

GEORGE A. AKERLOF

University of California, Berkeley

WILLIAM T. DICKENS

Brookings Institution

*Unfinished Business in the
Macroeconomics of Low Inflation:
A Tribute to George and Bill
by Bill and George*

THE RETIREMENT OF GEORGE Perry and William Brainard as editors of the *Brookings Papers* gives us an opportunity to say publicly what we have often said privately: the *Brookings Papers* is an important national institution. It is important because it has set the right tone for U.S. macroeconomic policy. George and Bill are both Keynesians, not just in the narrow tradition of IS-LM models with Phillips curves, but also in their broader methodological approach to macroeconomics. The *Brookings Papers* has always reflected their view that macroeconomics should be a pragmatic and judicious mixture of theory and common sense, informed by statistical analysis. That of course reflects the methodology of Keynes, throughout his life and especially in *The General Theory*. We think that U.S. macroeconomic policy has benefited enormously from such a balanced pragmatic-empirical approach.

Anyone who doubts the benefits of such an approach to macroeconomics needs only look to the north—to Canada—where doctrinaire use of an extreme form of natural rate theory in the 1990s led policymakers to push inflation too low, resulting in an unusually wide unemployment gap

We thank Robert Akerlof and Janet Yellen for helpful comments. George Akerlof is grateful to the Canadian Institute for Advanced Research and to the National Science Foundation under Research Grant SES 04-17871 for invaluable financial support.

relative to the United States. This is just one example where a more empirical, more nuanced macroeconomics—such as presented for the last thirty years in the *Brookings Papers*—has implications for national welfare. Macroeconomic policy, as no one appreciated better than Keynes, often even makes the difference between prosperity and depression.

It will be a hard task for the new editors to take the place of George and Bill, and we wish them luck. Neither this Bill nor this George has ever envied George and Bill their difficult job. They have edited the *Brookings Papers* the hard way. The easy way is to trawl the conferences and the economics department hallways for the best of what already exists. Instead, for the most part, George and Bill recruited people to push forward their pragmatic agenda for macroeconomic research. The influence of George and Bill, even in the late stages of producing these papers, is clear to all members of the Brookings Panel. We have continued to be amazed at how George and Bill could take the meeting drafts, which were not always in the best of shape, and quickly turn them into interesting, readable gems. In short, they have done the impossible: admittedly with some very good help from their authors, they have produced something like eight to twelve significant and relevant new papers on macroeconomics per year, year after year.

Some Background on This Paper

This volume in honor of George Perry and William Brainard is an opportunity to reflect back on the two Brookings Papers we wrote with Perry (with comments and occasional help from Brainard) on the macroeconomics of low inflation.¹ Although they took somewhat different approaches, these two papers had similar policy implications. Both found a significant cost of permanently low inflation in terms of permanently high unemployment.

Each paper examined the effects of a different type of money illusion. The earlier paper, “The Macroeconomics of Low Inflation,” examined the consequences of downward nominal wage rigidity, or resistance to nominal wage cuts. “Near-Rational Wage and Price Setting” examined the consequences of people thinking in nominal rather than in real terms when

1. Akerlof, Dickens, and Perry (1996, 2000).

inflation is very low. When this happens, a trade-off will emerge between inflation and unemployment—not just in the short run, but also in the long run when actual inflation and expected inflation are equal.

Although the existence of such a trade-off was previously well known, its magnitude was not. Both papers found surprisingly large long-run increases in unemployment from permanent reductions in inflation to zero. For example, in the benchmark simulation in the first paper, with limited nominal wage cuts for continuing workers, we found that a permanent reduction in annual inflation from 2 percent to zero would increase unemployment by 1.5 percentage points; when we estimated a benchmark Phillips curve using time-series data, unemployment increased permanently by 2.6 percentage points. In the benchmark “ignored expectations” simulations, a decline in annual inflation from 2.6 percent to zero yielded a 1.5-percentage-point increase in unemployment.² The preferred parameter estimates for our empirical model imply that moderate inflation will reduce unemployment by between 1.9 and 3 percentage points from what it would be with no inflation at all.³ The exact size of these trade-offs depends on the specific values of the parameters or the specific equations we estimate. But literally hundreds of simulations and estimates point to the same conclusion. There is a significant permanent trade-off between inflation and unemployment when inflation is low.

Why did these papers find such significant trade-offs? And why did they obtain these results so robustly? Both papers rely on one crucial assumption. Here we shall explain what that assumption is, why it is so crucial, and why, now as then, we think it is well justified. We welcome this opportunity not just to reminisce about the fruitful collaboration we had with George Perry, our coauthor on these papers as well as our editor, and with Bill Brainard, about a decade ago. By putting our results in a new perspective, we shall also complete some important unfinished business.

Two Calculations

The models presented in both papers posit a notional wage. This is the wage that firms would pay in the absence of any money illusion. In the

2. Akerlof, Dickens, and Perry (2000, p. 18).

3. Akerlof, Dickens, and Perry (2000, table 1, p. 32).

first paper, the departure of actual from notional wages was due to nominal wage rigidity; in the second, it was due to failure, at low levels of inflation, to incorporate inflationary expectations into the wage bargain. These manifestations of money illusion give rise to a wedge. At each level of employment this wedge causes a tendency for real wages to be permanently higher than the notional wages that emerge from the real wage bargains by an amount $S(\pi)$, where π is the rate of inflation. This wedge persists even in the long run when its effects are fully anticipated. Since $S(\pi)$ declines as inflation rises, equilibrium unemployment will rise as inflation falls.

We can then decompose the analysis into two parts. The first part asks: How large is $S(\pi)$? Or, equivalently, what is the seemingly exogenous tendency for real wages to be lower because inflation is π rather than zero? The second part asks: What change in unemployment results from this shift $S(\pi)$? Or, equivalently, what is the S -multiplier?

It seems natural that if S is expressed in percentage terms, the effects on employment can be obtained by multiplying S by the elasticity of demand for labor. This, in turn, suggests low estimates of the effects of S on unemployment. Daniel Hamermesh's 1996 survey reports that microeconomic studies of the elasticity of demand for labor find it to be in the range of 0.2 to 0.9.⁴ This would then suggest that the effect of inflation on unemployment is in the range $0.2 S(\pi)$ to $0.9 S(\pi)$.

It turns out, however, that this method of calculation is not correct for several reasons, one of which is that it ignores changes in the aggregate supply of labor. In the theoretical models presented in our two papers, we showed that one could approximate the impact of changing rates of inflation on equilibrium unemployment by dividing $S(\pi)$ by the slope of the short-run price Phillips curve. (To be precise, that slope is the change in price inflation with respect to the rate of unemployment in the short run, when expectations about inflation are fixed.) That use of the slope of the Phillips curve raises the following question: With a wedge of $S(\pi)$, how much additional unemployment induces actual prices to equal expected prices? It turns out that the S -multiplier is the inverse of the slope of the price Phillips curve. Conventionally, it is believed that such slopes are approximately $\frac{1}{2}$; so multiplying the difference $S(0) - S(\pi)$ by 2 gives the reduction in unemployment from permanently increasing inflation from 0

4. Hamermesh (1996).

to π . Thus, our models predict impacts on unemployment as much as ten times as high as from the labor demand approach.

This then takes us to our unfinished business. Why did we not obtain estimates that were very much lower than we did? Where is Waldo? That is, where is the elasticity of demand for labor? Here we shall construct a model that shows where the elasticity of demand for labor is hiding, namely, as a term in the price Phillips curve. It turns out that the slope of that Phillips curve reflects two factors. The first of these is the relationship between (notional) wages and employment in the wage setting equation that can be thought of as the aggregate elasticity of supply of labor. The second is, as intuition dictates, the elasticity of demand for labor.

But here the theory and our derivation turn out to be useful because “the demand for labor” is an ambiguous term. For example, one simple *microeconomic* definition of the demand for labor is employment by an *individual firm or individual industry* as a function of its own wage, with aggregate income and the prices and wages of other firms or industries fixed. The conventional *macroeconomic* demand for labor asks a different question: How will employment change for *all* firms if the real wages for all of them, and aggregate income, are changing? This is the appropriate demand concept in our formula for the *S*-multiplier. We shall argue that this elasticity, in contrast to the conventional microeconomic estimates of the elasticities, is sufficiently large to be fully consistent with our high values for the *S*-multiplier.

The micro and the macro concepts of the demand for labor will also differ depending on whether capital is assumed to be fixed or variable. As we shall see, for the macroeconomic definition relevant for the *S*-multiplier, the elasticity of demand, which is high when capital is fixed, is still higher when capital is variable. The consistency with our previous estimates is thus unaffected by considering capital as either fixed or variable.

The Theory

We adapt a model from Charles Bean’s survey of European unemployment.⁵ Thus, all of the assumptions are rather standard.

5. Bean (1994).

In the usual theory of the Phillips curve, the rate of inflation is the sum of two components. By definition, the rate of inflation is (approximately)

$$(1) \quad p_t - p_{t-1},$$

where p_t denotes the log of the price level at time t . Next $p_t - p_{t-1}$ can be decomposed into two terms:

$$(2) \quad p_t - p_{t-1} = (p_t - p_t^e) + (p_t^e - p_{t-1}),$$

where p_t^e is the log of the expected price level at time t (with those expectations made at $t - 1$). The second bracketed term in equation 2 is (approximately) the rate of inflation expected at time $t - 1$ (since p_{t-1} is known at $t - 1$); we shall denote it as π_t^e . The first bracketed term is the (log of the) desired price of a representative firm relative to the expected general price level. According to equation 2, then, inflation in excess of expectations depends upon the extent to which this desired price exceeds the expected price. As can be seen, inflation will be equal to expectations if this desired price equals the expected price.

A Phillips curve then develops from the determination of relative prices. It is the result of two considerations. The first is the markup of prices over variable costs. The second is what those variable costs might be. We begin by describing the price markup.

The Price Markup

We shall assume that the firm produces output according to the following production function:

$$(3) \quad Q_t = F(N_t, K_t),$$

where Q_t is output, N_t is employment, and K_t is the capital stock. We assume that the firm faces a demand function for its product of the form

$$(4) \quad D_t = A(P_t/\bar{P}_t)^{-\beta} Y_t,$$

where D_t is the demand for the firm's product, Y_t is aggregate real demand, P_t is the price charged by the individual firm for its product, and \bar{P}_t is the average price level for the whole economy. β , of course, is the price elasticity of demand for the product of this representative firm. Note that we have made demand for the product of this firm propor-

tional to aggregate income. Like demand for the average firm, its demand will be proportional to aggregate income for a given relative price of its product.

This yields a familiar equation, which says that the firm hires labor until the marginal revenue product is equal to the real wage:

$$(5) \quad \partial F / \partial N = MW_i / P_i,$$

where M is the markup factor, which will be $\beta / (\beta - 1)$; W_i is the nominal wage; and P_i is the price of the firm's product.

Variable Costs: Wages

We write the equation for the notional real wage then as

$$(6) \quad w_i - p_i^e = a - bu_i,$$

where u_i is the rate of unemployment. In a longer paper we would derive this relationship from microeconomic assumptions as our 1996 and 2000 papers did, but such a relationship can be derived from a wide range of bargaining, efficiency wage, or search models. We assume that all firms treat this wage as the exogenously given supply price for labor.

The Effects of Money Illusion

To consider the consequences of either money wage rigidity or ignored inflationary expectations, we must then describe the additional effect on wage outcomes. We shall illustrate here with money wage rigidity. Our paper on downward wage rigidity described how censoring of negative wage changes caused a tendency for wages to be higher than the notional relation in equation 6. This censoring occurs because of workers' reluctance to take cuts in money wages (and perhaps also firms' reluctance to give them). Such stickiness causes a tendency for money wages to be higher than they would be otherwise by an amount S , which depends on inflation π . The lower is π , the more frequent (and larger) will be the truncations of notional wage changes at zero. Our 1996 paper described the recursive nature of S ; it also described how S will vary with the steady-state value of π . In steady states with higher inflation, there will be less truncation, and therefore S will take on lower values: $S = S(\pi)$, with $S'(\pi) < 0$. $S(\pi)$ represented here in percentage terms. So the presence of

S leads to a shift in the expected real wage relative to its notional value in equation 6:

$$(6') \quad w_t - p_t^e = a - bu_t + S(\pi_t).$$

The Phillips Curve

We next take the Taylor series expansion of the logarithmic transformation of the markup equation 5. In addition, we express labor demand in terms of the unemployment rate, u_t , rather than in terms of the employment rate. This restates equation 5 in log form as

$$(5') \quad p_t - p_t^e = g - \alpha u_t + m + w_t - p_t^e,$$

where g is a constant from the Taylor series expansion.

Note that when prices equal price expectations, equation 5' can be interpreted as an implicit labor demand function. Since the equation is in log form, $-1/\alpha$ is the elasticity of demand for labor. That is precisely the *macroeconomic* elasticity of the demand for labor that results from the changing demand for labor according to equation 5 as the real wage changes for all firms.

Combining equation 5' with the notional wage equation 6' yields the real price that producers think they are setting as

$$(7) \quad p_t - p_t^e = G - (b + \alpha)u_t + S(\pi_t),$$

where $G = g + a$.

In turn, combining equation 7 with equation 2 yields a price Phillips curve of the form

$$(7') \quad \pi_t = G - (b + \alpha)u_t + S(\pi_t) + \pi_t^e.$$

Relating Unemployment to S

The rate of unemployment that will maintain the steady-state rate of inflation is then

$$(8) \quad u(\pi) = H + S(\pi)/[b + \alpha],$$

where H is a constant equal to $G/[b + \alpha]$.

For us the interesting parameter is the coefficient on $S(\pi)$. That determines the extent to which lower S will allow lower unemployment to be maintained at higher levels of inflation. And that coefficient is $1/[b + \alpha]$.

This coefficient has an intuitive interpretation. The parameter $-b$ is the slope of the macroeconomic supply of labor curve in $(w - p, u)$ space, and α is the slope of the labor demand curve in that space. $S(\pi)$ can be considered to be a wedge between the actual wage that firms pay (which is the basis of their labor demand) and the notional wage. Of course, the increase in unemployment is equal to the wedge divided by the sum of the absolute values of the slopes of the two lines.

One can look at what is happening another way. When $S(\pi)$ rises as π falls, one or both of two things must happen to keep firms' desired prices equal to expected prices. Either unemployment must grow to reduce the notional wage, or the marginal product of labor must rise, as it might with falling employment. The parameter b reflects the first effect and the parameter α the second.

In our two previous papers we assumed that the elasticity of output with respect to labor was 1. In that case equation 5 implies a constant real wage or an infinitely elastic macro demand for labor; α equals zero; and the unemployment effects of nominal rigidity depend only on the slope of the notional wage equation. But, more generally, the macro elasticity of demand for labor will matter as well. In the extreme, if the supply of capital is fixed and the elasticity of substitution of capital for labor is zero, the macro elasticity of demand for labor will be zero, the denominator of equation 8 will be infinite, and the S -multiplier will be zero.

Elasticity of Labor Demand

What, then, is a realistic value for α ? In general, $-\alpha$ is the elasticity of the marginal product of labor with respect to labor. Consider first what α would be if capital were fixed. With Cobb-Douglas production it is the capital coefficient (which is equal to the income share of capital with perfect competition in the factor market). More generally, with any production function with constant returns to scale, it is $(1 - NF_N/Y)/s$, where F_N is the derivative of F with respect to N , and s is the elasticity of substitution between capital and labor. The numerator is capital's share if there is perfect competition.

Typical values of the share of capital range from 0.2 to 0.4; the elasticity of substitution is normally assumed to be 1 or perhaps a bit less. These imply values of α that will not preclude a slope for the price Phillips curve of $1/2$. Thus, with conventional values of the parameters of the production function, equation 8 is consistent with S -multipliers that are quite large.

At business cycle frequencies it is natural to assume that capital is fixed, but we are concerned with the value of the S -multiplier in the long run, where there is a permanent trade-off between inflation and unemployment. For these longer-run considerations we should consider the values of α when capital is variable. Such variability can occur for two reasons: first, domestic saving will increase or decrease with the labor input as output and employment rise or fall; also, capital is potentially available from abroad. Both of these factors will cause the supply of capital to be more elastic and the long-run value of α to be less than with fixed capital. If in the long run there is an infinitely elastic supply of capital at a fixed cost, there will be no diminution of the marginal product of labor from the lower unemployment at all. Indeed, α will be exactly zero and the $S(\pi)$ multiplier will be equal to $1/b$.

There is another reason why one might suspect that α is quite small or even negative. We have assumed that firms mark up over marginal costs, as they should if they are profit maximizers. However, firsthand accounts of pricing suggest that firms typically include overhead in their costs when computing prices. If they do this, α will be equal to the elasticity of the average product of labor with respect to labor, and from Okun's law we know that that elasticity—at least in the short run—is positive. This was the major reason why the benchmark simulations of both of our previous papers were based on production functions without a declining marginal product of labor. Output was instead assumed proportional to labor input. As a result, in these simulations the S -multiplier was $1/b$, unaffected by the considerations of labor demand that are the unfinished business of this paper.

A high macro elasticity of demand for labor is also consistent with a number of empirical observations. At business cycle frequencies, real wages show very little variation in response to changes in employment. From a longer-term perspective, the U.S. economy has absorbed repeated waves of immigrants with little apparent impact on wages.⁶ Similarly, the

6. Friedberg and Hunt (1995).

French economy showed little impact on wages from the Algerian repatriation of the 1960s,⁷ as did the Israeli economy when faced with a wave of immigration from Russia more recently.⁸ Rachel Friedberg and Jennifer Hunt conclude from a survey of such shocks that a 10 percent increase in the labor force is associated with at most a 1 percent decline in the wages of native workers. If the marginal product of labor were very sensitive to changes in employment (implying a large value for α), one would expect to see large wage impacts from such changes.

Finally, we showed above how the coefficient on unemployment in the price Phillips curve is equal to 1 divided by our multiplier. We also noted that typical estimates for the United States place the value of that coefficient at 0.5. This, too, suggests values for α of considerably less than 1.⁹

Micro versus Macro Estimates of Labor Demand

The micro and the macro estimates differ because they make different assumptions regarding what is fixed. Both concern how a shift in wages will shift the quantity of labor demanded when the firm also adjusts its own price optimally. The macro elasticity assumes that the wage shift occurs at all firms; it also assumes that all other firms adjust their prices and that, as they do so, aggregate income will change. This macro elasticity can be inferred directly from equation 5 by seeing how the demand for labor varies as the wage varies relative to the aggregate price level, \bar{P} . But aggregate income varies, and the price charged by the individual firm, P , corresponds exactly to \bar{P} . In contrast, the micro elasticity assumes that the wage shift occurs only at the individual firm; it further assumes that aggregate income is constant. The demand for labor varies according to equation 5, but the demand equation, equation 4, plays a major role. Now aggregate income is constant, and equation 4 relates P to \bar{P} , with income assumed constant.

Why, then, are the macro estimates larger than the micro estimates? We have seen that the major difference between the estimates is that in

7. Hunt (1992).

8. Cohen and Hsieh (2000).

9. In a more elaborate model it is possible for the coefficient of unemployment in the price Phillips curve to be smaller than the denominator of the $S(\pi)$ multiplier, if markups are large and if the elasticity of the marginal product of labor with respect to labor is large. It seems unlikely, however, that both conditions hold.

Table 1. Elasticities of Demand for Labor under Different Assumptions^a

<i>Elasticity</i>	<i>Fixed capital</i>	<i>Variable capital^b</i>
Macro elasticity ^c	$-1/\alpha$	$-\infty$
Micro elasticity ^d	$-\beta/[1 - \alpha + \alpha\beta]$	$-(1 - \alpha)\beta - \alpha$

Source: Authors' model described in the text.

a. The production function is Cobb-Douglas with a share of capital of α in all cases.

b. Capital is freely available at a fixed cost.

c. Labor is demanded up to the point where its marginal revenue product equals the real wage.

d. Firms or industries assume that the prices of all other firms and aggregate output are fixed. The elasticity of demand for their own product is β .

the micro demand P/\bar{P} varies, and income is held constant, whereas in the macro demand P/\bar{P} remains constant, and income varies. The macro elasticity will be higher than the micro elasticity because it includes the feedback that occurs when one firm reduces its output in response to an increase in costs, thus reducing aggregate income and the demand for the products of all the other firms. Thus, an increase in costs produces a larger change in output at the macro level than at the level of a firm or industry. As a result, labor demand elasticities at the firm or industry level are smaller than those at the macro level, and high S -multipliers are not contradicted by low elasticities of firm and industry demand, such as those reported by Daniel Hamermesh.¹⁰

Table 1 compares four elasticities of the demand for labor. With a Cobb-Douglas production function with a capital share of output of α , the macro elasticity of demand for labor with fixed capital will be $-1/\alpha$. That will be a relatively large number, say, -4 , if α is equal to 0.25. In contrast, if capital is variable and there is an infinitely elastic supply at a fixed international cost of capital, the elasticity of the demand for labor will be infinite.

These macro elasticities are very different from the micro elasticities. We assume as before that the demand for the firm's product is given by equation 4, with an elasticity of demand for its product with respect to its own price of β ; the firm takes the prices of other firms and aggregate income as fixed. If capital is fixed, the elasticity of demand for labor by the firm will be $-\beta/(1 - \alpha + \alpha\beta)$. With variable capital and a fixed capital cost, the elasticity becomes $-(1 - \alpha)\beta - \alpha$. Note that with both fixed and

10. Hamermesh (1996).

variable capital the macro elasticities of demand are larger than the corresponding micro elasticities.

Summarizing Theory and Practice

The simulations of inflation-unemployment trade-offs in our two earlier papers assumed that the macroeconomic elasticity of labor demand was high—that is, that the marginal product of labor rises slowly as the quantity of labor declines. Because this assumption appears to run counter to conventional estimates of the microeconomic elasticity of labor demand, these simulations might seem misleading. However, there is good reason to think that the macroeconomic elasticity would be much larger than the microeconomic elasticity, and a wide array of macroeconomic data is consistent with this logic.

The Second Act

Our 1996 paper described the effects of inflation due to a shock, S , to real wages resulting from money wage rigidity. That shock rises as inflation falls. Our 2000 paper allowed for another possible source of such a shock. There, as inflation approaches zero, wages or prices are set increasingly with only nominal terms in mind. This is consistent with the use of money as a unit of account: when inflation is low, people think in nominal rather than in real terms. That paper calculated $S(\pi)$ from the ignorance of inflationary expectations in wage equations. Similar effects will also occur if price setters, rather than wage setters, ignore inflationary expectations. The previous analysis regarding the value of the S -multiplier carries over from the wage-rigidity model to the expectations-ignorance model. In both cases the effects of S on equilibrium unemployment were of significant size mainly because of the implicit assumption that the macro elasticity of labor demand is very high.

Estimating S

Beyond finding economically important values of the S -multiplier, our 1996 paper estimated a significant trade-off between inflation and unemployment for another reason. Our estimates of S , not just of its multiplier,

were high relative to those found by other investigators of sticky money wages. It is useful here, just by way of reminder, to indicate why we seemingly deviate from other investigators by eliminating important biases in their estimates of S .

The standard alternative method for estimating S is to posit that the distribution of wage changes would be symmetric in the absence of money wage rigidity, and to use the counterfactual distribution to compute the impact of having the negative money wage changes “swept up” to zero.¹¹ The asymmetry of the wage distribution does provide a good diagnostic test for the existence of money wage rigidity,¹² but three biases arise in using this asymmetry to calculate the value of S . The first two of these come from underestimation of the wage change truncation.¹³ Wage change distributions (including those used to estimate S) typically have very large and frequent reporting errors. It is easy to see that such reporting errors will seriously affect the amount of asymmetry in the data. For example, if there is no true negative wage change, all of the observations below zero will be spurious. Reconstruction of the wage change distribution to restore symmetry then seriously underestimates the amount of truncation. Since we wrote our paper, many authors have provided estimates of nominal rigidity that correct for this measurement error. All suggest that failure to take measurement error into account leads to significant understatements of the extent of downward nominal rigidity in nearly all data sets.¹⁴

The second difficulty with this method of estimating S arises because the effect of wage truncation on wages will be cumulative: previous periods’ truncations may still be affecting this period’s wages. In our paper we developed estimates of the truncation that were independent of reporting error. We also developed a method of recursion that allowed us to take account of this cumulative effect of wage truncation. This factor is not important in normal times but becomes very important as the average rate of wage inflation approaches zero.

This brings us to our third cause for the underestimation of S . Wage asymmetries, of course, have been measured at historic levels of inflation,

11. Card and Hyslop (1997).

12. See also Kahn (1997).

13. See McLaughlin (1994).

14. See Dickens and others (2006) for a review and error-corrected estimates of downward nominal rigidity for a wide range of countries.

which have almost universally been significantly in excess of zero. But the theory and our own simulations say that wage change truncation increases nonlinearly as inflation falls toward zero. Our explicit modeling of this nonlinearity constitutes a third reason why our estimates of S were so high for inflation that is close to zero.

Conclusion

In this paper we have completed some unfinished business. We have given an intuitive explanation for why our estimates of the effects of low inflation on unemployment are quite large. We have also justified the reasons for our choice of the parameter responsible for this conclusion. In particular, we have shown how the macro demand for labor will enter the long-run trade-off between inflation and unemployment. And we have further discussed why very low estimates of such a demand for labor can be ignored in assessing the trade-offs between inflation and unemployment.

References

- Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." *BPEA*, no. 1: 1–59.
- _____. 2000. "Near-Rational Wage and Price Setting and the Long-Run Phillips Curve." *BPEA*, no. 1: 1–44.
- Bean, Charles R. 1994. "European Unemployment: A Survey." *Journal of Economic Literature* 32, no. 2 (June): 573–619.
- Card, David, and Dean Hyslop. 1997. "Does Inflation 'Grease the Wheels of the Labor Market'?" In *Reducing Inflation: Motivation and Strategy*, edited by David H. Romer and Christina D. Romer, pp. 71–121. University of Chicago Press.
- Cohen, Sarit, and Chang-Tai Hsieh. 2000. "Macroeconomic and Labor Market Impact of Russian Immigration in Israel." Princeton University and Tel Aviv University (April).
- Dickens, William T., and others. 2006. "The Interaction of Labor Markets and Inflation: Micro Evidence from the International Wage Flexibility Project." Brookings (December).
- Friedberg, Rachel M., and Jennifer Hunt. 1995. "The Impact of Immigrants on Host Country Wages, Employment and Growth." *Journal of Economic Perspectives* 9, no. 2 (Spring): 23–44.

- Hamermesh, Daniel S. 1996. *Labor Demand*. Princeton University Press.
- Hunt, Jennifer. 1992. "The Impact of the 1962 Repatriates from Algeria on the French Labor Market." *Industrial and Labor Relations Review* 45, no. 3 (April): 556–72.
- Kahn, Shulamit. 1997. "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review* 87, no. 5 (December): 993–1008.
- McLaughlin, Kenneth J. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34, no. 3: 383–414.

General Discussion

Lawrence Summers suggested that although no one has explicitly rejected the Akerlof-Dickens-Perry analysis, no central bank has explicitly adopted that analysis in making monetary policy. The idea that the wage mechanism behaves differently at 1 percent annual inflation than at $2\frac{1}{2}$ percent inflation does not appear to be an important factor in central bank decisionmaking. Insofar as central banks have adopted higher inflation targets, this has occurred not because of concern about unemployment, but because of concern about the implications of deflation for real interest rates.

George Perry replied that the Akerlof-Dickens-Perry ideas actually did get some attention at the Federal Reserve and other central banks. A common theme in the literature before the mid-1990s was that zero inflation is ideal, even though people have understood the potential negative consequences of deflation for a long time. Therefore the apparent decision by the Federal Reserve not to pursue zero inflation may reflect some concern about the trade-off with unemployment.

William Dickens added that the Bank of New Zealand had expressly considered this research when making its decision to increase its inflation target. Further, when many economists were advocating zero inflation in the 1990s, he, Akerlof, and Perry argued that this option should be taken off the table, and this has since happened. In fact, all inflation-targeting central banks currently target inflation in the 1 to 2 percent range.

