Comment and Discussion

James H. Stock: From 1993 through mid-1998, the U.S. economy experienced several years of low unemployment and low and falling inflation that is nothing short of extraordinary. What should one make of this experience? Is it transitory, simply good luck, or has the economy changed in a fundamental way? Is the death of the Phillips curve, so long proclaimed, finally a reality?

Although it is widely known that the NAIRU is measured with considerable imprecision, the recent experience is surprising. For example, Staiger, Watson, and I have estimated that in 1989 the NAIRU, based on the GDP deflator, was 6.3 percent, with a 95 percent confidence interval of 5.0 to 7.4 percent.[1] The United States has now been at or below the lower end of this confidence interval for some time, yet inflation remains quiescent.

In his previous work, Robert Gordon has argued that the unemployment-based Phillips curve has been a trusty and stable relation, at least through the early 1990s. In the current paper, he turns his attention to the events of the past five years. His approach to this puzzle is, sensibly, to ask what went wrong with the constant NAIRU, circa-1992 Phillips curve. He considers four sets of factors that could have contributed to the good inflation performance, given recent unemployment: first, traditional supply shocks, in particular food and energy prices and import prices, which have been included in Gordon's empirical work at least since 1982; second, some new supply shocks, in particular medical prices and computer prices; third, recent measurement improvements

1. Staiger, Stock, and Watson (1997, table 1).

in the CPI; and fourth, an otherwise unexplained drop in the NAIRU. His empirical strategy is to decompose the error that he would have made had he been asked, at the end of 1992, to forecast the average inflation rate for the four quarters ending in 1998:2, given future values of the unemployment rate, food and energy prices, and his other control variables, but not the future shocks to the Phillips curve.

Gordon concludes that much of the fall in inflation—approximately 1.4 percentage points—is attributable to favorable traditional supply shocks (declining real food and energy prices and import prices), but that declines in the NAIRU have also been an important factor, explaining approximately 0.5 percentage point of the surprisingly good inflation performance. Although the new, lower NAIRU might persist, in all likelihood the favorable supply shocks will not. Thus Gordon's explanation of the Goldilocks economy echoes Goldilocks's fate. On the one hand, the bears do come home, rudely interrupting her sleep and ending her consumption windfall: there is, it turns out, no free lunch. On the other hand, she escapes with no more than a fright, presumably to pursue future policies that depend less heavily on favorable supply shocks.

My comment has three parts. First, I raise some disagreements with aspects of Gordon's conceptual framework, especially the introduction of his new supply shocks. Second, I reestimate his decomposition, altering aspects of his analysis with which I disagree. Third, I return to the claims of proponents of the new economy in light of this evidence and provide some further evidence that the puzzle is not about inflation, but about labor markets. Because of space constraints, this discussion exclusively considers inflation as measured by the GDP deflator.

Part of Gordon's explanation for the recent good news on inflation involves his three new supply shocks: declining real computer prices, moderation in health care inflation, and improvements in methods for measuring inflation. The third of these so-called supply shocks is not a supply shock at all, and it is confusing for Gordon to call it so. If one were modeling river flooding and the units of measurement of river depths had switched from feet to meters, one would not call the resulting break in the series a climate shock; rather, one would just adjust the series so that the units were comparable over time. Similarly, because methods for measuring inflation have changed, one needs to make an adjustment so that historical and current values of inflation are in the

same "units."[2] Gordon deserves credit for being the first in this literature to incorporate such an adjustment, but this adjustment has little to do with supply shocks.[3]

Although it is common sense to adjust for known improvements in measurement, the argument for singling out falling computer prices and moderating health care inflation is less clear. Returning to the flooding example, this is a bit like saying that St. Louis suffered flooding because the Mississippi was especially deep although the Missouri was not—which is no explanation at all. Nor does it explain low inflation to observe that two domestic sectors had falling or stable prices. Over any period, there always will be some sectors that contribute positively to inflation and others that contribute negatively. To justify subtracting the contributions of two such sectors requires arguing that these contributions are qualitatively different from those of the other sectors, in other years, that have been left in. One systematic approach to eliminating outlier sectors is to use a trimmed-mean CPI, or perhaps the median CPI. This issue is more than semantic. Gordon "explains" the decline in the NAIRU in terms of these two price shocks; but if *real* health care costs start to rise again, would one expect the NAIRU to rise as well? Probably not, because (as I argue below) the decline in the NAIRU seems to be linked to recent developments in labor markets.[4]

As a basis for further analysis, I have recomputed Gordon's decomposition in his table 8, using a specification that makes these and other minor changes. The motivation for one regressor—the change in pro-

2. This feet to meters view also leads to a different method from Gordon's for performing the dynamic simulations. Rather than treating measurement changes as a shock with dynamic consequences, the dynamic simulation should be performed in base-quarter (1992:4) units, and then the forecasts should be converted to current-quarter units.

3. The river depth analogy is imperfect, because the change in how inflation is measured affects real activity, slightly, through real expenditures that are indexed to the CPI. But these effects are arguably negligible for the purposes of explaining the current inflation-unemployment puzzle.

4. This raises the question of whether one should control for the traditional supply shocks of real food and energy prices and import price effects. One argument for retaining these supply shocks is that they are largely external to the U.S. economy; another is that they have been included in empirical Phillips curve specifications at least since the early 1980s and therefore are not subject to the ex post identification bias alluded to above. Whether these arguments are compelling is a matter of judgment, but for the purposes of this discussion I retain them in my specifications.

ductivity deviation—is unclear, so this has been omitted. Also, the sample period is different, and lagged inflation enters differently. The resulting specification is similar to that of Staiger, Watson, and I.[5] The results are presented in the first column of table B1. When a constant NAIRU of 6.0 is assumed, the forecast error for the constant NAIRU, zero supply shock, dynamic forecast of GDP inflation in 1998:2 is essentially the same as that of Gordon: -2.79 versus -2.75 (that is, $1.01 - 3.76$) in table 8. The TV-NAIRU used in table B1 is computed using the same value of τ_η as Gordon does, and the decline of inflation explained by the fall in the NAIRU, 0.63 percentage point, is close to Gordon's estimate of 0.53 percentage point. The contributions of the traditional supply shocks in tables B1 and 8 are also similar. Evidently, Gordon's decomposition is robust to these changes in specification. Interestingly, the root mean-squared error of the dynamic simulation in table B1 is less than Gordon's (0.47 versus 0.60 in table 3), so the changes in the specification actually improve the performance of the Phillips curve in the mid-1990s.

This decomposition is subject to several sources of sampling uncertainty: uncertainty about the constant NAIRU through 1992; uncertainty about the TV-NAIRU and the estimate of τ_η; and uncertainty about the other coefficients in the regression. The second column of table B1 illustrates the effect of the uncertainty about the constant NAIRU. For the specification described in the notes to the table, the estimated constant NAIRU is 6.3 percent. Using this value, the dynamic simulation forecast error is larger, and the amount of surprisingly low inflation explained by a drop in the TV-NAIRU approximately doubles, from 0.6 percentage point to 1.3 percentage points. A larger value of τ_η results in a lower estimate of the current TV-NAIRU and further increases the amount explained by a drop in the NAIRU. A proper treatment of all sources of uncertainty would, I suspect, result in a very large sampling uncertainty associated with this decomposition. This example also emphasizes the importance of the decline in the NAIRU to Gordon's story.

What can one conclude from this about the claims of proponents of the new economy? Gordon characterizes the new economy view as one in which the NAIRU has fallen. However, a more radical interpretation,

5. Staiger, Stock, and Watson (1997).

Table B1. Decomposition of Dynamic Simulation Forecast Error for Four-Quarter Inflation, 1998:2[a]

Percent change (annual rate)

	TV-NAIRU, 1992:4 NAIRU 6.0 percent (1)	*TV-NAIRU, 1992:4 NAIRU 6.3 percent (2)*	*New economy (3)*	*Capacity utlization rate (4)*	*Housing permits (5)*
Actual inflation, 1998:2	1.01	1.01	1.01	1.01	1.01
Predicted inflation, 1998:2[b]	3.80	4.47	3.80	2.72	2.76
Dynamic forecast error	−2.79	−3.46	−2.79	−1.71	−1.75
Decomposition of error					
Contribution of traditional supply shocks					
Food and energy	−0.50	−0.50	−0.50	−0.40	−0.53
Import prices	−0.59	−0.59	−0.59	−0.60	−0.51
Subtotal	−1.09	−1.09	−1.09	−1.00	−1.04
Other identified explanations					
CPI measurement[c]	−0.44	−0.44	−0.44	−0.44	−0.44
Change in NAIRU[d]	−0.63	−1.30	. . .	. . .	. . .
Change in slope	. . .	. . .	−0.68	. . .	. . .
Subtotal	−1.07	−1.74	−1.12	−0.44	−0.44
Total explained error	−2.16	−2.83	−2.21	−1.44	−1.48
Unexplained error	−0.63	−0.63	−0.58	−0.27	−0.27
Summary statistic					
Simulation root mean-squared error[e]	0.47	0.47	0.29	0.30	0.39
Simulation standard error	1.37	1.37	1.37	1.49	2.03

Source: Author's calculations.

a. All entries pertain to four-quarter inflation; inflation is measured by the annual percentage change in the GDP deflator. For each column, a base regression is estimated by ordinary least squares from 1963:3 to 1992:4, using earlier observations for initial conditions. Data are quarterly, and the regressions use the specification in Gordon's table 3, with the following modifications: the change in productivity deviation is dropped; the Nixon "on" and "off" variables are combined into a single price control variable; and lagged inflation enters as six lags of the quarterly change in inflation, where the coefficients are unrestricted.

b. Assumes a constant NAIRU and constant real food, energy, and import prices, with the following assumptions on the NAIRU: for column 1, the constant NAIRU is assumed to be 6.0 percent; for column 2, the in-sample estimate, 6.3 percent; for column 3, 6.0 percent; for columns 4 and 5, the constant "nonaccelerating inflation" rates of capacity utilization and housing starts are taken to be the in-sample estimates.

c. From 1992:4 to 1998:2, changes in measurement methodology reduced annual measured CPI inflation by 0.63 percentage point, which—using Gordon's factor of 0.7—corresponds to a 0.44 percentage point per year reduction in measured GDP deflator inflation.

d. For columns 1 and 2, these dynamic simulations are repeated, allowing the NAIRU to vary. The TV-NAIRU is estimated over the full sample, 1963:3 to 1998:2, using $\tau_\eta = 0.092$ (approximately Gordon's value).

e. The simulation root mean-squared error is the dynamic simulation forecast error that incorporates all the "explanations" relevant for that column.

which I believe is more in keeping with new economy rhetoric, is that the traditional Phillips curve has simply ceased to be relevant. There is no longer a link between the unemployment gap and changes of inflation; according to this view, the slope coefficient in the Phillips curve is now zero. Gordon's evidence does not address this interpretation.

This hypothesis is investigated in the third column of table B1. The circa-1992 Phillips curve is estimated through 1992:4, and the dynamic simulation proceeds like that in the first column, except that from 1993 onward the slope of the Phillips curve (more precisely, the sum of coefficients on the unemployment gap) is set to zero. The results are remarkable. No fancy econometrics here: simply setting this slope to zero makes the same contribution as using a TV-NAIRU. Indeed, the simulation root mean-squared forecast error is far smaller than those in the first two columns of table B1 and in table 3.

Performing an econometric test of the stability of this coefficient involves some subtleties. It is tempting simply to test for a break in 1992:4 but this would be misleading, because 1992:4 was after all chosen by preliminary data analysis as a likely candidate for a break date. That is, the natural t test has data-snooping bias. One way around this is to use a test for a break at an unknown date, such as the Quandt likelihood ratio test. Doing so yields a striking finding: the hypothesis of stability of the slope coefficient is rejected at the 5 percent level, and the break date is estimated to be 1993:1! When a regression is estimated incorporating this break date, the coefficient through 1992:4 is negative and statistically significant, but the coefficient for the post-1992 period is slightly positive and statistically insignificant.

One is therefore left with two competing interpretations of the recent unemployment-inflation experience, both of which are consistent with the data. Either the Phillips curve is stable, except for a shift in the NAIRU, or the Phillips curve is now defunct, and there is no link from the unemployment gap to changes in inflation. The implications of these two explanations are quite different, to put it mildly. But which is correct?

It is the job of a discussant to raise questions, not to answer them. Nonetheless, it is possible to provide some insights by examining other indicators of aggregate activity as predictors of inflation. Gordon looks at capacity utilization; another such indicator is housing authorizations (building permits). The fourth and fifth columns of table B1 report

results for specifications identical to that in the first column, except that the unemployment rate is replaced by each of these indicators, respectively, and the "natural rates" of capacity utilization and housing starts are held constant at pre-1993 estimated values. The results are striking. These circa-1992 equations have the same forecast error for 1998:2—after adjusting for improvements in CPI measurement and the actual paths of the traditional supply shocks, the forecast error is a mere −0.27 percentage point. Moreover, their root mean-squared forecast errors are even smaller than for the unemployment specifications that resort to a falling NAIRU and are comparable to the specification with a zero slope. These relations do not appear to exhibit a break in the early 1990s: tests for a change in the sum of the coefficients on these variables in 1993:1 are not significant at the 10 percent level when evaluated using the Quandt likelihood ratio test critical values. Finally, note that favorable supply shocks play the same role for these specifications as for the specifications based on the unemployment rate. Were it not for these favorable supply shocks, inflation today would be what it was in 1994.

I suspect that the results in the fourth and fifth columns are typical of what one would find using other proxies for activity. Recently, Mark Watson and I have examined the performance of a host of other indicators for forecasting inflation over the past two decades.[6] The details of our exercise differ from the analysis here: our data are monthly, the focus is on one year ahead forecasts, and the methodology is to simulate real-time forecasting by recursively updating the estimated models. Yet the overall message is similar. During the 1990s, the unemployment rate is a poor predictor of inflation; in addition to housing starts and capacity utilization, inflation is predicted well by industrial production, employment growth (as opposed to the unemployment rate), manufacturing and trade sales, and new composite activity indexes.

All of the above suggests that the relation between inflation and many real economic indicators has been stable in the 1990s. This, in turn, constitutes evidence against the new economy view that the link between inflation and aggregate activity has been weakened or broken, and is consistent with the view that the unemployment-based Phillips curve has been stable except for a drop in the NAIRU. Thus the puzzle

6. Stock and Watson (1998a).

is not why inflation has been so low given the unemployment rate, but rather, why the unemployment rate has been so low given inflation.

When viewed this way, the answer does not lie in an investigation of special factors that have held down price inflation, such as computers and medical prices. It suggests, instead, a closer look at labor markets and what changes, if any, have occurred in these markets in the 1990s but not elsewhere. Such factors might include the role of information technology in facilitating job searches and changes in the welfare system in the United States. These results also underscore the fact that the unemployment-based Phillips curve is but one of several tools that economists should use when forecasting inflation.

In summary, Robert Gordon has provided a simulating paper that continues his careful and important research into the relations between wages, prices, and the unemployment rate. This work emphasizes the importance of good fortune—in the guise of favorable supply shocks—in explaining the current state of the U.S. economy. Significant questions remain, however, about what special features of labor markets have led to instability in forecasting relations based on the unemployment rate but not on other aggregate indicators.

General discussion: Participants discussed economic developments that could account for a downward shift in the natural rate of unemployment. Benjamin Friedman noted several changes in the labor market that Lawrence Katz has frequently emphasized. There are a million more people in prison now than ten years ago in the United States, and they are probably drawn disproportionately from the ranks of people who might otherwise be unemployed. Today, a large fraction of new jobs involve sitting at keyboards, which implies that jobs are much more flexible across companies and industries than in the past. Finally, about one out of seven job openings today is filled through temp firms. Friedman recalled that a well-functioning employment service has long been advocated as the way to reduce the natural rate and reasoned that we have gotten the benefits of an employment service, but through private temp firms rather than the public sector.

William Dickens was skeptical about the importance of some of these labor market changes, noting that the growth in the prison population did not begin in 1993 and that other changes, such as people with poor employment prospects going off welfare and into the labor force, would

have made the current natural unemployment rate go up. Nevertheless, he thought it possible for the United States to sustain unemployment in the 4 percent range without inflation. His own bivariate vector autoregressions with prices and wages showed that price, rather than wage, innovations explained inflation, not only in the supply shock periods of the 1970s but also in the 1960s and 1980s. He speculated that firms may have planned for capacity additions on the expectation that the Federal Reserve would target the prevailing estimates of the natural unemployment rate. If capacity utilization was responsible for price pressures, the natural rate becomes a self-fulfilling prophecy. A series of favorable supply shocks have now allowed the economy to break this self-fulfilling prophecy, and capacity utilization is low enough to permit lower unemployment rates than would have been predicted from history.

William Brainard noted that the original motivation behind Phillips's idea was that unemployment is a good measure of pressure in labor markets. Since labor costs are so important, a tight labor market with rising wages would be closely associated with rising inflation. He reasoned that the whole NAIRU idea is damaged if that mechanism turns out not to be central to the inflation process. James Duesenberry suggested using wage equations that include vacancies as an explanatory variable. If wages are driven by how difficult it is to fill vacancies, a shift in the Beveridge curve in the direction of less unemployment for given vacancies would be consistent with the idea that the labor market functions more effectively today. He noted that unemployment has historically been used in price and wage equations due to the absence of reliable data on vacancies, but this procedure would miss apparently important recent labor market developments.

Edmund Phelps commented that analysis of the natural rate should make use of the rich literature on labor markets, much of it from Europe, and should employ real explanatory variables such as real prices, real policies, and institutions. He saw good evidence that the natural rate declined over the past ten years and that a sharp fall occurred in the past five years. In particular, he cited the downward trend in the proportion of the labor force who exhibit relatively high unemployment rates and the construction boom that dates from late 1993. He reasoned that the capital goods boom might have lowered the natural rate through two channels. First, it might have raised employment when it pulled up real prices in the newly attractive capital goods sector; and

second, even after capital has shifted to that sector from the consumer goods sector, construction is relatively labor intensive.

Charles Schultze observed that because Gordon's NAIRU for wages was much higher than his NAIRU for prices, his results implied the implausible outcome that the trend labor share would increase indefinitely. He also questioned the formulation of Gordon's wage equation—in which the dependent variable is nominal wage growth minus the trend growth rate of productivity—because it assumes an immediate adjustment of wages to changes in the productivity trend. Finally, he disagreed with the paper's conclusion that the fall in inflation in the 1990s did not overturn the natural rate framework. He thought it was more accurate to conclude that *if* the natural rate has consistently been a valid explanation for the relationship between inflation and unemployment, the paper provides an estimate of that rate and how it has changed over time.

George Perry pointed out that Gordon's estimated NAIRU was above the actual unemployment rate in 1962. Given that in 1962 policymakers and economists outside the government agreed on the importance of getting the economy moving because unemployment was so high, he felt this cast serious doubt on estimates using Gordon's framework. In the alternative, Stock suggested that the experts were wrong in 1962, because they were extrapolating previous favorable experience. Brainard remarked that since the NAIRU is a smoothed two-way estimate generated via Kalman filtering, the estimated NAIRU for the 1960s is influenced by the data in the entire sample, which puts a lot of faith in the particular way Gordon has allowed change to occur. More generally, he observed that the structure behind Gordon's simple econometric framework could have changed in many ways, affecting any of the estimated coefficients: one would get a different perspective on the evolution of the economy by allowing time variation in the sum of the coefficients on unemployment or prices, rather than only in the NAIRU. And the data probably would not distinguish well among such alternative ways in which the structure might have changed.

Robert Hall recalled that in 1994 Gordon had predicted an upward jump in the rate of inflation, based on the Phillips relationship as it existed at the time. This understandable error, he noted, called attention to the fact that the Phillips equations are forecasting, not structural, relationships: there is no attempt at identification, and no structural

interpretation can be given of the findings. For example, the estimated low forecasting power of wages for prices in Gordon's paper failed to diminish Hall's confidence that imposing a 10 percent increase in labor costs on businesses would promptly raise prices. He believed it an important finding of the paper that the joint behavior of prices, wages, unemployment, and capacity utilization has changed. This knowledge is important, even though the mechanisms behind the change are unclear. He pointed out that the paradoxical behavior of prices today is nothing new, if one takes a broader historical and cross-country perspective. Neither the collapse of U.S. prices in the 1920s nor the end of four big inflations described by Sargent could be explained by a Phillips-type relationship. He concluded that the behavior of prices has been different in different episodes and may not be predictable.

Hall also drew attention to recent theoretical models of the labor market that seem completely orthogonal to Phillips curve research. In the modern Diamond-Mortensen-Caballero framework, the wage for each job match is always at its theoretical equilibrium, and Ricardo Caballero has stressed the ability of this framework to explain variations in the natural rate of unemployment, both over time and across countries, in a way that has nothing to do with any concept of adjustment process. Being completely adjustment free, this framework is dramatically at odds with the ideas in Gordon's paper and any Phillips curve research.

References

Franz, Wolfgang, and Robert J. Gordon. 1993. "German and American Wage and Price Dynamics: Differences and Common Themes." *European Economic Review* 37(4): 719–62.

Gordon, Robert J. 1971. "Inflation in Recession and Recovery." *BPEA, 1:1971*, 105–58.

———. 1975. "Alternative Responses of Policy to External Supply Shocks." *BPEA, 1:1975*, 183–206.

———. 1977. "Can the Inflation of the 1970s Be Explained?" *BPEA, 1:1977*, 253–77.

———. 1982. "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." In *Workers, Jobs, and Inflation*, edited by Martin N. Baily. Brookings.

———. 1983. "'Credibility' vs. 'Mainstream': Two Views of the Inflation Process." In *Inflation: Prospects and Remedies, Alternatives for the 1980s*, edited by William D. Nordhaus. Washington: Center for National Policy.

———. 1984. "Supply Shocks and Monetary Policy Revisited." *American Economic Review, Papers and Proceedings* 74(2): 38–43.

———. 1990. "U.S. Inflation, Labor's Share, and the Natural Rate of Unemployment." In *Economics of Wage Determination*, edited by Heinz König. Berlin: Springer-Verlag.

———. 1994. "Inflation and Unemployment: Where is the NAIRU?" Paper prepared for the Meeting of Academic Consultants to the Board of Governors of the Federal Reserve System. Washington, December 1.

———. 1995. "Estimating the NAIRU as a Time-Varying Parameter." Paper prepared for the Panel of Economic Advisors to the Congressional Budget Office. (November 16).

———. 1997. "The Time-Varying NAIRU and Its Implications for Economic Policy." *Journal of Economic Perspectives* 11(1): 11–32.

Gordon, Robert J., and Stephen R. King. 1982. "The Output Cost of Disinflation in Traditional and Vector Autoregressive Models." *BPEA, 1:1982*, 205–42.

King, Robert G., and Mark W. Watson. 1994. "The Post-War U. S. Phillips Curve: A Revisionist Econometric History." *Carnegie-Rochester Conference Series on Public Policy* 41: 157–219.

Levy, Mickey D. 1997. "Slaying the NAIRU Myth," *Jobs and Capital* 6(Summer): 15–19.

Phelps, Edmund S. 1978. "Commodity-Supply Shock and Full-Employment Monetary Policy." *Journal of Money, Credit, and Banking* 10(2): 206–21.

Seskin, Eugene P. 1998. "Annual Revision of the National Income and Product Accounts." *Survey of Current Business* 78(August): 7–35.

Shimer, Robert. Forthcoming. "Why Is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual 1998*, edited by Ben S. Bernanke and Julio J. Rotemberg. MIT Press.

Sims, Christopher A. 1987. "Comment." *BPEA, 1:1987*, 117–20.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment, and Monetary Policy." *Journal of Economic Perspectives* 11(1): 33–49.

Stock, James H. 1998. "Monetary Policy in a Changing Economy: Indicators, Rules, and the Shift towards Intangible Output." Paper prepared for Bank of Japan Conference on Monetary Policy in a World of Knowledge-Based Growth, Quality Change, and Uncertain Measurement. Tokyo, June 18–19.

Stock, James H., and Mark W. Watson. 1998a. "Forecasting Inflation." Unpublished paper. Harvard University, Kennedy School of Government and the Woodrow Wilson School.

———. 1998b. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model." *Journal of the American Statistical Association* 93(March): 349–58.

U.S. Bureau of Labor Statistics. 1997. "The Experimental CPI Using Geometric Means (CPI-U-XG)." Unpublished paper. Department of Labor (April).