# Predicting Real Growth Using the Yield Curve

by Joseph G. Haubrich and Ann M. Dombrosky

Joseph G. Haubrich is a consultant and economist and Ann M. Dombrosky is a financial reports analyst at the Federal Reserve Bank of Cleveland.

## Introduction

The yield curve, which plots the yield of Treasury bonds against their maturity, is one of the most closely watched financial indicators.[1] Many market observers carefully track the yield curve's shape, which is typically upward sloping and somewhat convex. At times, however, it becomes flat or slopes downward ("inverts," in Wall Street parlance), configurations that many business economists, financial analysts, and other practitioners regard as harbingers of recession (see figure 1).

A recent article in *Fortune* labeled the yield curve "a near-perfect tool for economic forecasting" (see Clark [1996]). In fact, forecasting with the yield curve does have a number of advantages. Financial market participants truly value accurate forecasts, since they can mean the difference between a large profit and a large loss. Financial data are also available more frequently than other statistics (on a minute-to-minute basis if one has a computer terminal), and such a simple test as an inversion does not require a sophisticated analysis.

In this *Review,* we examine the yield curve's ability to predict recessions and, more generally, future economic activity. After comparing the

curve's forecasts with the historical record, we judge its accuracy against other predictions, including naive forecasts, traditional leading indicators, and sophisticated professional projections. This article builds on a wide range of previous research, but, taking an eclectic approach, differs from the earlier work in a variety of ways. These differences show up mainly in the way we judge forecast performance.

Like the important early work of Harvey (1989, 1991, 1993) and Hu (1993), we use out-of-sample forecasts and compare yield curve forecasts with other predictions (including professional forecasts), but we extend our data set to the mid-1990s. In addition, we consider how adding the yield curve improves (or reduces) the accuracy of other forecasts. In this, we follow Estrella and Hardouvelis (1991), who do not, however, use out-of-sample forecasts. Finally, building on the recent work of Estrella and Mishkin (1995, 1996), we consider how well the yield curve predicts the *severity* of recessions, not just their probability, and compare the forecasts with a wider range of alternatives.

■  **1**  Yield curve reports appear in the "Credit Markets" section of *The Wall Street Journal* and the "Business Day" section of *The New York Times.*

## FIGURE 1

### Yield Curves[a]



Percent, weekly averages
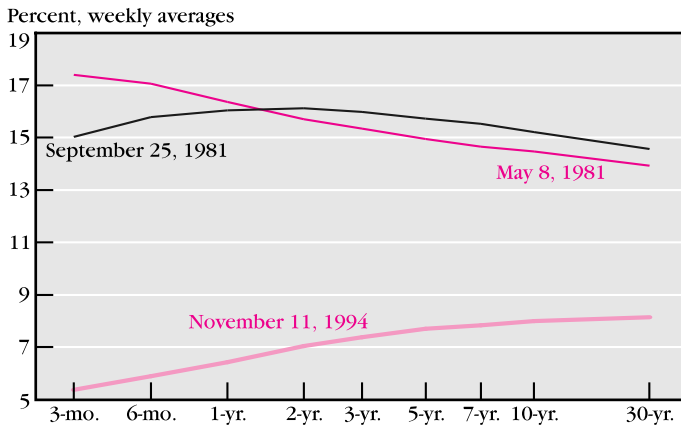
a. Three-month and six-month instruments are quoted from the secondary market on a yield basis; all other instruments are constant-maturity series.
SOURCE: Board of Governors of the Federal Reserve System.

The most distinguishing feature of this paper, however, is that it documents the decline in the yield curve's predictive ability over the past decade (1985–95) and discusses possible reasons for this phenomenon. By some measures, the yield curve should be an even better predictor now than it has been in the past.

Widespread use of the yield curve makes assessing its accuracy a worthwhile exercise for economists. But policymakers, too, need an accurate and timely predictor of future economic growth. The ready availability of term-structure data (as opposed to, say, quarterly GDP numbers) ensures a timely prediction, but accuracy is another question. Central bankers have an added incentive to understand the yield curve, since the federal funds rate and the discount rate are themselves interest rates. Uncovering the "stylized facts" about the curve can help the Federal Reserve to understand the market in which it operates.

With sophisticated macroeconometric models and highly paid professional forecasters, is there any place for a simple indicator like the yield curve? Aside from the knowledge gained about the curve itself, there are several reasons to answer that question affirmatively. Simple predictions may serve as a check on more complex models, perhaps highlighting when assumptions or relationships need rethinking. Agreement between predictions increases confidence in the results, while disagreement signals the need for a second look. A simple, popular indicator also

provides some insight into market sentiment. Of course, it's always a good idea to check whether the expensive and complicated forecasts actually do perform better.

After first reviewing some basics about the yield curve and the reasons it might predict future growth, we look at the actual relationship and compare predictions from the yield curve to those generated by naive statistical models, traditional indicators, professional forecasters, and an econometric model.

## I. Interest Rates and Real Economic Activity

**W**hile our main goal is a rather atheoretical assessment of the yield curve's predictive power, forecasts based on the yield curve are on a sounder economic footing than those based on hemlines or Superbowl victories. The best way to see this is to start with a simple theory of the term structure, called the *expectations hypothesis.*

Under this theory, long-term interest rates are the average of expected future short-term rates. If today's one-year rate is 5 percent and next year's one-year rate is expected to be 7 percent, the two-year rate should be 6 percent ($[7 + 5] \div 2 = 6$). More generally, the expectations hypothesis equates the yield (at time $t$) on an $n$-period bond, $Y_{nt}$, and a sequence of one-period bonds:[2]

$$Y_{nt} = E_t(Y_{1,\,t}\,Y_{1,\,t+1}\,Y_{1,\,t+2}\,...\,Y_{1,\,t+n-1}).$$

If low interest rates are associated with recessions, then an inverted term structure—implying that upcoming rates will be lower—predicts a recession.[3]

One possible reason for expecting low interest rates in recessions might be termed the *policy anticipations hypothesis.* If policymakers act to reduce short-term interest rates in recessions, market participants who expect a recession would also expect low rates. The yield curve picks up the financial market's estimate of future policy.[4]

■ **2** See chapter 7 of Mishkin (1989) for a fuller treatment of this issue. Mishkin also points out the main flaw in the expectations hypothesis: The term structure normally slopes up, but interest rates do not trend up over time. Campbell (1995) offers a useful discussion of related points.

■ **3** For a classic documentation of this pattern, see Kessel (1965).

■ **4** Rudebusch (1995) takes this approach.

Another possibility is that current monetary policy may shift both the yield curve and future output. For example, tight monetary policy might raise short-term interest rates, flattening the yield curve and leading to slower future growth. Conversely, easy policy could reduce short-term interest rates, steepen the yield curve, and stimulate future growth. The yield curve predicts future output because each of these shifts follows from the same underlying cause: monetary policy. Taking this logic one step further, monetary policy may react to output, so that the yield curve picks up a complex intermingling of policy actions, reactions, and real effects.

In these explanations, the yield curve reflects future output indirectly, by predicting future interest rates or future monetary policy. It may also reflect future output directly, because the 10-year interest rate may depend on the market's guess of output in 10 years.

The expectations hypothesis certainly marks the beginning of wisdom about the yield curve, but only the beginning. The 30-year bond may have a high interest rate not because people expect interest rates to rise, but because such a bond must offer a high return to get people to hold it in the first place. (This is commonly called the risk premium, though for some theories that may be a misnomer.) Investors may dislike wide swings in prices as market expectations about the distant future change over time. Conversely, there may be reasons why some people would rather hold a 30-year bond than a one-year bond. For example, they may be saving for retirement and prefer the certain payoff on the longer-term note (this is sometimes called the *preferred habitat hypothesis*).

The risk premium provides another reason why the yield curve may be a useful predictor: The premium itself holds information. As a simple example, consider that recessions may make people uncertain about future income and employment, or even about future interest rates. The risk premium on a longer-term bond reflects this. In conjunction with changes working through the expectations hypothesis, the yield curve may take some very strange twists indeed, becoming inverted, humped, or even u-shaped.[5]

These explanations provide an additional motivation for investigating yield curve predictions. They also hint at the many important issues that transcend the yield curve's predictive power. It matters, for instance, if the curve reacts to future policy, to movements in output, or to some combination of the two. But these considerations fall by the wayside if the yield curve is not an accurate predictor of future economic activity. In this article, we concentrate on that more basic issue, leaving determination of the underlying causes for another day.

## II. Data and Computation

There are many ways of using the yield curve to predict future real activity. One common method uses inversions (when short rates exceed long rates) as recession indicators. Is it possible, however, to predict the magnitude as well as the direction of future growth? Does a large inversion predict a severe recession? Does a steep yield curve predict a strong recovery?[6] Operationally, this means relating a particular measure of yield curve "steepness" to future real growth. In taking this route, we follow and build on the related work of Estrella and Hardouvelis (1991).

Obtaining predictions from the yield curve requires much preliminary work. Three principles guided us through the many decisions that were required: Keep the process simple, preserve comparability with previous work, and avoid data snooping. Thus, we avoided both complicated nonlinear specifications and a detailed search for the "best" predictor.

To begin with, there is no unambiguous measure of yield curve steepness. A yield curve may be flat at the short end and steep at the long end. The standard solution uses a spread, or the difference between two rates (in effect, a simple linear approximation of the nonlinear yield curve).[7] This means choosing a particular spread, in itself no trivial matter. Among the 10 most commonly watched interest rates (the federal funds rate and the three-month, six-month, and one-, two-, three-, five-, seven-, 10-, and 30-year Treasury rates), 45 possible spreads exist.[8]

An additional problem is that there are several types of yield curves or term structures. In fact, it sometimes helps to draw a distinction between the yield curve and the term structure. The yield curve is the relation between the

---

■  **5**   Stambaugh (1988) makes this point.  For a less technical description, see Haubrich (1991).

■  **6**   Other approaches also exist. For example, Harvey (1988) examines whether the term structure predicts changes in consumption.

■  **7**   Frankel and Lown (1994) is one of the few papers that considers nonlinear measures of steepness.

■  **8**   If there are $n$ rates, there are $n/2\,(n-1)$ spreads. This is the classic formula for combinations.  See Niven (1965), chapter 2.

yield on Treasury securities and their maturity. The term structure is a particular yield curve—that for zero-coupon Treasury securities. The term structure is theoretically more interesting. It answers the question, "How much would I pay for one dollar delivered 10 years from to-day?" The problem is that a zero-coupon Treasury security rarely matures in exactly 10 years. What we actually observe in the market are prices (and thus yields) on existing Treasury securities. These may not mature in precisely 10 years (or whatever maturity you choose), and they often have coupon payments. That is, a 10-year Treasury note pays interest semiannually at a specified coupon rate, so its yield is not the yield called for in the term structure.

Finding the desired interest rate almost always involves estimation of some kind. Calculating the theoretically pure term structure is often quite difficult, as it must be estimated from coupon bonds of the wrong maturity, all subject to taxation. (Using zero-coupon bonds may help, but this approach introduces problems of its own, as the market is thinner and the tax treatment of coupons and principal differs.) This means that the pure term structure is not available in real time, when the Federal Reserve must attempt to discern the course of the economy. To avoid stale data, we must turn to the more "rough-and-ready" yield curve. Even here, the problem of matching maturities arises. Fortunately, the Treasury Department publishes a "constant-maturity" series, where market data are used to estimate today's yield on a 10-year Treasury note, even though no such note exists.[9]

For our study, we use data from the Federal Reserve's weekly H.15 statistical release ("Selected Interest Rates"), which compiles interest rates from various sources. For the spread, we chose the 10-year CMT rate minus the secondary-market three-month Treasury bill rate. In addition to allowing a comparison with the work of Estrella and Hardouvelis (1991), Harvey (1989, 1993), and Estrella and Mishkin (1995, 1996), choosing only one spread enables us to minimize the problems of data snooping (Lo and MacKinlay [1990]) and the associated spuriously good results.[10] That is, trying every single spread would produce something that looked like a good predictor, but it very likely would be a statistical fluke akin to Superbowl victories and hemlines. We then convert the bill rate, which is published on a discount rate basis, to a coupon-equivalent yield so that it is on the same basis as the 10-year rate.[11]

Also following Estrella and Hardouvelis, we use quarterly averages for the spread. This smoothes the anomalous rates that appear at the turn of each month.[12] A priori there is no presumption that GDP should correlate better with a particular date's spread than with the quarterly average.[13]

As our measure of real growth, we use the four-quarter percent change in real (fixed-weight) GDP. GDP is, of course, the standard measure of aggregate economic activity, and the four-quarter forecast horizon answers the "what-happens-next-year" type of question without embroiling us in data snooping issues regarding the optimal horizon choice.

Our sample period runs from 1961:IQ through 1995:IIIQ. This covers various inflationary experiences, episodes of monetary policy tightening and easing, and several business cycles and recessions. Included are five recessions, inflation rates from 1 percent to more than 13 percent, and a federal funds rate ranging from under 3 percent to over 19 percent.

Our basic model, then, is designed to predict real GDP growth four quarters into the future based on the current yield spread. Operationally, we accomplish this by running a series of regressions (detailed below) using real GDP growth and the interest rate spread lagged four quarters (for example, the interest rate spread used for 1961:IQ is actually from 1960:IQ).

The next step involves comparing the yield curve forecasts with a sequence of increasingly sophisticated predictions using other techniques. We start with a naive (but surprisingly effective) technique which assumes that GDP growth over the next four quarters will be the same as it was over the last four. (That is, the growth rate is a random walk.) We then regress real GDP growth against the index of leading economic indicators (lagged four quarters). This enables us to make a comparison with another simple and popular forecasting technique.

---

■ **9** See Smithson (1995) for a good description of constant-maturity Treasuries (CMTs).

■ **10** Work by Knez, Litterman, and Scheinkman (1994) suggests that using more or different rates would not capture much additional information.

■ **11** The three-month CMT rate was not published before May 1995. To keep the data consistent, we use the secondary-market three-month Treasury bill rate throughout. For such a short rate, the differences are minimal.

■ **12** Park and Reinganum (1986) document this calendar effect.

■ **13** We also reworked the results using data for the last week of each quarter for the 1963–95 period. The findings were comparable, although the predictive power of the spread decreased somewhat.

**B O X  1**

**Forecasting Equations**

Yield spread: in-sample

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \alpha + \beta\ spread_t$$

Yield spread: out-of-sample

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \alpha + \beta\ spread_t$$

Naive

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \frac{RGDP_t - RGDP_{t-4}}{RGDP_t}$$

Leading indicators

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \alpha + \beta\ index_t$$

Lagged GDP

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \alpha + \beta\ \frac{RGDP_t - RGDP_{t-4}}{RGDP_t}$$

Lagged GDP plus yield spread

$$\frac{RGDP_{t+4} - RGDP_t}{RGDP_t} = \alpha + \beta\ \frac{RGDP_t - RGDP_{t-4}}{RGDP_t} + \gamma\ spread_t$$

We next look at two additional forecasts generated by statistical procedures. We regress real GDP growth against its own lag (again, four quarters) and against its own lag and the 10-year, three-month spread.

The final, and most sophisticated, alternative forecasts we consider come from the Blue Chip organization and DRI/McGraw–Hill (hereafter referred to simply as DRI). We first compare the results of our model with forecasts from *Blue Chip Economic Indicators,* beginning with the July 1984 issue.[14] We use the one-year-ahead Blue Chip consensus forecasts for real GDP (or real GNP when GDP forecasts are unavailable), labeled "percent change from same quarter in prior year." These forecasts are taken from the Blue Chip newsletters corresponding to the first month of each quarter.

We next compare our results with predictions from DRI, reported in various issues of its *Review of the U.S. Economy.* DRI generates these forecasts from an econometric model. Although we tried to collect forecasts for the same period as our Blue Chip forecasts (that is, from issues corresponding to the first month of each quarter), our DRI data set is missing two points: 1985:IIIQ and 1987:IIQ. We use forecasts for real GDP (or real GNP when GDP forecasts are unavailable) one year ahead.

Box 1 summarizes the regressions used to forecast future real GDP growth.

## III.  Forecast Results

**D**oes the yield curve accurately predict future GDP? First, look directly at the data. Figure 2 shows the growth of real GDP and the lagged spread between the 10-year and three-month Treasury yields. A decline in the growth of real GDP is usually preceded by a decrease in the yield spread, and a narrowing yield spread often signals a decrease in real GDP growth. A negative yield spread (inverted yield curve) usually precedes recessions, but not always. For example, the yield spread turned negative in the third and fourth quarters of 1966, but no recession occurred for the next three years. (The recession that began in late 1969 was preceded by two quarters of a negative yield spread.) The latest recession, which occurred in 1990–91, was preceded by a yield curve more accurately described as flat than inverted.

Figure 3 plots the same data in a different form. It shows a scatterplot, with each point representing a particular combination of real

■  **14**  *Blue Chip Economic Indicators* is a monthly collection of economic forecasts by a panel of economists from some of the top firms in the United States (the so-called Blue Chip companies). The Blue Chip consensus forecast for real GDP is the average of about 50 individual forecasts. See Zarnowitz and Lambros (1987) for evidence that the consensus forecast predicts much better than individual forecasts, and Lamont (1994) for a possible explanation.

## FIGURE 2

### Real GDP Growth and Lagged Yield Spread

Percent yield
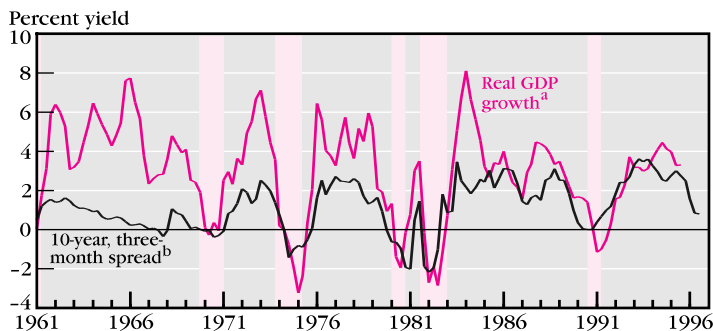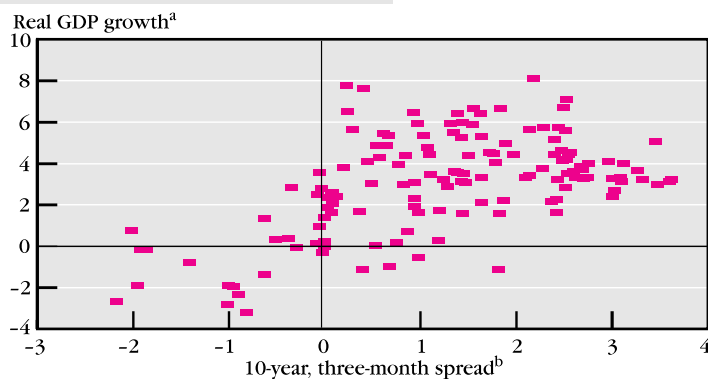


a. Four-quarter percent change.
b. Lagged four quarters.
NOTE: Shaded areas indicate recessions.
SOURCES: Board of Governors of the Federal Reserve System; and U.S. Department of Commerce, Bureau of Economic Analysis.

## FIGURE 3

### Scatterplot of Real GDP Growth and Lagged Yield Spread

Real GDP growth[a]



a. Four-quarter percent change.
b. Lagged four quarters.
SOURCES: Board of Governors of the Federal Reserve System; and U.S. Department of Commerce, Bureau of Economic Analysis.

GDP growth and the lagged yield spread. Even a casual look at the results reveals that the relationship between the two variables is usually positive; that is, positive real GDP growth is associated with a positive lagged yield spread, and vice versa.

Plotting the data gives a strong, albeit qualitative, impression that the yield spread predicts future real activity. We desire a more quantitative prediction, one that says more than "The yield curve is steep; looks like good times." To

generate the GDP predictions, we ran an in-sample regression, using the entire sample to generate each predicted data point. This is the sort of comparison Estrella and Hardouvelis (1991) make, and our results, presented below, confirm their assessment that the 10-year, three-month spread has significant predictive power for real GDP growth:[15]

Real GDP growth = 1.8399 + 0.9791$spread$
$\qquad$ (3.89) $\quad$ (4.50)

$$R^2 = 0.291, \ D-W = 0.352.$$

The yield spread emerges as statistically and economically significant, translating almost one for one into expected future growth. Thus, a spread of 100 basis points (1 percent) implies future growth of 2.8 percent (we derive this as 1.8 + 0.98 × 1). The $R^2$ indicates that much variation remains to be explained. Figure 4 plots our in-sample real GDP predictions versus actual real GDP growth.

The in-sample results are somewhat misleading, as the coefficients depend on information not available early in the sample. Figure 5 plots another series of predicted real GDP growth, this time generated from an out-of-sample regression. Each data point in this chart is based on a regression using only the data (yield spreads) before the predicted data point. That is, the predicted GDP growth rate for, say, 1980:IQ is based on the data sample from 1961:IQ through 1979:IVQ. Hence, this regression generates a true forecast because it uses available data to predict future (out-of-sample) real GDP growth.[16]

The predicted GDP series from the in-sample and out-of-sample regressions are broadly similar and generally follow the actual GDP data. The root mean square error (RMSE) of the predictions is 2.04 for the in-sample and 2.10 for the out-of-sample forecasts. It is not surprising that the in-sample regression performs slightly better. If we calculate the RMSE over the last 10 years of our data set (1985:IIIQ to 1995:IIIQ), the in-sample regression (RMSE 1.07) again does better than the out-of-sample regression (RMSE 2.09).

■ **15** The *t* statistics are Newey–West corrected with five lags. This offsets the bias created by overlapping prediction intervals (a serious problem in this case, as indicated by the Durbin–Watson statistic).

■ **16** We also corrected the regression for a more subtle problem. Because first-quarter GDP numbers become available only after the first quarter ends, we should not use those numbers in a regression. (That is, as of the first quarter, we still don't know the four-quarter growth rate over last year's first quarter.)
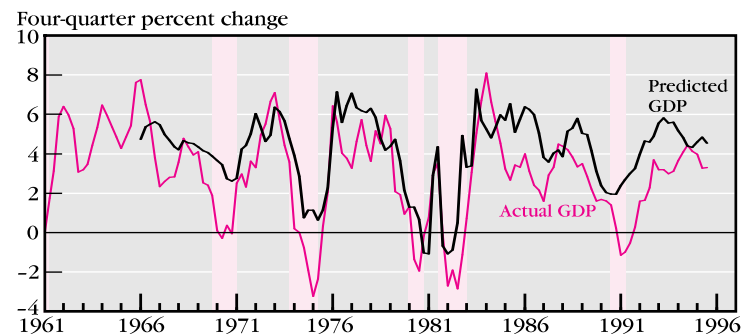
## FIGURE 4

### Real GDP Predictions: In-Sample
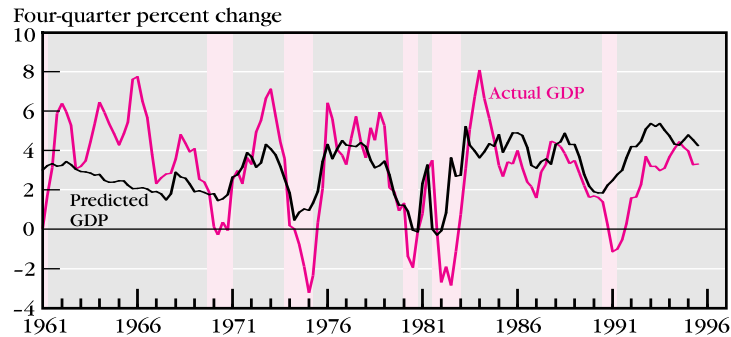
Four-quarter percent change



NOTE: Shaded areas indicate recessions.
SOURCES: U.S. Department of Commerce, Bureau of Economic Analysis; and authors' calculations.

## FIGURE 5

### Real GDP Predictions: Out-of-Sample

Four-quarter percent change



NOTE: Shaded areas indicate recessions.
SOURCES: U.S. Department of Commerce, Bureau of Economic Analysis; and authors' calculations.

## TABLE 1

### RMSE of Forecasts

|  | 1965:IVQ–1995:IIIQ | 1985:IIIQ–1995:IIIQ |
| --- | --- | --- |
| Yield spread: in-sample | 2.04 | 1.07 |
| Yield spread: out-of-sample | 2.10 | 2.09 |
| Naive | 3.15 | 1.66 |
| Leading indicators | 2.37 | 1.36 |
| Lagged GDP | 2.49 | 1.46 |
| Lagged GDP plus yield spread | 2.19 | 2.06 |
| DRI[a] | n.a. | 0.63 |
| Blue Chip | n.a. | 1.13 |

a. DRI forecasts are missing two quarters.
SOURCE: Authors' calculations.

We use this RMSE criterion to compare the yield spread forecasts with those derived from other techniques. The results are reported in table 1.

For the entire sample, the yield curve emerges as the most accurate predictor of real economic growth. Furthermore, adding the yield spread to a lagged GDP regression improves the forecast, while adding lagged GDP to the yield spread worsens the forecast. For in-sample regressions, adding variables never hurts, but it quite commonly reduces the performance of the out-of-sample regressions.

Curiously, the 1985–95 subsample completely reverses the results. The yield spread becomes the least accurate forecast, and adding it to lagged GDP actually worsens the fit. The leading indicators emerge as the best of the "low-cost" forecasts, and the two professional services do markedly better than the rest. In part, this may reflect the simple specifications used in the regression forecasts: With more lags, a simple regression forecast is often better than a sophisticated model (see Chatfield [1984], chapter 5). But the change is even more significant than that. Using unpublished data, Harvey (1989) finds that over the 1976–85 period, the yield spread performs as well as or better than seven professional forecasting services (including DRI but not Blue Chip).

The dramatic drop in forecasting ability may result from several factors. It certainly reflects a changing relationship between the yield curve and the economy. The coefficients in the term-spread regression demonstrate this. At the beginning of the sample, using only 20 data points (five years of quarterly data), the coefficient on the term spread is –0.14 (statistically insignificant). Midway through the sample, after 70 data points, the coefficient is 1.48, and for the whole sample, 0.98. While one advantage of an out-of-sample procedure is that it allows the coefficients to change, the influence of the first 20 years may force the "wrong" coefficient on the last 10. It is also quite reasonable that the relationship between the yield curve and the real economy might have changed over 30 years. Advances in technology, new production processes, changes in market organization or in the way the market reacts to new information, or even shifts in Federal Reserve policy (the famous "Lucas critique") might have altered the relationship between the yield curve and real activity.

Evidence suggests that both the timing and the size of the relationship between the yield curve and real activity has in fact changed. If
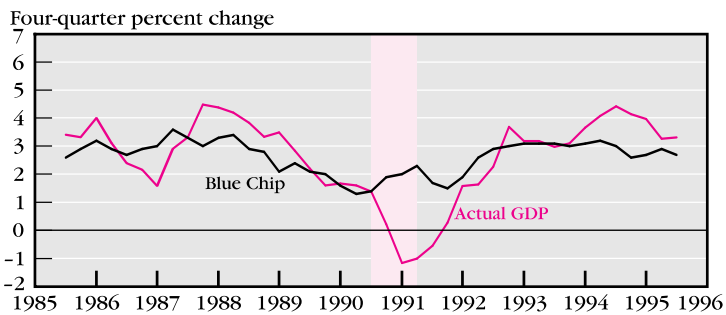
## T A B L E  2

**Correlations between Lagged Yield Spread and Real GDP Growth**

| Lag | 1960:IQ–1995:IIIQ | 1960:IQ–1985:IIQ | 1985:IIIQ–1995:IIIQ |
|-----|-------------------|------------------|---------------------|
| 0 | 0.041 | 0.075 | 0.143 |
| 1 | 0.206 | 0.252 | 0.358 |
| 2 | 0.366 | 0.439 | 0.538 |
| 3 | 0.481 | 0.585 | 0.641 |
| 4 | 0.540 | 0.666 | 0.684 |
| 5 | 0.505 | 0.624 | 0.708 |
| 6 | 0.423 | 0.523 | 0.736 |
| 7 | 0.331 | 0.398 | 0.721 |
| 8 | 0.218 | 0.251 | 0.660 |

SOURCE:  Authors' calculations.

## F I G U R E  6

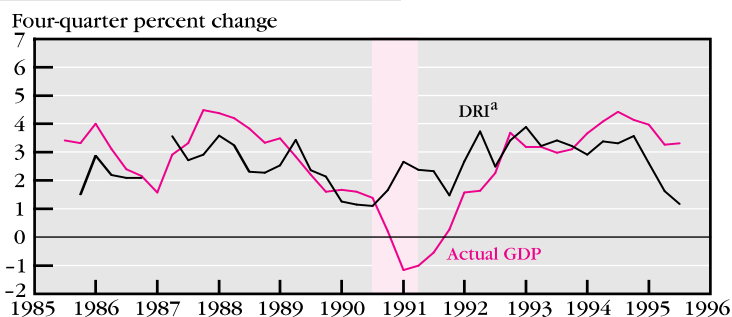**Real GDP Predictions:  Blue Chip**

Four-quarter percent change

NOTE:  Shaded area indicates recession.
SOURCES:  U.S. Department of Commerce, Bureau of Economic Analysis;
and *Blue Chip Economic Indicators,* various issues.

## F I G U R E  7

**Real GDP Predictions:  DRI**

Four-quarter percent change

a. 1985:IIIQ and 1987:IIQ data are missing.
NOTE:  Shaded area indicates recession.
SOURCES:  U.S. Department of Commerce, Bureau of Economic Analysis;
and DRI/McGraw–Hill.

we look at the correlations between the yield spread and real GDP growth at different lags (table 2), we see that the recent period has higher correlations between lags of yield spreads and real GDP growth. We also see that the largest correlation for the early (and total) period, 0.666, occurs at a lag of four quarters, exactly the lag used in our regressions. In the recent period, however, despite a higher correlation at four lags, the highest correlation (0.736) is reached at six quarters. The correlations drop off more slowly in the latter period as well. This accounts for our somewhat paradoxical conclusion: Despite better correlations between the yield curve and real GDP—with an in-sample RMSE that beats even the Blue Chip forecast—regressions using past data are less reliable predictors.

It is somewhat instructive to make a more detailed comparison with the sophisticated forecasts. The Blue Chip predictions can be considered out-of-sample because each one is based on data available at the time of the prediction. Figure 6 plots the Blue Chip consensus forecasts against actual real GDP growth. The Blue Chip forecasts appear much smoother than GDP, as they consistently underpredict real GDP when economic growth is high and overpredict real GDP when economic growth is negative.

The DRI forecasts are plotted in figure 7. These appear broadly similar to the Blue Chip forecasts, although the DRI series is more volatile. Like the Blue Chip forecasts, the DRI forecasts generally underpredict GDP when economic growth is high and overpredict GDP when economic growth is low. Our out-of-sample forecasts based on the yield spread overpredict GDP growth for all but one quarter during the 1985:IIIQ–1995:IIIQ period (corresponding to the Blue Chip and DRI data sets).

## IV.  Conclusion

**D**oes the yield curve accurately predict real economic growth? Answering this seemingly simple question requires a surprising amount of preliminary work. Much of this paper is devoted to refining the initial question to confront the realities of the financial marketplace.

Fortunately, the answer is less complex, if somewhat nuanced. The 10-year, three-month spread has substantial predictive power, and in this we confirm a variety of earlier studies. Over the past 30 years, it provides one of the best (in our sample, *the* best) forecasts of real growth four quarters into the future. Over the past decade, it has been less successful: Indeed,

the yield curve was the *worst* forecast we examined. This shift seemingly results from a change in the relationship between the yield curve and real economic activity—one that has become closer, but nonetheless has made regressions based on past data less useful.

An interesting topic for future research would be to examine whether simple fixes, such as a rolling regression model or more lags, could improve the recent performance of the yield curve. Certainly the simple yield curve growth forecast should not serve as a replacement for the consensus predictions of the Blue Chip panel or the DRI econometric model. It does, however, provide enough information to serve as a useful check on the more sophisticated forecasts and to encourage future research into the reasons behind the yield curve's worsening performance.

## References

Campbell, J.Y. "Some Lessons from the Yield Curve," *Journal of Economic Perspectives,* vol. 9, no. 3 (Summer 1995), pp. 129–52.

Chatfield, C. *The Analysis of Time Series: An Introduction,* 3d. ed. New York: Chapman and Hall, 1984.

Clark, K. "A Near-Perfect Tool for Economic Forecasting," *Fortune,* July 22, 1996, pp. 24–26.

Estrella, A., and G.A. Hardouvelis. "The Term Structure as a Predictor of Real Economic Activity," *Journal of Finance,* vol. 46, no. 2 (June 1991), pp. 555–76.

Estrella, A., and F. S. Mishkin. "Predicting U.S. Recessions: Financial Variables as Leading Indicators," National Bureau of Economic Research, Working Paper No. 5379, 1995.

_____ , and _____ . "The Yield Curve as a Predictor of U.S. Recessions," Federal Reserve Bank of New York, *Current Issues in Economics and Finance,* vol. 2, no. 7, June 1996.

Frankel, J.A., and C.S. Lown. "An Indicator of Future Inflation Extracted from the Steepness of the Interest Rate Yield Curve along Its Entire Length," *Quarterly Journal of Economics,* vol. 109, no. 2 (May 1994), pp. 517–30.

Harvey, C.R. "The Real Term Structure and Consumption Growth," *Journal of Financial Economics,* vol. 22, no. 2 (December 1988), pp. 305–33.

_____ . "Forecasts of Economic Growth from the Bond and Stock Markets," *Financial Analysts Journal,* September/October 1989, pp. 38–45.

_____ . "The Term Structure and World Economic Growth," *Journal of Fixed Income,* June 1991, pp. 7–19.

_____ . "Term Structure Forecasts Economic Growth," *Financial Analysts Journal,* vol. 49, no. 3 (May/June 1993), pp. 6–8.

Haubrich, J.G. "Wholesale Money Market," in *The New Palgrave Dictionary of Money and Finance.* New York: Stockton Press, 1992, pp. 798–800.

Hu, Z. "The Yield Curve and Real Activity," *International Monetary Fund Staff Papers,* vol. 40, no. 4 (December 1993), pp. 781–806.

Kessel, R.A. "The Cyclical Behavior of the Term Structure of Interest Rates," National Bureau of Economic Research, Occasional Paper No. 91, 1965.

Knez, P.J., R. Litterman, and J. Scheinkman. "Explorations into Factors Explaining Money Market Returns," *Journal of Finance,* vol. 49, no. 5 (December 1994), pp. 1861–82.

Lamont, O. "Macroeconomic Forecasts and Microeconomic Forecasters," Princeton University, unpublished manuscript, 1994.

Lo, A.W., and A.C. MacKinlay. "Data-Snooping Biases in Tests of Financial Asset Pricing Models," *Review of Financial Studies,* vol. 3, no. 3 (1990), pp. 431–67.

Mishkin, F. S. *The Economics of Money, Banking, and Financial Markets,* 2d. ed. Glenview, Ill.: Scott, Foresman, and Company, 1989.

Niven, I. *Mathematics of Choice: How to Count without Counting.* Washington, D.C.: The Mathematical Association of America, 1965.

**Park, S.Y., and M.R. Reinganum.** "The Puzzling Behavior of Treasury Bills that Mature at the Turn of Calendar Months," *Journal of Financial Economics,* vol. 16, no. 2 (June 1986), pp. 267–83.

**Rudebusch, G.D.** "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure," *Journal of Monetary Economics,* vol. 35, no. 2 (April 1995), pp. 245–74.

**Smithson, C.** "ABC of CMT," *Risk,* vol. 8, no. 9 (September 1995), pp. 30–31.

**Stambaugh, R.F.** "The Information in Forward Rates: Implications for Models of the Term Structure," *Journal of Financial Economics,* vol. 21, no. 1 (May 1988), pp. 41–70.

**Zarnowitz, V., and L.A. Lambros.** "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy,* vol. 95, no. 3 (June 1987), pp. 591–621.