

REAL-TIME FORECASTING

Dean Croushore
Professor of Economics and Rigsby Fellow, University of Richmond
Visiting Scholar, Federal Reserve Bank of Philadelphia

JEL codes: E01, E17

Introduction

This paper will discuss real-time forecasting in a macroeconomic policy context. I will begin by talking about the Survey of Professional Forecasters, a survey of private-sector forecasters. Next, I will discuss research on real-time data analysis and its importance in forecasting. Finally, I will discuss real-time forecasting in the 1990s.

In a policy environment, such as the one I faced for 14 years at the Federal Reserve, you have three basic choices for developing forecasts in a real-time forecasting environment. One possibility, used by many policy analysts, is simply to use forecasts made by others, such as Macroeconomic Advisors. After all, forecasting firms devote considerable resources to forecasting, so why not trust their forecasts? An alternative is to look at surveys of forecasters, such as the Survey of Professional Forecasters. This gives you a range of forecasts and you can base your decisions on the median forecast, which is usually a better forecast than the forecast provided by any individual forecaster. The third possibility is to create your own forecasting model. This gives you the ability to tweak the forecast to your own needs and to specify your own baseline underlying the forecast. You can do some simple things like I did at the Fed, such as forecasting GDP for the current quarter based on the employment data that are released early in the month. Or you can run time-series models of your own specification that often hold their own against much larger, more sophisticated models. Or, you could buy a large model

forecasting software program, such as the one provided by Macroeconomic Advisors, and then modify some of their assumptions to your own liking to produce your own forecast based on their model structure. Unless you have many resources at your disposal, however, you are unlikely to want to produce a large model forecast on your own. You are unlikely to do better than others, and almost certain to produce much worse forecasts, unless you have a large number of economists working on it, such as the dozens that work on forecasting at the Federal Reserve Board.

The major concern that you should have about all these forecasting models is the role of judgment in the outcome of the forecasting exercise. The more you study forecasting, the more you realize how much impact judgment has—there is no such thing as a pure model forecast. First, there is judgment in determining what model to use. Second, there is judgment about the underlying key parameters of the model: how do you determine the natural rate of unemployment, or the growth rate of GDP, or the equilibrium real interest rate, which are generally not determined within a model? Those factors tend to drive the forecast much more than you might think.

The Survey of Professional Forecasters

My own involvement with forecasting began in 1990 with a research paper in which I wanted to get data on inflation expectations. I used a survey that was run by the American Statistical Association and National Bureau of Economic Research, known as the Economic Outlook Survey. The survey contained much useful information, and I was impressed that the survey was begun in 1968, and was the longest quarterly survey of forecasters in existence. But shortly after using the survey myself, I read an announcement that the survey was folding because of lack of

interest and no organization that was willing to run the survey. As an economist at the Federal Reserve, I thought the survey was incredibly useful—it gave great insight into the expectations of the country’s leading forecasters. I was determined that the survey should not die, and so I contacted Robert Allison of the National Bureau of Economic Research and Victor Zarnowitz of the University of Chicago about the possibility of the Federal Reserve Bank of Philadelphia taking over the survey. They were both enthusiastic about having an institution like the Fed running the survey, and so we took over, missing only one quarterly survey in the transition.

After taking over the survey, I, along with my co-researcher Leonard Mills, began to rehabilitate the survey. We renamed it the Survey of Professional Forecasters (SPF). We increased the number of participants (the participant list was down to 13 forecasters when we took it over), tightened procedures for production of the survey results, and added questions to increase the value of the survey to researchers and to policymakers. Today, the survey is used to provide forecasts to policymakers before the meetings of the Federal Open Market Committee, as well as providing a solid database of historical forecasts for use by macroeconomic researchers. The Philadelphia Fed’s web site¹ provides complete details on the survey’s history and all the individual responses to each survey, as well as write-ups for each survey and both median and mean data across forecasters for each macroeconomic variable included in the survey.

If people wish to use a forecast, or a survey of forecasts, for making decisions, they would like to feel confident that the survey provides valid forecasts. So, much research has been done on the accuracy of forecasts from surveys, including the Survey of Professional

¹ <http://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/>

Forecasters. Two standard tests of the accuracy of forecasts are tests of unbiasedness (that forecast errors have zero mean over long periods) and efficiency (that forecast errors are uncorrelated with information known when the forecast was made). If forecasts are unbiased and inefficient, then people are likely to find them useful. If forecasts are biased or inefficient, then it should be possible for someone to improve on the forecasts in real time.

SPF forecasts generally pass the tests of unbiasedness—forecasts are unbiased in long samples. However, over short periods, the forecasts might have persistent errors. Figure 1 provides an example of SPF forecasts of inflation based on the GDP inflation rate, and how the forecasts compare with the measure of the GDP inflation rate that is released one month after the end of the quarter.² In the short run, the forecasts sometimes exhibit patterns, in which forecast errors persist for some time. But, as Croushore (2010) points out, the forecasters adapt fairly quickly to structural changes in the economy that lead to short-run persistence of forecast errors, and before long the errors disappear. If the forecasts were perfectly accurate, all the points would lie on the 45-degree line in Figure 1. Although many points are off the 45-degree line, on average over the 35 years of data shown here, the plotted points lie fairly symmetrically around that line.

Most research also shows that the SPF forecasts pass tests for efficiency. However, there are exceptions. Some of the exceptions found in the literature are not valid because although they show that the forecast errors are correlated with another variable, they don't use the data that the forecasters had at the time they made their forecasts. Instead, those studies use revised data that the forecasters would not have known, as so their tests are not really tests of efficiency.

² See Croushore (2010) for more details on the data shown here.

Ball and Croushore (2003) show that there is a tendency for forecast errors to be correlated with changes in monetary policy. As Figure 2 shows, SPF output forecast errors are negatively associated with changes in the real federal funds interest rate. When the Fed tightens monetary policy (and thus the real federal funds rate increases), forecasters reduce their forecasts for output growth, but not by enough. As a result, output growth falls more than expected by the forecasters, and thus there is a negative relationship between output forecast errors and changes in the real federal funds rate.

For the most part, though, despite the Ball-Croushore findings, the forecast errors in the SPF tend to be small. The survey's forecasts are generally better than simple univariate time-series models at short horizons, as Stark (2010) notes in a recent review of the forecast accuracy of the SPF. However, there are some imperfections in the survey forecasts, especially for long horizons and with respect to the survey's efficiency in responding to changes in monetary policy.

Real-Time Data

In evaluating forecasts of macroeconomic variables, researchers must be aware of data revisions. Some researchers are not careful about this issue, so they grab data from the current database and perform tests on forecasts as if the data in their database were the same as the data that were available to researchers in real time. This is a dangerous and invalid practice. Many papers have been written that show that some new model provides better forecasts than the SPF; but in most cases the forecasting advantage comes because the researcher is comparing forecasts from a model using a recent data set to the forecasts made by the SPF forecasters using a completely different data set. Of course, the two sets of forecasts are not comparable.

To be able to compare forecasts made with a new model to the SPF forecasts in a legitimate manner, one would need to have at one's disposal a real-time database, showing what the data looked like at the time the SPF forecasters were making their forecasts. This is, in fact, the purpose of the Real-Time Dataset for Macroeconomists that Tom Stark and I developed in the late 1990s and early 2000s; see Croushore and Stark (2000, 2001) for details. The data set was developed at the Federal Reserve Bank of Philadelphia and new variables are continuously being added to the data set, based on work that continues at the Fed and work that my students have completed at the University of Richmond. A database of this nature needs good institutional support, as it is a club good, being excludable but non-rival. The Federal Reserve is a natural institution for supporting such projects, as it falls under the domain of providing macroeconomic data to the public for no charge.

Following the success of the Real-Time Dataset for Macroeconomists, other real-time databases have been developed all over the world. In the United States, the Federal Reserve Bank of St. Louis developed the ALFRED database, keeping successive vintages of the FRED database and making the data available in a convenient form. The Bureau of Economic Analysis has also kept all the vintages of their Excel files containing National Income and Product Account data available since 2002. The OECD now has a large real-time dataset containing data for all the countries in the OECD, and the Euro Area Business Cycle Network recently made a real-time database available for all the countries in the euro area. The United Kingdom, New Zealand, and Japan also maintain their own real-time databases.³

³ See my web page at <https://facultystaff.richmond.edu/~dcrousho/data.htm> for links to all of these datasets.

Some government statistical agencies in some countries have been reluctant to help researchers develop real-time datasets. They fear that if data revisions are examined by researchers, the statistical agencies will be subject to criticism because of systematic revisions. But research on data revisions is not intended to be critical of those agencies. The research findings might help the agencies strengthen their procedures to avoid having predictable revisions, for example. Economists understand that data agencies have limited resources and cannot produce perfect data releases given their constraints. The goal of research is to help people understand the limitations of the data and to explore the implications of those revisions for structural macroeconomic modeling, forecasting, and policy analysis. In addition, data revisions often reflect new information that cannot be known any earlier. For example, tax returns give the government statistical agencies much better data on income for the preceding year than they had during that year; so GDP and income statistics are improved dramatically. Or, take the example of inflation measures: by construction, the consumer price index is not revised (except for changes in the seasonal pattern), whereas the personal consumption expenditures price index is revised; yet the latter is a much superior measure of inflation precisely because the revisions reflect changes in weights applied to different sectors that provide a more accurate view of the economy.

The typical structure of real-time data sets is shown in the data matrix in Table 1, which illustrates real-time data on real U.S. output. Each column reports a data vintage, that is, the date at which the data are observed. So, the column labeled Nov 1965 tells you what someone in November 1965 would have observed at the time. Each row shows the data for a date for which real output is measured. Thus the upper left value of 306.4 (in billions of real dollars) is the value for real output in the first quarter of 1947, as someone in November 1965 would have observed

in the government's database. As you move across a given row in the table, you see how data are revised. For example, in November 1965, the first release of the data on real output for 1965Q3 was 609.1 (as before, in billions of real dollars). That number was revised to 613.0 in the data set of February 1966, and remained at that level in the data set of May 1966. The large increase seen in later vintages of the data of 3636.3 is not because of revisions to data, but because of changes in the base year from 1958 in the vintages from 1965 and 1966, to a base year of 2005 in vintages in 2009 and 2010. Moving down the main diagonal of the table, the last recorded observation in each column shows the initial release of the data for each date: 609.1 for 1965Q3, 621.7 for 1965Q4, 633.8 for 1966Q1, 13014.0 for 2009Q3, 13155.0 for 2009Q4, and 13254.7 for 2010Q1.

If data revisions were small and random, we would not worry about them affecting the structural modeling, forecasting, and policy analysis. But a look at the revisions should convince you that the revisions may be large and consequential. For example, Figure 3 shows the revisions to real output growth from the first quarter of 1977. From its initial release at 5.2%, it was revised up a few months later to 7.5%, and ultimately was revised up as high as 9.6%. But later it was revised down as low as 4.7% in the benchmark revision of July 2010. So the revisions can be large and can occur even three decades years after the first release of the data for a particular date.

You might think that such large revisions are rare and affect output growth for just one quarter, but even in the long run, data revisions can be large. For example, if you average real output growth over five-year periods, you will find large revisions which could potentially affect your view of long-run economic growth. As Table 2 shows, however, the 5-year growth can be revised by as much as 0.6 percentage points, which is a large revision for a growth rate that is as

low as 1.9%. For example, the growth rate in the first half of the 1970s was initially released as 2.1%, but by 1999 was revised up to 2.6%, nearly a 25% increase in the growth rate.

In modeling data revisions, a key question for which an answer is needed for modeling or forecasting is whether the data revisions can be modeled as providing news or reducing noise. Data revisions that provide news are those that come about when the government's data releases are optimal forecasts of later releases. Under that situation, data revisions will not be predictable in advance from data known (by anyone) at the time the data are released. Providing such data revisions requires the government statistical agency to not report its sample information alone, but to use judgment and forecasting models to optimally forecast the values of missing data, so that the data release is an optimal forecast of later data releases. But often data releases are not constructed in this manner, but rather fill in the missing source data with forecasts in such a way that the data release is not an optimal forecast of later data releases, usually because today's data release is correlated with other data known at the time. In such a situation, future data revisions will be predictable, and data revisions reduce measurement error, but each data release is not an optimal forecast of future data releases. For example, we know that there is a strong correlation between GDP data and employment data. If the government reports the GDP data based on its sample of the elements of GDP, but ignores the information contained in the employment data, then its data release will contain measurement error and data revisions reduce noise. If I were a forecaster interested in predicting future data revisions, I would look at the data on employment and form a forecast of future GDP releases using a model combining the data from the GDP release and the employment release. Such a forecast would be a better forecast of future releases of GDP than the government's release of GDP data.

This discussion raises a key question: should the government use the limited source data that it knows, combined with other information such as data on employment and industrial production to form an optimal forecast of GDP? Or should the government follow a simple rule to fill in its missing data on forecasts of GDP and produce a noisy measure, ignore data on other variables? You might think that the first method would be preferable, which seems intuitively clear. But the danger is that once you start forecasting with extraneous variables, with forecasting being more art than science, the data releases for GDP will become very subjective. As a government statistical agency, you then open yourself up to criticism that you are manipulating the data for political means. On the other hand, if you follow standard and well established procedures for filling in missing data, you avoid any possibility of people thinking that you are manipulating the data for political reasons because anyone can replicate your results, even though your data releases are not optimal forecasts of later data releases. In this situation, noise trumps news.

Forecasting

How does a forecaster produce optimal forecasts in real time? First, any forecaster needs good data; a forecast is only as good as the data used to generate the forecast. The literature on forecasting mainly focuses on model development--trying to build a better forecasting model, especially comparing forecasts from a new model to other models or to forecasts made in real time.⁴

⁴ In this section, I will report the main results in two survey papers, Croushore (2006 and forthcoming).

The first question in this literature is: are data revisions large enough to affect forecasts in a meaningful way? We have seen that data revisions may be substantial, but what is the impact of those revisions on forecasts? Stark and Croushore (2002) suggest three ways this can occur: (1) by changing the data that are input into a model; (2) by changing the coefficient estimates of the model; and (3) by changing the structure of a model, such as changing the number of lags that provide the model's best fit. Stark and Croushore (2002) illustrate these ideas using "repeated observation forecasting" (ROF), which uses different real-time data vintages for the same sample period to see how forecasts change as the vintage of the data changes. By running these ROFs allowing changes in the lags in the model, allowing coefficient estimates to change, and changing vintages, we can observe all three ways in which forecasts change. To isolate which reason is the main cause of changes in the forecasts, we can run another set of ROFs that keeps the number of lags unchanged; a comparison of the baseline result and this one reveals the importance of changes in the lag structure in the model. To isolate the effect of changes in parameter estimates, we can keep the parameter estimates fixed and generate forecasts based on the different vintages of data, to see how much the forecasts are affected by parameter changes. Everything else must be due to changes in the data input into the model.

Looking at the literature on forecasting with real-time data reveals no broad, general tendency. For some variables and for some forecasting methods, the forecasts change significantly because of data revisions. However, other variables and forecasting methods are more robust to data revisions, as Croushore (forthcoming) shows.

Because data revisions have mixed effects on forecasts, we might ask: is there an optimal method for forecasters to adjust their forecasts in the face of data uncertainty? This is a much more difficult question to answer and there have not been very many research papers that have

tried to tackle it. A few papers seem to find some degree of ability to improve forecasts by accounting for data revisions (see Koenig et al., 2003), but the predictability of data revisions is fairly small, and larger forecasting gains may be found by pursuing other aspects of modeling other than modeling data revisions. Some researchers model data revisions with time-series models, but the evidence in Croushore and Stark (2001) suggest that this will be problematic because benchmark revisions cannot be characterized as ARMA models, and benchmark revisions are the most significant type of data revisions. This explains why sophisticated filtering methods and state-space models often fail to deliver improvements in forecasts.

In summary, forecasters face difficult issues in forecasting in real time in the face of data revisions. It is not clear that the payoff to optimally handling data revisions exceeds the benefits of working on other aspects of forecast modeling, especially if data revisions are difficult to predict, as is generally the case.

Application: Forecasting in the 1990s

To illustrate real-time forecasting problems, I will demonstrate with a real-life example from my own forecasting experience at the Federal Reserve, showing forecasts from the Survey of Professional Forecasters and showing the effects of data revisions. In the 1990s, the tech boom provided an unexpected burst of productivity, increasing output growth, reducing the unemployment rate, and causing forecasters to rethink key aspects of their models. In this section, I will look at the forecast errors made in the 1990s and show how forecasters eventually caught up to the change in productivity growth, but it took some time.

In my analysis, I will look at one-year-ahead forecasts of various macroeconomic variables. Each of the SPF forecasts is made in the middle of the quarter, shortly after the first

release of the GDP data for the previous quarter. After that, data revisions occur once a year for data over the past three years, and benchmark revisions every five years or so cause significant data revisions for many years' worth of data.

Forecasts for real GDP growth made in the mid-1990s were wrong for many years in a row (Figure 4). The forecasters had expected real GDP growth to be just a bit over 2% for those years, and it turned out to be double that number. Forecasters were slow to realize that the tech boom had brought a persistently high growth rate of productivity, which translated directly into higher GDP growth. After persistent forecast errors from 1996 to 1998, the forecasters began to raise their forecasts for GDP growth in 2000. By 2001, the forecasts called for GDP growth close to 3%, just in time for the tech bubble to burst, driving GDP growth substantially lower, as the United States experienced a mild recession.

With GDP growth much faster than the forecasters expected, you might think that inflation would be higher than the forecasters thought, but in fact the opposite was true (Figure 5). Because the source of the increase in GDP growth was productivity growth, this was a classic supply shock, causing faster real GDP growth and slower growth of the price level. So the forecasters were again persistently incorrect in their inflation forecasts from the early 1990s until the end of 1999. They kept thinking that output was above potential output and so they kept thinking that inflation would rise in the future; but in fact, the forecasters had pegged potential output as too low, and inflation fell almost continuously throughout the decade.

For the most part, the forecast errors on real GDP growth translated into errors in unemployment forecasts (Figure 6). The stronger-than-anticipated growth of GDP meant that the unemployment rate would decline more than was forecasted, to be sure. But the forecasters were also very unsure of what the natural rate of unemployment was. Several years after the economic

recovery from the 1990-1991 recession, they thought the unemployment rate might have bottomed out at 5.5%. But the tech boom kept the demand for workers growing throughout the decade, and the forecasters were wrong almost continuously throughout the decade. They gradually ratcheted down their view of the natural rate of unemployment, but were always behind the curve, until the 2001 recession.

In describing these forecast errors, I have used a data set with vintage November 2001, which could be deceptive because of data revisions. In real time, the forecasters did not see the “Actual” line that I have shown in these charts. Rather, they observed early releases of the data, which might have looked quite different. So, Figures 4, 5, and 6 look pretty bad, as the forecasts clearly made severe and persistent errors. But, if you only knew what they knew at the time they made their forecasts, the forecast errors would not have looked as bad, which is why the forecasters were slow to adjust their methods. For example, Figure 7 shows the same data as Figure 4 for real GDP growth forecasts, but adds in a line labeled “Real Time Actuals” showing at each date what the forecasters thought the last data point looked like when they made their forecasts. Of course, because these are one-year ahead forecasts, they are always a year behind, so they still appear to make persistent forecast errors. But the “Real-Time Actual” line is generally closer to the forecast line than is the “Actual” line. The point is that the forecasters did not know how severe their forecast errors were in real time; they only realized it much later.

Another way to see how slow the forecasters were to change their outlook can be seen in their long-term forecasts for real output growth. The SPF forecasters are asked to provide a forecast for average real GDP growth for the next ten years. As Figure 8 shows, the forecasters seemed to be reluctant to change their views about real GDP growth in the future, despite the persistent real GDP growth rates of about 4% throughout the second half of the 1990s. In fact,

the forecasters had lowered their forecasts of real GDP growth over the coming decade in 1996, just as the productivity boom was starting. They finally changed their view in 2000 and 2001, just as a mild recession was beginning.

Conclusions

Forecasters face a difficult task in real time. As we have seen, data revisions can wreak havoc with forecasts. As the example from the 1990s shows, when structural change occurs in the economy, it may take forecasters a long time to adjust their models. That situation is exacerbated when data are revised and the initial releases of the data are much different than the later data, as was the case with real GDP growth in the second half of the 1990s. The key to good forecasting is probably the use of judgment, rather than technical expertise. In the 1990s, some forecasters recognized the permanent (or at least persistent) shift in productivity growth, including Fed chairman Alan Greenspan. The SPF forecasts took a long time to catch up to the productivity boom, but some individual forecasters performed much better.

If you want to become a real-time forecaster, you should think about major elements of your model, such as the growth rate of potential output (and the growth rate of productivity) and the natural rate of unemployment. If you can make a better guess about changes in these variables over time, you can outperform forecasters who have greater technical expertise. But every forecaster, no matter how talented, will have trouble dealing with data revisions, which are largely unforecastable, and which can make forecast errors surprisingly large.

References

- Ball, Laurence, and Dean Croushore. 2003. "Expectations and the Effects of Monetary Policy." *Journal of Money, Credit, and Banking* 35 (4): 473-484.
- Croushore, Dean. 2006. "Forecasting with Real-Time Macroeconomic Data." In: Graham Elliott, Clive W.J. Granger, and Allan Timmermann, eds., *Handbook of Economic Forecasting*, Amsterdam: North-Holland.
- Croushore, Dean. 2010. "An Evaluation of Inflation Forecasts from Surveys Using Real-Time Data." *B.E. Journal of Macroeconomics: Contributions* 10 (1), Article 10.
- Croushore, Dean. Forthcoming. "Forecasting with Real-time Data Vintages." In: Michael P. Clements and David F. Hendry, eds., *Oxford Handbook of Economic Forecasting*, Oxford: Oxford University Press.
- Croushore, Dean, and Tom Stark. 2000. "A Funny Thing Happened on the Way to the Data Bank: A Real-Time Data Set for Macroeconomists," *Federal Reserve Bank of Philadelphia Business Review* (Sept./Oct.): 15–27.
- Croushore, Dean, and Tom Stark. 2001. "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105: 111–130.
- Koenig, Evan, Sheila Dolmas, and Jeremy Piger. 2003. "The Use and Abuse of 'Real-Time' Data in Economic Forecasting," *Review of Economics and Statistics* 85: 618-628.
- Stark, Tom. 2010. "Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters." Working paper. Philadelphia, PA: Federal Reserve Bank of Philadelphia.
- Stark, Tom, and Dean Croushore. 2002. "Forecasting with a Real-Time Data Set for Macroeconomists." *Journal of Macroeconomics* 24 (December): 507–31.

Table 1: Data Matrix for Real Output (billions of real dollars)


Vintage	Nov 1965	Feb 1966	May 1966	...	Nov 2009	Feb 2010	May 2010
Date							
1947Q1	306.4	306.4	306.4		1772.2	1772.2	1772.2
1947Q2	309.0	309.0	309.0		1769.5	1769.5	1769.5
1947Q3	309.6	309.6	309.6		1768.0	1768.0	1768.0
.							
.							
.							
1965Q3	609.1	613.0	613.0		3636.3	3636.3	3636.3
1965Q4	NA	621.7	624.4		3724.0	3724.0	3724.0
1966Q1	NA	NA	633.8		3815.4	3815.4	3815.4
.							
.							
.							
2009Q2	NA	NA	NA		12901.5	12901.5	12901.5
2009Q3	NA	NA	NA		13014.0	12973.0	12973.0
2009Q4	NA	NA	NA		NA	13155.0	13149.5
2010Q1	NA	NA	NA		NA	NA	13254.7

Source: Real-Time Dataset for Macroeconomists, Federal Reserve Bank of Philadelphia

Table 2: Five-Year Average Annual Growth Rate of Real Output Across Vintages									
Vintage Year	1975	1980	1985	1991	1995	1999	2003	2009	2010
5-year period									
49Q4 to 54Q4	5.2	5.1	5.1	5.5	5.5	5.3	5.3	5.3	5.3
54Q4 to 59Q4	2.9	3.0	3.0	2.7	2.7	3.2	3.2	3.2	3.2
59Q4 to 64Q4	4.1	4.0	4.0	3.9	4.0	4.2	4.2	4.3	4.3
64Q4 to 69Q4	4.3	4.0	4.1	4.0	4.0	4.4	4.4	4.4	4.4
69Q4 to 74Q4	2.1	2.2	2.5	2.1	2.3	2.6	2.6	2.6	2.6
74Q4 to 79Q4	NA	3.7	3.9	3.5	3.4	3.9	4.0	3.9	3.9
79Q4 to 84Q4	NA	NA	2.2	2.0	1.9	2.2	2.5	2.5	2.5
84Q4 to 89Q4	NA	NA	NA	3.2	3.0	3.2	3.5	3.6	3.5
89Q4 to 94Q4	NA	NA	NA	NA	2.3	1.9	2.4	2.5	2.5
94Q4 to 99Q4	NA	NA	NA	NA	NA	NA	3.9	4.0	4.1
99Q4 to 04Q4	NA	NA	NA	NA	NA	NA	NA	2.2	2.4
04Q4 to 09Q4	NA	NA	NA	NA	NA	NA	NA	NA	1.2

Source: Author’s calculations from Real-Time Dataset for Macroeconomists, Federal Reserve Bank of Philadelphia

Figure 1
SPF Forecasts Versus One-Quarter-Later Actuals



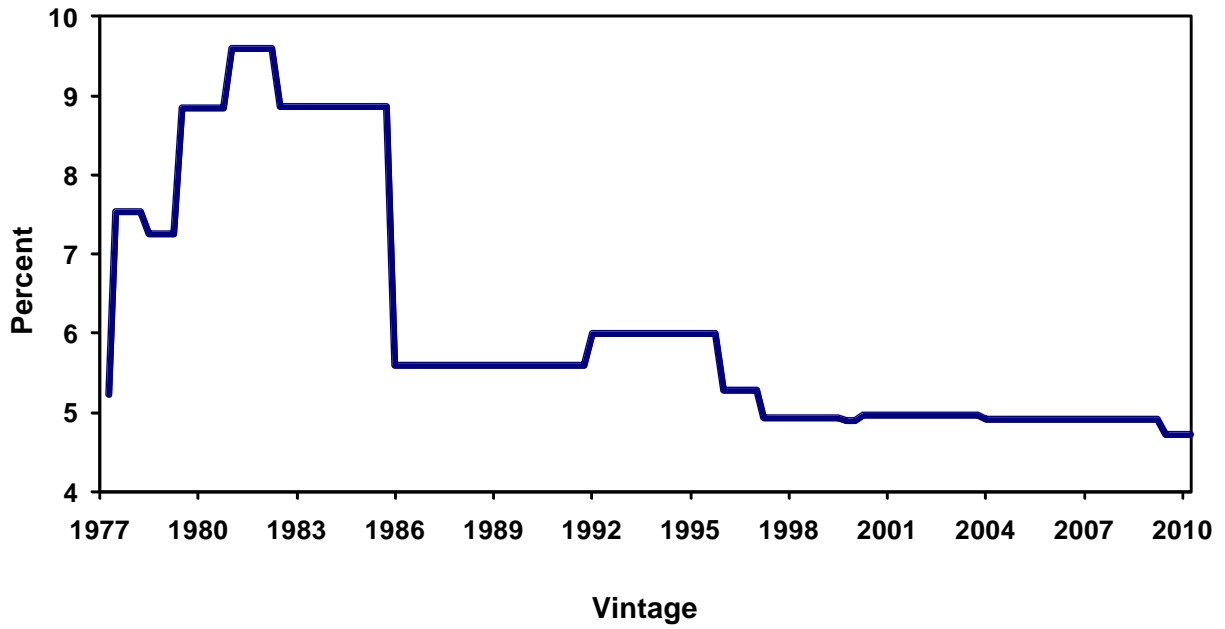
Source: Author's calculations from data from Survey of Professional Forecasters and Real-Time Dataset for Macroeconomists, Federal Reserve Bank of Philadelphia

Figure 2
Output Forecast Errors and Change in Real Fed Funds Rate




Source: Author's calculations from data from Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia and FRED database, Federal Reserve Bank of St. Louis

Figure 3
Real Output Growth for 1977Q1
(as viewed from the perspective of 133 different vintages)



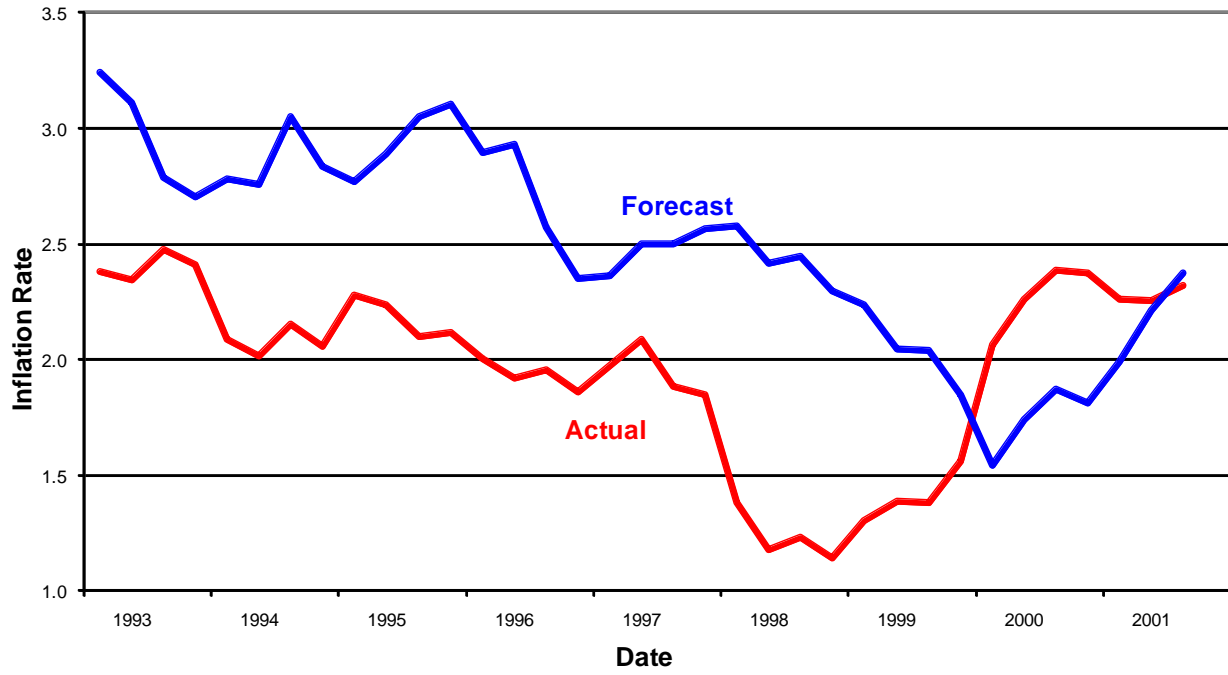
Source: RealTime Dataset for Macroeconomists, Federal Reserve Bank of Philadelphia

Figure 4
Real GDP Growth Forecasts and Actuals



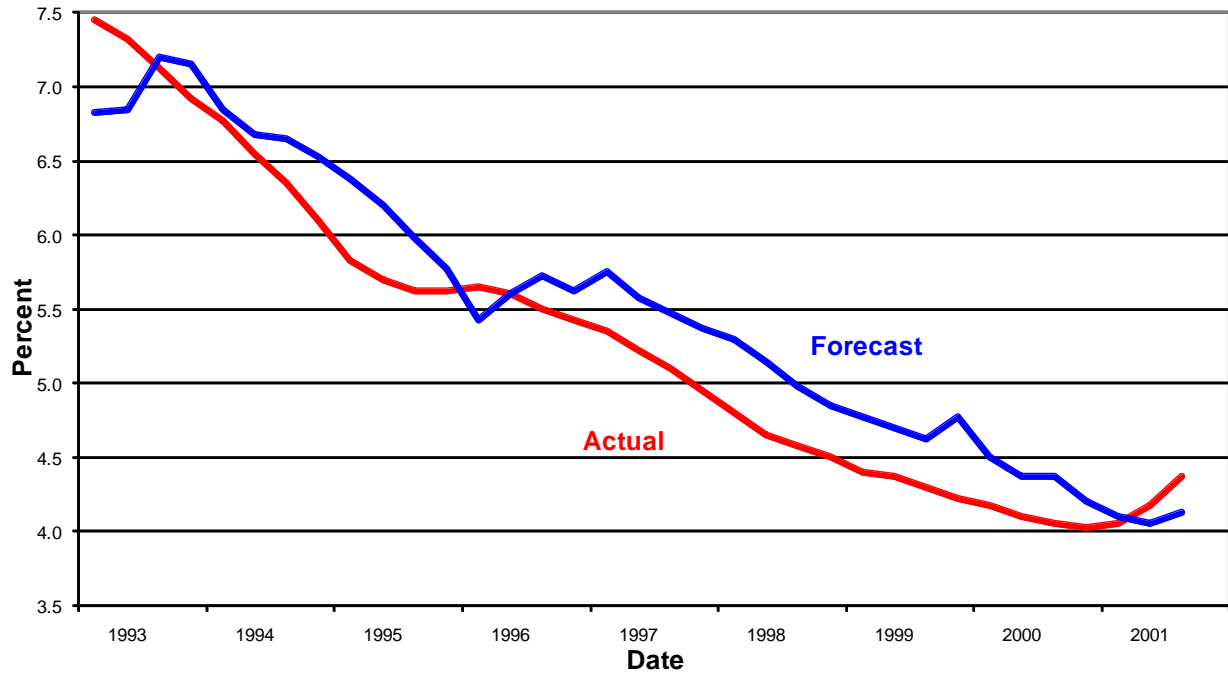
Source: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia and FRED Database, Federal Reserve Bank of St. Louis

Figure 5
Inflation (GDP Price Index) Forecasts and Actuals



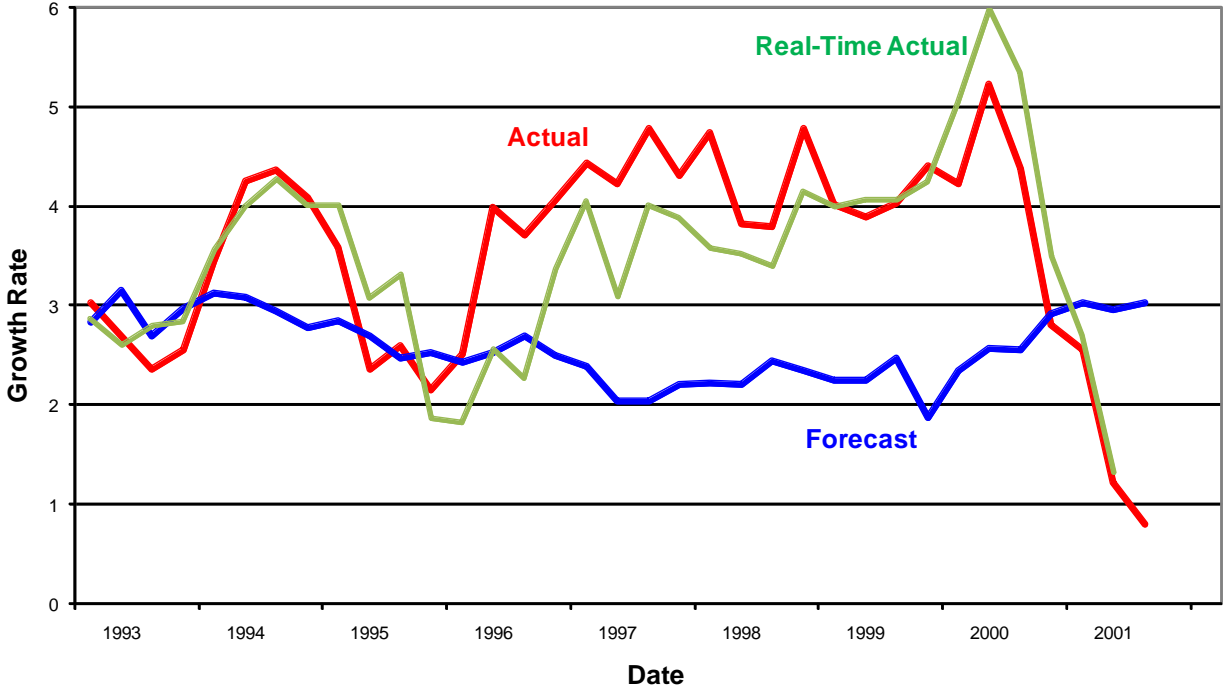
Source: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia and FRED Database, Federal Reserve Bank of St. Louis

Figure 6
Unemployment Rate Forecasts and Actuals



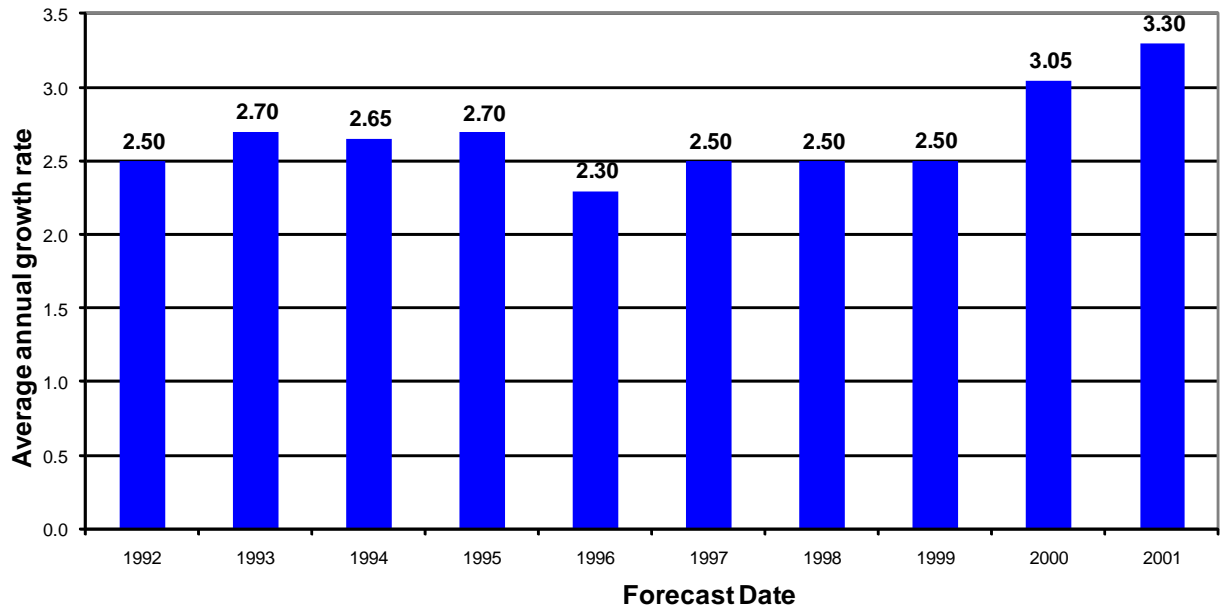
Source: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia and FRED Database, Federal Reserve Bank of St. Louis

Figure 7
Real GDP Growth Forecasts and Real-Time Actuals



Source: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia and FRED Database, Federal Reserve Bank of St. Louis

Figure 8
Real GDP: 10-Year Forecasts



Source: Survey of Professional Forecasters, Federal Reserve Bank of Philadelphia