Comments on Charles Wyplosz's "Do We Know How Low Inflation Should Be?"

William T. Dickens
The Brookings Institution

A little inflation may be a good thing in a modern economy. In particular, it may be possible to maintain lower rates of unemployment with low inflation than with zero or very low inflation.[1] In two recent articles, George Akerlof , George Perry and I (1996,2000) (hereafter, ADP) have presented evidence for two specific mechanisms by which inflation may affect the equilibrium level of unemployment. We modeled those mechanisms, and argued the case that large permanent reductions in unemployment may be obtained by moving from either a high or very low rate of inflation to a moderate rate (2-4% in the United States). The two mechanisms we examined were nominal rigidity in wage setting, and near rationality in the use of inflationary expectations in price and wage setting.

In these papers we estimated Phillips Curve relations from which we deduced the magnitude of the effects of nominal rigidity and near rationality on the long-run relationship between inflation and unemployment. But these empirical exercises do little to verify the *nature* of this relationship because that was simply assumed when we accepted the dictates of our theory in setting up the specification we estimated. In this paper Charles Wyplosz takes a very different approach to roughly the same empirical problem. He estimates models of unemployment in which he allows the NAIRU, or natural rate of unemployment, to vary with the rate of inflation in a very general way. Such an approach has advantages and disadvantages relative to the approach my

[1] For example see Tobin (1972, p11); Samuelson and Solow (1960, p.182) and Schultze (1959, p.134).

colleagues at The Brookings Institution and I have taken. Below I discuss Wyplosz's results and present some of my own. Taken together, our results suggest the possibility that very low rates of inflation may cause unemployment to be higher than it would be at moderate rates of inflation in the Euro zone.

Is Very Low Inflation the Best Policy?

Several other papers being presented at this meeting discuss the potential costs of inflation in depth. I would like to call attention to several arguments for choosing low rates of inflation over very low or zero inflation. At least since Keynes there has been concern that the prevalence of nominally denominated debt may make expenditures very sensitive to the level of nominal income, and increase systemic financial risk in times of deflation. Very low or zero inflation may make deflation unavoidable in a contraction – particularly if monetary policy becomes less effective when nominal interest rates approach the floor of zero. In addition, it has been argued that worker resistance to nominal wage cuts prevents real wage adjustments. My colleagues and I demonstrated the potential importance of this last argument in our 1996 *Brookings* paper.

Further, in a paper just published this Summer, George Akerlof and George Perry and I (2000) have developed a new argument for moderate inflation. At the center of the case for very low inflation is the belief that the equilibrium level of unemployment either does not change or increases with the rate of inflation. A core assumption of that natural rate model is that economic agents everywhere and always form expectations about future inflation and raise wages and prices in anticipation of it. In our recent paper we present evidence that real world price and wage setters simply do not behave that way – at least not all the time. In particular, we argue that when inflation is low a significant

fraction of price and wage setters probably ignore or underweight anticipated inflation in setting future prices. The prices they set will then lag a fixed distance behind what they would be if inflation was fully anticipated and the overall price level will be lower relative to the money supply than it would be at zero inflation. As inflation increases, the cost of such near rational behavior increases and more agents will fully anticipate inflation in wage and price setting as the natural rate model predicts. When they do, prices will tend to rise relative to the money supply. As a result, the real money supply (and therefore real demand) will be larger in times of low inflation than either high inflation or zero inflation. Near rationality at low inflation creates a trade-off between inflation and real activity that central banks can exploit to the advantage of the economy.

Table 1

Inflation and Unemployment in Europe In the 1960s

	Average Inflation	**Average Unemployment**
	60-68	**60-68**
United Kingdom	3.4	1.9
France	3.6	1.7
Germany	2.3	1.1

If these arguments are right then moderate inflation ought to be better for economies than very low rates of inflation. While there is substantial evidence that very high rates of inflation are related to poor economic performance, there is no persuasive evidence that growth is lower at three percent inflation than at zero or one percent inflation. In fact, the 1960s were a period of high growth and very low unemployment

despite average rates of inflation above the levels desired by the European Central Bank and other authorities that have announced inflation targets.

While this experience fits the story that moderate inflation is better than low inflation, some have pointed to the recent experience of the United States, or the past history of Japan and other countries, as evidence that low unemployment is compatible with very low inflation. There are two things to keep in mind in evaluating these arguments. First, examples from times prior to the institution of modern systems of industrial relations aren't probative. Considerations of nominal rigidity were irrelevant in economies based largely on craft and agricultural production where job tenures tended to be short and self-employment was common. Second, nominal rigidity is not a problem in an economy that is experiencing rapid nominal wage growth *either* because of moderate inflation *or* rapid productivity growth. The recent success of the United States in maintaining both low inflation and unemployment has probably been due in large part to its upturn in productivity growth. High productivity growth probably also helps explain Japan's remarkable labor market prior to the 90s.

But is there evidence to support this view for Europe, and if so, what is the rate of inflation that minimized unemployment for the Euro? These are the questions that Wyplosz poses.

Wyplosz's Approach vs. ADP

Wyplosz takes three very general approaches to estimating the long-run relation between inflation and unemployment. In the first, he estimates a nearly standard accelerationist Phillips Curve. Following Phelps and Zoega (2000) he includes a measure of Tobin's Q in the regression. To allow non-linearities in the long run Phillip's Curve

Wyplosz introduces a polynomial in a long centered moving average of expected inflation. Expectations are deduced from the difference between long-term nominal interest rates in a country and the world (US) real long rates. Wyplosz's second approach is to include the polynomial of the long moving average in an equation to predict unemployment directly. The expected value for unemployment in that equation is interpreted as the natural rate. The third approach extends the second by combining data from four countries to estimate the long-run conditional expectation of unemployment given the rate of inflation.

This essentially atheoretic approach to estimating how the equilibrium rate of unemployment might vary with the rate of inflation is complimentary to the approach that my Brookings colleagues and I have taken. Wyplosz's approach has several advantages over our approach. First, to the extent that there are other influences on the shape of the long-run Phillips Curve besides those envisioned in our models, his method could pick them up while they would only lead to problems in estimating our model. A further advantage of his approach is that it could provide a much stronger test of the effects of nominal rigidity and near rationality than what my colleagues and I have done. While we assumed the functional form, his estimates test it. Informally one might inspect his estimated long-run Phillips Curves for the characteristic shape implied by the theory, or one could formally test whether the addition of his polynomials to equations such as those estimated by my colleagues and I significantly improve their fit.

There are also disadvantages to the approach that Wyplosz takes. For one, the long-run Phillips Curve may not be stable. As discussed above, increasing the rate of productivity growth will reduce the rate of inflation at which rigidity effects are felt

allowing the attainment of lower rates of inflation for the same rate of unemployment. Wyplosz includes a measure of Tobin's Q in his equations in part to capture such effects. But, a changing value of Tobin's Q would shift the position of the long-run Phillips Curve left or right as his equations are specified, it would not shift it up or down. In contrast, a model of nominal rigidity would predict that declining productivity growth would shift the nearly horizontal section of the long-run Phillips Curve at low inflation up. Examples of this can be seen in my Figures 2-5.[2]

A second drawback to Wyplosz's approach relative to the approach that my colleagues and I have taken is there may be information in the short-run behavior of inflation and unemployment that would help identify a structural model. That information would be ignored in Wyplosz's approach. For example, near rationality implies that past inflation will be less predictive of future inflation when inflation rates have been low than when they have been high. Alternatively, nominal rigidity implies that periods of very low inflation or deflation should affect the economy like a supply shock producing either higher unemployment or a burst of inflation in their wake.

By using a very flexible functional form in the context of a method that ignores some identifying variance Wyplosz's approach necessarily has low power. Our approach is more likely to reject the hypothesis of the natural rate, but will be biased if the true model is not the one we envision. Wyplosz will be less likely to reject the natural rate model if it is false, but will be less vulnerable to most forms of misspecification.

Given the complementary nature of these two approaches it is worthwhile to see what they both have to say about the unemployment-inflation relation in Europe. Of

[2] Also, Wyplosz uses a centered moving average for his polynomial term that captures the long-run relationship between inflation and unemployment while most theories would suggest that it is lagged values

course both methods are subject to the usual problems that bedevil all attempts to estimate Phillip's curves. The Phillip's curve itself is not a structural model though it can be derived from several different models. As such its identification is always an issue. In particular, the assumption that unemployment is effectively exogenous to price setting at the frequency at which the relationship is estimated (typically a quarter or a year) is always questionable, and is particularly problematic in open economies subject to external supply shocks. In addition, the question of how to model expectations is vexing. Still, the method is one that is often used to inform monetary policy and is a frequently used device in academic research. Thus I reconsider Wyplosz's analysis and conclusions in the next section, present my own work on applying the ADP model to other countries besides the U.S. in the section after that, and conclude with a discussion of what lessons I think policy makers should take from these two exercises.

Reconsidering Wyplosz's Analysis

Wyplosz approaches the estimation of the effects of inflation on the long-run unemployment rate in three ways. First, he estimates Phillips curves where the natural rate is a non-linear function of the long-term moving averages of the inflation expectations and the q ratio. Second, he estimates for each of his four countries the mean of the unemployment rate conditional on q, a time trend, and his polynomial in the MA of inflationary expectations. Finally, he pools data for all four countries and estimates the conditional mean of unemployment assuming that changes in the rate of inflation have the same effect on unemployment in all countries. I find the results of the last two exercises interesting and important, but the first unconvincing.

---

of inflation that should be relevant for current unemployment.

Wyplosz presents his Phillips curve results as evidence of non-linearties in the long-run relationship between inflation and unemployment. He notes that the coefficients on the unemployment rate are not very precisely estimated and he therefore eschews computing the implied natural rates. There are other reasons not to perform this exercise as well. In all four countries the implied equilibrium unemployment rates corresponding to at least some inflation rates are negative for most years. In Germany, it is negative for values of inflation less than .7% and greater than 3.5% for almost all values of q. In France, the coefficient on unemployment is positive so that the equilibrium unemployment rate is negative for values of inflation above 1%. In the Netherlands the unemployment rate is negative below 1% and above 8.5%. In Switzerland it is negative for inflation rates above 4% in most years.

For France and Germany these results are largely nonsensical, but in Switzerland and the Netherlands most observations fall in the range for which the implied equilibrium unemployment rate is positive. These results illustrate something important about Wyplosz's method. When one fits a polynomial (or for that matter any model), unless the functional form is motivated by a theory there is no reason to believe that the results have any meaning outside the range for which there is data. Even within the range of the data, the shape of the polynomial will be determined to give the best fit where most of the observations are. In both of these countries there are observations in the range in which the implied equilibrium unemployment rate is negative, but only a few. What this probably means is that the gain in fit afforded in the ranges of inflation in which most observations fall from twisting the polynomial this way were greater than the loss of having a few observations fit very poorly.

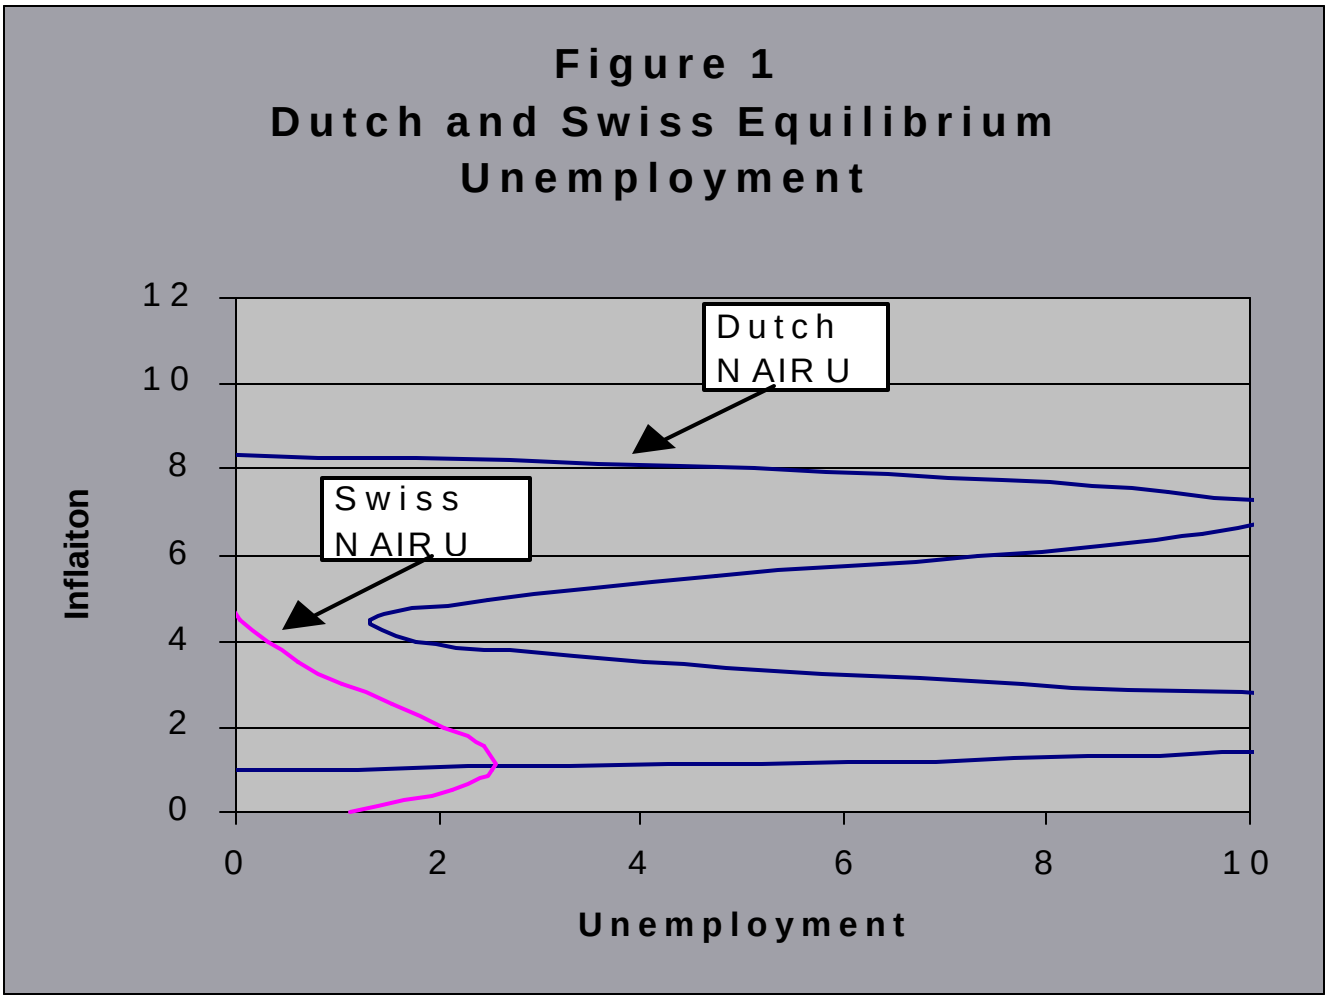

For the Netherlands and Switzerland Wyplosz notes that the coefficients on inflation in the polynomial are all positive and interprets this as indicating that only sand effects are present in these countries. In my Figure 1, I have plotted the Phillips curves implied by Wyplosz's estimates for these two countries. In both cases the estimated values of q (the exponential decay term multiplying the polynomial) are sufficiently negative to produce significant declines in unemployment with increasing inflation within the range in which most observations fall. In fact, with the exception of the what is going on at extremely low and high rates of inflation in these two countries the results are reminiscent of my colleagues and my findings. However, none of Wyplosz's Phillips curve estimates include controls for many changes going on in these countries labor

markets over the sample period and the results for the conditional unemployment rate suggest this may be crucial.

Table 2 and Figures 3 and 5 in Wyplosz's paper present the results for the conditional expected unemployment rate. Note that in all four countries the trend is positive, large and statistically significant. As the author suggests, the most likely explanation for this is that there are important left-out features of the labor markets in these countries that are affecting the unemployment rate. Given that there is no account taken of these in the Phillips curve analysis those results are suspect at best. However, the time trend may be an adequate control, and to the extent it is, there is more reason to take seriously the results of Wyplosz's second and third approaches to the data.

The author draws three conclusions from his results on country-specific conditional expected values. First, for nearly all of the range from 1 to 5% inflation in every country for every specification, unemployment declines with inflation. Second, in nearly every specification there is a small range of inflation rates – mostly below 1% and always below 2% – for which unemployment is increasing with inflation. Finally, while in most specifications a constant expected unemployment rate can be rejected, with the exception of France the confidence intervals for the conditional mean at relevant inflation rates are broad enough so that for all practical purposes a constant rate cannot be ruled out.

The last conclusion nearly negates the first two, and that is why Wyplosz undertakes his third exercise. Since the shapes of relationships are so similar across countries, the loose fit in any one country should not be the last word. Under these

circumstances it does make sense to pool the data the way he does and to estimate how the expected unemployment rate will vary with inflation across countries.

The results of this analysis are presented in Table 3 and Figure 4 in Wyplosz's paper. They suggest a strong negative relation between inflation and unemployment over nearly the entire range of post-war inflation rates. Unemployment rates are at a maximum in the range the ECB has chosen as its target for inflation. However, recall my earlier comments about the difficulties of interpreting the results of curve fitting exercises for those parts of the curve at the periphery of the sample. As a quick look at Wyplosz's Figure A1 will confirm, these four countries collectively have not had that much experience with inflation above 6% or below 2%. The curves will thus reflect only what is going on between these two values. Making the polynomial fit well in this range may require it to fit poorly outside the range. I suspect that explains these results.

In particular, I am very suspect of the finding of increasing unemployment with inflation at very low rates of inflation. First, I can think of no theory that would motivate such sharply increasing unemployment rates in the range or 0-1% inflation that would also accommodate large decreases as unemployment climbed above that rate. Most stories about how inflation might cause increasing unemployment would imply constant or increasing effects as inflation increased. Second, as Wyplosz's Figures 3 and 4 show, the relationship between inflation and unemployment is very imprecisely estimated in this interval and in no case could one rule out steadily declining unemployment with increasing inflation. In fact, in 2 of the 8 specifications presented in Wyplosz's Figure 5 a declining rate is estimated for this range. This imprecision and sensitivity to specification is to be expected given how few values of expected inflation fall in this range –

particularly since nearly all of those that do are during periods of above-normal unemployment. The best conclusion is that we simply can't tell from these data what the long-run relationship between inflation and unemployment is in this range.

Results for the ADP Model Outside the United States

If it is problematic to estimate the relationship between inflation and unemployment at very low rates of inflation using a flexible functional form, can more progress be made using a functional form dictated by theory? If the ADP model was found to fit as well outside the U.S. as it does inside and to suggest an important trade-off between inflation and unemployment at low rates of inflation then one might have some confidence in these findings even if there was little experience with very low rates of inflation. A good model should be able to forecast beyond the range of data on which it is estimated. With this hope I present some preliminary attempts to fit the ADP model for countries outside the United States.

Table 2 presents estimates of the important parameters of the ADP Phillips curve for Canada, the United Kingdom, France and Germany. Typical results for the United States are also presented. The form of the ADP Phillips curve is

$$\pi_t = a + \Phi(D + E\pi^2_{Lt-1})\pi^e_t - bu_{Lt} + gX_t + S_t(\sigma, g_t, \pi_{t-1}, \pi_{t-2}, \ldots) + \varepsilon_t$$

where π is the rate of inflation, Φ is the cumulative standard normal density function, u is unemployment, X is a vector of dummy variables and controls for supply shocks, S is the term we add to account for nominal rigidity (described in Akerlof, Dickens and Perry 1996) which is a function of the standard deviation of desired wage change (σ), trend productivity growth (g), and past values of inflation. ε is assumed to be an i.i.d. error, the subscript Lx denotes a weighted average of lagged values starting with period x and

going back and the superscript e denotes the expected rate (modeled in all of these specifications as a weighted average of past values).

Table 2

Parameter Estimates for ADP Phillips Curves for Several Countries
(standard errors in paranethesis)

Parameters	United States	Canada	United Kingdom	France	Germany
Constant (a)	.024 (.012)	.025 (.008)	.000 (.018)	-.009 (.007)	-.085 (.021)
Sum of Coefficients of Unemployment (b)	-.43 (.18)	-.33 (.10)	-.35 (.16)	-.21 (.06)	-.36 (.05)
Constant in Coefficient on Expectations (D)	.02 (.60)	-.68 (.59)	-1.36 (1.53)	-.46 (.47)	-.49 (.38)
Coefficient of π^2 in Coefficient of Expectations (E)	.117 (.071)	.123 (.048)	.118 (.114)	.116 (.054)	.248 (.118)
Standard Deviation of Desired Wage Change (in Percent) (σ)	2.2 (1.4)	2.7 (1.3)	7.4 (3.1)	7.7 (1.6)	17.6 (2.6)

The general specification and the method of estimation are described in Akerlof, Dickens and Perry (2000). The specific functional forms and the data used are described in the appendix.

On first examination, it seems that the results for other countries resemble those for the United States. Some aspects look quite promising. Without any time trends or controls for changes in labor market institutions we get a strong negative relationship between the level of unemployment and inflation in every country. If the coefficient on expected inflation is one, as the natural rate hypothesis implies, then the constant term in the cumulative normal multiplying expected inflation should be a large positive number. In all examples presented (except the United States), the value is negative -- though small positive values could not be ruled out. Estimated values suggest that the coefficient on expected inflation is considerably less than 1 when inflation is low, as the theory of near-

rationality in price and wage setting suggests. In addition, in every country there is considerable variation in the value of the cumulative normal term over the sample as the coefficient on the weighted average of lagged inflation is large and typically statistically significant.

Besides the U.S., the one country where the coefficient of the square of lagged inflation in the cumulative normal is not statistically significant is the U.K.. In many of the specifications that we estimated for the U.S. this parameter was not significantly different from zero when the term for nominal rigidity was included, but significant when that term was excluded. This is the case for this specification for the U.K. as well.

One place where we start to see an important difference between the U.S. and other countries is in the last line of Table 2. The larger the standard deviation of desired wage changes the higher the level of inflation at which nominal rigidity becomes important for unemployment. Were nominal rigidity no problem for an economy, this parameter should be estimated to be zero. Instead we find that in Europe this parameter is estimated to be much larger than in the U.S. or Canada which, taken literally, would mean that nominal rigidity is much more of a problem for Europe. However, the first line of Table 2 shows the constant term to be zero or negative for all the European countries. This implies the impossibility of a negative asymptote for unemployment as inflation rises and calls into question the validity of the exercise.

Figures 2-5 show the implied long-run Phillips curves for Canada, the U.K., France, and Germany. Each figure shows two long-run Phillips curves corresponding to different rates of productivity growth – that which prevailed in the 60s and the average

rate since 1980. As noted above, changes in the rate of productivity growth can have substantial effects on the long-run Phillips curve when nominal rigidity is a factor.

The Canadian Phillips curves in Figure 2 look very similar to those we estimated for the United States (Akerlof, Dickens and Perry 2000). Nominal rigidity causes unemployment to rise substantially at very low rates of inflation. Near rationality allows

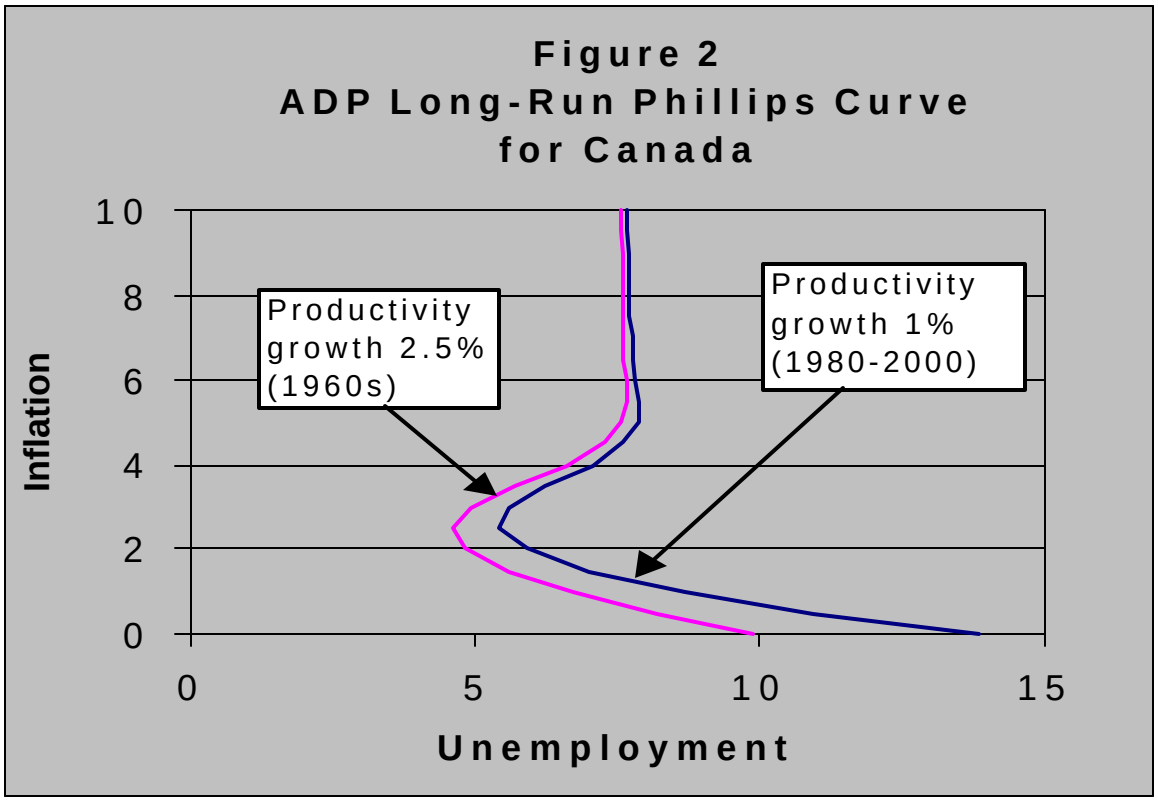

unemployment to fall noticeably below the natural rate of about 7.5% to a low of about 5.5% at an inflation rate between 2 and 3%. I have not done the exhaustive specification search we did for the United States for any of these countries, but the several specifications we have estimated for Canada all yield qualitatively similar results. The ADP Phillips curves for the U.K. in Figure 3 looks somewhat like those for Canada except that unemployment is predicted to fall steadily for inflation rates above 5%. There is a local minimum for the unemployment rate at a rate of inflation between 3 and 4%.

The value is higher than for Canada and the United States because the effects of nominal rigidity are estimated to be important at higher rates of inflation due to the greater variability of wage change. Again, limited robustness checks suggest that ADP Phillips curves estimated for the U.K. all have roughly this form. In particular, all have the minimum unemployment rate occurring at very high rates of inflation and most have negative asymptotes.

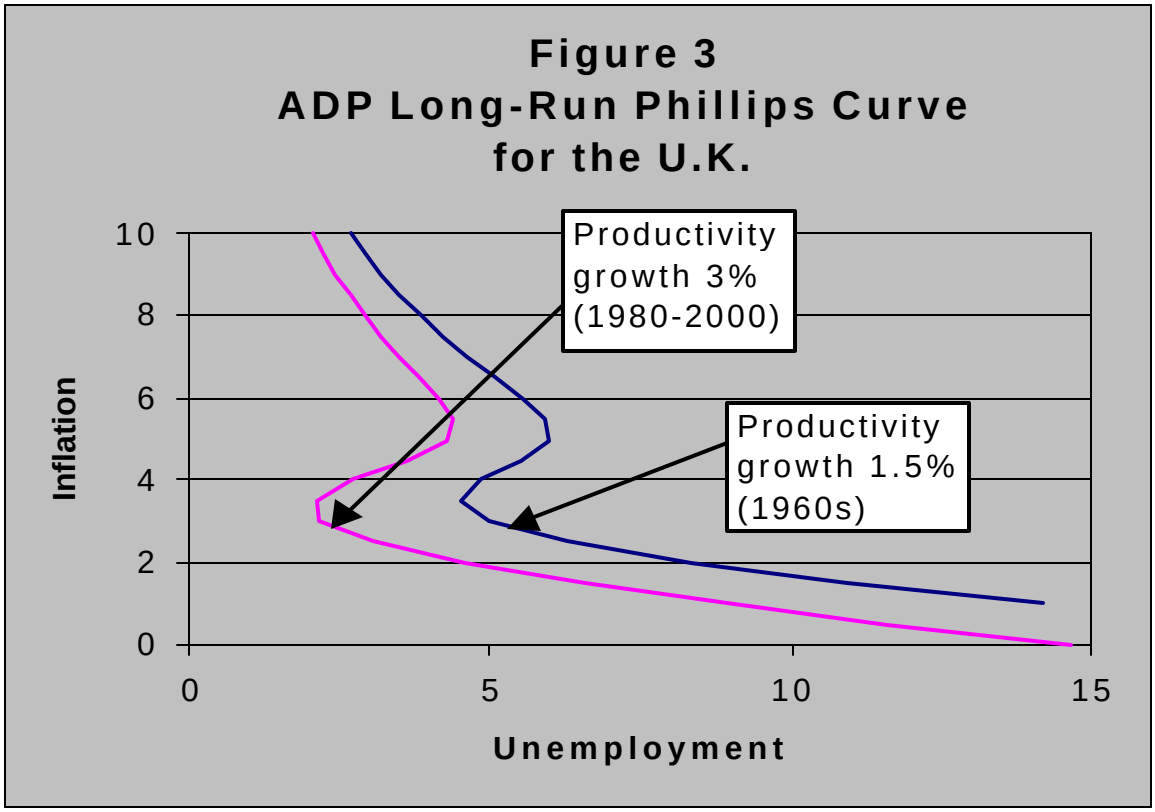

The French Phillips curves shown in Figure 4 are even more problematic. Only at very high rates of productivity growth do we see the local minimum for unemployment caused by near rationality. As with the U.K., I have run limited robustness checks and find similar results for a range of different specifications. In particular, I have attempted adding a time trend to the French estimates. While the time trend is estimated to be large

and statistically significant in any Phillips curve that does not include the term for nominal rigidity, the inclusion of that term causes it to fall to nearly zero and to become statistically insignificant. However, the negative asymptote for unemployment as inflation increases is also characteristic of most specifications we estimated. Two specifications (out of 15) yielded positive asymptotes, but negative values for the minimum unemployment rate for at least some values of trend productivity growth observed during the sample period. Also, the lag structure for unemployment suggests that changes in unemployment are very important for the inflation process in France possibly indicating the existence of insider-outsider problems.

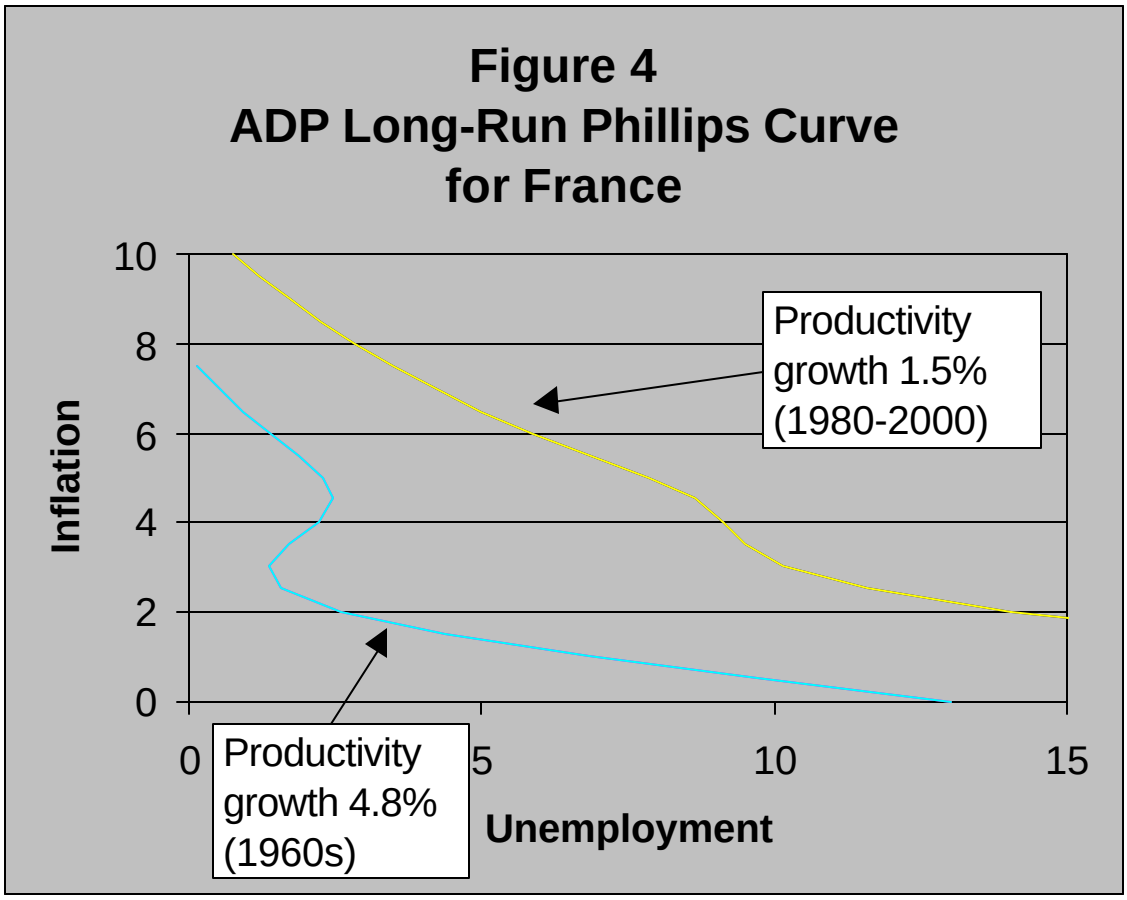

Figure 5 presents typical curves for Germany. The very large value of the standard deviation of wage change estimated for Germany, combined with the negative intercept, results in long-run Phillips curves which are virtually horizontal and that shift up as productivity growth slows. I have only estimated five specifications for Germany, but all have this character.

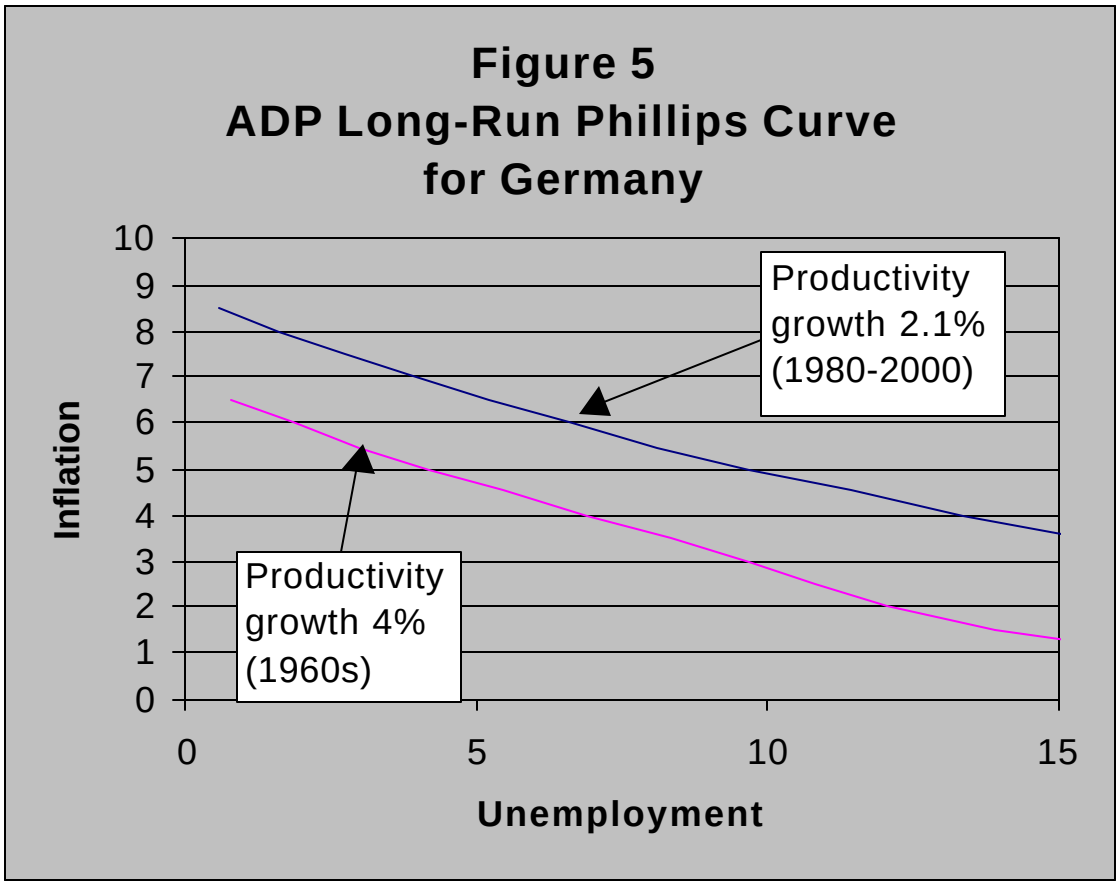

How seriously should we take these results? With respect to the findings for near rationality, it important to remember that the short-run impact of expectations on the rate of inflation is at least as important for the identification of this effect as is the position of the long-run Phillips curve. Similarly, our approach to estimating the effects of nominal rigidity has implications for the short-run behavior of inflation, though I suspect that most

of the identification in these estimates comes from the long-run implications. Thus even though the models estimated for Europe cannot reflect the true long-run relationship between inflation and unemployment, these estimates still provide evidence of the importance of nominal rigidity, and particularly near rationality.

The estimates for the standard deviation of wage change are particularly surprising and suspect. One might expect the corporatist wage setting practices of Germany and France to lead to lower variation in wage change – not higher. However, it has been argued that Europeans are much less mobile than North Americans and this may mean that shocks to local labor markets require larger wage changes to equate labor supply and labor demand. This might help explain an anomaly in the wage curve literature that Blanchflower and Oswald (1994) describe. In Europe there is a negative relationship between wage levels and unemployment in nearly all cross-sectional data sets. In the United States the relationship is only negative with panel data when geographic dummies are included. Otherwise it is usually positive. This suggest that labor moves to equate expected incomes in the medium run in the United States as in the Harris-Todaro model, but not in Europe. It would be useful to examine panel data for these countries to examine the distribution of wage changes directly and to compare the distribution to that in similar U.S. data.

Conclusion

The evidence presented by Charles Wyplosz and the ADP Phillips curves I have estimated for Europe certainly don't make an iron clad case against a target range for inflation of 0-2%, but they should give the ECB reason for concern. Wyplosz's consistent finding that unemployment declines at rates of inflation above 2% combined with the

evidence I have presented for the relevance of near rationality, and possibly nominal rigidity, in countries outside the United States suggests that targeting an inflation rate below 2% could be very costly.

Certainly there is more to the European unemployment problem than too low a rate of inflation. Wyplosz's finding of a significant trend in his models of the conditional mean of unemployment, and the inability of the ADP model to tell a coherent story about unemployment at high rates of inflation are suggestive of this. Real wage rigidities as well as other labor market problems could be contributing to the inability of these models to fit European data. It is even possible that these problems cause the models Wyplosz and I have estimated to be misspecified so that the results are meaningless. However, what if near rationality and nominal rigidity are problems in the Euro zone? Targeting too low a rate of inflation could indefinitely delay Europe's full recovery from its high rates of unemployment even if structural reforms are successful in removing other impediments.

The Brookings Papers on Economic Activity (1) pp1-44.

---------- (1996) “The Macroeconomics of Low Inflation,” *The Brookings Papers on Economic Activity* (1) pp1-59.

Bagnoli, Philip (1997) "Productivity Growth: Discussion And Twelve Sector Survey", Brookings Discussion Papers in International Economics, No. 132, (July).

Blanchflower, David G., and Andrew J. Oswald (1994) *The Wage Curve*, MIT Press.

Phelps, Edmund and Gylfi Zoega (2000) “Structural Booms: Productivity Expectations and Asset Valuations,” unpublished paper

Samuelson, Paul A., and Robert M. Solow (1960) “Problems of Achieving and Maintaining a Stable Price Level: Analytical Aspects of Anti-inflation Policy,” *American Economic Review, Papers and Proceedings*, 50(2) pp.177-94.

Schultze, Charles L. (1959) “Recent Inflation in the United States,” Study Paper 1. Joint Economic Committee, 86 Cong. 1 sess. (September).

Tobin, James, (1972) “Inflation and Unemployment,” *American Economic Review* 62(1) pp1-18.

Appendix

Functional Forms for ADP Phillips Curves for Several Countries

	Canada	United Kingdom	France	Germany
Number of Lags on Unemployment	11	7	3	11
Lag Structure of Inflationary Expectations (p_e)	Unrestricted	Piecewise Linear	Unrestricted	Unrestricted
Lag Structure for Inflation (p)	Piecewise Linear	Geometric Decay	Piecewise Linear	Piecewise Linear
Coefficient of Inflationary Expectations (p_e)				

*The specification of each term and the method of estimation are described in Akerlof, Dickens and Perry (2000). For these specifications the only supply shock control was an dummy for the first two quarters in1973 and the third quarter in 1978 through the third quarter of 1979 as a control for the oil price shocks of those years. Dummies were used instead of a food and fuel price series or import price deflators to accommodate both the price changes and the policy responses to those price changes (such as rationing). The dummies were not significant in Canada and were not used in the specification presented here though they were used in other specifications. Results for all Canadian specifications were substantively similar to results obtained in the U.S.

Data Sources

The Consumer Price Index series used for each country was taken from the International Monetary Fund's "International Financial Statistics" CD-ROM. Unemployment data for France, Canada and the United Kingdom was taken from OECD's *Main Economic Indicators*, and for all years for which an unemployment rate was not published by the OECD, a rate was constructed using the OECD's registered unemployment and total employment series. Unemployment rates for Western Germany are from the German Federal Employment Services Institute for Employment Research. The OECD unemployment rate series for France and Canada were multiplied by a scaling factor, so that the series did not differ significantly from the Bureau of Labor Statistics' "Unemployment Rates: Approximating U.S. Concepts" series in *Comparative Civilian Labor Force Statistics: Ten Countries, 1959-1999*. The productivity trend was constructed by combining sources presented in Table II of Bagnoli (1997).