Not Your Father's Recovery?

Kenneth R. Beauchemin

There has been much talk about a disappointing recovery in the wake of the Great Recession—that this time it is much slower. Comparing features of this recovery to past recoveries casts some doubt on that view. The comparison is made using a scaled-down version of the sophisticated and powerful models that real forecasters actually use. Applying it to real GDP growth, unemployment, inflation, and the federal funds rate suggests that the recovery looks consistent with past recoveries—at least so far.

Everything should be made as simple as possible, but not simpler.

—Albert Einstein

If all goes according to the usual business-cycle dating procedures, a committee at the National Bureau of Economic Research (NBER) will soon convene to declare that what has come to be known as the "Great Recession" came to an end in June 2009. As we all know, it was a doosie. Beginning in December of 2007, it will have lasted 19 months—the longest downturn since the Great Depression. It will also go down as the deepest, as the United States shed 4.1 percent of gross domestic product (GDP) from peak to trough. The unemployment rate more than doubled, rising from 5.0 percent in December 2007 to a peak of 10.1 percent in October 2009.

Now that we are a year into the recovery, and the annual July "benchmark" revisions to the National Income and Product Accounts are complete, the time is ripe for an assessment of the recovery so far. Popular opinion strongly suggests that it has been "substandard." But what is the standard by which the strength of a recovery can be measured? Some point to a historical tendency for deep recessions to be followed by rapid recoveries and vice versa.

This belief looks to be validated in figure 1, which plots the peak-to-trough output loss in each of the eleven NBER-dated postwar recessions (horizontal scale) versus the growth in real GDP in the four quarters following the final quarter of each recession (vertical scale). The way the data points are scattered suggests that output loss and subsequent growth are positively related, as indicated by the upward-sloping line that "best fits" the ten episodes prior to the Great Recession. By this standard we are indeed being cheated. The line predicts 9.7 percent growth from the third quarter of 2009 to the second quarter of 2010, but we got only 3.2 percent instead!

Figure 1. The Severity of the Recession and the Strength of the Recovery
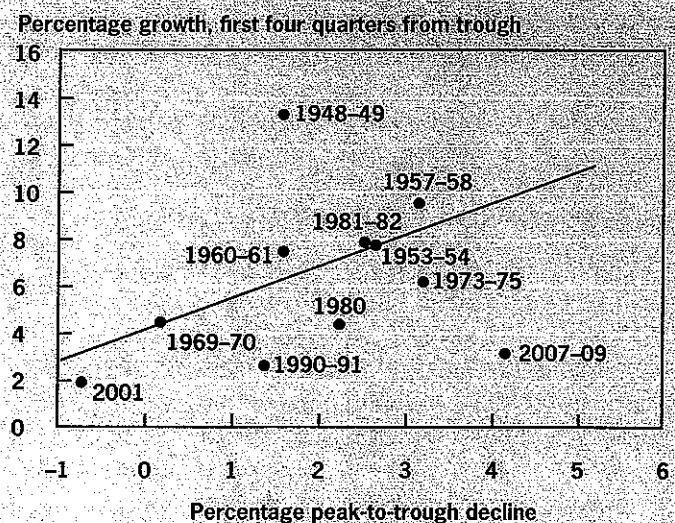
Time for a collective groan? Only if we accept the line-of-best-fit in figure 1 as the appropriate historical standard for economic recoveries to live up to. And that line is essentially being offered up as a forecasting model, but at best it is a "rule of thumb." As humans we excel at recognizing patterns, but some times we are too good. Figure 1 may be a case in point. When the Great Recession is added to the mix—which by the way increases the sample size by a whopping 10 percent—the "pattern" seems to fade away. Fortunately, there are simple forecasting techiques to help place our pattern-recognizing instincts in check.

Vector Autoregressions and the Force of History

An economy is a complex, dynamic, and stochastic system. It is complex because there are lots of relevant quantities which are all tied up with each other in a seamless web of market and nonmarket interactions. Everything affects everything else. As a result, economies are complicated and not likely to yield their secrets to collections of scatterplots. An economy is also dynamic in that its current and past conditions have strong implications for all of its possible future conditions. And economies have stochastic or random elements. There are uncertainties in life that must be tolerated by the people and institutions that comprise an economy, as well as those who try to predict its future course.

Macroeconomic forecasters have in their toolkit a powerful class of models called vector-autoregressive (VAR) models, which reflect these basic properties. The "vector" part of the name simply means that they forecast multiple economic concepts (or variables) at once. The "autoregression" means that a forecast of any one variable depends on the past historical (or forecast) values of all of the variables in the model—including its own. Like actual economies, VARs are complex, dynamic, and stochastic.

But in another important sense, they are not. In its simplest form, a VAR is what economists call a "reduced-form" model. They generate an extremely useful set of correlations that summarize the historical data, but contain little structure motivated or produced by economic theory. And structure is useful for interpreting a forecast. But that very lack of structure imbues VAR models with enough flexibility to capture the "force of history" in a forecast; too much structure can degrade forecasting ability. To paraphrase Mark Twain, history does not repeat itself, but it rhymes. A VAR model is good at picking up on the rhymes.

A VAR model forecast is essentially an average of past business cycles, and as such, it is the most natural candidate to produce our view of what a standard recovery from the Great Recession would look like. In what follows, we will put ourselves in the shoes of a macroeconomic forecaster approximately one year ago, when the U.S. economy was on the verge of a recovery that was not yet apparent. We

will specify a simple VAR model—too simple by professional forecasting standards but one adequate to our task. To estimate the model's parameters, we will use only the data through the second quarter of 2009, the presumed trough of the recession. We will then simulate the model from the third quarter of 2009 to the second quarter of 2010 to establish a historical standard for the recovery, one year out. We then compare the simulated recovery to what has actually transpired in the past year to gauge how extraordinary or ordinary the recovery has actually been so far.

Recovering from the Great Recession

Although VAR models and some of their more complicated cousins can handle a fairly large number of variables, we will construct one containing four variables that are of major interest to Federal Reserve policymakers: real GDP (growth rate), core personal consumption expenditures (PCE) inflation, the unemployment rate, and the federal funds rate.[1] (See table 1 for variable definitions.) To conduct the experiment, the model is fitted to data from the first quarter of 1959 through the second quarter of 2009.[2]

Figure 2 presents the results of our experiment. Each panel in the figure contains the retrospective forecast and the actual data values for each of the variables. They also contain the 90 percent probability bands to help quantify the degree of uncertainty surrounding the forecast. If the recovery-period data wander outside of these bands, we can be pretty sure that something out of the ordinary has been going on, as there is only one chance in ten that the pure force of history would have sent them out there.

It is evident from the real GDP panel that economic growth over the past year appears to have been somewhat substandard. Although it started out roughly in line with historical norms, it has weakened in the last two quarters. Nevertheless, real GDP growth continues to stay within the probability bands. Note though, that the rule-of-thumb suggested by figure 1 is far too demanding: the 3.2 percent actual growth rate over the previous four quarters compares much more favorably to the 5.8 percent rate expected by the VAR model, than the 9.7 percent rate predicted by the rule of thumb.

The behavior of the unemployment rate during the recovery looks truly anomalous. While the VAR model expected the unemployment rate to have fallen to an average of 8.0 percent last quarter, the actual rate remained stubbornly high at 9.7 percent for two consecutive quarters, and it is currently flirting with the upper probability band marking the 95th percentile. A fair bit of the back-story here is doubtlessly the tremendous productivity growth recorded during the past year, which has allowed businesses to stretch existing labor, finding efficiencies where they could, with workers who were quite happily employed despite the extra effort they were required to expend. But that just pushes the puzzle into a deeper realm.

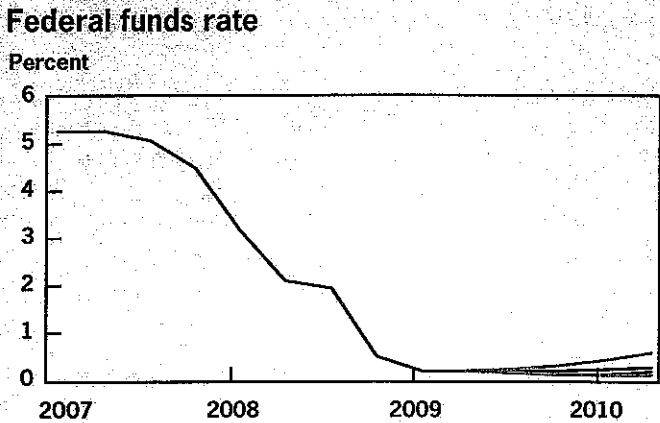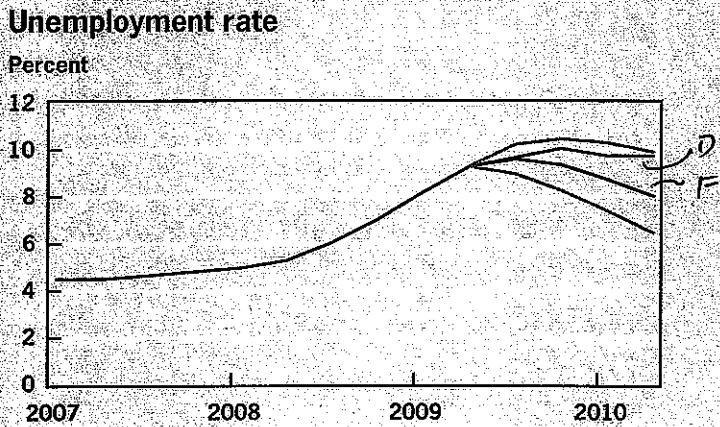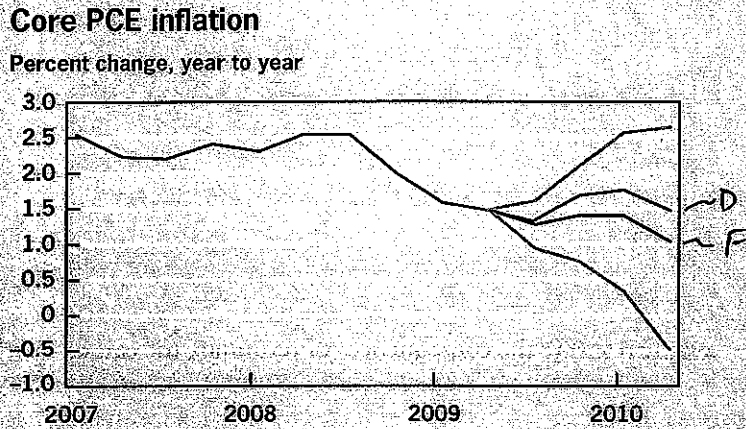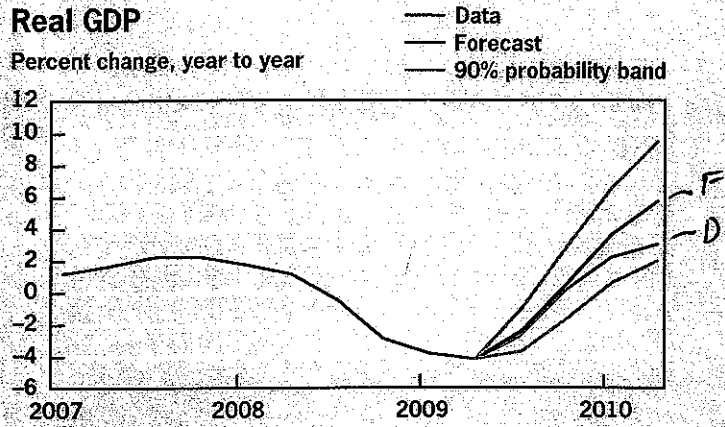## Figure 2. Model's Forecast Using Data from 1959:Q1–2009:Q2

Real GDP
Percent change, year to year

— Data
— Forecast
— 90% probability band

12
10
8
6
4
2
0
-2
-4
-6

2007 2008 2009 2010

F
D

Core PCE inflation
Percent change, year to year

3.0
2.5
2.0
1.5
1.0
0.5
0
-0.5
-1.0

2007 2008 2009 2010

D
F

Unemployment rate
Percent

12
10
8
6
4
2
0

2007 2008 2009 2010

D
F

Federal funds rate
Percent

6
5
4
3
2
1
0

2007 2008 2009 2010

Figure 3. Model's Forecast Using Data from 1983:Q1–2009:Q2

Real GDP
Percent change, year to year

— Data
— Forecast
— 90% probability band

6
4
2
0
-2
-4
-6

2007 2008 2009 2010

D
F

Core PCE inflation
Percent change, year to year

3.0
2.5
2.0
1.5
1.0
0.5
0

2007 2008 2009 2010

D
F

Unemployment rate
Percent

12
10
8
6
4
2
0

2007 2008 2009 2010

D
F

Federal funds rate
Percent

6
5
4
3
2
1
0

2007 2008 2009 2010

Table 1. Definitions

Real GDP	The price-adjusted value of all goods and services produced domestically in a specified period of time (such as a quarter).
Core PCE inflation	The rate of change in the core PCE index, which measures the average price of goods and services consumed in the U.S. excluding food and energy. The Federal Reserve relies on the fourth-quarter percent change in the index as its principal measure of U.S. inflation.
Unemployment rate	Percentage of the U.S. labor force that is unemployed and looking for a job.
Federal funds rate	The interbank interest rate charged on overnight loans; the principal policy instrument of the Federal Reserve.

Turning to inflation, the year-on-year change in the core PCE index has not gone too differently from the retrospective forecast, with the actual data showing somewhat higher inflation than would have been predicted by the model a year ago—1.5 percent realized versus the 1.0 percent rate expected. The higher inflation rate is mildly curious given that the unemployment rate revealed considerably more resource slack than would be historically consistent.

Finally, the experiment also indicates that the continuation of the near-zero federal funds rate policy in the previous year was very much in line with historical forces. At first blush, this seems puzzling considering that the policy is without precedent in U.S. experience. However, the policy had already taken hold at the time of the retrospective forecast, and there is a strong tendency for central banks, including the Federal Reserve, to change overnight lending rates infrequently. This tendency leads to a smoothing of interest rates, which the VAR model can deftly pick up on.

The Force of a More Modern History
Next we consider the same experiment, except that the VAR model will fit only the data in the latter part of the historical period. Instead of using the data all the way back to 1959, we instead use just the data from 1983 onward. Why throw out the earlier data? It is information after all. Nevertheless, it is always a concern that the structure of the economy changes over time, and if this is so, the more recent data may better reflect how the economy currently works. We still have the previous results to compare to, and that alone is informative.

Economists have offered a number of good reasons to downplay the earlier data. For example, it is conjectured that the change in Federal Reserve operating procedures that helped banish the high inflation rates of the 1970s has led to substantial changes in inflation dynamics. Others have detected a significant change in labor market behavior, finding evidence that the 1981-1982 recession marked the last recession in which a large portion of layoffs were temporary. Since then, changes in rehiring behavior have led to what have become known as "jobless" recoveries.

Figure 3 shows the consequences of rerunning the earlier experiment with the more current sample. In terms of real GDP growth, the 3.2 percent four-quarter growth rate now exceeds the expectations set down by the VAR simulation of only 2.4 percent growth, but nevertheless falls comfortably in the 90 percent probability bands. The unemployment rate is better aligned with its historical standard but is still higher than would be expected. Actual core PCE inflation lines up nearly perfectly with the VAR simulation, and the near-zero federal funds rate path remains historically consistent.

Conclusion
There has been much talk about a disappointing recovery in the wake of the Great Recession—that this time it is much slower. These simple exercises cast some doubt on that view. Taking the two experiments together, real GDP growth is not so different than we might have reasonably expected a year ago. The relationship between output growth and the unemployment rate looks somewhat anomalous, with an unemployment rate remaining higher than history alone would suggest, but not resoundingly so. Core PCE inflation turned out roughly as anticipated over the previous year, as has the Federal Reserve policy of a near-zero federal funds rate.

I have looked at the recovery from 30,000 feet. From that height, with some inevitable slippage and nuance, the recovery looks consistent with past recoveries—at least so far. Closer to the ground, recoveries will all look very different, with harder-hit sectors taking longer to recover or becoming permanently smaller. Credit will be tighter for longer in some sectors than in others. The list goes on and on. The idiosyncrasies should not be ignored, but neither should the long view.

Footnotes

1. More specifically, the VAR includes the natural log first differences of real GDP and the core PCE price index, and the natural logs of the unemployment rate and the federal funds rate.

2. Since many economic time series undergo frequent revision (for example, real GDP and the core PCE deflator in the present case), the data currently available for the post-2006 period are only approximately equal to the data available to a researcher at the close of the second quarter last year. I will avoid this complication, and assume that the present data set—or the current data "vintage"-was the one available last summer.

Recommended Reading

James H. Stock and Mark W. Watson. 2001. "Vector Autoregressions," *Journal of Economic Perspectives*, vol. 15, 101-15.