and Z. Wiener (1998) "Term Structure of Interest Rates" *Mathematica in Education and Research* 7, pp. 13–21.

■ Brace, A., D. Gatarek and M. Musiela (1997) "The Market Model of Interest Rate Dynamics," *Mathematical Finance* 7, pp. 127–147.

■ Cohen, R.D. (Nov. 2002) "The Relationship Between the Equity Risk Premium, Duration and Dividend Yield," *Wilmott Magazine*, pp. 84–97.

■ Diebold, F.X., Piazzesi, M., Rudebusch, G. (2005) "Modeling Bond Yields in Finance and Macroeconomics" *American Economic Review P&P*, pp. 415–420.

■ Esterella, A. (2005) "The Yield Curve as a Leading Indicator: Frequently Asked Questions" Federal Reserve Bank of New York Web Article, pp. 1–23.

■ Evans, C.L. and D. Marshall (2001) "Economic Determinants of the Nominal Treasury Yield Curve," Unpublished paper, Federal Reserve Bank of Chicago.

■ Fabozzi, F.J., *Bond Markets, Analysis and Strategies*, 3rd Ed., Prentice-Hall Int., 1996.

■ Fabozzi, F.J., *Investment Management*, 2nd Ed., Prentice-Hall Inc., 1999.

■ Fisher, M. (Q1 2001) "Forces That Shape the Yield Curve," *Economic Review*, Federal Reserve Bank of Atlanta, pp. 1–15.

■ Kessel, R.A. (1965) "The Cyclical Behaviour of the Term Structure of Interest Rates," *National Bureau of Economic Research Occasional Paper No. 91*.

■ McCulloch, J.H. and Kwon, H.C. (1993) "U.S. Term Structure Data, 1947–1991" Ohio State University Working Paper # 93-6 [see also <http://www.econ.ohio-state.edu/jhm/ts/mccckwon/mccull.htm>].

■ Mitchell, W.C. (1913) *Business Cycles*, Berkeley, University of California Press.

■ Siegel, A.H. and C.R. Nelson (1988) "Long-term Behavior of Yield Curves," *Journal of Financial and Quantitative Analysis* 23, pp. 105–110.