Breaking the Yield Spread

Erik Swanson

Breaking the Yield Spread

Abstract

The yield spread has long been considered a reliable indicator for predicting future recessions. This paper uses data on the ten-year Treasury Bond and three-month T-Bill rates from the second quarter of 1953 until the third quarter of 2007 to update the a probit model to proposed by Estrella and Mishkin (1995). I find that the relationship between the spread and recessions appears to be potentially unstable in recent years and use estimates of cointegrating vectors to assess whether the information content of the yield spread has changed. Analysis shows that the ten-year bond rate no longer tracks the three-month rate as it has in the past, implying that the yield spread may no longer carry useful information about the future state of the macro economy.

Introduction

Among economists, there is a strong interest in looking at the ability of financial variables to predict real economic variables, such as future economic growth (Estrella and Mishkin 1995). Typically, growth forecasters use a large number of economic variables in making their forecasts of future activity, but, if financial variables carry information on the future of the economy, they could also make an important contribution to growth forecasts (Estrella and Mishkin 1995). The reason for this is the contemporaneous availability of data on financial variables; when forecasters try to make predictions based on variables such as previous quarter's growth rate or monetary base, they must wait for these variables to become available from their sources. Financial variables, however, have levels determined by market activity and thus are instantly accessible by anyone.

A number of financial variables could conceivably have predictive power, but one in particular, the yield spread, has been shown in previous research to be a good indicator of future economic activity. The yield spread is the difference between two different interest rates – usually the rate of a long term bond minus the rate on a short term bond. The inclusion of the yield spread in the early experiments in prediction was largely an accident of availability – a number of authors simply looked at a variety of easily available financial variables in an attempt to get a good fit for predicting future growth.

The first section of this paper will review previous literature in the field. Section two gives one possible motivation for the relationship between the yield spread and future growth using the Taylor rule and the Pure Expectations Hypothesis. Since the model presented here does not yield an estimable equation, section three deals with using the yield spread to predict the occurrence of recessions using a probit model. My results are consistent with those in earlier work and the updated data clearly predict the recession which occurred in 2001. There is, however, a recent spike in probability which may or may not indicate a recession on the immediate horizon. Noting that there may have been a change in the information content of the yield spread, section four explores the evolution of the relationship between the ten-year and three-month bond rates which comprise my yield spread.

Literature Review

Papers on the use of the yield spread as a predictor of economic growth began appearing in the late 1980s. Stock and Watson (1989) included two yield spreads in their proposed Index of Leading Indicators – the spread between the return on 6- month commercial paper and the 6-month Treasury Bill, and the spread between the return on the 10-year Treasury and the 1-year Treasury. Stock and Watson indicate that both of these variables are among their most effective predictors of future growth. They leave these spreads out, however, when they compose an index meant to predict the likelihood of a recession in future quarters.

Stock and Watson also bring up an important flaw in their index: the potential for overfitting. If too many explanatory variables are included in their regression, then, while the fit may be very strong for in-sample data, it is likely that index will not make strong predictions when used on out-of-sample data. This very phenomenon drove the eventual discontinuation of the Stock and Watson index of leading indicators. While the index was updated several times to better fit past data, it was never successful at making strong out-of-sample predictions. A model which uses fewer parameters than the Stock and Watson index would most likely have more robust estimates for those parameters and thus provide better forecasts of future growth.

While Stock and Watson (1989) focused on predictions on a short time horizon – at most two quarters – Estrella and Hardouvelis (1991) explored the usefulness of the term structure in predicting performance farther out. Using as their measure of spread the difference between the 10-year treasury rate and the 3-month T-Bill rate, they find that the spread has a strong predictive power for real growth for horizons up to four years.

The results of Estrella and Hardouvelis (1991) were confirmed by Plosser and Rouwenhorst (1994), who also found that the slope of the term structure predicts real output growth in the US. They add to the literature by suggesting that the term spread can be decomposed into several elements – the slope of the long-term term structure, the slope of the short-term term structure (both using forward rates) and the current short term rate. To perform this decomposition, Plosser and Rouwenhorst compute implied forward rates based on different points along the yield curve – by taking a weighted difference of the current k-period bond rate and the current k-1 period bond rate, they find the implied forward rate on a k-1 period bond in the future. Using the same method, they find the implied forward rate on a one-period bond. These two forward rates allow them to break the term structure down into three parts. The first is the difference between the implied forward rate on the k-1 period bond and the implied forward rate on the one-period bond, which should carry information about the future state of the term-structure. The second is the difference between the implied one-period forward rate and the current one-period rate, which, under the expectations hypothesis, is an unbiased predictor of future change in the one-period rate. The third is the one-period rate itself.

Using this method, Plosser and Rouwenhorst conclude that the most information on future growth is held in the long end of the term structure (the first element of the decomposition), and that most of the results of full term structure predictions can be seen in predictions based on the current short term rate and the long term spread. By examining whether proxy variables for past and future monetary policy push out the predictive power of the spread (i.e., including them in regressions with the term spread decompositions forces the parameter estimates on the term spread terms to become statistically insignificant), Plosser and Rouwenhorst conclude that the spread contains information other than about expectations for future monetary policy, and thus cannot be discounted as redundant predictors.

In 1995, Estrella and Mishkin changed the discussion. Instead of attempting to actually forecast future real growth, they instead tried to predict the occurrence of recession, a binary dependent variable. While other studies has performed linear regressions, Estrella and Mishkin used a probit model to see if a number of financial variables were useful in predicting whether or not the economy would be in recession. Of the variables they examined, the ones which performed best were the yield spread (defined as the average quarterly difference between the rate on the 10 year treasury and the 3 month t-bill) and the quarterly average change in the NYSE. Between these two, the performance of the NYSE was particularly strong over short time intervals of one and two quarters ahead, while the yield spread was the most accurate indicator between three and six quarters ahead.

In addition to looking at a probit model to predict recession as opposed to a model which predicts level of future growth, Estrella and Mishkin also looked at the out-of-sample performance of their indicators. Of all the variables examined (which included both Stock and Watson's index and the Conference Boards Index of Leading Indicators), they found that the most accurate predictor out-of-sample was a composite of the NYSE and yield spread, with this composite deriving much of its short term power from the NYSE and its longer term power from the yield spread. Estrella and Mishkin show that, among financial variables and existing indexes of leading indicators, the simple yield spread is the best predictor of a recession four to six quarters in the future.

Estrella and Hardouvelis (1991) had shown that the relationship between real growth and the yield spread was not necessarily policy invariant – changing monetary policy regimes could cause the relationship to change. They also point out that the relationship may not be stable over time. In order to address these issues, Estrella, Rodrigues, and Schich (2000) explored the issue of stability of the relationship between the slope of the yield curve and real output. They look at both linear models with endogenous growth in industrial production and probit models attempting to predict a recession, both with yield spread as the explanatory variable. While there is significant evidence of instability in predicting change in industrial growth, in the case of the probit models there is no evidence of instability in predictive power for any time horizon or with any maturity combination defining the yield spread in either the United States or Germany.

Feroli (2004) continued the discussion of policy and the effects of the yield spread, modeling the effects of monetary policy changes on the predictive power of the yield spread. He finds that the yield spread is related to the output gap, but that the relationship is a function both of how strongly the monetary authority targets the output gap and of how strongly the monetary authority insists on smoothing interest rate movements. He also notes a break in the predictive power of the yield curve, contrary to the work of Estrella, Rodrigues, and Schich (2000), with the start of the Volcker monetary policy regime in 1979.

Modeling the Influence of the Yield Spread

The relationship between the yield spread and the output can be modeled in a number of ways. The following model is originally presented in Feroli (2004). Begin by looking at the Pure Expectations Hypothesis and the Taylor Rule:

$$i_t^N = \frac{1}{N} E_t\left(i_t^1 + i_{t+1}^1 + i_{t+2}^1 + \cdots + i_{t+N-1}^1\right) \qquad (1)$$

$$i_t^1 = \rho i_{t-1}^1 + (1-\rho)\left[i^* + \beta\left(E_t \pi_{t,k} - \pi^*\right) + \gamma E_t x_{t,q}\right] \qquad (2)$$

The Pure Expectations Hypothesis (1) is a long-standing theory in economics to describe the shape of the term structure of interest rates. It says that the yield on a long term bond is merely an average of expected yields on short term bonds over the life of the long bond. Thus, investors are indifferent between holding, say, a ten-year bond and holding ten one-year bonds in succession, rolling the interest over. This most basic form of the expectations hypothesis does not allow for a term premium – that is, a consistently higher rate paid on a long term bond to compensate investors for what they perceive to be higher risk associated with the increased holding period.

The Taylor rule is a monetary policy rule which describes how the monetary authority makes its policy decisions. There are two major concerns for the monetary authority, inflation and the output. Thus, the Taylor rule gives weights to how closely the monetary authority wishes to target each of these. Based on the chosen targets and the weights, the Taylor rule indicates where the interest rate should be set. This version of the rule also includes an interest smoothing term, which represents the emphasis the monetary authority places on keeping interest rates stable.

In this model ρ represents the important the monetary authority places on keeping interest rates stable. The term i^* represents the desired interest rate of the monetary authority when the output gap and inflation are at their desired levels. The terms k and q represent how far into the future we are looking. The output gap is given by x and inflation by π. The parameters β and γ represent the weights the monetary authority puts on deviation from targeted inflation and output.

To simplify, we let k and q equal 0 (which is equivalent to assuming that the monetary authority is primarily concerned with were inflation and output are *now* as opposed to in the future) and iterate equation 2 one period forward, then we look at the expectations hypothesis for a two period bond minus a one period bond and substitute to solve for the output gap:

$$E_t x_{t+1} = \frac{2}{(1-\rho)\gamma} S_t^{2,1} - \frac{1}{\gamma}\left[i^* + \beta\left(E_t \pi_{t+1} - \pi^*\right) - i_t^1\right] \qquad (3)$$

If we now make the assumption that the expectations in the market are rational – that is, take the expected value at time t to be equal to be an unbiased predictor of the true value in the future - and allow for error terms in both expected inflation and the expected output gap, we get:

$$x_{t+1} = \frac{2}{(1-\rho)\gamma} S_t^{2,1} - \frac{1}{\gamma}\left[i^* + \beta\left(\pi_{t+1} - \pi^*\right) - i_t^1\right] + \varepsilon_t^y + \frac{\varepsilon_t^\pi}{\gamma} \qquad (4)$$

This equation leaves us with the result that the output gap one period in the future depends on the current level of the yield spread and the inflation term. The yield spread term is what we focus on here. The coefficient on this term depends on both how interested the central bank is in smoothing interest rates and how strongly they are targeting the output gap. If the bank focuses more strongly on targeting the output gap, the influence of the spread (actually of both terms) will be diminished. This is a logical result because, if the bank cares more about the spread, it should do a better job of controlling it. There is also a negative relationship with the interest smoothing parameter – if the bank decides that it wants to smooth the path of interest rate movements more carefully, then the spread will have a greater influence on the output gap.

The second term here represents a problem for forecasting. Because the future output gap depends on inflation in the same period, it is not possible to make predictions for future output using this entire equation. Also, as Feroli (2004) notes, the right hand side contains endogenous variables and thus cannot be determined by OLS. Thus, this model is useful in that it motivates a relationship between future output and the yield spread.

Predicting Recessions

Ultimately, the goal of looking at financial indicators of macroeconomic performance is to produce a metric by which we can quickly judge, by looking at consistently and immediately available information, the future performance of the macro economy. While it is possible to capture expected future inflation by looking at the rate of return on Treasury Inflation Protected Securities, previous econometric analysis, such as Estrella and Mishkin (1995), has shown that strong results can be seen looking only at the yield spread and ignoring the omitted variable bias.

These results become particularly strong when, instead of focusing on predicting real output, we instead attempt to predict recessions. It makes intuitive sense that this method is more successful – the simple yield spread is a relatively blunt instrument for predicting a variable as nuanced as changes in real output. Predicting a binary variable like recessions (as defined by the NBER) allows for a stronger fit. This procedure follows the work of Estrella and Mishkin (1995). I will attempt to estimate the following equation:

$$\Pr(RECESSION_{t+k} = 1) = \Phi(\alpha + \beta SPREAD_t) \qquad (5)$$

where Φ is the cumulative distribution function of the normal distribution, t is the time of the observation, and k is the number of quarters lagged. This specification uses a constant and *SPREAD* in quarter t, weighted by β, to predict the probability of a recession in quarter t+k.

For the term *SPREAD*, we use the difference between the rate on the 10-year treasury bond and the 3-month T-Bill. The data run from 1953, quarter 2 until 2007, quarter 3. While this data is available in monthly form, we use quarterly data instead to be consistent with Estrella and Mishkin (1995). Also, quarterly data has less noise and thus will produce smoother results in estimation. For comparison, results of regressions on monthly data are presented in Table 2.

Here we follow Estrella and Mishkin (1995) in looking at regressions for time horizons from one to six quarters. We only use the data until 2006, quarter 1, since this is the latest observation for which we can get a six quarter lag of the recession indicator. In order that the results be consistent across the regressions, Table 1 shows regressions using only this data.

The measure Pseudo R-Squared is meant to be analogous to the measure R-Squared in Ordinary Least Squares Regression. It is calculated by taking the difference between the absolute value of the log likelihood in the model predicted with only a constant term and the absolute value of the log likelihood of the full model, weighted by the absolute value of the log likelihood of the constant-only model. Thus is represents the percentage improvement in the predictive power of the full model over the predictive power of the model using a constant only. To compare with the linear form of R-Squared, which measures the percentage of variation in the dependent variable which is explained by the independent variable, Pseudo R-Squared also represents how much of the variation in the binary dependent variable is captured by the independent variables. The difference is, a linear version of R-Squared cannot be computed in a probit model, since the errors do not have a similar interpretation. Thus, Pseudo R-Squared makes use of the available information – the likelihoods computed in fitting the model, to explain how well the variation is captured. This measure of Pseudo R-Squared is different from the one used by Estrella and Mishkin (1995), which used a different measure proposed by Estrella (1995). Their measure takes the form:

$$pseudo-R^2 = 1-\left(\frac{\log L_U}{\log L_C}\right)^{-\frac{2}{n}\log L_C} \quad (5)$$

which they say better mirrors the linear version of R-Squared. Here L_U and L_C are the likelihoods calculated for the unrestricted and constant only models respectively, and *n* is the number of observations. In this paper, however, we will use the measure discussed above for the sake of simplicity.

These results mirror Estrella and Mishkin (1995) well, especially considering the much wider range of our data. The coefficients on the spread term in all the regressions are significant at the 1% level, so all of these models predict better than a constant term alone. The strongest predictions are made at the three and four quarter horizons, but the best is at three quarters. This differs from Estrella and Mishkin (1995), who found the strongest relationship at the 4 quarter horizon when looking at the spread alone. Figure 1 shows the recessions as defined by the NBER overlaid with the predicted probability of recession in each quarter as determined by the model.

Figure 1:

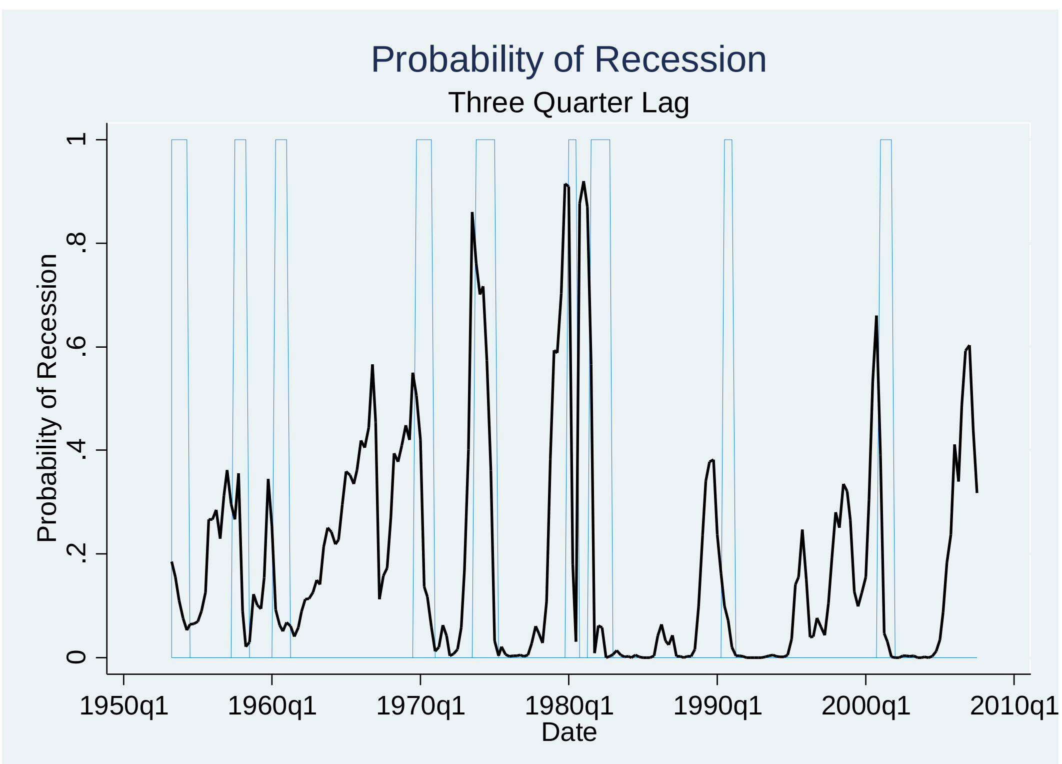

It is clear looking at this graph that every recession since 1970 has been accurately predicted by the yield spread. These results are very strong in-sample, but, perhaps even more importantly, they are preserved out-of-sample. For instance, we can look at the same graph, but based on a regression which uses the data available only through 1994, quarter 4. This graph can be seen in Figure 2.

Figure 2:

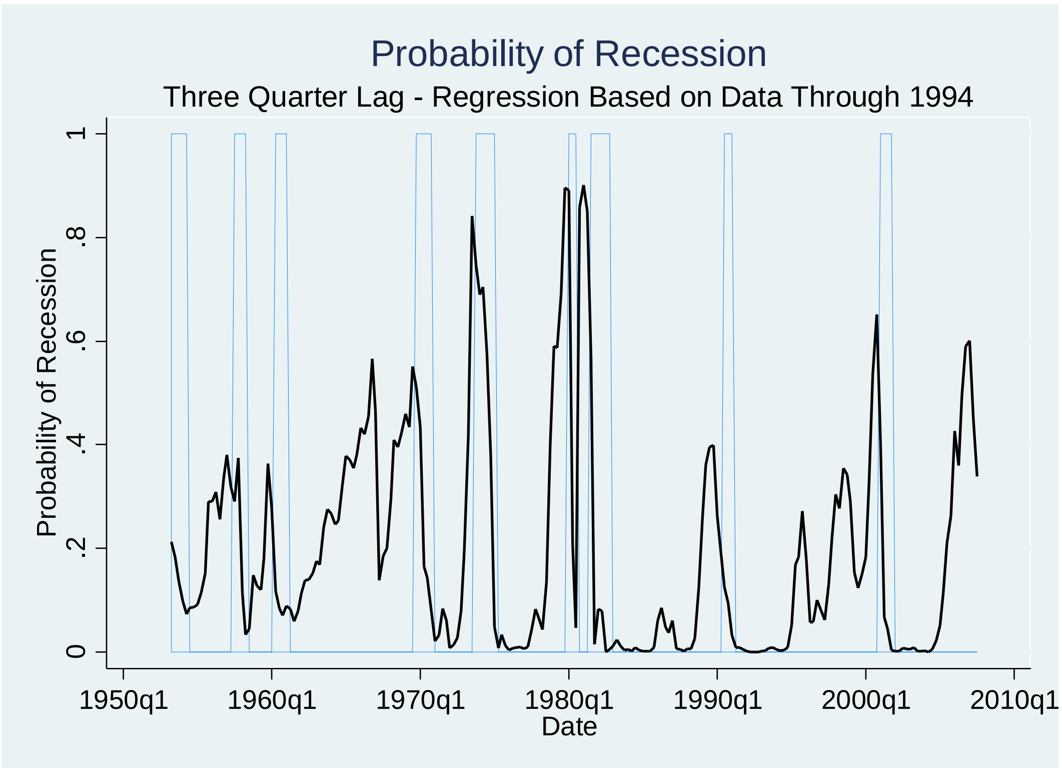

The closeness of this fit is remarkable – the second graph predicts the 2001 recession with nearly the same accuracy as the first, despite being 7 years out-of-sample. The peak probability in both cases occurs in the 4th quarter of 2000 (the quarter before the recession started), with the in-sample model giving a probability of a recession 3 quarters ahead at 0.6597 and the out-of-sample model giving a probability of 0.6520. The ability of the spread to predict out-of-sample consistently over time is explained in more detail in Estrella and Mishkin (1995).

It is impossible to look at these graphs and not notice the sharp spike on the right side. Both graphs peak at a 2007, quarter 1 with a probability of recession of over 60%. This means that the model is strongly predicting a recession starting the 2007, quarter 4. There isn't another example of such a high predicted probability that did not actually lead to a recession since 1966. The two small false-positives in the 1990s were not of nearly the magnitude of the current prediction. Currently, other economic signals are mixed. The Federal Reserve recently cut interest rates, but has not signaled an intention to do so again in the future. Oil prices are soaring and hurting markets. Yet the preliminary estimates for real 3rd quarter growth in 2007 are a robust 3.9% - which would certainly seem to imply that a recession will not start with the present quarter. If we assume that a recession does occur in the next few quarters, then the model is still functioning well – at this point it is safe to say that a predicted probability of recession over 50% is a *very* strong indicator that one will happen.

The more interesting case is to assume that a recession does not occur in the near future – if this is the case, what has changed? One possible reason for the ability of the spread to predict recessions to break down would be if the long rate were no longer reflecting the same information as it has historically. Looking at a plot of the rate on the 10-year treasury against the rate on the 3-month T-Bill, we see that the long rate has not been tracking the short rate recently nearly as closely as it has historically. This plot can be seen in Figure 3.

Figure 3:

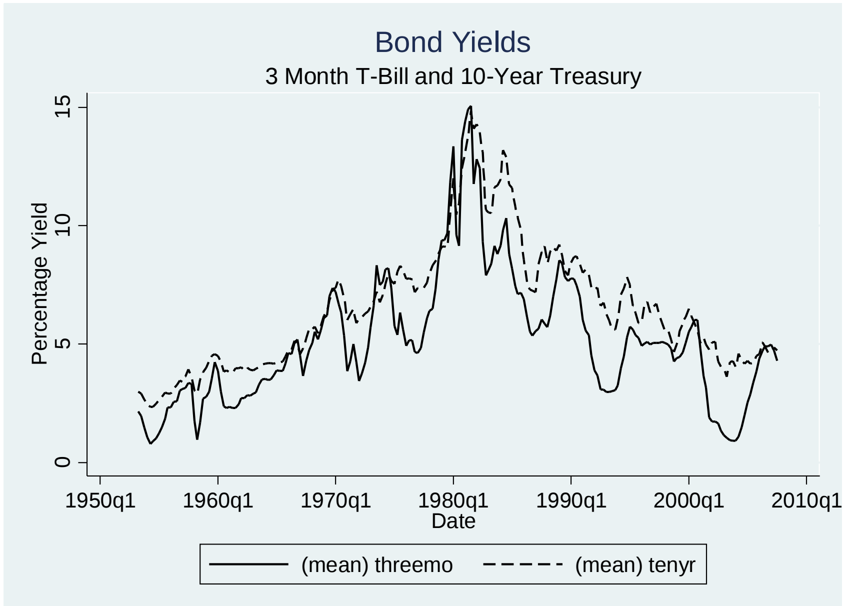

We can see from the graph that, during the most recent period of tightening, the 10-year rate did not track the short rate up. The 10-year rate appears to be to be varying less in general since about 2000. If the relationship between the 3-month rate and the 10-year rate has changed, then the fundamental information carried in the spread has changed.

Testing for Cointegration

Based on the economic model, which starts by assuming the Pure Expectations Hypothesis, the movements of the ten-year rate and the three-month rate should be related over time. One implication of the expectations hypothesis is that the expected yield from holding bonds of any two durations must be the same over the period of the shorter bond. Thus, it can be shown that a high yield spread implies rising long-term rates over the life of the short-term bond. Long-term rates must rise because a rise in interest rates causes the prices of lower-yielding bonds to fall, thus inflicting the holder of the long term bond with a capital loss and equalizing his return to that of the holder of the short-term bond. Short-term rates must rise over the life of the long-term bond because the rate on the long-term bond is merely an average of expected short-term rates over its lifetime. Whether or not these implications of the expectations hypothesis are borne out empirically is a question for a different paper; it suffices here to say that, by considering the expectations hypothesis as a basis for the relationship between the yield spread and real growth, the expectations hypothesis also clearly motivates a relationship between the movements of interest rates at different maturities.

One way to examine the relationship between the ten-year and three-month rate and how it may have changed is to look at whether there is a cointegrating relationship between the two rates. If the cointegrating relationship has changed, then the information content of the yield spread has changed and thus the model to predict recessions may not function as well as it has in the past. When looking at cointegration, I will examine trends in monthly data as opposed to quarterly data. Monthly data provides many more observation, which will make the results of my estimation more statistically robust.

Two non-stationary time series variables are defined as cointegrated if there exists a cointegrating vector $(1, -\beta)$ which forms a linear combination of the variable which is stationary. A time series variable is non-stationary if it is autocorrelated – if its movements now can be explained by its movements in previous time periods – and if the coefficients on its autoregressive terms sum to greater than one. A non-stationary variable is characterized by the fact that the effect of shocks do not diminish over time, so they do not revert to a mean.

It is easy to show that both the three-month rate and the ten-year rate are non-stationary by looking at the output of the augmented Dickey-Fuller (ADF) test. The null hypothesis of the ADF test is that the series has a unit-root and is thus non-stationary. It follows its own distribution, the critical values for which are reported in Table 3. Cointegration also requires that the two series be integrated of the same rank. A series is I(n) – integrated of order n - if it is non-stationary but becomes stationary when taking *n* differences. The ADF results for first differences of both the ten-year and three-month rates are reported in Table 3. Based on the ADF statistics, the ten-year and three-month rates are both I(1) series. They are non-stationary on their own, but stationary in first differences. Thus I can consider whether or not they are cointegrated without further modification.

There are two common methodologies for testing for cointegration. The first is Engle-Granger test. This test seeks to estimate the cointegrating relationship based on a linear regression. To carry out the test, one must first establish that the two series are integrated of the same order (as shown above for the ten-year and three-month rates), then run an ordinary least-squares regression of one series on the other. If the residuals of this regression are stationary under and the Engle-Granger ADF test (which is the same as a standard ADF test, but with more extreme critical values), then the series are said to be cointegrated. The estimate of the coefficient β in the linear regression can be considered as part of the cointegrating vector $(1, -\beta)$, but it is not a very reliable estimate.

The residuals from the regression of the ten-year rate on the three-month rate have an ADF statistic of -3.605, which exceeds the 5% critical value of the Engle-Granger ADF of -3.41 in absolute value. Thus I conclude from the Engle-Granger test that the series are cointegrated in the full samples. The Engle-Granger test, however, is not always considered to be reliable. It relies upon repeated uses of the ADF test, along with the special critical values for the Engle-Granger ADF test, which are not robust to minor differences in specification. If one fails to correctly identify the form of the time series (whether or not there is a time trend, or a constant, or a number of other possible specifications), then the critical values are not reliable for making statistical inference. Thus I will use a different method to further these results.

This more reliable method to estimate cointegrating relationships is the Johansen method. It is based on the estimation of a Vector Error Correction Model (VECM). If we consider two series X and Y to be cointegrated, then the VECM will take the following form:

$$\Delta Y_t = \beta_{10} + \beta_{11}\Delta Y_{t-1} + \cdots + \beta_{1p}\Delta Y_{t-p} + \gamma_{11}\Delta X_{t-1} + \cdots + \gamma_{1p}\Delta X_{t-p} + \alpha_1 (Y_{t-1} - \theta X_{t-1}) + \varepsilon_{1t} \qquad (6)$$

This is the same as a standard vector autoregression (VAR) model, except for the term $Y_{t-1} - \theta X_{t-1}$ at the end, which represents the cointegrating relationship between series X and Y. To fit this model, I first need to establish a value for p, or the maximum number of lags to be included in the VECM. Because the data here is monthly, I considered using up to 24 lags to get two full year cycles. Based on a series of criteria which recommend optimum lag lengths, I am choosing to use three lags here. This choice is based on the minimal Schwartz's Bayesian information criterion. Other information criterion recommended longer lag lengths, but it would be difficult to accommodate these longer lags here because of the small number of observations, particularly in subsamples of the data.

Using three lags, I first calculate the Johansen trace statistic. This test considers whether the number of cointegrating relationship is *r* or if it is more than *r*. It is impossible to have more cointegrating relationships that the number of series being considered minus one, so here it is only possible that one cointegrating relationship exists. The test is based on the estimated characteristic roots (or eigenvalues) of the estimated matrix of coefficients in the VAR(p) model. The results for this test are reported in Table 4. Because the trace statistic for rank 1 is less than the 5% critical value, this test implies the existence of a single cointegrating relationship between the ten-year and three-month rates.

To estimate the actual value of this cointegrating relationship, I use the VECM model above, where Y is the ten-year rate and X is the three-month rate, also allowing for a constant to exist in the cointegrating vector. The resulting estimation of the cointegrating vector is seen in Table 5. Using this cointegrating vector to create a new series and running a probit model of monthly recessions on it actually yields slightly better results than a model using the simple yield spread.

The presence of a constant in this vector poses a problem for my assumption of the Pure Expectations Hypothesis. It means that the ten-year rate should be higher in the long run than the three-month rate, which implies a "term premium" on the longer bond. The term premium is the extra interest which a bond must pay in order to make investors assume what they perceive as additional risk in holding long term bonds. The expectations hypothesis, however, implies that there should be no term premium – investors are indifferent between holding one long bond or a series of short bonds over their preferred maturity horizon. Forms of the expectations hypothesis with a time-invariant term premium do exist, and would allow for a motivation of the influence of the spread similar to the one considered in the economic model above. Thus, I conclude that, while this result provides evidence contradictory to one of the assumptions of my economic model, it does not invalidate the result of that model in motivating the relationship between the yield spread and future real growth.

Taking the cointegrating vector in Table 5 to define the true relationship between the ten-year and three-month rates across the whole time period, it is possible to see if the estimates of cointegrating relationships in smaller subsamples are consistent with this whole. Specifically, I am interested in seeing if the cointegrating relationship has been consistent with its long run pattern in the recent past. First, however, I will consider other possible locations of a break.

The most obvious place to test for a split is at the point of the monetary policy regime change that came in the third quarter of 1979. The results for the cointegrating vectors in these timeframes can be seen in Table 5. Based on a two-sample t-test, there is not a significant difference between the estimates on the coefficient of the three-month rate in these periods (the value of the constant is not particularly important). Thus, I cannot conclude that there is a significant difference between the cointegrating relationships in these two periods.

Another possible break is at the start of the monetary policy regime of Alan Greenspan. Estimates for these periods are also reported in Table 5. Again, a two sample t-test cannot reject the null that the cointegrating relationships are the same in the two-periods. However, an F-test shows that the variances of the parameter estimates in the two subsamples are significantly different at the 1% significance level. While some of this is due to the smaller sample size (242 vs. 409) in the post-Greenspan period, it nevertheless appears that the relationship between the rates has become significantly more variable.

It may be particularly illustrative to look at an even more recent period, starting say in the beginning of 1993, since the relationship between the yield spread and recessions has been more erratic during this period. This matches up to the start of a new political administration and is also early enough that the minor false positives indicated by the model in the mid-1990s are captured in the sample period. Estimates of the cointegrating vectors in this period are shown in Table 5. A two-sample t-test rejects the null hypothesis that the estimates of the cointegrating relationship are the same at the 1% significance level. Thus I conclude that there has been a change in the cointegrating relationship between the two rates since at least 1993. Based on the results of this test, I cannot pinpoint the time of the change, but can only note that there is a difference before and after January 1993.

Finally, I will consider the period since the end of the most recent recession, that is, since December 2001, as the ten-year rate seems to be tracking the three-month rate particularly poorly during this period. The results are reported in Table 5. No significant cointegrating relationship exists on this time horizon – a t-test shows that the estimate for the relationship is not significantly different from zero. This result is corroborated by the Johansen trace statistic, which finds that there exist no cointegrating vectors. Thus, I conclude that in the recent past the relationship between the ten-year rate and the three-month rate has deteriorated almost completely. This result has significant consequences for the reliability of predicting recessions based on the yield spread as it indicates the information content of the spread is significantly different than it has been in the past.

Conclusion

Historically, the yield spread has been a very good predictor of recessions, especially at the three and four quarter horizons. This fact is extremely useful in that it provides an easy and immediately available way to observe predicted downturns in future growth. Using a data set which is longer in both directions than that used by Estrella and Mishkin (1995), I find results consistent with theirs – good Pseudo R-Squared statistics at a range of time horizons.

The relationship, however, may not be as stable as Estrella, Rodrigues, and Schich (2000) had previously found. By analyzing estimates of cointegrating vectors produced by a vector error correction model, I find that, at least since 1993, there has been a statistically significant change in the relationship between the ten-year and three-month bond rates. When there is a stable relationship by which the rates move together, deviations from this relationship provide insight into the state of the macro economy. However, if, as my results imply, the relationship between the two rates is no longer stable, then deviations from their normal joint evolution will no longer be as telling as they have been in the past.

The recession model predicts a downturn starting in the final quarter of 2007. While that quarter is coming to a close, and apparently without evidence that it will have been the start of a recession, the economy is certainly on the brink of potential decline. Between the falling dollar and the fallout resulting from the rapidly declining housing sector, it is plausible that a recession may loom in the near future. If this is the case, then it appears the model is still predicting well. However, if this recession does not materialize, the failure of the model may be attributed to the changing relationship between the two interest rates. Regardless, now is not the time to abandon the yield spread when looking at the future macro growth of the US economy. Its historically strong performance demands that it continue to at least receive attention until a conclusive statement can be made about its unreliability.

References

Blough, S. R. Yield curve forecasts of inflation: A cautionary tale. *New England Economic Review, ,* 3.

Bonser-Neal, C., & Morley, T. R. Does yield spread predict real economic activity.. *Economic Review (01612387), 82*(3), 37.

Cook, T., & Hahn, T. Interest rate expectations and the slope of the money market yield curve. *Economic Review (00946893), 76*(5), 3.

Cooper, J. C. (2006). Don't watch this curve too closely. *Business week,* (3966), 22-22.

Dotsey, M. (1998). The Predictive Content of The Interest Rate Term Spread for Future Economic Growth. Federal Reserve Bank of Richmond *Economic Quarterly* 84, 31-51.

Estrella, A., & Hardouvelis. (1991). The term structure as a predictor of real economic activity. *The Journal o f Finance*, 46(2), 555-576.

Estrella, A. & Mishkin. (1995). Predicting US Recessions, Financial Variables as Leading Indicators. NBER Working Paper Series. 5376.

Estrella, A. & Mishkin. (1996). The Yield Curve as a Predictor of US Recessions. *Current Issues in Economics and Finance*, 2(7).

Estrella, A., Rodrigues, and Schich. (2000). How Stable is the Predictive Power of the Yield Curve? Evidence from Germany and the United States.

Estrella, A., & Trubin, M. R. (2006). The yield curve as a leading indicator: Some practical issues. (cover story). *Current Issues in Economics & Finance*, *12*(5), 1-7.

Feroli, Michael. (2004). Monetary Policy and the Information Content of the Yield Spread. *Topics in Macroeconomics*, 4(1), article 13.

Haubrich, J. G. (2006). Does the yield curve signal recession? (cover story). *Economic Commentary,* , 1.

Haubrich, J. G., & Meyer, B. (2007). The yield curve's predictive power. *Economic Trends (07482922),* , 12-13.

Kozicki, S. Predicting real growth and inflation with the yield spread. *Economic Review (01612387), 82*(4), 39.

Kozicki, S., & Sellon, G. Longer-term perspectives on the yield curve and monetary policy. *Economic Review (01612387), 90*(4), 5-33.

MacKenzie, M. (2006). Fresh data could steepen yield-curve inversion. *Wall Street Journal - Eastern Edition, 247*(47), C8.

Plosser, C. I., and K. G. Rouwenhorst. (1994). International Term Structures and Real Economic Growth, *Journal of Monetary Economics* 33, 133-155.

Pu Shen, & Corning, J. Can TIPS help identify long-term inflation expectations? *Economic Review (01612387), 86*(4), 61.

Sanford, J. (2005). Yield curve-onomics. *Canadian Business, 78*(20), 167-168.

Shoesmith, G. L. (2003). Predicting national and regional recessions using probit modeling and interest-rate spreads. *Journal of Regional Science, 43*(2), 373-392.

Stock, J.H., and M.W. Watson (1989). New Indexes of Coincident and Leading Indicators, In

O.J. Blanchard And S. Fischer, Eds., *NBER Macroeconomics Annual*, Cambridge MA: MIT Press.

Stoyko, S. (2006). Treasurys may remain steady as yield-curve inversion persists. *Wall Street Journal - Eastern Edition, 247*(42), C10.

Taylor, J. B. (1993). Discretion Versus Policy Rules in Practice, *Carnegie-Rochester Conference Series On Public Policy* 39, 195-214.

Wyss, D. (2006). Does the yield curve matter? *Business Week Online,* , 4-4.

Table 1:

Probit Model Results for Various Horizons (Quarterly)

	Number of Quarters Lagged					
	1	2	3	4	5	6
Pseudo R-Squared	0.1141	0.2622	0.336	0.2792	0.2264	0.1482
Coefficient	-0.4563	-0.8197	-1.0107	-0.85	-0.7319	-0.55
T-Statistic	-4.47	-5.85	-6.18	-5.96	-5.51	-4.74
Constant	-0.3837	-0.1255	-0.0432	-0.1696	-0.2703	-0.3933

Table 2:

Probit Models for Various Horizons (Monthly)

	Number of Months Lagged							
	5	6	7	8	9	10	11	12
Pseudo R-Squared	0.1728	0.211	0.2518	0.274	0.2827	0.2735	0.267	0.259
Coefficient	-0.574	-0.659	-0.759	-0.808	-0.829	-0.806	-0.791	-0.778
T-Statistic	0.066	0.071	0.079	0.082	0.083	0.082	0.081	0.081
Constant	-0.456	-0.404	-0.349	-0.331	-0.328	-0.347	-0.36	-0.374

Table 3:

Augmented Dickey-Fuller Tests for a Unit Root

	ADF Stat	1% Critical Value	5% CV	10% CV	Reject?
Ten-Year Rate	-1.569	-3.43	-2.86	-2.57	No
Three-Month Rate	-2.155	-3.43	-2.86	-2.57	No
∆Ten-Year Rate	-18.445	-3.43	-2.86	-2.57	Yes
∆Three-Month Rate	-18.012	-3.43	-2.86	-2.57	Yes

Table 4:

Johansen Trace Statistic Tests for Cointegration

Maximum Rank	Trace Statistic	5% Critical Value
0	29.069	15.41
1	3.4671	3.76

Table 5

Estimates of Cointegrating Vectors

Beta	Full Sample	Pre-Volcker	Post-Volcker	Pre-Greenspan	Post-Greenspan	Pre 1993	Post 1993	Post Dec 2001
Ten-Year	1	1	1	1	1	1	1	1
Three-Month	-1.0782	-1.0228	-1.057	-1.1004	-0.8626	-1.0957	-0.508	-0.1802
Std. Error	0.0865	0.1398	0.1141	0.068	0.2404	0.0783	0.1818	0.0737
Constant	-0.9263	-0.616	-1.5403	-0.6091	-2.2397	-0.8005	-3.274	-4.0118