9-2000

A Funny Thing Happened on the Way to the Data Bank: A Real-Time Data Set for Macroeconomists

Dean D. Croushore
University of Richmond, dcrousho@richmond.edu

Tom Stark

Recommended Citation

Croushore, Dean and Tom Stark. "A Funny Thing Happened on the Way to the Data Bank: A Real-Time Data Set for Macroeconomists" *Business Review - Federal Reserve Bank of Philadelphia* 5 (September/October 2000): 15-27.

A Funny Thing Happened On the Way to the Data Bank:
A Real-Time Data Set for Macroeconomists

*Dean Croushore and Tom Stark**

In October 1999, the U.S. government dramatically revised its data series on real gross domestic product, the best measure of the economy's total output. The new data showed that the economy had been growing somewhat faster over the previous decade than had been reported earlier. When data are revised, economists face unique problems when forecasting, studying the economy, and analyzing economic policy.

For example, economists are constantly trying new methods of forecasting the economy. An economist develops a new forecasting method by taking data about the economy, such as real output, unemployment, interest rates, and inflation rates, then relating those variables to each other through a set of equations that make up an economic model. The economist then looks at how well the model explains movements of the data in the past and how well it forecasts future

*Dean Croushore is an assistant vice president and economist in the Research Department of the Philadelphia Fed. He is head of the department's Macroeconomics section. Tom Stark is an economic analyst in the Research Department.

movements of the data. But substantial data revisions, like those in October 1999, throw a monkey wrench into the development of economic models. A key problem is that the data now being used to develop forecasting models can differ from the data used prior to October 1999.

Data revisions also cause problems when economists analyze past decisions about changes in policy, especially monetary policy. Many economists write articles about how monetary policy has been conducted in the past. They look at today's economic data and argue that monetary policy was tightened or loosened, that is, interest rates were increased or reduced, in response to, say, changes in real output or changes in the inflation rate. But often the data they're looking at have been revised dramatically and look nothing like the data that monetary policymakers were confronted with at the time the decision about interest rates was made.

Because of problems like these, economists need a data set containing only the observations that were known at each point in time. Such a data set would answer questions such as: What data were available to the Federal Reserve when it met to discuss monetary policy issues in February 1974? If an economist were to prepare a forecast of output growth or inflation using a new model and using data that were known in October 1987, what would the forecast be?

These types of questions can be answered only by constructing a data set that shows what the data looked like at different points in time. Doing so has been the subject of a project that the Federal Reserve Bank of Philadelphia has undertaken over the past seven years. The project required a painstaking collection of data series as they appeared in printed documents from the past. The result is a real-time data set for macroeconomists.

The data set is quite large, as you might expect, and will get larger over time as we add variables to it. Research using the real-time data set is in its preliminary stages, but it generally shows that: (1) the results of certain types of forecasting methods are very sensitive to revisions in the data, while other methods are more stable; (2) estimates of how monetary policymakers react to data are sometimes quite different when real-time data are used; and (3) the results of empirical research in macroeconomics sometimes change significantly when revised data are used. In addition, the data set can be used to study the process of data revision, which may itself be important.

THE DATA SET

The real-time data set was constructed to reflect, at each date, exactly what the macroeconomic data looked like at that date. We use the term vintage to describe each different date for which we have data as they looked at the time.

For example, suppose we were to look at the growth rate of real output for the first quarter of 1977. The first time real output for that quarter was reported, the national income and product accounts showed that real output grew 5.2 percent—that's the reading in our May 1977 vintage of the real-time data set. Today, when we look at the national income and product accounts, the growth rate of real output for the first quarter of 1977 is listed as 5.0 percent.[1] You can pick any vintage between May 1977 and today and look in our data set to see the value of real output for the first quarter of 1977 as recorded in that vintage.

Currently, the data set consists of 23 quarterly variables, including quarterly observations of 10 monthly variables. The variables include nominal output, real output and all of its components, measures of the money supply, measures of bank reserves, and the unemployment rate; for a complete list, see *Variables Included in the Real-Time Data Set*.

There is a new vintage of the data set every three months, beginning in November 1965. The

[1]When we say "today," we mean May 2000, when this article was written.

data included in each vintage are those an economic analyst would have had available in the middle of each quarter. Thus, the vintages correspond to data as they existed on November 15, 1965; February 15, 1966; May 15, 1966; August 15, 1966; November 15, 1966; and so on. For most variables, each vintage contains all the historical data (back to the first quarter of 1947) available at the time.

The data set is posted on the Internet at www.phil.frb.org/econ/forecast/reaindex.html. The web page contains links to the data itself, research papers that describe the data in more detail and use the data in a variety of empirical exercises, a bibliography of research papers that deal with real-time data issues, complete documentation on the data set, a description of changes in the data set, and a note on data we need to complete the data set, in case anyone can tell us of their whereabouts.

As you can imagine, this type of data set would have been easy to create if only someone had collected the data as time went on. We've been collecting some of these data since 1991. But gathering the bulk of the data set required us to go back into historical documents (mainly the *Survey of Current Business* for data from the national income and product accounts) and manually enter the data into a computer spreadsheet.[2]

Two major problems occurred in constructing this data set. First, the historical documentation sometimes did not make clear the exact date on which the data were available. Since we want this data set to include only those observations that would have been available to someone on a particular date, it's especially important not to include observations that were published after

[2]Much of this work was done by interns from Princeton University and the University of Pennsylvania, as well as research assistants at the Federal Reserve Bank of Philadelphia. We thank all of them for their hard work and dedication to this monumental task. We especially wish to acknowledge one student, Bill Wong, whose contributions were particularly notable.

Variables Included in the Real-Time Data Set

Quarterly Observations:

Nominal GNP (before 1992)
or GDP (1992 and after)

Real GNP (before 1992)
or GDP (1992 and after)

- Consumption and its components:
  - Durables
  - Nondurables
  - Services
- Components of investment:
  - Business fixed investment
  - Residential investment
  - Change in business inventories (Change in private inventories after August 1999)
- Government purchases (government consumption and gross investment, 1996 and after)
- Exports
- Imports

Chain-Weighted GDP Price Index

Monthly Observations:
(quarterly averages of these variables are also available in the quarterly data sets)

Money Supply Measures:

- M1
- M2

Reserve Measures:
(data from Board of Governors):

- Total reserves
- Nonborrowed reserves
- Nonborrowed reserves plus extended credit
- Monetary base

Civilian Unemployment Rate

Consumer Price Index

3-month T-bill Rate

10-year T-bond Rate

the date in question. Consequently, we spent a lot of time trying to determine exactly when data were available. Whenever there was doubt about the timing, we didn't include the data until we were sure about the date on which it had been made available to the public. We have prepared complete documentation, describing in detail all the source data, what was included, and what wasn't.

The second major problem was verifying the accuracy of the data that we typed into our spreadsheets. With such a huge data set, the opportunity for data-entry errors is large. To minimize the chance of errors in the data set, we did a large number of checks to ensure that components added up to totals; for example, total consumption spending should equal consumption spending on durables plus consumption spending on nondurables plus consumption spending on services.[3] We also plotted many of the variables to see if there were numbers that didn't make sense or that contained typos. We're confident that the data set contains few errors; any errors that remain are likely to be small.

DATA REVISIONS

One important use of the data set is to characterize how data are revised. Many data series are revised on a regular basis because the government issues preliminary numbers before all the underlying information is available. For example, the Bureau of Economic Analysis (BEA), the government agency that issues the gross domestic product (GDP) data, releases its first report on the nation's GDP near the end of the month following the end of a quarter; that release is called the advance report. But at the time of the advance report, the BEA doesn't yet have complete information, so it makes projections about certain components of GDP from incomplete source data. As time goes on, the source data become more complete. But it usually isn't until the *following year* that better information, such as income-tax records and economic census data, is available. So the GDP data undergo a continual process of revision. The data for the first quarter of 2000 were first released on April 27, 2000; they were revised on May 25, 2000, again on June 29, 2000, and yet again on July 28, 2000. Some time will pass before the first quarter observation is revised again, generally in July of each of the following three years. Thus, the data for the first quarter of 2000 will change in July 2001, July 2002, and July 2003. Each revision will be based on more complete information, so the data should become more reliable over time.

In addition to this regular schedule of revisions, the government periodically (about every five years) makes major changes, called benchmark revisions, to the data for the national income and product accounts. The most recent of these (as of this writing) occurred in October 1999. Benchmark revisions incorporate new source data and may also include changes in definitions of variables or changes in methodology. The changes are necessary, in part, because our economy is constantly changing: different types of products enter the market and different accounting methods need to be used. For example, in the benchmark revision of October 1999, the BEA changed the way it classified computer software purchased by businesses and government. Formerly treated as an office expense, such software is now treated as an investment, which is more logical because software lasts many years. The October 1999 revisions raised the average growth rate of real output over the previous two decades.

Other benchmark revisions include changes in methodology that improve the quality of the data. In the benchmark revision of January 1996, for example, the method of calculating real output was changed from a fixed-weight to a chain-weight method. Why? Because economic research had shown that the chain-weight method

[3]Prior to 1996, the components of real output added up to real output, but that's not true under the chain-weighting method used since 1996.

was an improvement over the fixed-weight method, which tended to distort calculations of real output growth in the distant past. The chain-weight method eliminates this problem.[4]

How Large Can Revisions Be? To get an idea of the size of revisions, let's return to our example of the growth rate of real output for the first quarter of 1977. Earlier, we noted that in the May 1977 vintage, the growth rate was 5.2 percent, but in the May 2000 vintage, it was 5.0 percent. That difference of just 0.2 percentage point hides quite a wild ride (Figure 1). We began at 5.2 percent in May 1977, but in the August 1977 vintage, the growth rate for the first quarter of 1977 was revised to 7.5 percent, the result of the annual revision of the data that incorporated new information. In August 1978, the growth rate was revised down slightly to 7.3 percent as more new information, including data from tax returns, was incorporated into the accounting process. Then in August 1979, the availability of even more new data caused the growth rate for the first quarter of 1977 to be revised up to 8.9 percent. Note that, even two and a half years after the fact, the raw data on real output were still being modified, as more and more records became available.

But variation in the growth rate of real output for the first quarter of 1977—from 5.2 percent to 7.5 percent to 7.3 percent to 8.9 percent—is minor compared to what happened after that. A benchmark revision of the national income accounts in late 1980 caused the growth rate to rise to 9.6 percent. A minor change in August 1982 brought the growth rate back down to 8.9 percent. Yet another benchmark revision in late 1985 drove the growth rate, as recorded in our February 1986 vintage, all the way down to 5.6 percent. It remained there until late 1991, when another benchmark revision nudged it back to 6.0 percent. In February 1996, it changed to 5.3 percent. Then, in May 1997, 20 years after the fact, the growth rate was revised again, this time

[4]For more details on chain weighting and what it means, see the article by Steven Landefeld and Robert Parker.

FIGURE 1: Real Output Growth for 1977Q1

(as viewed from the perspective of 93 different vintages)

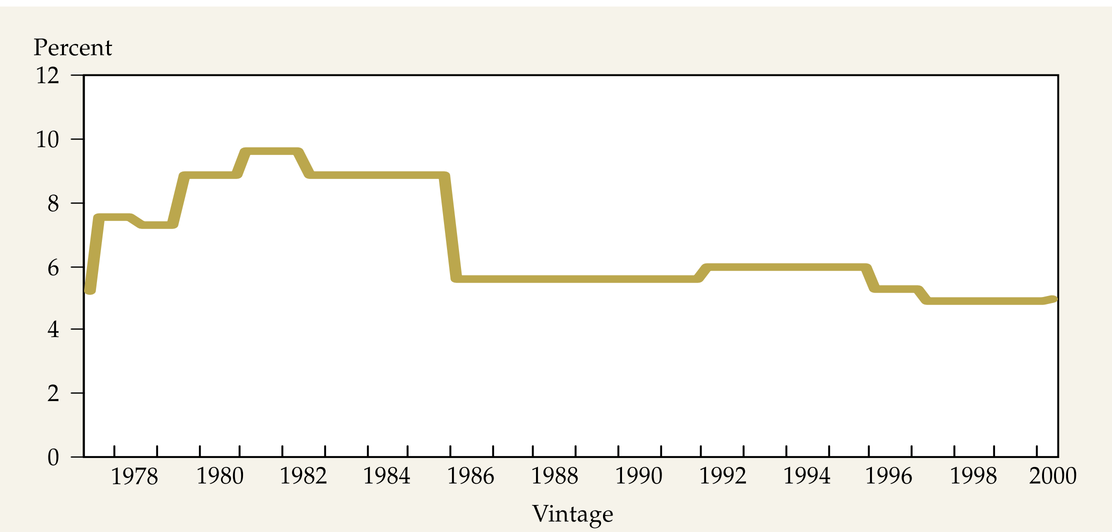

down to 4.9 percent, as the output data were changed to be consistent with newly available data on wealth. In early 2000, the growth rate was revised up slightly to 5.0 percent.

These changes in the measure of the growth rate of real output in a particular quarter are fairly dramatic. It's particularly interesting that the numbers changed so much from their initial values long after the fact, especially the decline in the growth rate from 8.9 percent to 5.6 percent in the February 1986 vintage.

Another perspective on the size of revisions can be gained by examining a chart that shows the relative frequency of revisions of a given size to the growth rate of real output (Figure 2).[5] The revision represents the difference between the annualized growth rate of real output as reported in the BEA's advance report and the growth rate for that quarter in the latest vintage of data at the time this article was written. Each bar in the chart shows the percentage of times (on the vertical axis) a revision of a particular size occurs (shown by the ranges on the horizontal axis). For example, the tallest bar on the chart shows that just over 25 percent of the time, the total revision to quarterly real output growth from its initial release to the latest available data ranged from a decline of 0.5 percent to an increase of 0.5 percent annually. You can see that many of the revisions aren't too far from zero, but a few are quite large, either positive or negative.

How Big Are Revisions Over Longer Periods? The example above showed that data revisions in a particular quarter can be fairly substantial. But we know there's a lot of volatility in quarter-to-quarter growth rates of real output and not nearly as much over longer periods. The same may be true of revisions to the growth rates. Consequently, we examine the extent to which data revisions affect five-year growth rates.

If we examine data on nominal output growth, real output growth, and inflation over

[5] This figure shows the revisions for all quarters from the third quarter of 1965 to the second quarter of 1999. The labels associated with the ranges shown on the horizontal axis are rounded to one decimal place.

FIGURE 2: Relative Frequency of Data Revisions

Quarterly Growth Rate of Real Output

Size of Revision from Advance Report to Latest Vintage

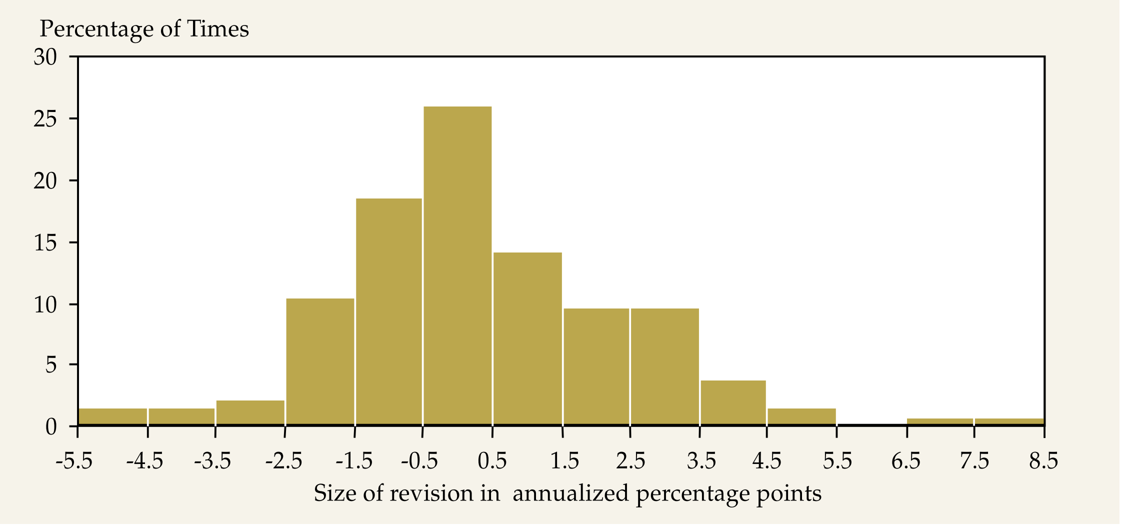

five-year periods, we see that even long after the fact, the five-year growth rates can change (Table).[6] For example, inflation averaged 7.7 percent from 1975 to 1979 according to the 1995 benchmark vintage of the data, but only 7.2 percent according to the 1999 benchmark vintage of the data. Real output growth (the inflation-adjusted growth rate of output) from 1955 to 1959 was as low as 2.7 percent in the 1995 benchmark vintage of the data, but as high as 3.2 percent in the 1999 benchmark vintage.

Thus, even five-year average growth rates may be substantially different across vintages of the data, though revisions are much smaller than for quarterly data. Even nominal output (the dollar value of output), which is easier to measure than real output and inflation, gets revised long after the fact, thanks to changes in how output is defined.

Another way to see how large data revisions may be is to look at a time-series plot that compares the data as they appeared in the BEA's advance report to how they stand today. Since we've already seen that revisions to quarterly growth rates are very volatile and revisions to five-year growth rates are smoother but still substantial,

Table: Averages Over Five Years For Benchmark Vintages

Annualized percentage points

Vintage Year:	'75	'80	'85	'91	'95	'99
Period	**Nominal Output Growth Rate**					
1950 to 1954	7.9	7.9	7.9	8.1	8.0	8.0
1955 to 1959	5.6	5.6	5.7	5.7	5.7	5.7
1960 to 1964	5.6	5.5	5.6	5.6	5.7	5.6
1965 to 1969	8.0	8.1	8.2	8.3	8.2	8.2
1970 to 1974	8.6	8.8	8.9	9.1	9.0	9.1
1975 to 1979	NA	11.1	11.2	11.3	11.4	11.4
1980 to 1984	NA	NA	8.5	8.2	8.5	8.6
1985 to 1989	NA	NA	NA	6.5	6.7	6.7
1990 to 1994	NA	NA	NA	NA	5.2	5.1
	Real Output Growth Rate					
1950 to 1954	5.2	5.1	5.1	5.5	5.5	5.3
1955 to 1959	2.9	3.0	3.0	2.7	2.7	3.2
1960 to 1964	4.1	4.0	4.0	3.9	4.0	4.2
1965 to 1969	4.3	4.0	4.1	4.0	4.0	4.4
1970 to 1974	2.1	2.2	2.5	2.1	2.3	2.6
1975 to 1979	NA	3.7	3.9	3.5	3.4	3.9
1980 to 1984	NA	NA	2.2	2.0	1.9	2.2
1985 to 1989	NA	NA	NA	3.2	3.0	3.2
1990 to 1994	NA	NA	NA	NA	2.3	1.9
	Inflation					
1950 to 1954	2.6	2.7	2.7	2.5	2.4	2.6
1955 to 1959	2.6	2.6	2.6	2.9	2.9	2.4
1960 to 1964	1.4	1.5	1.5	1.6	1.6	1.3
1965 to 1969	3.6	3.9	3.9	4.1	4.1	3.7
1970 to 1974	6.3	6.5	6.2	6.8	6.5	6.3
1975 to 1979	NA	7.1	7.0	7.5	7.7	7.2
1980 to 1984	NA	NA	6.1	6.1	6.4	6.2
1985 to 1989	NA	NA	NA	3.3	3.6	3.4
1990 to 1994	NA	NA	NA	NA	2.9	3.1

[6]The vintages chosen in this table are the last vintages of the data set prior to a benchmark revision: November 1975, November 1980, November 1985, November 1991, November 1995, and August 1999.

FIGURE 3: Revisions to Real Output Growth

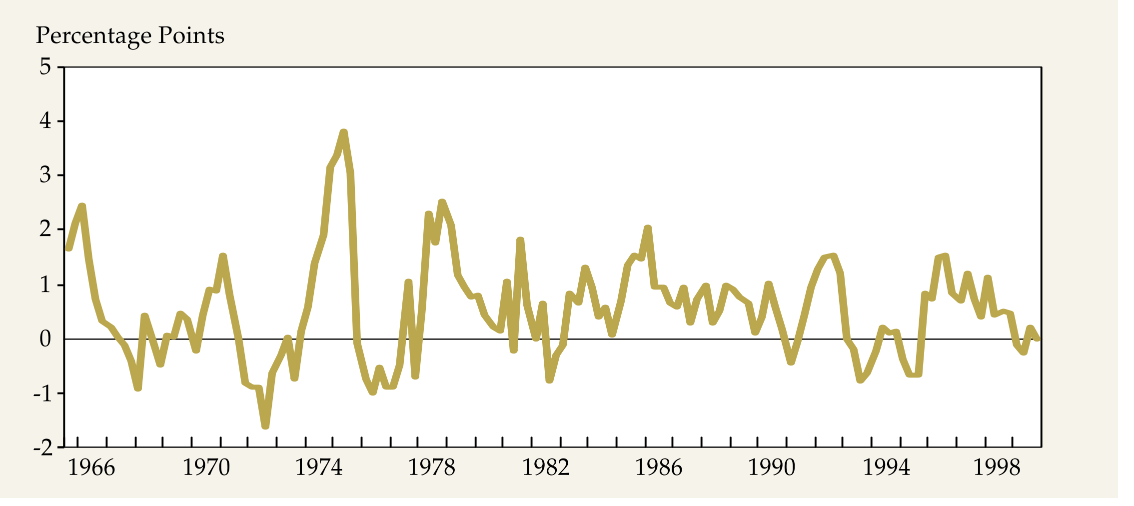

we'll take a look at real output growth over one year (Figure 3). The figure shows the differences between the growth rates of real output as they appear in one recent vintage (November 1999) and the growth rates of real output as each was first reported in the BEA's advance report. As you can see, the one-year growth rates are often revised dramatically—by over 3.5 percentage points in one instance.

Do Data Revisions Change Our Perception of Recessions? An important aspect of data revisions is how they affect our view of business cycles, in particular, the severity of recessions. Our sense of the severity of recessions, measured by the average rate of growth of real output, often changes when data are revised.[7] For example, in our November 1991 vintage, the average growth rate of real output in the recession that lasted from the third quarter of 1990 through the first quarter of 1991 was -1.0 percent. But the recession appeared worse when real output data were revised downward; in the August 1992 vintage, the average growth rate of real output was -2.8 percent. However, later still, the recession appeared less severe, when the average growth rate of real output was revised to -1.8 percent (in the November 1999 vintage).

IS RESEARCH IN MACROECONOMICS SENSITIVE TO DATA REVISIONS?

The real-time data set can also be used to examine research in macroeconomics to see if results are sensitive to the vintage of data being used; that is, do the results change significantly if a researcher uses a different vintage? In a recent paper, we examined a number of different empirical studies and found that some hold up very well, but other results change when different vintages of the data are used. These tests for the sensitivity of results are helpful to macroeconomic researchers who need to know if they can draw general conclusions from their results.[8]

We examined the 1990 paper by Finn Kydland and Ed Prescott, which showed the relationship

[7]Note that in a recession, many sectors of the economy turn down together, so the growth rate of real output is usually negative.

[8]See our 1999b research paper for more examples and more details.

of a number of economic variables to real output.[9] Kydland and Prescott used some simple statistics to show the relationships between different macroeconomic variables. The article is important because its results are one standard by which macroeconomists decide whether their business-cycle models fit the facts well enough to be useful.

The main statistic Kydland and Prescott looked at was the correlation statistic, which measures the degree to which variation in one variable is associated with variation in another variable. A negative correlation would mean that when one variable rises or falls, the other usually moves in the opposite direction. A positive correlation would mean that when one variable rises or falls, the other one usually moves in the same direction. The correlation can never be greater than 1 or less than -1, and the closer the correlation is to 1 (or to -1), the closer is the association between the two variables.

Kydland and Prescott found that the price index had a negative correlation with real output of –0.55, which means that the price index and real output *generally* move in opposite directions. Using a more recent vintage of the data, we find that the correlation is now slightly more negative: –0.66. Kydland and Prescott found that the correlation between output and consumer spending was 0.82; in today's vintage data it's 0.88. They found that the correlation between the M2 measure of the money supply and real output was 0.46; in today's vintage data it's 0.48. Looking at many other variables yielded similar results, so we conclude that the results of Kydland and Prescott hold up quite well.

We also examined a 1989 paper by Olivier Blanchard and Danny Quah, who used a small model of the economy to examine how a shock to the demand for goods and services (such as a war, which increases government purchases sharply) or a shock to the supply of goods and services (such as a dramatic increase in oil prices) affected the economy.[10] While most of Blanchard and Quah's empirical results hold up fairly well when we look at different vintages of the data, in one case they don't. When we examine how a demand shock affects the unemployment rate, we find that in more recent vintages of the data, there's a much larger effect (Figure 4).[11] Each line in the figure corresponds to a different vintage of the data and shows how the unemployment rate responds over time to a demand shock. When we use the February 1988 vintage, the unemployment rate drops immediately, then declines even more for several quarters until the end of the third quarter after the shock. Then the rate gradually returns to its starting point. But the impact of a shock to demand on the unemployment rate is bigger when we use the November 1993 vintage of the data and gets dramatically bigger when we use the February 1998 vintage. So although most of Blanchard and Quah's results weren't affected by the choice of vintage, the vintage strongly affected their estimate of the impact of a shock to demand on the unemployment rate. Evidently the statistical technique used in that study is sensitive to data revisions.

From these and other studies we examined, we concluded that most empirical work in mac-

[9] The data were adjusted by a statistical procedure to remove long-term trends, in order to focus on their movement over the business cycle.

[10] A shock is a sudden and surprising change to supply or demand.

[11] The figure shows the response of the unemployment rate to a demand shock that increases demand enough to lower the unemployment rate by one percentage point if no other variable in the model responds to the shock in the period in which the shock occurs. (In technical terms, the demand shock shifts the equation in the model describing demand, by changing the intercept term for unemployment by one percentage point.) The opposite effect on the unemployment rate would occur if there was a decrease in demand.

FIGURE 4: How the Unemployment Rate Responds To a Shock That Increases Demand

(as viewed from the perspective of the 3 vintages shown)

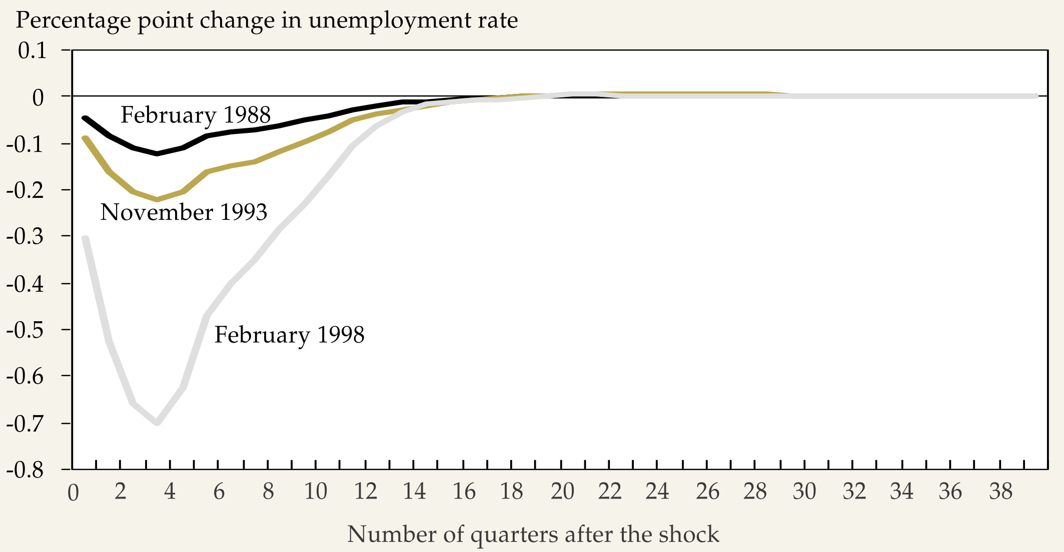

roeconomics holds up fairly well when the vintage of the data is changed, but some empirical methods, like that used by Blanchard and Quah, are more sensitive to vintage than others.

POLICY ANALYSIS

The real-time data set also helps economists understand policy actions. An economist studying past economic policies is probably doing so in light of the data as they exist today. But today's data have been revised extensively and may be quite different from the data that policymakers had available to them when they made their decisions. But if the economist has a real-time data set, she can see exactly what the economy looked like to policymakers when they made their decisions.

Consider the situation in early October 1992. Today's data tell us the economy was in pretty good shape in late 1992. Real output grew 4.3 percent in the first quarter, 4.0 percent in the second quarter, and 3.1 percent in the third quarter. But if you read accounts from that time, policymakers were clearly worried about whether the economy was recovering from the recession, and they were contemplating actions to stimulate the economy. Why were policymakers so worried? According to the data available to them, the economy had grown just 2.9 percent in the first quarter (less than today's revised number of 4.3 percent shows) and 1.5 percent in the second quarter (much lower than today's 4.0 percent). Statistics for the third quarter had not yet been released, but forecasts suggested that economic growth had not picked up much from the second quarter's anemic 1.5 percent. In addition, a number of monthly indicators pointed to a decline in the economy. (Later, many of these indicators were also revised up significantly.) Thus, it would be hard for an economist today to understand the policy concerns of the past without knowing the data policymakers were looking at.

Using the data that policymakers had before them would seem to be especially important if we were trying to model how policymakers act, a research area some economists have been interested in recently.[12]

USING REAL-TIME DATA FOR ANALYZING FORECASTS

The real-time data set can be used in a variety of ways to evaluate forecasts. Its main use, however, is likely to be in constructing new forecasting models. Sometimes an economist creates a new forecasting model using today's data, then claims that had this model been used in the past, it would have generated better forecasts than those generated by the models forecasters were using at the time. But such a claim isn't valid because past forecasters didn't have the same data to work with as today's economists do. To properly compare forecasts, an economist needs to work with a real-time data set, feed the proper vintages of the data into the forecasting model, and then see if the forecast is better.

A Simple Model with One Variable. To illustrate this idea, we've generated a simple forecasting model that uses only the history of real output to generate forecasts of future real output. We ran a simulation exercise comparing two procedures: (1) using today's data vintage and pretending that such data were available earlier; and (2) feeding data from the real-time data set into the model to generate forecasts. The first method is the technique an economist is forced to use in the absence of a real-time data set. Doing so assumes that the data aren't too different from what would have been available to a forecaster at the time. But as we've seen, that's not true. The second method uses the data available to a forecaster at the time a particular forecast was made.

The simulation exercise amounts to reconstructing what a forecaster would have done in real time. Consider a forecaster in February 1975 who wanted to forecast real output growth for the coming year. Data on real output through the fourth quarter of 1974 were available to her. For illustrative purposes, we assume that she used a very simple model to forecast future real output based on its history.[13] Using our real-time data set, we know exactly what data were available to her (our February 1975 vintage data), and we generate a forecast for the growth rate of real output over the next four quarters. The forecast turns out to be 1.3 percent. Then, imagine that three months go by, and we repeat the exercise, this time using the May 1975 vintage data. Again, we forecast real output over the next four quarters, and we find that the forecast is –3.0 percent (that's a recession forecast, with the economy's real output declining 3 percent from one year to the next). We continue this way, taking subsequent vintages of our data set one at a time, until we include very recent data, generating a new forecast with each new vintage of data. We call these forecasts real-time forecasts, since they're based on real-time data. We want to see how different these forecasts are from forecasts generated using today's data (the latest available data at the time we did our study) instead of real-time data. So we repeat the same exercise, but we use just the data available today in the same type of procedure.

To compare these two sets of forecasts, we can plot them against each other to see how different they are (Figure 5). The plot shows the forecasts based on the real-time data on the horizontal axis and the forecasts based on today's data on the vertical axis. If the forecasts were unaffected by whether we had real-time data, they'd all be

[12] See the paper by Dean Croushore and Charles Evans for an example of recent research in this area.

[13] We're using a time-series model called an autoregressive model with a four-quarter lag structure. For more details on these methods and the results, see our 1999a paper.

FIGURE 5: Two Real Output Growth Forecasts From a Simple Model

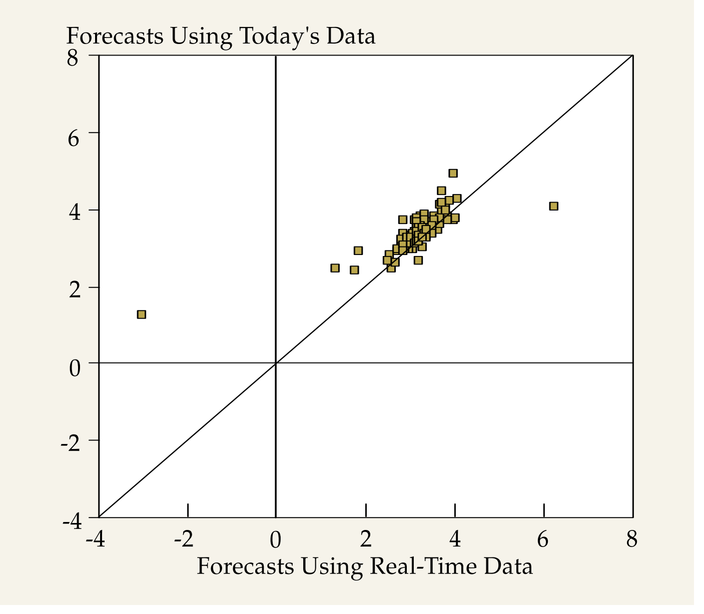

FIGURE 6: Two Real Output Growth Forecasts From a Complex Model

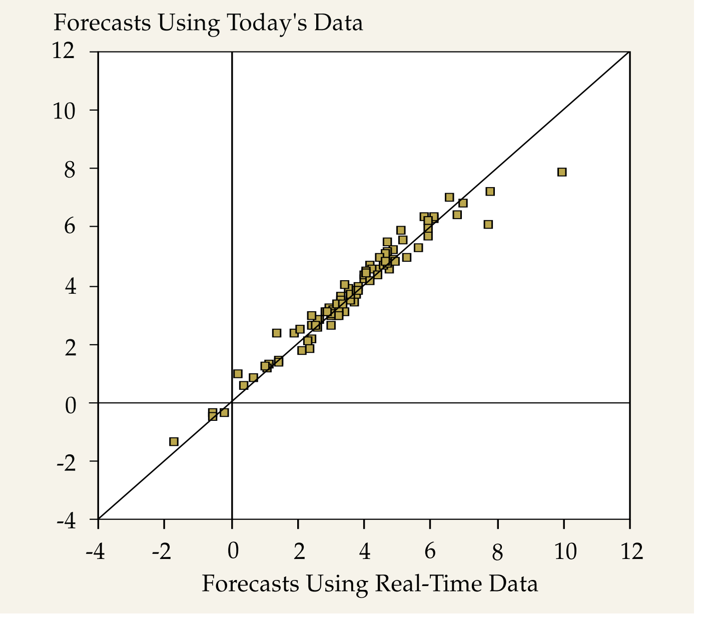

on the diagonal (45-degree) line that's drawn through the figure. Points on that line are those for which the forecast based on real-time data is identical to the forecast based on today's data. Though many points are on or near the diagonal, some points are far away from it.

Notice, for example, the point that's far to the left. That point came from the forecast made using real-time data available through May 1975, mentioned above, which forecasts a decline in real output of 3.0 percent. But revisions to the data over time caused the forecast using today's data to be much different—a 1.3 percent rise. Similarly, the point that's far to the right was from the forecast for the fourth quarter of 1976; the real-time forecast is for 6.2 percent growth, but the forecast using today's data is 4.1 percent.

Thus, in this simple model, revisions to the one variable being forecast cause the forecasts to diverge, in some cases by significant amounts.

A Complex Model with Many Variables. We can confirm the importance of using the real-time data set by performing a similar exercise using a complex forecasting model we've developed to forecast seven major macroeconomic variables, including real output, inflation, and interest rates.[14] Our tests have shown that this model provides dramatically better forecasts than the simple model used in the previous exercise. Repeating the same type of analysis used in the simpler model generates forecasts that aren't affected nearly as much by the choice of data vintage (Figure 6). The forecasts are generally quite close to the diagonal line, so that the real-time forecasts and the forecasts based on today's data are generally close to each other. Still, the forecasts diverge considerably from each other at certain dates. For example, the point furthest to the right is the forecast for the third quarter of 1976. The real-time forecast is 9.9 percent, but the forecast using today's data is 7.9 percent. In this model, the divergence between forecasts can arise because of revisions to any or all of the seven variables in the model, so figuring out the cause of the differences isn't easy. Nonetheless, the fact that differences arise indicates that data vintage matters for complex forecasting models as well as simple ones.

In both models, forecasts may be sensitive to the vintage of the data being used. For analyzing a new forecasting model, the best data set to use is the real-time data set.

[14]The model is a quarterly Bayesian vector error-corrections model. For more details, see the paper by Tom Stark.

SUMMARY

The real-time data set has a variety of uses, such as helping us understand how data are revised, testing the robustness of macroeconomic studies, analyzing policy actions and concerns, and developing forecasting models. It's our intention to keep adding variables to the data set over time and to maintain the data on the Internet for interested researchers. Though developing this data set was not easy, we hope it will prove valuable to economists and policymakers, regardless of their vintage.

REFERENCES

Blanchard, Olivier Jean, and Danny Quah. "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review* 79 (September 1989), pp. 655-73.

Croushore, Dean, and Charles L. Evans. "Data Revisions and the Identification of Monetary Policy Shocks," manuscript, December 1999.

Croushore, Dean, and Tom Stark. "Does Data Vintage Matter for Forecasting?" Federal Reserve Bank of Philadelphia Working Paper 99-15, October 1999a.

Croushore, Dean, and Tom Stark. "A Real-Time Data Set for Macroeconomists: Does the Data Vintage Matter?" Federal Reserve Bank of Philadelphia Working Paper 99-21, December 1999b.

Kydland, Finn E., and Edward C. Prescott. "Business Cycles: Real Facts and a Monetary Myth," Federal Reserve Bank of Minneapolis *Quarterly Review* (Spring 1990), pp. 3-18.

Landefeld, J. Steven, and Robert P. Parker. "BEA's Chain Indexes, Time Series, and Measures of Long-Term Economic Growth," *Survey of Current Business* (May 1997), pp. 58-68.

Stark, Tom. "A Bayesian Vector Error Corrections Model of the U.S. Economy," Federal Reserve Bank of Philadelphia Working Paper 98-12, June 1998.