Unfinished Business in the Macroeconomics of Low Inflation:

A Tribute to George and Bill by Bill and George

George A. Akerlof and William T. Dickens*

August 27, 2007

* 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.

The retirement of George Perry and Bill Brainard gives us an opportunity to say publicly what we have often said privately: *Brookings Papers* is an important national institution. It is important because it has set the right tone for US macro 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. *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 US macro policy has benefited enormously from being based on such a balanced pragmatic-empirical approach.

If we doubt the benefits of such an approach to macroeconomics we need only look to the North—to Canada—where doctrinaire use of an extreme form of natural rate theory in the 1990's led them to push inflation too low resulting in an unusual unemployment gap 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 *Brookings Papers*—has implications for national welfare. Macroeconomic policy, as no one appreciated better than Keynes, oftentimes 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 them their difficult job. George and Bill have edited *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. But, instead, for the most part, George and Bill recruited people to push forward their agenda, which

in turn was shaped by their vision of the important issues in macro policy. Doing so, 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. In addition to commissioning these papers, they have also edited them. We have continued to be amazed at how George and Bill could take first drafts, which were not always in the best of shape, and, quickly turn them into interesting, readable gems.

Introduction

This conference in honor of George Perry and Bill Brainard is an opportunity to reflect back on the two papers we wrote with George Perry (with comments and occasional help from Bill Brainard) on the macroeconomics of low inflation (Akerlof, Dickens and Perry (1996, 2000)). Although they took somewhat different approaches, these two papers had similar policy implications. They both found a significant cost of permanently low inflation in terms of permanently high unemployment.

These papers respectively examined the effects of two types of money illusion. "The Macroeconomics of Low Inflation" examined the consequences of downward nominal wage rigidity—resistance to nominal wage cuts. "Near Rational Wage and Price Setting" examined the consequences of people thinking in nominal terms—rather than in real terms—when inflation is very low. If people adopt nominal thinking when inflation is low, there will be a trade-off between inflation and unemployment. This trade-off will occur at low inflation even in the long-run.

While the existence of such a trade-off is obvious in these papers, its magnitude is not.

These two papers both found surprisingly large increases in unemployment from reductions in inflation to zero. For example, in the benchmark simulation with limited nominal wage cuts for continuing workers we found that a permanent reduction in inflation from 2 percent to zero would increase unemployment by 1.5 percentage points; with the benchmark estimated Phillips Curve, unemployment increased permanently by 2.6 percentage points. In the benchmark "ignored- expectations" simulations, a decline in inflation from 2.6 percent to zero yields a 1.5 percentage increase in unemployment (Akerlof, Dickens and Perry, 2000, p. 18). The preferred parameter estimates for our empirical model imply that moderate inflation will reduce unemployment by 1.9 to 3 percentage points from what it would be with no inflation at all (Akerlof, Dickens and Perry, 2000, Table 1, p. 32). 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 obtain such significant trade-offs? Why did they obtain these results so robustly? These papers rely on one crucial assumption. Here we shall explain what that assumption is, why it is so crucial, and also why, now as then, we think that this assumption is well justified. We welcome this opportunity, not just to reminisce about the fun times we had with George Perry, and also with Bill Brainard, while we were writing these papers about a decade ago. By putting our results in a new perspective, we shall also complete some important unfinished business.

Two Calculations

In both of our models (Akerlof, Dickens, and Perry 1996, 2000) we posit a notional

wage. This is the wage that would be paid by firms in the absence of any money illusion. In our first paper (1996) the departure of actual from wages was due to nominal wage rigidity; in the second paper (2000) it was due to failure, at low levels of inflation, to incorporate inflationary expectations into the wage bargain. These money-illusion considerations cause a tendency for real wages to be higher by an amount $S(\pi)$, which declines as inflation (π) goes from zero to low to moderate rates.

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 of the analysis 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 that the effects on employment can be obtained by multiplying S by the elasticity of demand for labor. Furthermore, this suggests low estimates of the effects of S on unemployment. Daniel Hamermesh's (1993) survey suggests that microeconometric studies of the elasticity of demand for labor find it to be in the range of .2 to .9. This would then suggest that the effect of inflation on unemployment is in the range: .2 $S(\pi)$ to .9 $S(\pi)$.

It turns out that the preceding method of calculation is not correct for several reasons – at least one of which is that it ignores changes in the aggregate supply of labor. In the theoretical models presented in our two papers supply effects were all that mattered. We also showed that one could approximate the impact of changing rates of inflation on equilibrium unemployment in our models by dividing $S(\pi)$ by the slope of the price Phillips Curve. That use of the slope of the Phillips Curve asks the following question: with a tendency for real wages to rise by $S(\pi)$, how

much additional unemployment would be necessary so that expected prices equal actual 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 ½, so multiplying $S(\pi)$ by approximately two gives a good approximation to the results of our models. Thus our models predict impacts on unemployment as much as ten times as high as 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? We shall here construct a model that shows where the elasticity of demand for labor is hiding, 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 factor that is reflected in the slope of the price Phillips Curve 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 industry* as a function of its own wage, with the prices and wages of other firms fixed. The conventional *macroeconomic* demand for labor asks a different question. It asks how employment will change for all firms if the real wage changes for all of them. This is the appropriate demand concept in our formula for the S-multiplier. We will argue that this elasticity, as opposed to microeconomic elasticities, is quite large – even infinite.

The Theory

We adapt a model from Charles Bean's survey of European unemployment (1994). Thus all of the assumptions are rather standard.

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) $p_t - p_{t-1}$

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

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

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

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

The Mark-Up. We shall assume that the firm produces output according to the production function:

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

where Q_t is output, N_t is employment and K_t is the capital stock.

We assume that the firm has a demand function for its product of the form:

(4) $$D_t = BY_t (P_t / \overline{P}_t)^{-\beta},$$

where D_t is the demand for the firm's product, Y_t is overall aggregate real demand, P_t is the price charged by the individual firm for its product. $\overline{P}_t$ is the average price level in the whole economy. β, of course, is the elasticity of demand for the product of this representative firm. Note that we made demand for the product of this firm proportional to income. Like the demand for the average firm, its demand will be proportional to income, for a given relative price of its product.

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

(5) $\partial F/\partial N = M\ W_t/P_t$

where M is the mark-up factor, which will be $\beta/(\beta - 1)$; W is the nominal wage; and P_t is the firm's price.

Variable costs: wages.

We shall write the equation for the notional real wage then as:

(6) $w_t - p^e_t = a - b\ u_t$,

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

The Effects of Money Illusion

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

(6') $w_t - p^e_t = a - b\, u_t + S(\pi_t)$.

The Phillips Curve: We take the Taylor series expansion of the log of the mark-up equation, expressing labor demand in terms of the unemployment rate u_t, rather than the employment rate to get

(5') $p_t - p^e_t = g - \alpha\, u_t + m + w_t - p^e_t$.

Note that when prices equal price expectations (5') can be interpreted as an implicit labor demand function with $-1/\alpha$ as the elasticity of demand for labor.

We combine this with the notional wage equation (6'). Together they yield the notional real price as:

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

where $G = g + a$.

In turn, using (2), this yields a price Phillips Curve of the form

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

Relation between unemployment and S

The rate of unemployment which will maintain the steady state rate of inflation, is then

(8) $u(\pi) = H + S(\pi)/[b + \alpha]$, 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 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 macro 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 firms' actual wage (which is the basis of their labor demand) and their notional wage. Of course, the increase in unemployment is equal to the wedge divided by the sum of the slopes of the two lines.

Alternatively, we 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 while the parameter α reflects the second.

In our two papers we assumed that the elasticity of output with respect to labor was one. In that case equation (5) implies a constant real wage or an infinitely elastic macro demand for labor, the term 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 there is a fixed supply of capital and a zero elasticity of substitution of capital for labor, the macro elasticity of demand for labor will be zero, the denominator of (8) will be infinite, and the S-multiplier equal to zero. What is a realistic value for alpha?

Choice of α

In general, minus alpha is the elasticity of the marginal product of labor with respect to labor. Let's first consider what alpha would be if capital were fixed. With Cobb-Douglas production it is the capital coefficient (which is equal to the share of capital with perfect competition in the product market). More generally with any CRS production function it is $(1-NF_N/Y)/s$ where s is the elasticity of substitution of capital and labor. The numerator is capital's share if there is perfect competition.

Typical values of the share of capital range from .2 to .4; the elasticity of substitution is normally assumed to be one, or perhaps a bit less. These imply values of alpha that are in line with a slope for the price Phillips curve of ½. Thus with conventional values of the parameters of the production function, (8) is consistent with S multipliers that are quite large.

But in the long-run capital will not be fixed. To consider a long-run Phillips Curve we should consider the values of alpha when capital is variable. Capital is variable for two reasons: first, domestic savings 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 alpha 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 alpha will be exactly zero. In this case the $S(\pi)$ multiplier would be equal to 1/b, as we assumed in our two earlier papers.

There is another reason why we might suspect that alpha 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, first hand accounts of pricing suggest that firms typically include overhead in their costs when computing prices. If they do this then alpha 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-- at least in the short run -- is positive.

A high macro elasticity of demand for labor is consistent with a number of empirical observations. At business cycle frequencies, real wages show very little variation in response to changes in employment. Taking a longer term perspective, the U.S. economy has absorbed repeated waves of immigrants with little apparent impact on wages (Friedberg and Hunt 1995). Similarly the French economy showed little impact on wages from the Algerian repatriation (Hunt 1992) as did the Israeli economy when faced with a wave of immigration from Russia (Cohen and Hsieh 2000). Friedberg and Hunt (1995) conclude from a survey of such shocks that

a 10% increase in the labor force is associated with, at most, a 1% decline in the wages of natives. If the marginal product of labor was very sensitive to changes in employment (implying a large value for alpha) then we 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 one over our multiplier. We also noted that typical estimates for the US place the value of that coefficient at .5. This too suggests values for alpha of considerably less than one.[1]

Micro vs. Macro Estimates of Labor Demand

Why do the micro estimates of labor demand appear to be so very much less than the macro estimates? Typically the micro demand for labor is estimated from industry or firm behavior. In either case there is something missing that is present in the macro demand for labor. At the macro level when one firm reduces its output in response to an increase in its costs, that reduces aggregate income and the demand for the products of all the other firms. With all firms reducing their output in response to an increase in costs the aggregate effect becomes important. At the firm (or industry) level there is no perceptible feedback from the reduction in output to the reduction in demand for the firm's (or industry's) product. Thus an increase in costs produces a much 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 very much smaller than those at the macro level. Thus high S-multipliers are not contradicted by the low elasticities of firm and industry demand such as those reported by Hamermesh (1993).

[1] 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 mark-ups are large and if the elasticity of the marginal product of labor with respect to labor is large. It seems unlikely

Summary of Theory and Practice.

In principle there is the possibility that the calculations of inflation-unemployment trade-offs in our two earlier papers were off-base because we were unduly skeptical of the rising marginal product of labor as unemployment rises. Such a rising marginal product of labor would be indicated by very low macro elasticities of demand for labor. But, in practice we believe that our choice of a very low value for alpha was appropriate and in keeping with a wide array of macro data.

The Second Act

Our first paper (Akerlof, Dickens and Perry (1996)) described the effects of inflation because of a shock, S, to real wages due to money wage rigidity. That shock rises as inflation falls. The second paper allowed for another possible explanation for S. There, as inflation falls, wage setting or price setting is done 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. This 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 α is zero, or, if not exactly zero, close to it.

---

that both conditions hold.

Estimation of S

Beyond our optimistic values of the S-multiplier, "The Macroeconomics of Low Inflation" estimated a significant trade-off between inflation and employment 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 estimations of S.

The standard alternative method for estimating S (Card and Hyslop (1997)) is to posit that the distribution of wage changes would be symmetric in the absence of money wage rigidity, and to use the counter factual distribution to compute the impact of having many wage changes "swept up" to zero.

The asymmetry of the wage distribution does provide a good diagnostic test for the existence of wage changes,[2] but there are two biases in using this asymmetry to calculate the value of S. The first two of these biases come from underestimation of the wage-change truncation.[3] Wage change distributions (including those used to estimate S) typically have very large and frequent reporting errors. It is easy to see that frequent 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

[2]See also Kahn (1997).

[3]See McLaughlin (1994).

we wrote our paper many authors have attempted to estimate the extent of nominal rigidity correcting for 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.[4]

The second difficulty with this method of estimating S occurs because the effect on wages of truncation 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 would allow us to take account of the cumulative effects of wage truncation on wages. 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. Those levels of inflation 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.

[4] See Dickens et al. 2006 for a review and error corrected estimates of downward nominal rigidity for a wide range of countries.

We have also justified the reasons for our choices of the parameter responsible for this conclusion. Especially, 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* 1:1996, 1-76.

Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "Near -Rational Wage and Price Setting and the Long-Run Phillips Curve," *BPEA* 1:2000, 1-60.

Brainard, William C., and George L. Perry. 2000. "Making Policy in a Changing World." In *Economic Events, Ideas, and Policies: The 1960s and After*, edited by William Brainard and George Perry. Brookings.

Bean, Charles R., "European Unemployment: A Survey," *Journal of Economic Literature*, 32: 2 (Jun., 1994), pp. 573-619.

Barro, Robert J. and Herschel I. Grossman, "A General Disequilibrium Model of Income and Employment," *American Economic Review,* 61: 1 (March, 1971), pp. 82-93.

Card, David and Dean Hyslop, "Does Inflation 'Grease the Wheels of the Labor Market'?"pp. 71-121, in David Romer and Christina Romer, eds., *Reducing Inflation: Motivation and Strategy.* Chicago: University of Chicago Press, 1997.

Cohen, Sarit and Chang-Tai Hsieh. "Macroeconomic and Labor Market Impact of Russian Immigration in Israel." Princeton University working paper. April 2000.

William T. Dickens, Lorenz Goette, Erica L. Groshen, Steinar Holden, Julian Messina, Mark E. Schweitzer, Jarkko Turunen, and Melanie Ward, "The Interaction of Labor Markets and Inflation: Micro Evidence from the International Wage Flexibility Project," Brookings mimeo December 2006.

Friedberg, Rachel M. "The Impact of Immigrants on Host Country Wages, Employment and Growth" *The Journal of Economic Perspectives*, Vol. 9, No. 2. (Spring, 1995), pp. 23-44.

Hamermesh, Daniel. *Labor Demand*. Princeton University Press, New Jersey. 1993.

Hunt, Jennifer. "The Impact of the 1962 Repatriates from Algeria on the French Labor Market." *Industrial and Labor Relations Review*, Vol. 45, No. 3. (Apr., 1992), pp. 556-572

Kahn, Shulamit, "Evidence of Nominal Wage Stickiness from Microdata." *American Economic Review.* 87:5, (December 1997), pp. 993-1008.

McLaughlin, Kenneth J., "Rigid Wages?" *Journal of Monetary Economics*, 34 (1994), pp. 383-

414.