An evaluation of real GDP forecasts: 1996–2001

Spencer Krane

Introduction and summary

Increases in real U.S. gross domestic product (GDP) averaged an annual rate of 3.2 percent between the fourth quarters of 1991 and 1995 (the solid line in panel A of figure 1), a relatively slow pace of growth considering that the economy was emerging from the 1990–91 recession. Output then surged in the second half of the decade, with current estimates showing real GDP rising at an average annual rate of 4.4 percent over the 1996–99 period. At the same time, inflation fell, with the rate of increase in consumer prices (measured by the Consumer Price Index, or CPI) moving from 5.4 percent in 1990 to an average of just 2.4 percent in the second half of the decade (solid line in panel B). The bars in the graphs show average forecasts of real GDP growth and CPI inflation made at the beginning of each year.[1] Between 1996 and 1999, average real GDP forecasts were in the range of 2.1 percent to 2.3 percent, while the CPI forecasts were in the range of 2.2 percent to 3 percent. Clearly, forecasters failed to predict the outstanding performance of the economy—they consistently underpredicted GDP growth and, though to a lesser degree, they overpredicted inflation.

At the turn of the millennium, forecasts for real GDP growth were in the range of about 3 percent to 3.5 percent. While not quite as robust as the actual rates of growth recorded during the second half of the decade, this still represented a solid gain in output and a step up from the projections made in that earlier period. Instead, in the second half of 2000, the expansion began to falter. The weakness intensified in early 2001, with the economy falling into recession in March. So again, forecasters failed to predict a major development in the economy.

How should we interpret these forecast errors? The economy is always being hit by shocks, and real GDP growth naturally fluctuates a great deal. Furthermore, recessions are irregular occurrences that can be generated by a variety of unforeseeable events. So, were the forecast errors during the 1996–2001 period unusual, or did they simply reflect the inherent difficulties in forecasting? If the errors were unusual, then why is this so? In particular, did forecasters change the way that they were constructing projections, or did the economy behave in an unusual manner? This article addresses these questions.

To do so, I first present a narrative account of the evolution of real GDP forecasts made during the 1996–2001 period. This narrative shows, qualitatively, that forecasters appeared to view most of the errors they were experiencing during the 1996–99 period as transitory and left GDP projections at a pace just somewhat below their benchmarks for longer-run growth. However, around the turn of the millennium, they boosted their projections for GDP growth, both for the long run and the nearer term. Indeed, they did so just around the time that the economy began to weaken.

This strategy clearly resulted in some large and, during 1996–99, persistent forecast errors for real GDP. I next show that, statistically, the 1996–99 errors were unusual—based on forecasters' track records, the odds of seeing such a string of underpredictions were quite small. The forecast errors in 2000 and 2001, though large in an absolute sense, were not so significant relative to the performance around earlier turning points in the economy.

Next, I examine whether the errors were influenced by some change in the way forecasters were making their projections. I use semiannual data back to the early 1980s to characterize the "typical" way that

Spencer Krane is a vice president and economic advisor at the Federal Reserve Bank of Chicago. *The author would like to thank Charlie Evans, Helen Koshy, Tina Lam, David Marshall, seminar participants at the Chicago Fed, and especially, Michael Munley for helpful comments and assistance. He also would like to acknowledge Aspen publishers for allowing use of the Blue Chip data.*

FIGURE 1

Real GDP growth and CPI inflation, 1990–2001

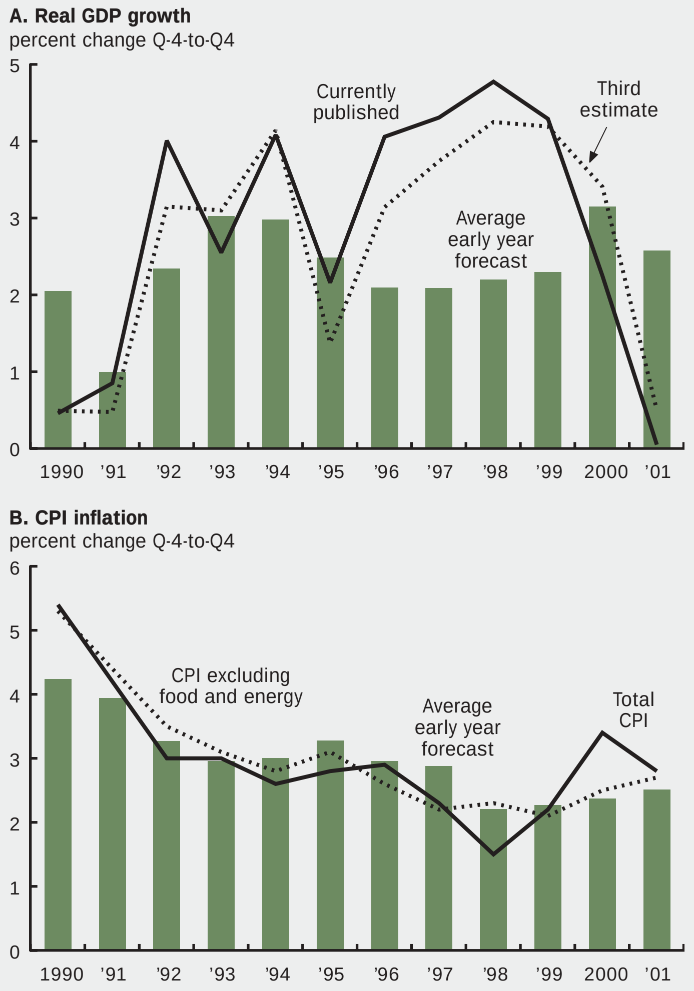

Sources: Federal Open Market Committee (FOMC), 1979–2001, *Federal Reserve Board Monetary Policy Reports to Congress*, Administration, 1979–2001, *The Budget of the United States Government*; submissions and mid-session reviews, and 1979–2001, *Economic Report of the President*; Congressional Budget Office (CBO), 1979–2001, *The Economic and Budget Outlook*, submissions and mid-year updates; Blue Chip, 1978–2001, *Blue Chip Economic Indicators*, various issues; Federal Reserve Bank of Philadelphia, *Survey of Professional Forecasters*; U.S. Bureau of Economic Analysis, *National Income and Product Accounts;* and U.S. Bureau of Labor Statistics, *Consumer Price Index.*

forecasters adjust projections for growth at various forecast horizons. I find that forecasters appear to view most shocks as being transitory—they may alter their near-term outlook in response to incoming data, but they generally do not change medium- and longer-term forecasts very much. This means that perceptions of longer-run trends—or potential GDP growth—provide an important anchor for projections more than a couple quarters out. As just noted, this characterization seems to describe the forecasts made between 1996 and 1999. Some other identifiable factors, such as recessions or shifts in economic policy, also have had a regular statistical influence on medium-term forecasts. However, such factors did not seem to be in play during the second half of the 1990s, while in 2001, forecasters appeared to react in a fairly typical fashion to the signals that the economy was weakening. Accordingly, forecasters probably did not behave unusually during the 1996–2001 period.

These results suggest that the forecast errors during this time likely reflect some unusual behavior in the economy. The final portion of this article discusses a couple of important candidates. First, during the second half of the 1990s, there was a marked and persistent pick-up in productivity growth, a rare development given the mature stage of the business cycle. Thus, the surprising step-up in actual GDP growth around mid-decade may have reflected the response of households and businesses to more robust underlying trends in productivity. Second, much of the downshift in overall economic activity in 2000 and 2001 reflected a surprisingly abrupt swing from boom to bust in business fixed investment. This swing seemed to accompany a rather sharp reassessment by financial markets and businesses of the earnings potential of certain investment projects, particularly in the high-technology area. To be sure, claims were made in the late 1990s that a high tech “bubble” had developed. But not only are such phenomena problematic to identify ex ante, predicting the timing and magnitude of any “bursting of the bubble” is virtually impossible. Indeed, at the turn of the millennium, even the more pessimistic forecasters thought that real GDP would rise at more than a 2 percent pace in 2000 and 2001.

Of course, the benefit of hindsight allows us to analyze history with some knowledge of the important shocks that hit the economy and of the responses of households and businesses to those events. Forecasters do not have this luxury. By their very nature, shocks are unknowable in advance. And once shocks begin to unfold, forecasters must make numerous judgment calls regarding their magnitude and persistence. If the surprises are unusual—such as those during the 1996–2001 period—history provides little guidance on how to make such judgments. Forecasting is further complicated by the fact that incoming data rarely provide a clear-cut reading on the course of events and because a good deal of time must pass before any persistent change in the economy can be identified with much statistical confidence. As a result, real-time forecasting is a much more difficult exercise than dissecting the performance of projections after the fact.

The data

The forecasters

For the sake of generality, I consider five widely cited public and private sector forecasts. The forecasts are best described as judgmental, although many are informed to varying degrees by econometric models.

Three important public agencies publish forecasts twice a year: The members of the Federal Open Market Committee (FOMC) and other District Federal Reserve Bank presidents present projections in their semiannual *Monetary Policy Reports* to Congress; and the Administration and the Congressional Budget Office (CBO) publish forecasts in conjunction with the submission and mid-session reviews of the President's Budget.[2] Many private-sector economists publish macroeconomic forecasts. I use two commonly cited averages—the consensus outlook published by *Blue Chip Economic Indicators* and the median projections from the Federal Reserve Bank of Philadelphia's *Survey of Professional Forecasts* (SPF) (see Croushore, 1993). Blue Chip forecasts are made each month, while the SPF is published quarterly. The current Blue Chip sample covers 52 forecasters, while the SPF covers about 35; the samples share about 15 respondents.

The forecasts

The variables projected, forecast periodicity, forecast horizon, and conditioning information vary among these forecasts. Notably, projections for the current year are available from each of these sources, but the FOMC projects the following year only in its mid-year report. Forecasts for quarterly data are available only for the Blue Chip and the SPF. All the forecasts include projections of real GDP, an inflation measure, the unemployment rate, and, with the exception of the FOMC, some interest rate. Wide sectoral detail, however, is available only for the SPF. All of the forecasters except the FOMC publish "long-run" projections, although the exact definition of "long-run" and the availability of these forecasts vary somewhat across forecasters and over time.

I often refer to "early year" and "mid-year" projections for real GDP growth. The early year forecasts all are published in February, though some (notably the Administration's) often are completed a couple of months earlier. The mid-year FOMC and Blue Chip forecasts are released in early July, the SPF in August, while the exact month that the Administration and CBO mid-session reviews are released varies through the summer. I also make use of Blue Chip forecasts made in March and August, the two months when long-term forecasts are collected. Current-year forecasts refer to projections made for the increase in real GDP between the fourth quarter of the previous year and the fourth quarter of the current calendar year. Half-year forecasts refer to annualized growth between the fourth quarter of the previous year and the second quarter of a year or from the second to fourth quarters of the same year. In addition, in December 1991 the U.S. Bureau of Economic Analysis (BEA) moved from using gross national product (GNP) to using GDP as the featured measure of aggregate output; I use the forecasts for GNP prior to 1991.

Reference data

When comparing forecasts to outcomes, one must decide which vintage of the *National Income and Product Accounts* (NIPA) to use for the "actual" values of GDP and its components. At various times, I present calculations based on different vintages of the NIPA in order to compare forecasts with the historical data in hand when a particular projection was made or to highlight other features of the data. For the most part, I construct forecast errors by comparing projections with the "third" or "final" estimates of the NIPA. When a comprehensive revision has occurred between the time a forecast was made and the third estimate is released, I adjust the forecast error or other data presentations for the average revision to GDP growth over the previous several years. This purges the analysis of the influence of the rebasing of GDP or major definitional changes that occur with comprehensive revisions but most likely were not incorporated in earlier forecasts.

Forecasting experience of the late 1990s and the 2001 recession

Below, I present a narrative account of the evolution of real GDP forecasts made during the high-growth period of the second half of the 1990s and around the 2001 recession. The discussion highlights the errors experienced during these periods and some apparent regularities in forecasting procedures that might help explain these errors. Table 1 presents the early year and mid-year forecasts for GDP growth over the 1996–2001 period. Table 2 shows forecasters' assumptions for the longer-run trends in GDP and productivity.

TABLE 1

Current-year real GDP forecasts, 1996–2001
(Q4-to-Q4 percent change in real GDP)

	1996	1997	1998	1999	2000	2001
Early year						
FOMC	2.1	2.1	2.4	2.8	3.6	2.3
Administration	2.2	2.0	2.0	2.0	2.9	3.2
CBO	2.1	2.1	2.3	1.8	2.9	2.6
Blue Chip	2.0	1.9	2.1	2.4	3.2	2.3
SPF	2.0	2.3	2.2	2.5	3.1	2.5
Mid-year						
FOMC	2.6	3.1	3.1	3.6	4.3	1.6
Administration	2.1	3.0	2.4	3.2	3.9	1.7
CBO	2.1	3.0	2.9	3.6	4.0	1.7
Blue Chip	2.6	3.2	3.0	3.5	3.9	1.8
SPF	2.8	3.1	3.0	3.2	4.2	1.5
Actual						
Third NIPA estimate	3.1	3.7	4.3	4.2	3.4	0.5
Currently published	4.1	4.3	4.8	4.3	2.3	0.1
H1 error						
Blue Chip	1.1	2.0	1.5	1.6	1.4	–0.6
SPF	1.4	1.2	1.3	0.5	1.9	–0.5
H2 revision						
Blue Chip	0.1	0.6	0.3	0.7	0.1	–0.8
SPF	0.1	0.5	0.3	1.0	0.2	–1.5

Notes: The *National Income and Product Account* (NIPA) estimate for the Q4-to-Q4 increase in real GDP in 1999 (published in March 2000) was 4.6 percent; the figure in the table is adjusted for the comprehensive revisions to the NIPA that occurred in December 1999. Currently published are the data published in the 2002 annual NIPA revision. H1 error and H2 revisions are percentage points, annual rate. Since mid-year Blue Chip forecast is from July, second-quarter data are not yet available; its first half error is based on actual for Q1 and July forecast for Q2.

Sources: Federal Open Market Committee (FOMC), 1979–2001, *Federal Reserve Board Monetary Policy Reports to Congress*; Administration, 1979–2001, *The Budget of the United States Government*; submissions and mid-session reviews, and 1979–2001, *Economic Report of the President*; Congressional Budget Office (CBO), 1979–2001, *The Economic and Budget Outlook*, submissions and mid-year updates; Blue Chip, 1978–2001, *Blue Chip Economic Indicators*, various issues; Federal Reserve Bank of Philadelphia, *Survey of Professional Forecasters* (SPF); and Actual: U.S. Bureau of Economic Analysis, *National Income and Product Accounts*.

Background—Forecasts during the early 1990s

The recovery from the 1990–91 recession was weak. Typically, the economy experiences a period of above-trend growth following a recession, as households and businesses catch up on postponed spending and inventories adjust to increases in demand. But, based on data in hand in mid-1992, output rose just 1.6 percent between 1991:Q1 and 1992:Q1, well below the average gain of roughly 5 percent recorded during the first year of the previous five expansions.

Many observers thought that "headwinds"—such as banks' efforts to meet capital standards and dislocations from the downsizing of the defense industry—were holding back the recovery. But even once these headwinds subsided, forecasters were not expecting much make-up for the lost growth. Instead, at 2.7 percent, the average of the early year forecasts for real GDP growth between 1992 and 1995 was just a bit above the generally prevailing views of the economy's long-run growth potential. And these forecasts were fairly accurate: Actual growth averaged 2.6 percent. At about 1.0 percent, the root mean squared forecasts errors (RMSE) of the forecasts were well below their longer-run averages (see table 3).

Forecasts during the second half of the 1990s

Given the relatively lackluster performance of the economy over the previous five years, forecasters entered the second half of the decade with modest expectations. In early 1996, real GDP was estimated to have increased less than 1.4 percent (annual rate) over the first three quarters of 1995.[3] Forecasters thought that some of this weakness would persist, and the early year projections for growth in 1996 were all close to 2 percent. Instead, according to the third NIPA estimates, real GDP rose 3.1 percent that year. Forecasters' early year projections for 1997 and 1998 were not much different from those in 1996—all looked for real GDP to rise between 1.9 percent and 2.4 percent. Some upped their projections three-tenths or four-tenths of a percentage point in 1999. But, in each year, output came in much stronger than expected: Real GDP rose 3.7 percent in 1997, 4.3 percent in 1998, and 4.2 percent in 1999.[4]

TABLE 2

Evolution of long-run forecasts
(percent change, annual rate)

	1996–98	1999	2000	2001	2002
Real GDP					
Administration	2.3–2.4	2.4	3.0	2.9	3.1
CBO	2.0–2.2	2.3	2.8	3.1	3.1
Blue Chip	2.3–2.5	2.5	3.1	3.4	3.2
SPF	2.3–2.5	2.5	3.1	3.3	3.0
Productivity					
Administration	1.2–1.3	1.3	2.0	2.3	2.1
CBO	1.1–1.5	1.8	2.3	2.7	2.2
Blue Chip	N.A.	N.A.	N.A.	N.A.	N.A.
SPF	1.3–1.5	1.6	2.4	2.5	2.1
		1991:Q4 1995:Q4	**1995:Q4 2000:Q4**		
Actual					
GDP		2.6	3.9		
Productivity		1.1	2.5		

Notes: Long-run forecasts are from early year Administration, Congressional Budget Office (CBO), and *Survey of Professional Forecasters* (SPF) forecasts and the March Blue Chip. Due to changes in reporting, the horizons used to determine the long run for the Administration and CBO forecasts vary somewhat over time. Actual data for 1991:Q4–95:Q4 are as published in March 1996; actual for 1995:Q4–2000:Q4 are as published in the 2002 annual NIPA revision. N.A. indicates not available

Sources: Federal Open Market Committee (FOMC), 1979–2001, *Federal Reserve Board Monetary Policy Reports to Congress*; Administration, 1979–2001, *The Budget of the United States Government*; submissions and mid-session reviews, and 1979–2001, *Economic Report of the President*; Congressional Budget Office (CBO), 1979–2001, *The Economic and Budget Outlook*, submissions and mid-year updates; Blue Chip, 1978–2001, *Blue Chip Economic Indicators*, various issues; Federal Reserve Bank of Philadelphia, *Survey of Professional Forecasters*; Actual: U.S. Bureau of Economic Analysis, *National Income and Product Accounts*; and U.S. Department of Labor, Bureau of Labor Statistics.

All told, the early year forecasts shown in table 1 underpredicted real GDP growth by between 0.9 and 2.4 percentage points during the 1996–99 period. Thus, the most obvious characteristics of these forecasts is that, in contrast to the 1992–95 period, the errors made during the second half of the decade were persistently positive and they were large.

These forecasts exhibit another interesting feature. The fact that forecasters did not make substantial changes to their GDP projections suggests that they thought the errors they were experiencing largely reflected transitory shocks or factors that would be offset by other developments. This view is supported by the mid-year forecasts shown in table 1. While these all generally looked for stronger growth than the early year projections, the differences largely reflect the incorporation of data in hand for the first half of the year. This can be seen using the quarterly forecasts made by Blue Chip and SPF. Table 1 presents the errors in the early year forecasts for real GDP growth in the first half of the year and the revisions made at mid-year to forecasts of second-half growth.[5] In 1996, 1998, and 2000, forecasters made large errors in the first half of the year but did not revise their second-half projections very much. Modest upward adjustments were made in 1997, but these still left the second-half forecasts below 2.7 percent. In 1999, the forecasters made more substantial upward revisions to their projections for growth in the second half of the year, pushing them above the 3 percent mark.

If most variations in GDP growth are viewed as transitory, then perceptions of longer-run trends in growth must be an important factor anchoring the annual GDP forecasts. Indeed, between 1996 and early 1999, the published assumptions for long-run growth were all in the range of 2 percent to 2.5 percent (table 2). And in each year, the early year forecasts for annual growth were generally just somewhat below these assumed longer-run trends. However, after four years of persistently strong growth and low inflation, in late 1999 and early 2000 forecasters began to boost their assumptions for long-run real GDP growth to around 3 percent. Thus, it probably is not a coincidence that around this time forecasters' mid-year projections also included a substantial upward revision to the projection of growth in the second half of the year.

Forecasts for 2000 and 2001

Forecasts made early in 2000 were looking for real GDP to rise between 2.9 percent and 3.6 percent that year—close to forecasters' updated perception of potential growth. In the event, growth in the first half of the year was quite robust. According to the estimates in hand at mid-year, real GDP advanced at an annual rate of 5.5 percent in the first quarter of the year and likely posted another healthy gain in the second quarter.

Forecasters again did not think this "extra" strength would persist. For example, the Blue Chip and SPF mid-year forecasts for growth in 2000:H2 were both about 3.3 percent, and forecasts made at this time for real GDP growth in 2001 averaged about that pace. But instead of simply settling down to trend, GDP growth surprisingly collapsed during the second half of 2000. According to the NIPA estimates available in March 2001, real GDP growth slowed to a 2.2 percent rate in 2000:Q3 and a 1 percent pace in 2000:Q4.[6]

In response, forecasters began to project slower growth, with most early-year forecasts for the increase in real GDP in 2001 running between 2.3 percent and 2.6 percent. By mid-2001, the economic picture had soured further, and forecasters marked their projections for growth down substantially. That said, the changes were not large enough. The mid-year forecasts were clustered in the range of 1.5 percent to 1.8 percent. Instead, according to revised estimates published in July 2002, real GDP changed little over the four quarters of 2001—and it fell at an average annual rate of 0.8 percent over the first three quarters of the year. Thus, despite the downward revisions, forecasters failed to predict the 2001 recession.

But, again, forecast errors are not the complete story. Notably, relative to the 1996–99 period, the projections for growth in 2001 were adjusted quickly. For example, the early year Administration forecast for 2001 was based on the data on hand as of the middle of 2000:Q4. It projected real GDP growth in 2001 would be 3.2 percent—about the same as the SPF and Blue Chip forecasts released in November 2000. However, over the next couple of months, the extent of the slowdown in the economy showed through more clearly in the monthly indicators of activity. As a result, the 2001 early-year FOMC, CBO, Blue Chip, and SPF forecasts—which were based on data available in late January or early February—all had been marked down to between 2.3 percent and 2.6 percent.

How unusual were the forecast errors during 1999–2001?

Clearly, forecasters made larger errors during the second half of the 1990s then they did during the first half of the decade. And while they reacted quickly to incoming information, they missed the sharp deceleration in activity in 2000 and 2001. But economic growth varies substantially over time, and the fluctuations are difficult to predict. Thus, one must ask whether these forecast errors were unusual or simply reflect inherent difficulties in forecasting.

The first two columns of table 3 show some sample statistics for the errors in the various forecasts calculated using data between 1980 and 1995. The mean errors for the early year forecasts are near zero, while their root mean square errors (RMSE) range between 1.3 percentage points and 1.7 percentage points. For reference, the standard deviation of real GDP growth over that period was about 2 percent. Furthermore, based on a simple regression of the current error on its lagged value, one cannot reject the hypothesis that the errors are uncorrelated across years. The mean errors for the mid-year forecasts also are near zero, and their RMSEs are between 0.9 and 1.3 percentage points.[7, 8]

In contrast, for every forecaster, *all* four early year forecasts made between 1996 and 1999 underpredicted real GDP growth. Furthermore, the errors were large: The average errors varied between 1.5 and 1.8 percentage points (table 3, column 3). For every forecaster, this average was greater than one RMSE of the forecast errors experienced during the 1980–95 period. The mid-year forecasts were only slightly better—they too, all underpredicted growth, with average errors between 0.7 and 1.2 percentage points.

How unusual were these errors in a statistical sense? Suppose that each year's forecast errors were drawn from independent *t*-distributions with means and variances as estimated using the 1980–95 data. (That is, *t*-distributions with means and standard errors as shown in the first two columns of table 3 and 16 degrees of freedom.) Because there is only about a one-in-six chance of experiencing a *single* draw greater than one standard deviation from these *t*-distributions, the odds of drawing *four consecutive* errors of this size from independent distributions are miniscule.[9] Indeed, none of the five forecasters ever made four consecutive same-signed errors in their early year forecasts during the 1980–95 period. And, on average, each forecaster experienced only two strings of *three* consecutive same-signed errors—and half of these occurred during recessions.

How unusual were the errors in 2000 and 2001? The far right column of table 3 indicates that, on average during 2000 and 2001, both the early year and mid-year forecasts overpredicted real GDP growth by nearly 1 percentage point. In addition, as noted earlier, the errors in the early year prediction of real GDP growth in 2001 were quite large, between 1.8 and 2.7 percentage points. However, the dynamics of an economy dipping into recession are quite different than one in expansion. Indeed, as we see in the fourth column, the average errors in 2000 and 2001 are not

TABLE 3

Forecast statistics for errors in current-year real GDP growth forecasts
(percentage points)

	1980–1995		Mean errors for:		
	Mean error	RMSE	1996–99	1980–82 1990–91	2000–01
Early year					
FOMC	–0.02	1.30	1.48	–0.69	–0.99
Administration	–0.29	1.67	1.78	–1.39	–1.11
CBO	–0.19	1.45	1.76	–1.22	–0.81
Blue Chip	–0.18	1.39	1.73	–1.29	–0.81
SPF	–0.11	1.44	1.58	–1.27	–0.86
Mid-year					
FOMC	0.02	1.22	0.71	–0.37	–0.99
Administration	–0.20	1.26	1.16	–0.78	–0.86
CBO	–0.26	0.92	0.93	–1.41	–0.91
Blue Chip	–0.07	1.30	0.76	–0.74	–0.91
SPF	0.14	1.01	0.80	–0.09	–0.89
Standard deviation of GDP growth	2.05		0.51	0.77	

Notes: RMSE are root mean square forecast errors. Errors and standard deviations of GDP growth are calculated using the third estimates of Q4-to-Q4 real GDP growth (adjusted for comprehensive NIPA revisions).

Sources: Federal Open Market Committee (FOMC), 1979–2001, *Federal Reserve Board Monetary Policy Reports to Congress*; Administration, 1979–2001, *The Budget of the United States Government*; submissions and mid-session reviews, and 1979–2001, *Economic Report of the President*; Congressional Budget Office (CBO), 1979–2001, *The Economic and Budget Outlook*, submissions and mid-year updates; Blue Chip, 1978–2001, *Blue Chip Economic Indicators*, various issues; Federal Reserve Bank of Philadelphia, *Survey of Professional Forecasters*; and U.S. Bureau of Economic Analysis, *National Income and Product Accounts*.

much different from those observed during the 1980, 1981–82, and 1990–91 recessions.

How unusual were the forecast procedures?

The results in the previous section suggest that the forecast errors during 1996–99 were drawn from a different distribution than they were, on average, during 1980–95. The question then arises whether this disparity reflects unusual behavior on the part of the forecasters or an unusual performance by the economy. This section addresses the first part of this question.

Typical evolution of GDP forecasts

How do GDP forecasts "typically" evolve over time? Given the qualitative descriptions above, restricting analysis to annual forecasts might hide some interesting reactions—or non-reactions—of higher-frequency forecasts to incoming data. Furthermore, longer-term projections appear to be an important part of the story. Only the private sector forecasts publish both quarterly and long-term forecasts. Accordingly, I analyze the Blue Chip consensus numbers released each March and October, the two months when respondents also are surveyed for long-term forecasts.[10] Note that the time gap between these months corresponds roughly with the interval between the early year and mid-year forecasts used above. And since the different annual forecasts track each other relatively closely, the patterns in these data likely generalize fairly well to the behavior of other forecasters. The appendix describes these data in more detail.

Given the periodicity of these forecasts, I consider semiannual time series of growth projections for half-year periods. Let $f_t^{gdp}(t+k)$ be the forecast made in period t for (annualized) real GDP growth in period $t + k$. For example, if t falls in the first half of the year (that is, the March forecasts) and $k = 1$, then $f_t^{gdp}(t+k)$ is the forecast for growth between the second quarter and the fourth quarter of the year. The available forecast horizons are $k = 0, 1, 2$. Let $f_t^{gdp}(lr)$ be the forecast of long-run growth made at time t. Alternatively, for any half-year period t, I have a sequence of three forecasts made in half-year periods— $t – 2$, $t – 1$, and $t -$ $f_{t-2}^{gdp}(t)$, $f_{t-1}^{gdp}(t)$, and $f_t^{gdp}(t)$—respectively. These latter forecasts are the bars plotted in the three panels of figure 2 (with the time grid identifying period t, the half-year being forecast). The solid line in each panel is the forecast of long-run growth, $f_{t-k}^{gdp}(lr)$, and the dashed line is actual half-year GDP growth (see appendix).

As we can see, in general, the one-year and one-half-year ahead forecasts do not differ substantially from the longer-run outlook (panels A and B). The standard deviations of the differences between these forecasts and the long-run projections are 0.7 percentage point and 0.8 percentage point, respectively;

for reference, these standard deviations are just about one-quarter the size of the average half-year growth forecast. However, at times, some large differences do open up. Some occur in the first half of the 1980s, when activity was projected to bounce back from the deep recessions in 1980 and 1981–82. Others are found during 1989 and 1990, when real GDP was correctly projected to grow well below trend. In contrast, forecasters' projections for growth in the current half-year period (panel C) often differ substantially from their long-term outlook. The standard deviation of the difference between $f_t^{gdp}(t)$ and $f_t^{gdp}(lr)$ is 1.5 percentage points, with differentials running as large as 3 to 4 percentage points during recessions and the recovery in 1983.

Figure 3 presents a couple of factors that may help explain the patterns in figure 2. Panel A plots $f_{t-2}^{gdp}(t)$ (bars), $f_{t-2}^{gdp}(lr)$ (solid line), and $f_{t-2}^{tbr}(t-2)-f_{t-2}^{tbr}(lr)$ (dashed line), the expected deviation of the real Treasury bill rate from its long-run average the year before the end of the forecast period. The figure suggests that high interest rates may have led forecasters to lower their year-ahead growth projections in the mid- and late 1980s. The converse appears to be true in 1993 and 1994. Panel B plots $f_{t-2}^{gdp}(t)$ and $f_{t-2}^{gdp}(lr)$ along with the most recent value of the Chicago Fed National Activity Index, or CFNAI (the dashed line), that would have been observed at the time the forecast was made. The CFNAI is a convenient way to summarize a large number of the regular monthly indicators that forecasters use to gauge the current pace of economic activity.[11] (Note that a CFNAI value of zero corresponds with the indicators growing at their long-run averages.) In general, there does not appear to be much correlation between the CFNAI and the longer-run forecasts, with the possible exception of a negative correlation when projecting a recovery from recession. In contrast, forecasts for the current semiannual period do appear to change substantially in conjunction with such data. As we see in panel C, $f_t^{gdp}(t)$ often deviates from $f_t^{gdp}(lr)$ in the direction indicated by the movements in the CFNAI. The largest deviations are found in and around recessions.

Quantifying the forecast processes

This section estimates a couple of simple regression models in order to provide some rough quantification of the patterns exhibited in figures 2 and 3.

The first model considers how forecasts for growth over half-year periods differ from the outlook for longer-term GDP growth. The regression is:

$$f_t^{gdp}(t+k)-f_t^{gdp}(lr)=a+b_1[f_t^{tbr}(t+k-2)-f_t^{tbr}(lr)]+b_2 CFNAI_t^{nr}(t-1)+b_3 CFNAI_t^{r}(t-1)+u_t^{gdp}(t+k),$$

where k = 0, 1, 2 and the regressors are

1) $f_t^{tbr}(t+k-2)-f_t^{tbr}(lr)$: the difference between the real Treasury bill rate and the long-run value expected to be in place one year before the end of the forecast period;
2) $CFNAI_t^{nr}(t-1)$: the most recent value of the CFNAI known at the time the forecast was made if it is greater than –0.7; and
3) $CFNAI_t^{r}(t-1)$: the most recent value of the CFNAI known at the time the forecast was made if it is less than –0.7.

The *nr* and *r* superscripts refer to "no recession" and "recession" CFNAI values. This dichotomy is to address the observation that forecasters may react differently to incoming data in and around recessions. The boundary point is taken from Evans, Liu, and Pham-Kanter (2002); as they discuss, historically, when the CFNAI falls below –0.7, there is about a 70 percent chance that the economy is in recession.

The second model considers forecast revisions; that is, how forecasters *change* their projection for a particular semiannual period in light of recent forecast errors or other information that they learn between time t – 1 and time t. For the change in the forecast for real GDP growth in the current half-year period t, the model is:

$$f_t^{gdp}(t)-f_{t-1}^{gdp}(t)=a+b_1 rev_t^{gdp}(t-2)+b_2 err_t^{gdp}(t-1)+b_3 err_t^{tbr}(t-1)+b_4 err_t^{CFNAI}(t-1)+u_t^{gdp}(t),$$

where

1) $rev_t^{gdp}(t-2)$: the revision made between period t – 1 and t in the published estimate of real GDP growth over half-year t – 2;
2) $err_t^{gdp}(t-1)$: the error in the forecast made at time t – 1 for real GDP growth over half-year t – 1 based on actual GDP data available in period t;
3) $err_t^{tbr}(t-1)$: the error in the forecast made at time t – 1 for the (quarterly) real T-bill rate at the end of half-year t – 1; and

FIGURE 2

Evolution of Blue Chip half-year real GDP growth forecasts

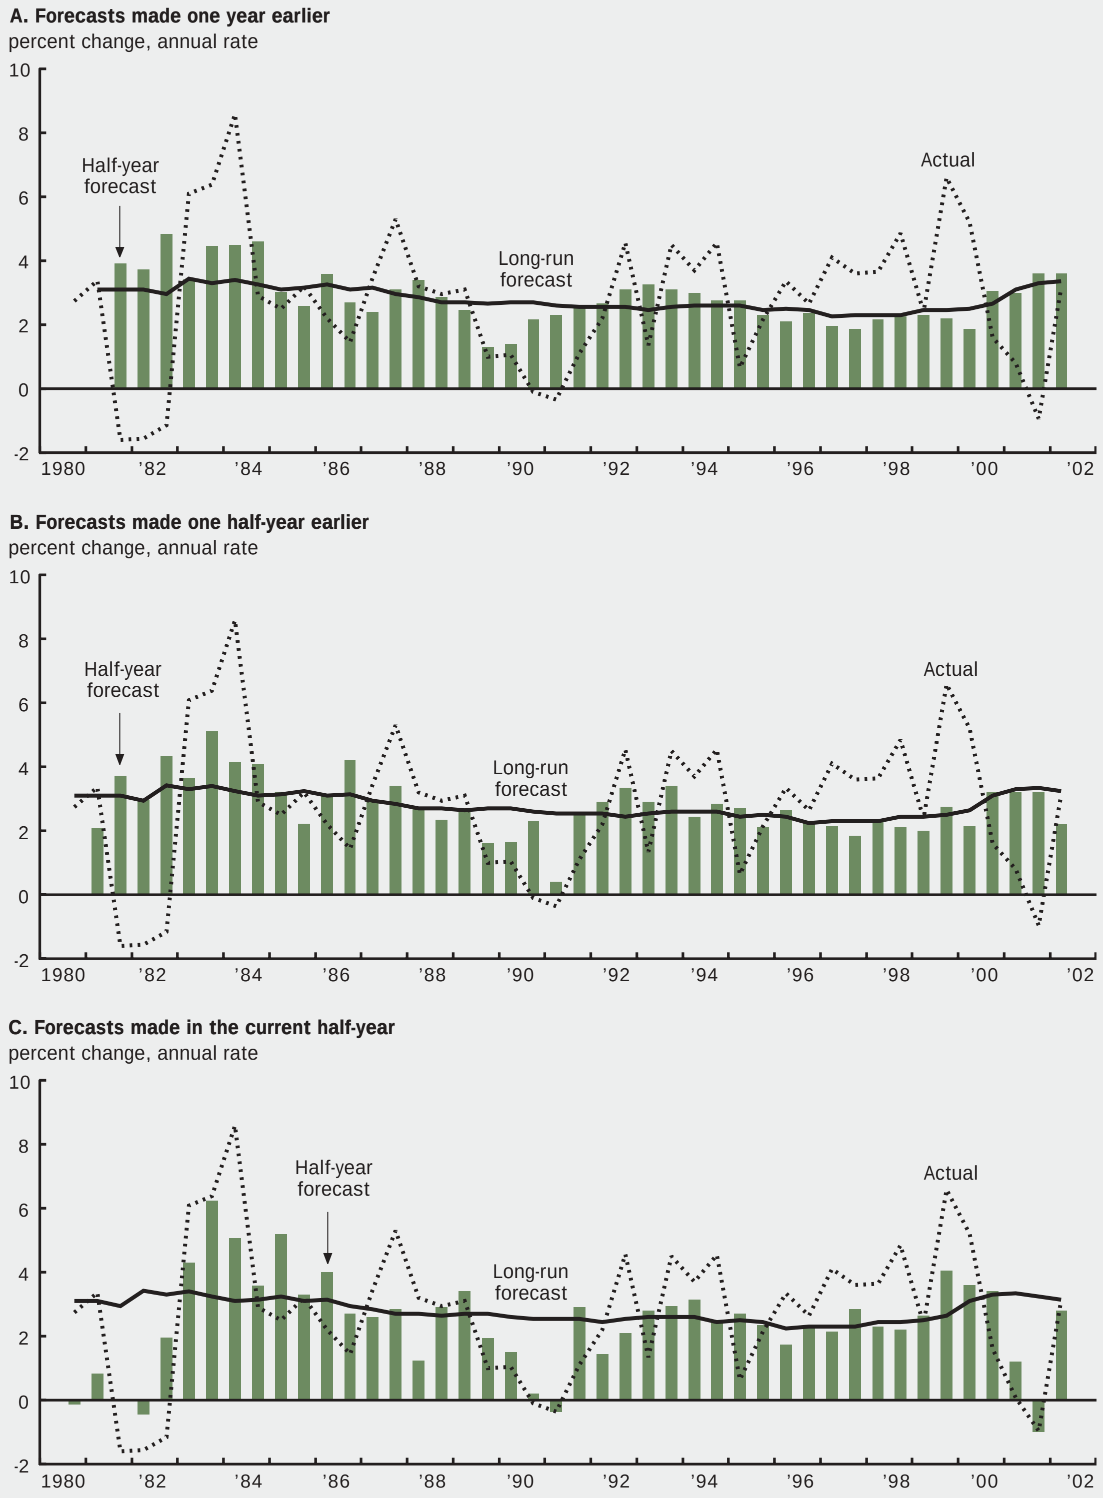

Sources: Blue Chip, 1978–2002, *Blue Chip Economic Indicators*, various issues; and U.S. Bureau of Economic Analysis, *National Income and Product Accounts*.

FIGURE 3

Interest rates, current activity, and Blue Chip half-year real GDP growth forecasts

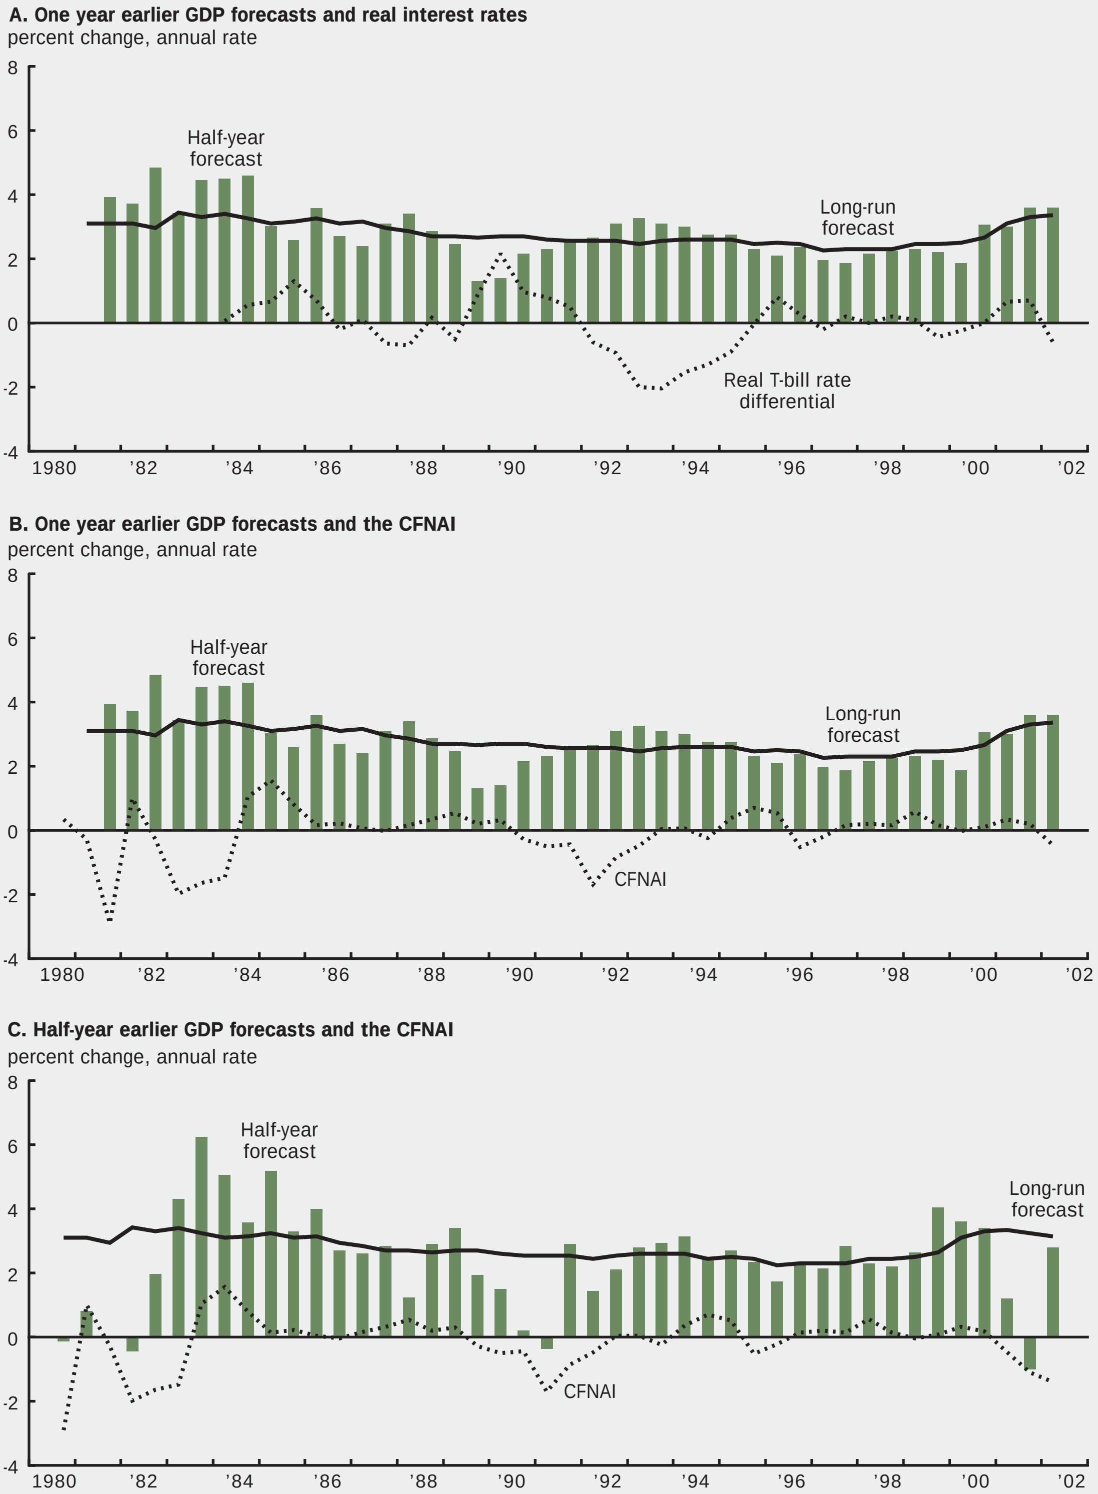

Note: To smooth inherent volatility, the three-month moving average of the CFNAI, which is designated CFNAI-MA3, is plotted in the figure. The real T-bill rate differential if the difference between the real T-bill rate and its long-run expectations (see appendix).

Sources: Federal Reserve Bank of Chicago, CFNAI; and Blue Chip, 1978–2002, *Blue Chip Economic Indicators*, various issues.

4) $err_t^{CNFAI}(t-1)$: the "shock" in the CFNAI learned at time *t*. This is the residual from a simple AR(2) model predicting the most recent value of the CFNAI that would be known at the time the period-*t* forecast of GDP is made.

I ran a similar equation for the real T-bill forecast. The equations for the period-*t* revisions in the longer-horizon forecasts (k = 1, 2, *lr*) are:

$$f_t^{gdp}(t+k) - f_{t-1}^{gdp}(t+k) = a + b_1 rev_t^{gdp}(t-2) +$$
$$b_2 err_t^{gdp}(t-1) + b_3 err_t^{tbr}(t-1) + b_4 err_t^{CFNAI}(t-1) +$$
$$\sum_{j<k} b_{5j} u_t^{gdp}(t+j) + \sum_{j<k} b_{6j} u_t^{tbr}(t+j) + u_t^{gdp}(t+k).$$

The extra terms in these regressions—the residuals from the shorter-horizon equations—test whether unaccounted for factors that generate revisions in forecasts for earlier time periods are expected to persist and affect growth in the farther out quarters. This is similar to tracking impulse responses in a vector autoregression.

The results for the first model are shown in table 4. As indicated by the $\bar{R}^2$ values, the interest rate deviations and the CFNAI explain more than 60 percent of the variation in the difference between $f_t^{gdp}(t)$ and $f_t^{gdp}(lr)$, but only about 20 percent of that in $f_t^{gdp}(t+2) - f_t^{gdp}(lr)$ and *none* in $f_t^{gdp}(t+1) - f_t^{gdp}(lr)$. As seen in the top row, a positive interest rate differential appears to be taken as a signal of strong activity in the near term, but causes forecasters to lower their one-year-ahead forecasts below $f_t^{gdp}(lr)$. The CFNAI terms indicate that current half-year forecasts are significantly raised or lowered relative to the long-term outlook in reaction to good or bad readings on incoming high-frequency indicators of activity. And the larger coefficient on $CFNAI_t^r(t-1)$ than $CFNAI_t^{nr}(t-1)$ indicates that the responses are bigger when the economy appears to be falling into recession. But the medium-term forecasts react little to the incoming data, the exception being that if the economy currently is in a recession, then forecasters will tend to predict a period of above-trend growth at the one-year-ahead horizon.

The results from the second model are shown in table 5. As shown by the $\bar{R}^2$ values, these factors explain about 40 percent of the variation in revisions to current and one-half-year-ahead forecasts but little of the changes to longer-run forecasts. Consistent with the first model, much of the explanatory power for the one-quarter-ahead revision comes from the shock to the CFNAI, but this shock has little predictive power for revisions to the out-quarter forecasts. None of the projections are revised much in response to the most recent GDP forecast error. And with the possible exception of the half-year-ahead forecast, the reactions to the portion of earlier GDP revisions not explained by the model are small. Errors and revisions in the outlook for the T-bill rate have at most a small influence on the GDP forecast revisions.

Together, these models suggest that projections of real GDP growth beyond the next couple of quarters usually do not vary far from forecasters' long-run growth outlook; the exceptions are when events such as recessions or changes in monetary policy come into play. Forecasters may make large revisions to near-term projections for real GDP growth in response to incoming high-frequency data, but the average responses to past GDP and interest rate forecast errors and revisions are small. These results suggest that forecasters think that most of the "shocks" revealed in incoming

TABLE 4

Explaining deviations in half-year forecasts of real GDP growth from the long-run forecast

		Forecasts for growth over the two quarters ending:		
		Current half year	Half year ahead	One year ahead
Regression on:				
$f_t^{tbr}(t+k-2) - f_t^{tbr}(lr)$		0.24	–0.12	–0.20
		(2.21)	(–1.22)	(–2.39)
$CFNAI_t^{nr}(t-1)$		1.46	0.30	0.12
		(4.37)	(1.04)	(0.54)
$CFNAI_t^{r}(t-1)$		2.01	–0.03	–0.42
		(6.32)	(–0.10)	(–2.26)
$\bar{R}^2$		0.61	0.00	0.18
Std. dev. of $f_t^{gdp}(t+k) - f_t^{gdp}(lr)$		1.30	0.70	0.58
1982–95:	Mean error	0.08	0.08	0.04
	RMSE	0.70	0.75	0.57
1996–99:	Mean error	–0.06	–0.21	–0.14
	RMSE	0.60	0.31	0.31
2000–02:	Mean error	–0.34	–0.11	0.02
	RMSE error	1.37	0.63	0.32

Notes: T-statistics in parentheses. Semiannual Blue Chip data, 1982:H2 to 2002:H1. RMSE are root mean square forecast errors.
Sources: Federal Reserve Bank of Chicago, CFNAI; Blue Chip, 1978–2002, *Blue Chip Economic Indicators*, various issues; and U.S. Bureau of Economic Analysis, *National Income and Product Accounts*.

TABLE 5

Explaining revisions to forecasts of real GDP growth

	Forecast for growth over the two quarters ending:			
	Current half year	Half year ahead	One year ahead	Long-run
Regression on:				
$rev_t^{gdp}(t-2)$	0.48 (1.78)	0.60 (4.41)	–0.24 (–1.58)	0.07 (2.37)
$err_t^{gdp}(t-1)$	0.08 (0.71)	–0.01 (–0.23)	–0.10 (–1.70)	–0.01 (–0.93)
$err_t^{tbr}(t-1)$	–0.21 (–0.64)	0.23 (1.55)	0.04 (0.26)	0.17 (2.77)
$\sum u_t^{gdp}(t+j)$		0.16 (1.92)	0.11 (0.46)	0.07 (1.20)
$\sum u_t^{tbr}(t+j)$		0.02 (0.14)	–0.06 (–0.12)	0.08 (0.58)
$err_t^{CFNAI}(t-1)$	1.48 (4.67)	0.18 (1.29)	–0.12 (–0.72)	–0.04 (–1.20)
$\bar{R}^2$	0.38	0.38	0.14	0.24
Std. dev. of $f_t^{gdp}(t+k) - f_{t-1}^{gdp}(t+k)$	1.29	0.57	0.55	0.10
1982–95: Mean error	–0.03	–0.03	0.05	–0.02
RMSE	0.83	0.45	0.47	0.08
1996–99: Mean error	0.36	0.22	–0.19	0.04
RMSE	0.61	0.34	0.44	0.07
2000–02: Mean error	–0.38	–0.15	0.13	0.04
RMSE	1.85	0.21	0.39	0.07

Notes: T-statistics in parentheses. Semiannual Blue Chip data, 1982:H2 to 2002:H1. RMSE are root mean square forecast errors.
Sources: Federal Reserve Bank of Chicago, CFNAI; Blue Chip, 1978–2002, *Blue Chip Economic Indicators*, various issues; and U.S. Bureau of Economic Analysis, *National Income and Product Accounts*.

monthly data or recent errors have only a transitory influence on real GDP growth or will be offset by other factors. Those shocks that are more persistent could be expected to elicit a policy response that would have an influence on output at a longer horizon. Indeed, in qualitative terms, the characterization of the GDP forecast process provided by these two simple models is consistent with the time-series evidence—such as that generated by structural vector autoregression (VAR) models—regarding the response of real GDP to various shocks (see appendix).

Were the forecasts in 1996–2001 unusual?

The above statistical description appears consistent with our earlier qualitative characterization of forecasts during the 1996–2001 period. Notably, as seen in figure 3, the Blue Chip one-year-ahead forecasts for real GDP growth in 1996–2000:H1 were a bit *lower* than the long-run projections. Thus, forecasters were not carrying earlier underpredictions or forecast revisions forward into higher projections for GDP growth in the out quarters. Indeed, forecasters were expecting other factors—such as external shocks from the Asian crisis in 1997 and the Russian default in 1998—to hold back growth. Not until 2000, when long-term forecasts were increased, do we see a boost in $f_{t-2}^{gdp}(t)$ and $f_{t-1}^{gdp}(t)$. Furthermore, the substantial downward revisions in $f_t^{gdp}(t)$ in 2000 and 2001 appear consistent with the declines in the CFNAI.

Supporting these qualitative descriptions, the errors in our simple equations describing the forecast process were not that different from those experienced prior to 1996. (Though given the quite weak explanatory power of these models, the analysis of errors only provides suggestive evidence.) As indicated by the average errors in the bottom portion of table 4, forecasts during 1996–99 were a bit lower than the first model predicts. Similarly, near-term forecasts were revised up a bit more than was typical (as shown in the bottom of table 5). However, in both cases, the differences are at most a few tenths of a percentage point on GDP growth and are not statistically significant. For the 2000–02:H1 period, the near-term forecasts are 0.3 to 0.4 percentage point lower than predicted by the models, but these errors are small relative to the revisions between $f_{t-1}^{gdp}(t)$ and $f_t^{gdp}(t)$ in 2000 and 2001.[12]

FIGURE 4

Nonfarm business productivity

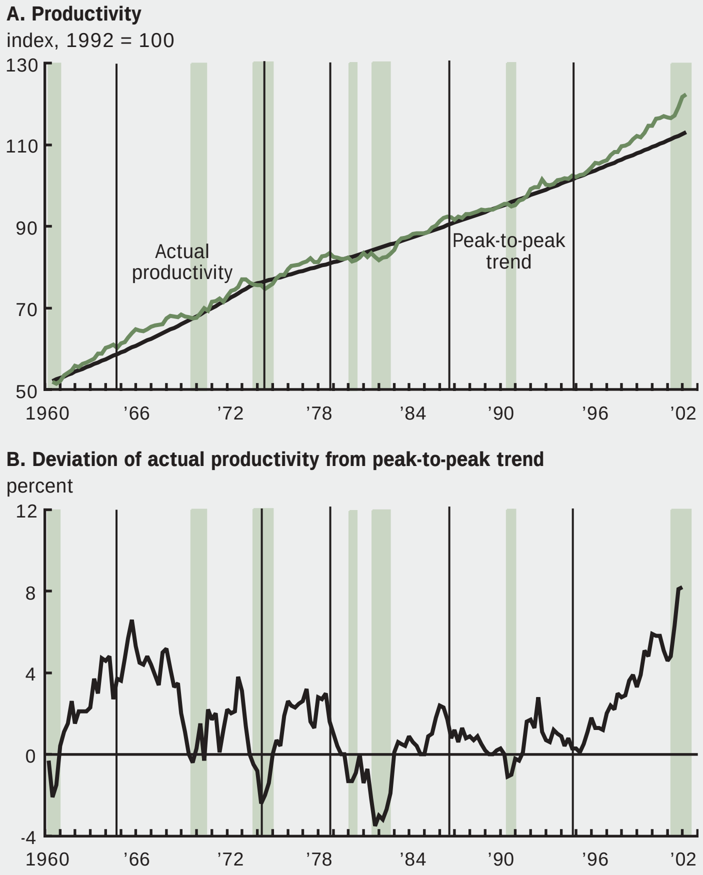

Notes: The vertical black lines denote the four-year mark after the end of the previous recession. The shaded areas indicate recessions as defined by the National Bureau of Economic Research.
Source: U.S. Department of Labor, Bureau of Labor Statistics.

was expected to run in the 1 percent to 1.5 percent range, about the same as the pace that had prevailed since the early 1970s. Demographic projections (not shown) showed the working age population rising about 1 percent per year, which was thought to translate into like-sized increases in hours worked. This left the projections for long-run GDP growth in the range of 2 percent to 2.5 percent.

Because long-run forecasts anchor the medium-term outlook, changes in productivity trends have important implications for the forecasting exercise. The colored line in panel A of figure 4 plots the level of productivity (output-per-hour) in the nonfarm business sector. The black line is the simple trend of productivity between business cycle peaks.[13] As we can see, productivity is quite cyclical—it typically falls during a recession (or period of weak growth) and rises sharply early in a recovery. But productivity rarely accelerates persistently during a mature business cycle. The vertical black lines in the figure denote the four-year mark after the end of the previous recession, while panel B plots the (percent) deviation in actual productivity from the peak-to-peak trend. As we can see, the only previous time that productivity remained well above trend four years into the expansion was during the late 1960s. But even then, the gap between actual and trend productivity was not increasing—that is, actual productivity growth was proceeding at its peak-to-peak trend. In contrast, in the mid-1990s, productivity growth picked up markedly and persistently outstripped earlier trends. The four-quarter increase in output-per-hour exceeded the 1.4 percent peak-to-peak trend that prevailed between 1980 and 1990 in *every* quarter between 1996:Q1 and the cyclical peak in 2001:Q1. The average growth rate of productivity over this period was 2.5 percent.

Even now, determining how much of the pick-up was transitory, though long-lived, and how much of it represented a permanently higher trend is a difficult task. Almost by definition, a change in the trend cannot be identified until we have observed a substantial amount of data following the break. Indeed,

Unusual behavior of the economy

Given that forecasters seemed to be conducting business as usual, the question is what economic developments made forecasting so difficult? It is beyond the scope of this article to catalog the vast number of factors—and forecasters' perceptions of them—that influenced the economy over 1996–2001. Instead, I focus on two related developments: the step-up in productivity growth and the boom and bust in business investment. Both of these were inherently difficult to predict. And both had important implications for GDP forecast errors during this period.

Acceleration in productivity

The trend in labor productivity is one of the fundamental determinants of long-run growth. As we see in table 2, in the mid-1990s, productivity growth

FIGURE 5

Capital services, investment, and productivity

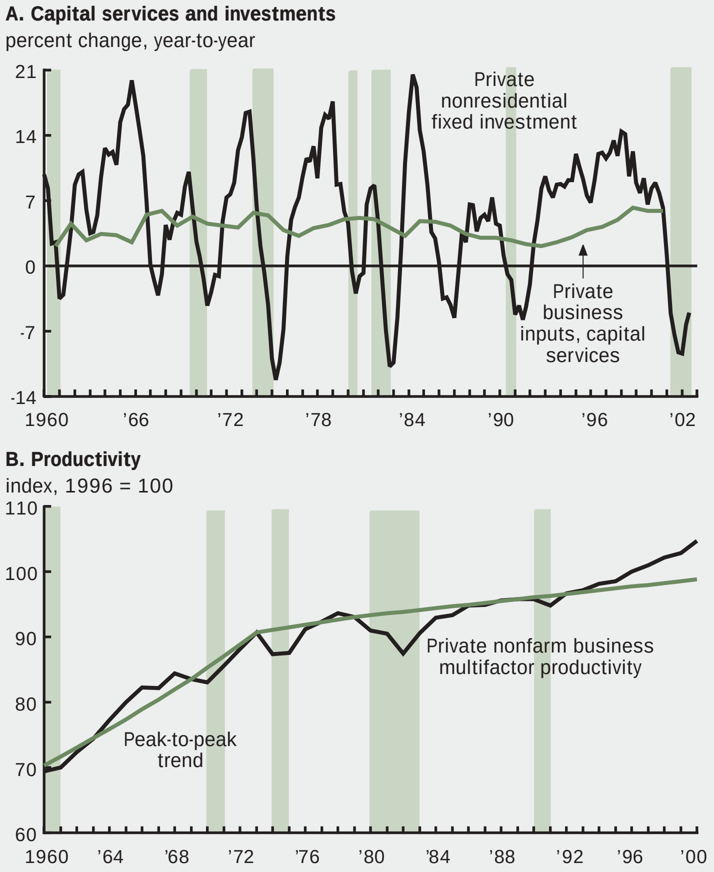

Note: The shaded areas indicate recessions as defined by the National Bureau of Economic Research. Panel A plots four-quarter percent changes in private nonresidential fixed investment and year-to-year changes in capital services.

Sources: U.S. Department of Labor, Bureau of Labor Statistics and U.S. Bureau of Economic Analysis.

as late as 1999, forecasters’ estimates of the economy’s longer-run trends in productivity growth remained between 1.3 percent and 1.8 percent.

Eventually, however, a confluence of corroborating evidence led forecasters to change their expectations. The fact that the high GDP growth was associated with low unemployment and subdued inflation indicated that the economy’s productive resources were not being strained. The economy also had proved unexpectedly resilient to external shocks. Furthermore, forecasters found themselves underpredicting every major component of domestic private demand, suggesting that the source of strength was some broad-based phenomenon as opposed to a sector-specific shock.[14] Finally, as discussed below, a good deal of the increase in productivity growth appeared to reflect sources that could prove to be persistent. To be sure, a great deal of uncertainty remained regarding how much of the pick-up in productivity reflected a permanently higher trend (see Gordon, 2000). But, by 2001, most of the forecasters had raised their assumptions for the trend growth in productivity to the 2.3 percent to 2.7 percent range. Correspondingly, they boosted long-run growth forecasts for real GDP growth to the 3 percent to 3.5 percent range. These long-run assumptions became a new anchor for nearer-term forecasts.

Increases in capital and information technology

One reason that forecasters changed their views of the trends in productivity is that some of the important factors underlying the gains were thought likely to be long-lived. In particular, a good deal of the step-up in growth that occurred in the second half of the 1990s reflected intensified capital deepening and developments in the information technology (IT) sector. Once in place, capital does not disappear, and the longer-run prospects for IT were quite optimistic.

Jorgenson and Stiroh (2000) and Oliner and Sichel (2000) both estimated that about half of the acceleration in productivity between the first and second halves of the 1990s was due to capital deepening—or an increase in the quality and quantity of capital used per hour worked. Surges in capital deepening often reflect cyclical weakness in hours. But this time the gains were due to a sustained pick-up in capital services, a measure of the productive input provided by the total business capital stock in the economy. Panel A of figure 5 shows the growth rates of aggregate capital services (colored line) along with business fixed investment (black line). Growth in capital services had edged down from 4.7 percent in 1985 to 2.1 percent by 1992, but a surge in investment in the 1990s boosted its growth to about 6 percent by the end of the decade.[15]

Indeed, the large gains in investment depicted in the figure account for a good deal of the pick-up in overall real GDP growth during the 1996–99 period. According to the July 2002 revised NIPA data, after increasing at an average annual rate of about 5 percent between the cyclical peak in 1990:Q3 and 1995:Q4, real business fixed investment (BFI) rose at about an 11 percent annual rate between 1995:Q4 and 2000:Q2. As a result, BFI moved from boosting real GDP growth by an average of about 0.5 percentage points per year during the first half of the 1990s to raising it between 0.8 and 1.5 percentage points per year during the second half of the decade.

Technology also was an important factor in the productivity acceleration. The studies cited above also estimate that between 60 percent and 100 percent of the increase in capital deepening reflected increases in the quantity and quality of high technology capital used by labor. Changes in technology also influence output per hour through other channels. Multifactor productivity refers to increases in output per hour that cannot be attributed to capital deepening or changes in labor quality. As we see in panel B of figure 5, multifactor productivity also exhibited an unusually sharp acceleration in the second half of the 1990s. Both Jorgenson and Stiroh and Oliner and Sichel calculated that improvements in the production of IT products made substantial contributions to this acceleration in multifactor productivity. And more recent estimates by these authors using up-to-date data point to even larger IT contributions to the acceleration in overall productivity in the second half of the 1990s.

Collapse of investment and decline in activity in 2000 and 2001

Even though forecasters boosted their views regarding the longer-run prospects for the economy, they expected several factors to moderate GDP growth in 2000 and 2001.[16] All told, forecasters believed that these factors would bring GDP growth down to its longer-run potential, but would not be sufficient to tip the economy into a recession.

However, as already noted, by their very nature, recessions are periods of unusual economic activity and are therefore hard to predict. This time, as shown in panel A of figure 5, the demand for capital equipment suddenly and surprisingly collapsed in the second half of 2000. In particular, in the high-tech area, bookings for capital equipment fell sharply, inventory–sales ratios backed up, and industrial production began to drop. Instead of the solid 10.5 percent annual rate increase projected by the SPF in August, BFI barely changed in the second half of 2000. In February 2001, the SPF forecast real BFI to increase 4.5 percent over the four quarters of the year; instead, according to the third NIPA estimates, it fell 9.4 percent. Similarly, the pace of inventory investment did more than just moderate; by 2001 firms were liquidating inventories at a sharp rate.

According to the July 2002 revised NIPA data, real BFI swung from double-digit gains to dropping at an average annual rate of 6.3 percent between 2000:Q2 and 2001:Q4. As a result, BFI reduced real GDP growth by 1.2 percentage points in 2001—a negative swing of 2 to 2.6 percentage points relative to its contributions to growth during the second half of the 1990s. Spending on high-technology equipment, which represents about one-third of total BFI, accounted for a good deal of this swing. Changes in inventory investment went from being, on balance, a neutral influence on GDP growth in 1999 and the first half of 2000 to reducing it by nearly 1.5 percentage points in 2001. In contrast, slower growth in all other sectors of the economy—with a share of about 85 percent—reduced real GDP growth by just about 1 percentage point between 2000 and 2001.

Investment and the adjustment of capital stocks

Thus, the forecast miss in GDP seemed to have been precipitated by a sudden swing in investment, followed by a sharp correction in inventories. Even though some stock adjustment had been anticipated, the extent of the drop-off clearly was underestimated. Why are such swings in investment so hard to forecast?

Some simple arithmetic regarding capital stocks and flows provides a useful way to frame the discussion. For any particular type of capital, call it the *i*th type,

$$K_t^i = K_{t-1}^i - \delta_t^i K_{t-1}^i + I_t^i,$$

or

$$I_t^i / K_{t-1}^i = g_t^i + \delta_t^i,$$

where I_t^i is investment, K_t^i is the end-of-period capital stock, δ_t^i is the depreciation rate, and g_t^i is the growth rate of this component of the capital stock. The simple arithmetic of this equation is: 1) if g_t^i and δ_t^i are relatively stable in the long run, then so will be I_t^i / K_{t-1}^i, meaning that investment and the capital stock will be growing at the same rate, g_t^i; and 2) to increase g_t^i, investment must grow faster than the capital stock for some period in order to boost I_t^i / K_{t-1}^i. Conversely, to lower g_t^i, investment will have to grow slower than capital for some time.

Suppose technological innovation makes some type of capital more productive, for example, a new chip makes computers more powerful. Businesses will want to raise the growth rate of computer capital to take advantage of the higher marginal value of the new computers. In order to do so, for some time *investment* in computers would have to *increase* at a higher rate than that of the computer capital stock. As the higher desired capital is achieved, growth in investment will fall. But to what rate? To the degree the innovation reflects a permanent change in the *growth rate* of technology, growth in both capital and investment will settle at a new higher g_t^i. To the extent that it is a *one-time step-up* in technology, growth will fall back to the original g_t^i.[17] The basic logic of this discussion extends to describing the behavior of aggregate investment and capital.

Gauging growth in investment and capital stock in 1996–2001

The arithmetic presented above indicates that in order to pin down the path for investment, forecasters—at least implicitly—have to make some judgment concerning the persistence of any observed pick-up in capital growth. Such decisions clearly were important during the 1996–2001 period. As we noted earlier, capital growth was spurred by the desire to incorporate advances in technology, boosting the growth in capital services, g_t, to around 6 percent by the end of the decade. The February 2000 SPF forecast projected that real BFI would increase about 8 percent that year—and some of this gain reflected spending that was thought to have been deferred due to Y2K. Thus, this forecast for the underlying rate of increase in real BFI was not far from the pace of growth in capital services. Such a projection produces a constant $I/K = g + \delta$, the equilibrium condition for stable growth in investment and capital.

In other words, it appears forecasters had come to believe that we had experienced a long-lived increase in the *rate* of advance in technology that should generate a persistent increase in the rate of growth in capital and in the investment spending to support this growth. But, given the magnitude of the swing in investment, it seems that forecasters overestimated where the growth rate of capital would settle over the medium term.[18] What happened to the determinants of capital stock growth that may have caused this miss?

Around this time, both the players in financial markets and the businesses making capital spending decisions appear to have reevaluated the earnings potential of certain investment projects. The deceleration in BFI was preceded in early 2000 by a decline in stock prices. In both the equity markets and investment, the retrenchments were particularly dramatic in the high-technology sectors—just as these sectors had led the surge on the upside. Whatever its root cause, such a reassessment clearly was a negative for new investment projects. And to the extent that expected payoffs to capital projects already undertaken were revised down, earlier investment may have pushed the capital stock to a level that, in retrospect, was too high. This would imply a period of below-trend growth in the capital stock and an even sharper retrenchment in investment in order to realign stocks with desired levels.

Could this reassessment have been predicted? To be sure, by conventional historical standards, equity valuation metrics—such as price–earnings ratios or dividend–price ratios—were at unprecedented levels in early 2000. And the high rates of investment had substantially pushed up growth in capital. Many commentators argued that these facts meant that the stock market was "overvalued" and that firms had overbuilt productive capacity. Based on these observations, one might have thought that a "bursting of the bubble" would lead to weak activity 2000 and 2001.

But actually forecasting such an event is problematic. Throughout the second half of the 1990s, stock market valuation metrics had been continuously attaining new historical records, and some observers had been continuously predicting market corrections (see for example, Campbell and Shiller, 2001). Yet equity markets kept moving up and investment surged further. Forecasters who may have lowered their earlier projections due to such reservations also would have underestimated the strength of the economy to an even greater degree than the consensus did in the late 1990s. Indeed, when it came to writing down numbers, even the more pessimistic forecasts did not predict outright declines in GDP in 2000 and 2001. In February 2000, the average of the lowest ten Blue Chip forecasts still had real GDP rising 2.2 percent that year, and this group even boosted their outlook to 3.3 percent in July. Even after the stock market declines and weak investment indicators during the second half of 2000, as of February 2001 the bottom-ten Blue Chip average forecast that real GDP would increase 1.2 percent that year. And in July, the pessimists still thought that output would rise at about a 1 percent annual rate in the second half of the year.

Conclusion: Implications for future forecasts

Because up-to-date estimates are not yet available (see note 15), we cannot look at the decomposition of productivity to see how growth in capital services or multifactor productivity has performed in recent quarters. However, as figure 4 shows, growth in total labor productivity has been very well maintained. Between the cyclical peak in 2001:Q1 and 2002:Q3, growth in output per hour has averaged a strong 4 percent annual rate.[19] This performance more resembles the cyclical patterns around the 1960 and 1969 recessions, when productivity trends appeared to be nearly 3 percent, than the behavior of output per hour around the recession between 1973 and 1990, when productivity trends were closer to 1.25 percent to 1.5 percent.[20]

A number of researchers have made rough estimates of what might be reasonable steady-state values to expect for growth in output per hour. As summarized in Oliner and Sichel (2002), the numerous scenarios considered in these papers produce a range of values between 1.3 percent and 3.2 percent, with point estimates largely between 2 percent and 2.8 percent. Thus, while a return to pre-1995 rates can not be ruled out, most analysts are guessing that the economy will experience higher productivity growth in the long run.

These estimates leave us with a relatively optimistic view about productivity trends going forward. In line with this perception, long-run forecasts for real GDP growth have not changed much over the past couple of years. The most recent assumption for long-run growth, made in October 2002 by the Blue Chip consensus, was 3.2 percent. Accordingly, despite the recession, and, to date, bumpy recovery, forecasters still are anchoring their cyclical projections for real GDP growth with solid trends in the underlying long-run pace of growth in economic activity.

NOTES

[1]The average forecasts plotted in figure 1 are the averages of the early year projections made by the Federal Open Market Committee (FOMC) and other Federal Reserve Bank presidents, the Administration, the Congressional Budget Office (CBO), the Blue Chip Consensus, and the median forecast from the Federal Reserve Bank of Philadelphia's *Survey of Professional Forecasters.*

[2]The Federal Reserve publishes a range and central tendency of forecasts made by the FOMC members and other Bank presidents. I use the middle of the central tendency as the FOMC point forecast. Other details regarding the data are available from the author upon request.

[3]These figures are the growth estimates that were available in mid-January 1996. Comprehensive revisions to the NIPA were published that month, but they covered data only through 1995:Q3; estimates for 1995:Q4 were delayed until March, after the early year forecasts for 1996 were made.

[4]The 1996, 1997, and 1998 figures are from the third NIPA estimates for growth in those years. The October 1999 comprehensive revisions to the NIPA added business expenditures on software to the estimates of business fixed investment. The BEA estimates this added 0.41 percentage point to average real GDP growth between 1992 and 1998. The 1999 growth figure cited above is the third estimate less 0.41 percentage point.

[5]First-half errors are estimated using the information available at the time of the mid-year forecast. The mid-year SPF forecasts are made in August, so the actual values used to calculate the first-half error are the first estimates of growth in the second quarter. The mid-year Blue Chip forecasts are made in July, before second-quarter data are available; the first-half Blue Chip "error" is thus calculated using the actual value for GDP in for the first quarter and the revision made between early year and mid-year in the forecast for second-quarter GDP growth.

[6]The revised NIPA estimates published in July 2002 paint a somewhat different picture of these developments. Real GDP growth during the first half of the year was revised down from 5.2 percent to 3.8 percent, and the increase in the third quarter is now estimated to be just 0.6 percent (annual rate). The estimate of real GDP growth in 2000:Q4 is still about 1 percent.

[7]A large literature exists that examines the performance of macroeconomic forecasts; see for, example, Berger and Krane (1984), McNees (1992, 1995), Romer and Romer (2000), Schuh (2001), and the references cited in these papers. Many papers conduct formal statistical tests of forecast efficiency. One criterion for efficiency is that forecast errors should be independent of information known at the time a forecast was made, which includes the lagged forecast error. Schuh rejects the efficiency of annual SPF forecasts of GDP, though the rejection is due to correlation with variables other than the lagged GDP forecast error.

[8]The 1980–95 period includes three recessions, 1980, 1981–82, and 1990–91. Excluding these years from the calculations, the mean errors of the early year forecasts are between 0.2 and 0.4 percentage point and the RMSEs are between 1.1 and 1.5 percentage points. For the mid-year forecasts, the means are between 0.1 and 0.3 percentage point and the RMSEs in the 0.6 to 0.8 percentage point range.

[9]Cumulative sum (CUSUM) plots also suggest a structural break in the distributions of the errors during this period. Recursive *t*-tests (see Harvey, 1989) using the 1996–99 errors easily reject that the errors have a zero mean when the tests are constructed using the standard deviation of the errors over the four-year period. However, the recursive *t*-tests only reject at between the 6 percent and 9 percent level if the standard deviation of the errors over the 1980–95 period is used. Finally, Schuh also finds that the average forecasts from the SPF, Blue Chip, and *Wall Street Journal* made statistically significant underpredictions of real GDP growth during the 1996–2000 period.

[10]The March and October Blue Chip surveys ask for forecasts of averages for GDP growth, inflation, and a number of other variables over two five-year intervals—one beginning two years from now and one beginning seven years from now. These rarely differ by more than one-tenth or two-tenths; I use their average as the long-run forecast. The Blue Chip is more useful than the SPF for this exercise, mainly because the latter publishes long-term forecasts just once a year and has been doing so for GDP only since 1992.

[11]The CFNAI is a weighted average of 85 monthly indicators in five broad categories: production and income, labor markets, consumption and housing, manufacturing and trade sales, and inventories and orders. The weights are chosen using principal component analysis and reflect the series' correlation with the (unobserved) common movement in all of the indicators. (See Fisher, 2000, and Evans, Liu, and Pham-Kanter, 2002.) To smooth through inherent volatility, the three-month moving average of the index often is used; this average is plotted in figure 3 and used elsewhere in this article.

[12]Similarly, Schuh concludes that the SPF forecasters were not behaving unusually during the 1996–2000 period. In addition, Schuh finds that the SPF forecasts fail to exploit certain statistical relationships among the forecast errors for different variables. He postulates that the large errors during this time may have in part reflected a confluence of macroeconomic factors—perhaps intensified by structural changes in the economy—that magnified the consequences of forecasters' failure to make efficient use of these relationships.

[13]Specifically, the trends connect the level of productivity between the business cycle peaks in 1960 and 1969, 1969 and 1973, 1973 and 1980, and 1980 and 1990; this last trend line is then extended through 2002:Q2.

[14]This statement is based on the SPF data, which include projections of personal consumption expenditures, residential investment, business fixed investment, government purchases (federal, state, and local), net exports, and inventory investment.

[15]Annual capital services data are published by the U.S. Bureau of Labor Statistics in conjunction with their multifactor productivity estimates. Jorgenson and Stiroh provide a description of why capital services measure the productivity of the capital stock. Note that the investment data in figure 5 are quarterly and are from the NIPA data available in late 2002. At that time, capital services and multifactor productivity data (shown in panel B of figure 5) were available only through 2000; furthermore, these data do not reflect the influence of the July 2002 annual revisions to the NIPA.

[16]First, monetary policy had been tightened—the federal funds rate had been raised 175 basis points between the spring of 1999 and the spring of 2000—and forecasters were expecting further increases in rates. Second, the price of imported oil had risen, which acts as a tax on U.S. energy consumers. Third, equity markets—which had been skyrocketing since late 1994—began to edge off in March 2000, so that the boost to spending from wealth effects was expected to wane. Furthermore, in order to adjust stocks to higher desired levels, outlays for housing, consumer durable goods, inventories, and business capital all had been increasing at high rates, and growth in these expenditures was expected to cool as the stock adjustment process ran its course. Finally, some drop in spending on high-tech equipment also was anticipated, following the temporary boost to outlays for these items in 1999 and early 2000 by firms addressing Y2K contingencies.

[17]Even if the advance is a permanent rise in the level, but not the growth rate, of technology, the *level* of the desired capital stock still is higher. Accordingly, the transition from the old to new time path for the capital stock will require some period of elevated capital stock growth and even higher investment growth. But once the new path is reached, g_t^i needs to fall back to its old value. Consequently, investment needs to grow *less* than the capital stock for some period to bring I_t^i / K_{t-1}^i back down to the original $g_t^i + \delta_t^i$.

[18]Even in "normal" times, investment is difficult to predict because of the large cyclical swings in its demand and the frictions caused by the costs of planning, installing, and operating new capital (see Oliner, Rudebusch, and Sichel, 1995). Y2K also complicated matters during the period, as firms first boosted high-tech investment in order to deal with potential problems and then delayed spending to avoid having to break in new equipment close to the January 2000 century date change. But the size of the errors noted above suggests that other factors also were in play during this period.

[19]Indeed, the strong performance of productivity may be one reason that the economy weathered the shock of the events of September 11, 2001, better than many predicted. Forecasters revised down their projections for real GDP a good deal immediately following the terrorist attacks, with the Blue Chip forecast from October 2001 projecting a 1 percent annual rate drop in real GDP in 2001:H2 and the SPF forecast made in November looking for a similar decline. According to the latest NIPA estimates, real GDP fell at an annual rate of 0.3 percent in 2001:Q3 but rose at a 2.7 percent pace in 2001:Q4. In the fall of 2001 the Blue Chip and SPF forecasts for growth in 2002 were in the 2.5 percent to 3 percent range; and as of December 2002, projections for growth in 2002 are in the 2.7 percent to 2.9 percent range.

[20]For example, if we assume that the cyclical trough occurred during 2001:Q4, then 2002:Q3 is three quarters after the trough. At this time, the level of productivity was 6.1 percent above its level at the 2001:Q1 peak. Three quarters after the troughs for the 1960 and 1969 recessions, productivity was 5.8 percent and 7.3 percent, respectively, above its value in the peak quarters. But three quarters after the troughs of the 1973, 1980, 1981, and 1990 recessions, productivity was just 1 percent to 4 percent higher than at the preceding cyclical peaks.

APPENDIX: DATA, TIMING CONVENTIONS, AND INTERPRETATIONS OF REGRESSION MODELS EVALUATING THE BLUE CHIP FORECASTS

In March, t signifies the first half of the year and the $k = 0$ forecast is for growth from the fourth quarter of the previous year to the second quarter of the current year. The $k = 1$ forecast is from the second quarter to the fourth quarter of the current year; the $k = 2$ forecast is from the fourth quarter of the current year to the second quarter of the following year. For October, t corresponds to the second half of the year; the $k = 0$ forecast is for second-to-fourth quarter growth, and so on. At the time the October Blue Chip is published, the most recent *National Income and Product Accounts* (NIPA) data are the third estimates for the second quarter of the current year; in March, the most recent data are the second estimates for the fourth quarter of the previous year. However, revisions between the second and third estimates of the NIPA usually are small, so that the statistical results probably are not substantially influenced by differences in the information sets available to the forecasters in March and October.

The actual data for gross domestic product (GDP) during the first half of the year are taken from the third NIPA estimates; the actuals for the second half of the year are the estimates published with the annual revisions made in the following summer. Both are adjusted for the average influence of any NIPA comprehensive revision that may have occurred between the time the forecast was made and the actual data were published. Given our timing conventions, $err_t^{gdp}(t-2)$ essentially reflects revisions to GDP that occur with comprehensive revisions to the NIPA that are larger than the adjustments described above. The results in table 5 indicate that forecasters apparently carry forward these influences in their medium-term forecasts.

The real Treasury bill forecast is constructed by taking the difference between the expected average nominal Treasury bill (T-bill) rate for a quarter and the expectation of long-run inflation. If t is in March, the interest rate differential is from the second quarter of the previous year; if t is in October, it is from the previous fourth quarter. The $t - 2$ value is used to account for lags between changes in interest rates and their influence on the real economy. The short-term Treasury bill forecasts were first available in 1982. Long-run T-bill forecasts were first made in 1983; I constructed a 1982:H2 value using other Blue Chip long-run interest rates.

The Chicago Fed National Activity Index (CFNAI) has been published only since 2000, so a real-time series is not available. Instead, I use the index as currently published. To account for publication lag, for March, I assume the forecasters knew the January value of the CFNAI; for October, I assume the latest available index was from August.

As noted in the text, in qualitative terms, the results from the regression models are consistent with the time-series evidence—such as that generated by structural vector autoregression (VAR) models—regarding the response of real GDP to various shocks. As an example, consider the results from Gali (1992). These show that a favorable one-standard-deviation supply shock increases real GDP by 2.8 percentage points (annual rate) in one quarter, and that a real-side demand shock boosts growth by 2 percentage points. But over the following four quarters, the supply shock raises average growth by just 0.4 percentage point, while the demand shock has little further effect. In contrast, a money supply shock has a 0.6 percentage point impact over the one to five quarter-ahead period. In addition, the short-lived shocks explain larger fractions of the GDP forecast error variance—at the one to five quarter horizon, the supply shock explains about two-thirds, demand shocks one-fifth, and the money supply shock about one-eighth of the variance. Of course, these figures are only illustrative, as such calculations are model-specific, notably with regard to the restrictions used to identify shocks.

REFERENCES

Berger A., and S. Krane, 1984, "The informational efficiency of econometric model forecasts," *The Review of Economics and Statistics*, Vol. 67, No. 1, pp. 128–134.

Campbell, J., and R. Shiller, 2001, "Valuation ratios and the long-run stock market outlook: An update," Cowles Foundation, discussion paper, No. 1295.

Croushore, D., 1993, "Introducing: The Survey of Professional Forecasters," *Business Review*, Federal Reserve Bank of Philadelphia, November/December, pp. 3–13.

Evans, C., C. Liu, and G. Pham-Kanter, 2002, "The 2001 recession and the Chicago Fed National Activity Index: Identifying business cycle turning points," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 26, No. 3, pp. 26–43.

Fisher, J., 2000, "Forecasting inflation with a lot of data," *Chicago Fed Letter*, Federal Reserve Bank of Chicago, No. 151.

Gali, J., 1992, "How well does the IS-LM model fit postwar U.S. data," *Quarterly Journal of Economics*, pp. 709–738.

Gordon, R., 2000, "Does the 'new economy' measure up to the great inventions of the past?," *Journal of Economic Perspectives*, Vol. 14, No. 4, pp. 49–74.

Harvey, A., 1989, *The Econometric Analysis of Time Series*, Cambridge, MA: The MIT Press.

Jorgenson, D., and K. Stiroh, 2000, "U.S. economic growth in the new millennium," *Brookings Papers on Economic Activity*, Vol. 1, pp. 125–211.

Jorgenson, D., M. Ho, and K. Stiroh, 2002, "Projecting productivity growth: Lessons from the U.S. growth resurgence," *Economic Review*, Federal Reserve Bank of Atlanta, Third Quarter, pp. 1–13.

McNees, S., 1995, "An assessment of 'official' economic forecasts," *New England Economic Review*, Federal Reserve Bank of Boston, July/August, pp. 13–23.

__________, 1992, "How large are economic forecast errors?," *New England Economic Review*, Federal Reserve Bank of Boston, July/August, pp. 26–42.

Oliner, S., G. Rudebusch, and D. Sichel, 1995, "New and old models of business investment: A comparison of forecasting performance," *Journal of Money, Credit, and Banking*, Vol. 27, pp. 806–826.

Oliner S., and D. Sichel, 2002, "Information technology and productivity: Where are we now and where are we going?," Federal Reserve Board, Finance and Economics, Discussion Series, No. 2002-29.

__________, 2000, "The resurgence of growth in the late 1990s: Is information technology the story," *Journal of Economic Perspectives*, Vol. 14, No. 4, pp. 3–22.

Romer C., and D. Romer, 2000, "Federal Reserve information and the behavior of interest rates," *American Economic Review*, Vol. 90, No. 3, pp. 429–457.

Schuh, S., 2001, "An evaluation of recent macroeconomic forecast errors," *New England Economic Review*, Federal Reserve Bank of Boston, January/February, pp. 35–56.