

Reflections on the Natural Rate Hypothesis

Joseph Stiglitz

Few concepts in economics with an acronym as unpleasant as NAIRU have attained as much public prominence in the past 10 years. Who, 10 years ago, would have thought that White House economic press conferences would begin with a question concerning the administration's estimate of the NAIRU? The issue typically posed is not whether there is a NAIRU, but what its level actually is.

I have become convinced that the NAIRU is a useful analytic concept. It is useful as a theory to understand the causes of inflation. It is useful as an empirical basis for predicting changes in the inflation rate. And, it is useful as a general guideline for thinking about macroeconomic policy. But these points are controversial within the profession; for example, half of the five papers in this symposium are openly hostile to the concept of the NAIRU. (I am including half of the Staiger, Stock and Watson paper in this calculation, for reasons that will become apparent as you read it.)

Let me begin my defense of the NAIRU by asserting that it is a very well defined concept. In fact, the main advantage of using the somewhat ugly term NAIRU instead of its more euphonious synonym, the natural rate, is that each time we use the term we are reminded of its meaning. The NAIRU is defined as the nonaccelerating inflation rate of unemployment; that is, the rate of unemployment consistent with an unchanging inflation rate. When unemployment is below the NAIRU, there is pressure for the inflation rate to rise; contrarily, when unemployment is above the NAIRU, there is pressure for the inflation rate to fall. In fact, economists should probably be calling it the NIIRU, for "nonincreasing inflation rate of unemployment," but this remedy for the slipped derivative probably sounds worse than the original disease.

I think of the theory behind the NAIRU essentially as a description about how the economy behaves out of equilibrium. When unemployment is below the

■ *Joseph Stiglitz is Chairman, Council of Economic Advisers, Washington, D.C.*

NAIRU, real wage demands are greater than the amount firms are willing to pay (at prevailing prices and price expectations). These incompatible wage-setting aspirations and price-setting plans are resolved through a wage-price spiral in which workers do not get the real wages they expect and firms do not get the real prices they expect. Consequently, inflation will be higher than was expected. Because inflation is pretty close to a random walk (as long as unemployment is not too far away from the NAIRU), today's inflation rate is a good proxy for the expectation of tomorrow's inflation rate; therefore, the disequilibrium will be resolved by a rising level of inflation. Equilibrium—defined in this case as inflation stable at a level equal to its expectation—is only achieved when unemployment rises to the NAIRU, making the behavior of wage setters compatible with the behavior of price setters. In other words, the natural rate hypothesis views changes in the inflation rate as a labor market phenomenon whose magnitude can be proxied by a particular measure of labor market slack: the unemployment rate.

In this paper, I want to elaborate on three criteria for evaluating whether the NAIRU is useful for academics and policymakers. First, does the deviation of unemployment from some natural rate provide a robust and useful way to predict changes in the inflation rate? This issue is an empirical one. It is a necessary condition for defending the natural rate hypothesis, but it is hardly sufficient. Thus the second criterion, can economists explain why the NAIRU changes over time? This is important if policymakers are to stay ahead of the curve, since in this rapidly evolving economy, it is important to understand the structural underpinnings of the coefficients we rely on for predictions. I am realistic about the fact that we will never be able to do this perfectly, so the third criterion asks whether the NAIRU is a useful way to frame policy discussions despite all of the uncertainty surrounding its precise level and direction of change. To this end we must ask if the Phillips curve relation is sufficiently robust, and if the theory of the NAIRU is sufficiently rich, to provide a fruitful tool for thinking about the perennial problem of balancing rapid growth with price stability in the face of uncertainty. The following sections discuss the natural rate hypothesis in the context of these three criteria: its ability to explain changes in the inflation rate, our ability to explain why the NAIRU changes, and the usefulness of these ideas for policy discussions.

The Empirical Success of the Phillips Curve

The most basic implication of the natural rate hypothesis is that changes in the inflation rate are (in large part) a labor market phenomenon whose magnitude can be proxied by a particular measure of labor market slack—the unemployment rate. I believe the data support the view that unemployment is a good predictor of the direction that inflation will move. One nonparametric way to approach this issue is to compare the direction of the change in inflation (as measured by the core experimental CPI) with the level of demographically adjusted unemployment.¹ Since 1960, inflation

¹ The demographically adjusted unemployment rate takes the actual unemployment rates for different demographic groups and weights them by their labor force shares in a given base year, in this case, 1993.

rose in 26 of the 32 quarters when the unemployment rate was below 5 percent, but inflation fell in 24 of the 27 quarters when unemployment was above 7 percent. This empirical regularity is very strong. It also suggests an upper bound on our uncertainty about the NAIRU: it looks as if it has been contained within 5 and 7 percent for the period from 1960 to the present. Note that this is the NAIRU for the demographically adjusted unemployment rate; mapping it into a concept that we can compare to the actual history of unemployment would widen the band somewhat.

This is an example of the proper test of the NAIRU hypothesis itself. In contrast, almost all other tests of the NAIRU assume structure for the quantitative relation between unemployment and inflation. We can generalize this relatively ad hoc procedure into more formal and powerful nonparametric tests of the natural rate hypothesis. At the CEA, we have done nonparametric tests that find very strong rejections of the proposition that changes in inflation are independent of the level of unemployment. In particular, the data show that the probability of inflation increasing over a given year is almost monotonically decreasing in the level of unemployment at the beginning of that year.

I want to go beyond these tests to discuss some tentative conclusions we have reached in our more structural research at the CEA, many of which are consistent with the empirical research in this area, as exemplified by some of the papers in this symposium. When we regress changes in the inflation rate on past values of unemployment, regardless of the precise specification or other variables included, we get *t*-statistics on lagged unemployment on the order of 4 or 5. These numbers are pretty standard for the literature; in this symposium, both the papers by Gordon and by Staiger, Stock and Watson find that the coefficient on unemployment is always significant at the 1 percent level. Such results make us just about as certain that unemployment has predictive power for changes in the inflation rate as we can ever be in empirical research.

Just saying that unemployment matters, however, is not enough. To defend my thesis that the NAIRU is a useful concept, I need to demonstrate how much it matters. In our regressions, we find that keeping the unemployment rate one percentage point below the NAIRU for one year will result in the inflation rate increasing by between 0.3 and 0.6 of a percentage point. That magnitude is definitely large enough to command the attention of policymakers. Finally, our analysis indicates that at least 20 percent of the variation in the inflation rate can be explained by unemployment alone. This figure serves as a reminder that the actual inflation process—and the policy decisions that must be made based on it—is much more complicated than a simple link between the NAIRU and inflation. However, economists should not readily surrender any parsimonious concept that can explain 20 percent of something important.

What do Economists Know about Changes in the NAIRU?

The second criterion for the NAIRU being a useful concept is that we can explain why it changes over time, predict those changes to some degree, and possibly even identify policies that might alter it. When Milton Friedman first proposed

the natural rate hypothesis in his presidential address to the American Economic Association in 1968, it sounded like royal edict had established the natural rate as another one of the universe's invariant constants. Today, there is general recognition that if a NAIRU exists, it must be changing over time.

A few years ago the typical estimates of the NAIRU ranged from 6.0 to 6.2 percent. If this natural rate remained operative today, the below 6 percent unemployment from August 1994 through the latter part of 1996 should have increased inflation. For example, if the natural rate were 6.2 percent and we assume that each year the unemployment rate is a percentage point below the NAIRU inflation rises by 0.5 of a percentage point, then the average unemployment rate of 5.6 percent from August 1994 to August 1996 would have led to a 0.6 point increase in the inflation rate. Instead, inflation, as measured by the yearlong increase in the core CPI, fell from 2.9 percent to 2.6 percent over that period.²

Through 1995 and 1996, inflationary pressures have been milder than in previous periods when unemployment was this low. For adherents to the NAIRU approach, the reason is clear: the natural rate of unemployment has fallen. In fact, the conclusion that the NAIRU is lower today is supported almost unanimously by econometric research in this area. Both papers in this symposium that estimate time-varying NAIRUs agree with our research at the CEA in showing that the natural rate has been falling steadily since at least 1984. Our own preliminary finding is that the NAIRU has fallen by around 1.5 percentage points since its peak in the early 1980s. I should caution, however, that the uncertainty surrounding this estimate, both the formal standard errors and the uncertainty over whether we are estimating the correct model, is very large.

The three main forces in this decline are the changing demographics of the labor force, the fact that productivity growth has become more in line with worker expectations and the general increase in the competitiveness of the labor and product markets. Accounting for demographic change is straightforward. Imagine that each demographic group has its own unchanging natural rate of unemployment: higher for teenagers than adults, higher for women than for men, and so on.³ Then imagine that the only source of changes in the NAIRU are changes in the proportion of these different groups in the labor force. If we assume that demographic changes had about the same effect on the NAIRU as they have had on observed unemployment, then about one-third of the 1.5 percentage point decline in the NAIRU can be attributed to demographic changes. The single most important demographic change is the aging of the baby boomers—we now have a more mature

² The weakness of inflationary pressures is also illustrated by the fact that in August 1996, wage growth was still lagging behind the sum of inflation and trend productivity growth. As noted in the text, inflation can be taken as 2.6 percent, and I consider trend productivity growth to be 1.1 percent, for a total of 3.7 percent. However, even after the fairly strong growth in earnings in the August 1996 unemployment report, average hourly earnings were growing at 3.6 percent annually. The broader-based Employment Cost Index shows that compensation increased only 2.9 percent over last year's level.

³ There are reasons why these groups might have different "natural rates." New entrants to the labor force (teenagers) typically are involved in more search, in order to find a good "match"; hence, there is likely to be a higher level of frictional unemployment.

labor force, with greater representation of age groups that have traditionally had lower unemployment rates.

The second explanation for the change in the NAIRU, which is linked to productivity growth, can be called the “wage-aspiration effect.” Although neither the level or rate of change of productivity has any long-run effect on the unemployment rate—witness the fact that unemployment has been about the same over the course of a century of massive productivity growth and large shifts in its trend growth rate—changes in the productivity growth rate can have temporary effects on the natural rate of unemployment. The intuition is that workers’ demands for increased real wages depend on their past rate of change, possibly because of the psychological fact that people get accustomed to a certain rate of increase in their standard of living. Thus, after a fall in the productivity growth rate, workers will initially demand real-wage growth based on their previous experience and thus faster than the increase in productivity, which puts upward pressure on the inflation rate and requires a higher level of unemployment for the economy to stay in equilibrium. But this increased NAIRU is only temporary, either because the productivity shock itself is temporary, or because workers will eventually moderate their demands in response to permanently lower productivity growth. Either way, the NAIRU can return to its preshock level. This effect, called the wage-aspiration effect, surely contributed to a rising NAIRU in the 1970s and early 1980s, in the aftermath of the productivity slowdown. However, workers have now had a lot of time to adjust their real-wage aspirations down to reflect the slower productivity growth, which should be helping the NAIRU rebound to its earlier, lower rate.

We may be able to make an even stronger statement than this. The official productivity series may seriously underestimate our actual productivity growth. If you calculate productivity growth from the income-side measures of the economy—as was Bureau of Labor Statistics policy until January 1996—then you find that productivity growth from the beginning of 1993 through mid-1996 has been 1.4 percent per year. If these GDI-based numbers turn out to be more accurate than the GDP-based numbers, then we can expect that the wage-aspiration effect has actually been working in reverse to lower the NAIRU. Our estimates at the CEA indicate the wage-aspiration effect has lowered the NAIRU by around 0.5 percentage point since the early to mid-1980s. However, I want to emphasize that there is a lot of uncertainty in this estimate.

The most likely suspects for the remaining decline in the NAIRU all fall under the heading of increased competitiveness of the product and labor markets. This is partly the consequence of the increased opening of markets at home and abroad through regulatory reform and trade-opening agreements. Although trade is still a relatively small part of our economy—around 10 percent—the fact that much of the U.S. manufacturing sector faces potential competition is enough to deliver wage restraint. Changes in labor markets, such as decreasing rates of unionization, have also had some salutary effects on inflation, regardless of their other impacts. Although quantifying these general notions of competitiveness and the institutional structure of the labor market is extremely difficult, it is not implausible to think that they account for the remaining 0.5 percentage point decline in the NAIRU since the early 1980s.

In addition to demography, wage aspiration and competitiveness, a fourth factor might currently be affecting the NAIRU and might affect it even more in the near future: hysteresis. The idea that sustained high unemployment will gradually raise the natural rate of unemployment is perhaps most familiar in the European context. The intuition behind the claim emphasizes that high sustained unemployment decreases both the work and job-search skills of the unemployed (outsiders) at the same time that those who remained employed (insiders) want to maintain wages at the expense of expanding employment. The reverse of these forces might be at work in the U.S. labor market today. Research has not conclusively shown that hysteresis is one of the forces that has lowered the NAIRU in the American economy. But if it does, then high unemployment is even worse than we thought, because it raises the NAIRU, and lower unemployment is even better than we thought, because it reduces the NAIRU.⁴

These trends may continue to reduce the NAIRU in the future. Demography will continue to lower the natural rate of unemployment as the current bulge of workers in the 35–49 age bracket move into the even lower unemployment 50 plus age bracket, although this effect will be partly offset by the increase in teenagers resulting from the so-called “baby-boom echo.” If hysteresis does work in reverse, the current spell of low unemployment should help to generate a lower NAIRU in the next few years. The other two factors affecting the natural rate are harder to predict, although the competitiveness of the economy seems likely to increase with the process of increased opening of international trade and continued reforms in the U.S. regulatory structure.

The NAIRU and Macroeconomic Policy

The third criterion that the natural rate hypothesis has to pass if it is to maintain its claim to an important place in the analytic toolkit of economists is that it be useful for discussing and analyzing policy even when there is uncertainty about its current and future level. In a sense, it is hard to think about macroeconomic policy without the concept of the NAIRU. Clearly, we would like to get unemployment as low as possible, without inflation accelerating. If there is no clear, systematic relation between inflation and unemployment, why wouldn't policymakers simply keep trying to push unemployment lower and lower?

Thus, policymakers must base their actions on some perceptions about the consequences of changed levels of economic activity on the rate of inflation. The more subtle question, put roughly, is, “Should policymakers ‘target’ the NAIRU?”

⁴ In proposing the concept of the natural rate of unemployment, Milton Friedman separated the demand-management choices of policymakers from the constraints on those choices. While the argument here agrees with Friedman that supply-side labor market factors like demography are of primary importance, the concept of hysteresis suggests that rather than macroeconomic policymakers being faced with a NAIRU given to them solely by the supply side of the economy, the natural rate also depends on the evolution of aggregate demand. In short, Friedman's dichotomy may be overstated.

To answer that, we need to know what are the consequences of basing policy on a mistaken estimate of the NAIRU.

There is a famous dictum of Alfred Marshall that the world is continuous, and I think this applies to macroeconomic policy as well: if the economy heads a little below the NAIRU, then we can expect a little more inflation. This view stands in opposition to the view, more common in nonacademic circles, that the NAIRU is like a precipice: take one step over it, and you fall into a spiral of rapidly accelerating inflation. The evidence simply does not support this view. In fact, the evidence suggests that we can go even further than Marshall and say that in this case the world is not only continuous but approximately linear: if you hold the unemployment rate below the NAIRU for a year, the inflation rate rises by about 0.3 to 0.6 of a percentage point per year. Contrary to the accelerationist view, not only does the economy not stand on a precipice—with a slight dose of inflation leading to ever-increasing levels of inflation—but the magnitude by which inflation rises does not increase when the unemployment rate is held down for a prolonged period of time.⁵ Thus, small mistakes have only small consequences.

What if macroeconomic policymakers do make the mistake of holding unemployment below the NAIRU for a period of time, and inflation rises as a result? What is required to reverse this type of a mistake? The traditional view was that the Phillips curve was convex—the costs of a disinflation in terms of lost employment were viewed as substantially higher than the benefits of greater employment from a similarly sized inflation. The underlying intuition was that contractions decreased output by creating more slack in the economy, thus raising unemployment, while expansions ran into capacity constraints, quickly resulting in higher prices with little employment gain. In this environment, a risk-neutral policymaker would be averse to experimenting with the unemployment rate. In contrast, if the Phillips curve is linear, then experimentation has zero expected cost—the cost of a deflation in terms of temporarily increased unemployment is exactly equal to the benefit of the lower unemployment that caused the inflation in the first place.

I want to go even further and suggest that there is some evidence that the Phillips curve might be concave. This concavity is consistent with the literature on asymmetric price adjustment, which shows that in monopolistically competitive markets, producers might adjust prices down to avoid being undercut by a rival but will be more reluctant to raise prices even in the face of generally rising prices. Our empirical research at the CEA has found that when we run Phillips curve regressions allowing for a kink at the NAIRU, we find that the best fit is with a concave function. Although this concavity is relatively mild, it is statistically significant. This particular nonlinearity, if in fact it holds up in future research, suggests that even risk-averse policymakers might want to engage in moderate experiments with the unemployment rate. If, hypothetically, they then discovered that the unemployment rate had fallen 0.5 percentage points below the NAIRU for a six-month period, then with a

⁵ One way to test this hypothesis is to estimate accelerationist models that nest nonaccelerationism as a restricted case. We have found that the coefficients on the acceleration terms are insignificant; in other words, the restricted model fits the data almost as well as the accelerationist model. This is true in a variety of specifications.

concave Phillips curve it would take less of a rise in unemployment, possibly 0.4 of a percentage point above the NAIRU for six months, to wring the consequent inflation back out of the system. If the hysteresis effect proves important, then these conclusions are reinforced.⁶

Conclusion

When I discuss the NAIRU with the press or in public, I am usually asked about what I think the Federal Reserve will or should do. The answer I and other members of the Clinton administration give is always the same: we do not comment on Fed policy. This stance has served the economy well by contributing to a sense of confidence that our economy is well managed. But there are also good economic reasons to avoid being pinned down to a particular NAIRU. The NAIRU can be moving. Unemployment explains only a portion of changes in inflation, and there are a variety of other economic goals besides simply fighting inflation. The analytic framework connecting NAIRU to inflation is only one of many ingredients that go into the formulation of macroeconomic policy. Moreover, there are broader concerns of macroeconomic policy: if we can understand better the determinants of the NAIRU and the inflationary process, perhaps we can better design policies that reduce the NAIRU and alter the tradeoffs.

The natural rate hypothesis passes all three tests I have outlined. Unemployment is an empirically successful way to predict changes in the inflation rate. Economists have good explanations for the now undeniable fact that the NAIRU has fallen in the last 15 years, and forces like demography and hysteresis will continue to put downward pressure on the natural rate. Finally, the natural rate provides a useful framework for thinking about policy questions even if there is considerable uncertainty about its exact magnitude. In particular, although no one knows exactly where the NAIRU is, these results suggest that in testing the waters, we do not risk drowning. If need be, we can always reverse course. But by experimenting, and showing some hesitation about restraining the economy through higher interest rates or other methods as the NAIRU draws nigh, we might learn a little more about the depth of the waters and possibly become better swimmers in the process.

■ *An earlier version of these remarks was delivered, along with preliminary drafts of the rest of the papers in this symposium, at a conference on "Perspectives on the Natural Rate of Unemployment," which was held at the Georgetown University Conference Center on September 13, 1996. Financial support for the conference was provided by the Andrew W. Mellon Foundation. Dr. Stiglitz wishes to thank Jason Furman for his help in preparing this paper.*

⁶ With hysteresis effects, policymakers face a complicated dynamic programming problem, in which policy today must take into account the impacts on the future NAIRU. This discussion has analyzed only some of the aspects of macroeconomic policymaking. A fuller discussion would take into account factors such as costs of adjustment and of variability in output and unemployment, and dynamic learning effects; for example, are there policies that can affect the degree of uncertainty about the value of the NAIRU or of policy tradeoffs?