Topic 6: The New-Keynesian Phillips Curve

The Phillips curve has been a central topic in macroeconomics since the 1950s and its successes and failures have been a major element in the evolution over time of the discipline. We will now discuss a popular modern version of the Phillips curve—known as the "New Keynesian" Phillips curve—that is consistent with rational expectations. We will start, however, with a brief review of the history of the Phillips curve relationship.[1]

The Phillips Curve

The idea that there should be some sort of positive relationship between inflation and output has been around almost as long as economics itself, but the modern incarnation of this relationship is usually traced to a late 1950s study by the LSE's A.W. Phillips, which documented a statistical relationship between wage inflation and unemployment in the UK. This "Phillips Curve" relationship was then also found to work well for price inflation and for other economies, and it became a key part of the standard Keynesian textbook model of the 1960s. As Keynesian economists saw it, the Phillips curve relationship of the form

$$\pi_t = \alpha - \gamma U_t \tag{1}$$

provided a menu of tradeoffs for policy-makers: They could use demand management policies to increase output and decrease unemployment, but this could only be done at the expense of higher inflation.

In his 1968 presidential address to the American Economic Association, Milton Friedman presented a sharp critique of the Keynesian Phillips curve. In particular, he criticized its neglect of expectations.[2] Friedman argued that the correct formulation of the inflation-unemployment tradeoff was an "expectations-augmented" Phillips curve of the form:

$$\pi_t = -\gamma(U_t - U^*) + \pi_t^e, \tag{2}$$

where inflation, π_t, is negatively correlated with deviations of the unemployment rate from its natural rate U^*, and where the entire curve is shifted up or down one-for-one with changes in π_t^e (the rate of inflation that agents had expected to prevail in time t). Friedman predicted that attempts to keep unemployment low at the expense of higher inflation would

[1] A more detailed review of this history is available in the Rudd-Whelan paper on the reading list, "Modelling Inflation Dynamics: A Critical Review of Recent Research." This is available at www.karlwhelan.com.

[2] Milton Friedman (1968). "The Role of Monetary Policy." *American Economic Review*, 58, 1-17.

just result in raised inflation expectations. Thus, the economy would not be able to sustain the low unemployment and would end up with higher inflation. Friedman's prediction of the stagflation that was to come in the 1970s was perhaps the most influential pieces of macroeconomic theorizing ever.

Friedman assumed that inflation expectations evolved over time as a result of actual past experience—that is, that expectations were formed *adaptively*. One simple model of this is to assume that $\pi_t^e = \pi_{t+1}$, so that inflation expectations are determined by what happened last period. In this case, the inflation relationship is of the form

$$\Delta\pi_t = -\gamma(U_t - U^*) \tag{3}$$

This is sometimes labelled an "accelerationist" Phillips curve because it implies that one can only keep unemployment low at the expense of an increasing inflation rate, and thus an accelerating price level. Empirically, when applied to monthly or quarterly data, this idea was usually implemented by assuming that inflation expectations were determined as a weighted average of past inflation rates. This implied an inflation equation of the form

$$\pi_t = \alpha - \gamma U_t + \sum_{i=1}^{N} \beta_i \pi_{t-i} \tag{4}$$

where the weights β_i were constrained to sum to one. This restriction means that the relationship can still be written in terms of first-differences as

$$\Delta\pi_t = \alpha - \gamma U_t + \sum_{i=1}^{N-1} \delta_i \Delta\pi_{t-i} \tag{5}$$

The model also introduced the concept of the NAIRU—the non-accelerating inflation rate of unemployment. This is the unemployment rate consistent with constant inflation and it is defined implicitly by

$$\alpha - \gamma U^* = 0 \Rightarrow U^* = \frac{\alpha}{\gamma} \tag{6}$$

Econometric Phillips curves of this sort tend to fit the data quite well. And their implied estimates of the NAIRU are still very much part of macroeconomic policy discussions today, with policy recommendations often made on the basis of whether unemployment is above or below this NAIRU level.[3]

[3]Note though the NAIRU terminology is actually a misnomer. If unemployment is below U^*, then inflation will be increasing, but not accelerating. The price level is what will be accelerating.

The fact that inflation depends on its own lagged values in this formulation also has important implications for monetary policy. Consider, for instance, a central bank that wants to reduce inflation from a high level. If this Phillips curve is correct, then it will be very difficult to reduce inflation quickly without a significant increase in unemployment. So, this Phillips curve suggests that gradualist policies are the best way to reduce inflation.

The Rational Expectations Challenge

This "demise" of the traditional Phillips curve, and the sense that it was due to inadequate modelling of expectations, was a major impetus for the rational expectations school of thought in the 1970s, led by Robert Lucas and Thomas Sargent. Lucas and Sargent also rejected the "accelerationist" reformulation of the Phillips curve because it relied on the assumption of adaptive expectations, which does allow for the idea that agents process information in an optimal manner.

In addition to being more precise about expectations formation, this school of economists relied more heavily on neoclassical "microfoundations" for macroeconomic models. Often, as well as rejecting the Phillips curve, these economists also questioned the whole basis for Keynesian economics, i.e. the assumption that monetary policy could systematically affect output even in the short-run.

The principal response of Keynesian economists to these theoretical critiques has been to attempt to build models that incorporate rational expectations and that provide a microeconomic justification for monetary policy having at least short-run effects. To explain why monetary policy might have effects on the economy, one needs a theory of why inflation is not just determined by some nominal anchor such as the money supply. The most common microeconomic rationale put forward has been *sticky prices*. With sticky prices, an increase in the money stock can produce a short-run increase in real spending power and thus can boost real output.

This modern approach, featuring rational expectations and some form of microfoundations, is known as *New Keynesian* macroeconomics. We will now describe one of the key New-Keynesian models, and explore its implications for the behaviour of inflation and output.

Pricing à la Calvo

There are lots of different ways of formulating the idea that prices may be sticky. Some of the best known formulations were those introduced in papers in the late seventies by John Taylor and Stanley Fischer.[4] These papers essentially invented New Keynesian economics. Here, however, we will use a formulation known as *Calvo pricing*, after the economist who first introduced it.[5] Though not the most realistic formulation of sticky prices, it turns out to provide analytically convenient expressions, and has implications that are very similar to those of more realistic (but more complicated) formulations.

The form of price rigidity faced by the Calvo firm is as follows. Each period, only a random fraction $(1-\theta)$ of firms are able to reset their price; all other firms keep their prices unchanged. When firms do get to reset their price, they must take into account that the price may be fixed for many periods. We assume they do this by choosing a log-price, z_t, that minimizes the "loss function"

$$L(z_t) = \sum_{k=0}^{\infty} (\theta\beta)^k E_t \left(z_t - p^*_{t+k}\right)^2 \tag{7}$$

where β is between zero and one, and p^*_{t+k} is the log of the optimal price that the firm would set in period $t+k$ if there were no price rigidity.

This expression probably looks a bit intimidating, so it's worth discussing it a bit to explain what it means. The loss function has a number of different elements:

- The term $E_t \left(z_t - p^*_{t+k}\right)^2$ describes the expected loss in profits for the firm at time $t+k$ due to the fact that it will not be able to set a frictionless optimal price that period. This quadratic function is intended just as an approximation to some more general profit function. What is important here is to note that because the firm may be stuck with the price z_t for some time, it will lose profits relative to what it would have been able to obtain if there were no price rigidities.
- The summation $\sum_{k=0}^{\infty}$ shows that the firm considers the implications of the price set today for all possible future periods.

[4] See, for instance, Stanley Fischer (1977), "Long-Term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy*, 85, 191-205, and John Taylor (1979), "Staggered Wage Setting in a Macro Model," *American Economic Review*, Papers and Proceedings, Vol. 69, 108-113.

[5] Guillermo Calvo, "Staggered Contracts in a Utility-Maximizing Framework" *Journal of Monetary Economics*, September 1983.

- However, the fact that $\beta < 1$ implies that the firm places less weight on future losses than on today's losses. A dollar today is worth more than a dollar tomorrow because it can be re-invested. By the same argument, a dollar lost today is more important than a dollar lost tomorrow.

- Future losses are actually discounted at rate $(\theta\beta)^k$, not just β^k. This is because the firm only considers the *expected* future losses from the price being fixed at z_t. The chance that the price will be fixed until $t+k$ is θ^k. So the period $t+k$ loss is weighted by this probability. There is no point in the firm worrying too much about losses that might occur from having the wrong price far off in the future, when it is unlikely that the price will remained fixed for that long.

The Optimal Reset Price

After all that, the actual solution for the optimal value of z_t, (i.e. the price chosen by the firms who get to reset) is quite simple. Each of the terms featuring the choice variable z_t—that is, each of the $\left(z_t - p^*_{t+k}\right)^2$ terms—need to be differentiated with respect to z_t and then the sum of these derivatives is set equal to zero. This means

$$L'(z_t) = 2\sum_{k=0}^{\infty} (\theta\beta)^k E_t \left(z_t - p^*_{t+k}\right) = 0 \tag{8}$$

Separating out the z_t terms from the p^*_{t+k} terms, this implies

$$\left[\sum_{k=0}^{\infty} (\theta\beta)^k\right] z_t = \sum_{k=0}^{\infty} (\theta\beta)^k E_t p^*_{t+k} \tag{9}$$

Now, we can use our old pal the geometric sum formula to simplify the left side of this equation. In other words, we use the fact that

$$\sum_{k=0}^{\infty} (\theta\beta)^k = \frac{1}{1-\theta\beta} \tag{10}$$

to re-write the equation as

$$\frac{z_t}{1-\theta\beta} = \sum_{k=0}^{\infty} (\theta\beta)^k E_t p^*_{t+k} \tag{11}$$

implying a solution of the form

$$z_t = (1-\theta\beta)\sum_{k=0}^{\infty} (\theta\beta)^k E_t p^*_{t+k} \tag{12}$$

Stated in English, all this equation says is that the optimal solution is for the firm to set its price equal to a weighted average of the prices that it would have expected to set in the future if there weren't any price rigidities. Unable to change price each period, the firm chooses to try to keep close "on average" to the right price.

And what is this "frictionless optimal" price, p_t^*? We will assume that the firm's optimal pricing strategy without frictions would involve setting prices as a fixed markup over marginal cost:

$$p_t^* = \mu + mc_t \tag{13}$$

Thus, the optimal reset price can be written as

$$z_t = (1 - \theta\beta) \sum_{k=0}^{\infty} (\theta\beta)^k E_t (\mu + mc_{t+k}) \tag{14}$$

The New-Keynesian Phillips Curve

Now, we can show how to derive the behaviour of aggregate inflation in the Calvo economy. The following derivation is a bit subtle, and you will not be asked to repeat it in the exam.

The aggregate price level in the Calvo economy is just a weighted average of last period's aggregate price level and the new reset price, where the weight is determined by θ:

$$p_t = \theta p_{t-1} + (1 - \theta) z_t, \tag{15}$$

This can be re-arranged to express the reset price as a function of the current and past aggregate price levels

$$z_t = \frac{1}{1 - \theta} (p_t - \theta p_{t-1}) \tag{16}$$

Now, let's examine equation (14) for the optimal reset price again. We have shown that the first-order stochastic difference equation

$$y_t = ax_t + bE_t y_{t+1} \tag{17}$$

can be solved to give

$$y_t = a \sum_{k=0}^{\infty} b^k E_t x_{t+k} \tag{18}$$

Examining equation (14), we can see that z_t must obey a first-order stochastic difference equation with

$$y_t \quad = \quad z_t \tag{19}$$

$$x_t = \mu + mc_t \tag{20}$$

$$a = 1 - \theta\beta \tag{21}$$

$$b = \theta\beta \tag{22}$$

In other words, we can write the reset price as

$$z_t = \theta\beta E_t z_{t+1} + (1 - \theta\beta)(\mu + mc_t) \tag{23}$$

Substituting in the expression for z_t in equation (16) we get

$$\frac{1}{1-\theta}(p_t - \theta p_{t-1}) = \frac{\theta\beta}{1-\theta}(E_t p_{t+1} - \theta p_t) + (1 - \theta\beta)(\mu + mc_t) \tag{24}$$

After a bunch of re-arrangements, this equation can be shown to imply

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1-\theta)(1-\theta\beta)}{\theta}(\mu + mc_t - p_t) \tag{25}$$

where $\pi_t = p_t - p_{t-1}$ is the inflation rate.

This equation is known as the *New-Keynesian Phillips Curve*. It states that inflation is a function of two factors:

- Next period's expected inflation rate, $E_t\pi_{t+1}$.
- The gap between the frictionless optimal price level $\mu + mc_t$ and the current price level p_t. Another way to state this is that inflation depends positively on *real marginal cost*, $mc_t - p_t$.

Why is real marginal cost a driving variable for inflation? Firms in the Calvo model would like to keep their price as a fixed markup over marginal cost. If the ratio of marginal cost to price is getting high (i.e. if $mc_t - p_t$ is high) then this will spark inflationary pressures because those firms that are re-setting prices will, on average, be raising them.

Real Marginal Cost and Output

For simplicity, we will denote the deviation of real marginal cost from its frictionless level of $-\mu$ as

$$\hat{mc}_t^r = \mu + mc_t - p_t \tag{26}$$

so we can write the NKPC as

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1-\theta)(1-\theta\beta)}{\theta}\hat{mc}_t^r \tag{27}$$

One problem with attempting to implement this model empirically, is that we don't actually observe data on real marginal cost. National accounts data contain information on the factors that affect *average* costs such as wages, but do not tell us about the cost of producing an additional unit of output. That said, it seems very likely that marginal costs are procyclical, and more so than prices. When production levels are high relative to potential output, there is more competition for the available factors of production, and this leads to increases in real costs, i.e. increases in the costs of the factors over and above increases in prices. Some examples of the procyclicality of real marginal costs are fairly obvious. For example, the existence of overtime wage premia generally means a substantial jump in the marginal cost of labour once output levels are high enough to require more than the standard workweek.

For these reasons, many researchers implement the NKPC using a measure of the *output gap* (the deviation of output from its potential level) as a proxy for real marginal cost. In other words, they assume a relationship such as

$$\hat{mc}_t^r = \lambda y_t \tag{28}$$

where y_t is the output gap. This implies a New-Keynesian Phillips curve of the form

$$\pi_t = \beta E_t \pi_{t+1} + \gamma y_t \tag{29}$$

where

$$\gamma = \frac{\lambda (1-\theta)(1-\theta\beta)}{\theta} \tag{30}$$

And this approach can be implemented empirically.

The NKPC, Monetary Policy and the Lucas Critique

The New-Keynesian approach assumes that firms have rational expectations. Thus, we can apply the repeated substitution method to equation (29) to arrive at

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k} \tag{31}$$

Inflation today depends on the whole sequence of expected future output gaps. Thus, the NKPC sees inflation as behaving according to the classic "asset-price" logic that we saw with the dividend-discount stock price model.

The NKPC model turns much of the standard reasoning about inflation, based on the accelerationist model, on its head. The idea that there is considerable inertia in inflation,

and hence that it is difficult to reduce inflation quickly, does not hold in this framework—indeed, according to the NKPC, there is no "intrinsic" inertia in inflation, in the sense that there is no structural dependence of inflation on its own lagged values. Thus, the NKPC has very different implications for monetary policy. This model implies there is no need for gradualist policies to reduce inflation. According to the NKPC, low inflation can be achieved immediately by the central bank announcing (and the public believing) that it is committing itself to elminating positive output gaps in the future: This can be seen from equation (31).

Advocates of the NKPC will concede that the accelerationist model fits the data reasonably well. However, they view this as a so-called reduced-form relationship, not a structural relationship. For advocates of the NKPC, lagged inflation terms enter econometric Phillips curves merely because they are proxying for expectations of future values of the output gap, which are what truly determine current inflation (see equation 31). Thus, the supposed inertial nature of inflation is in reality only a sort of statistical mirage, and so should not have any bearing on the conduct of policy.

The NKPC model suggests that accelerationist Phillips curves may be subject to the Lucas critique. The correlation between lagged and expected future values of inflation is likely to vary across monetary policy regimes: In periods when the central bank has little credibility, the public may formulate its inflation expectations based on actual recent inflation performance, rather than on the public statements of the central bank. By contrast, if the central bank maintains a credible inflation target, then recent lagged values of inflation may play only a small role in the formulation of expectations. This would imply that the accuracy of accelerationist regressions would change over time as policy changes, and thus that they are not useful tools for the analysis of the effects of policy.

Of course, whether these are good policy recommendations depends on whether or not the NKPC is a good model of the inflation-output relationship.

Testing the New-Keynesian Phillips Curve

The NKPC approach is intuitively appealing but does it provide a good model of empirical inflation behaviour. One way to check is to estimate equation (29) econometrically. One problem for this idea is that the equation contains the term $E_t\pi_{t+1}$ which can't really be observed. This problem is sometimes tackled by observing that

$$E_t\pi_{t+1} = \pi_{t+1} + \epsilon_t \tag{32}$$

In other words, expected inflation equals actual inflation plus a random error term. So the equation can be re-written as

$$\pi_t = \beta\pi_{t+1} + \gamma y_t - \beta\epsilon_t \tag{33}$$

This equation can then be estimated to give values for β and γ. One problem is that the error term here $-\gamma\epsilon_t$ is correlated with one of the variables and this results in biased coefficients. The usual solution is to estimate the equation with IV—run a first-stage regression of variables thought to be uncorrelated with the error and then include the fitted value in the second stage. However, when this is done, the estimates of γ the coefficient on the output gap tends to be negative, which is not what we would have hoped for.

Another way to assess the model is to examine whether the pure forward-looking relationship

$$\pi_t = \gamma \sum_{k=0}^{\infty} \beta^k E_t y_{t+k} \tag{34}$$

seems to fit the data. A number of empirical studies have found that it does not. For instance, look at the upper panel of Figure 2: This figure is from a Rudd-Whelan paper called (for reasons that will become apparent in a minute) "Does Labor's Share Drive Inflation?"[6] The solid line in this figure is actual US inflation. The dashed line is an estimate of the present discounted value of output gaps generated using an econometric forecasting model. The figure reveals a shockingly bad performance for the NKPC model. The present value of output gaps turns out to be *negatively correlated* with inflation. So, not only does inflation not appear to equal this present value, it does not even have the correctly-signed correlation.

One way to understand these failures is to note that the discount rate β should be close to one, so one can re-write (33) as

$$\Delta\pi_{t+1} = -\gamma y_t + \beta\epsilon_t \tag{35}$$

So the model predicts a negative relationship between the change in inflation and the output gap. But we discussed earlier the good fit of models that predict a negative relationship between the change in inflation and the unemployment rate and thus a positive relationship between the change in inflation and the output gap. So the model makes a prediction that

[6] This paper can be downloaded at www.karlwhelan.com and can also be found in the April 2005 edition of the *Journal of Money, Credit, and Banking.*

is empirically wrong and it shouldn't be surprising that the key coefficient thus has the "wrong" sign.

A final way to explain the model's failure is in terms of leading indicators. According to the model, inflation should be a positive leading indicator of future values of the output gap. In fact, high inflation tends to be a negative leading indicator for output.

The Labour Share Approach

Recall that the true "driving variable" in the NKPC is actually the ratio of marginal cost to the price level, or real marginal cost as we called it. The output gap version of the model was introduced because we don't observe a time series for real marginal cost but presume that it is procyclical. Galí and Gertler (1999) suggest that the problem here may be that "potential output" is hard to measure so empirical proxies for the output gap, based on detrended actual output, may be poor proxies for the real output gap. They suggest an alternative approach based on trying to construct a more direct proxy for real marginal cost. Specifically, this approach tries to measure real marginal cost directly by assuming that it can be proxied with real average variable cost, where it is also assumed that labour is the only variable factor of production. In this case, average nominal variable cost is $\frac{wL}{Y}$ where w is the wage rate and L is labour input. So this proxy for *real* marginal cost is $\frac{wL}{pY}$, which is also the labour share of national income.

Galí and Gertler argue that this produces a more sensible version of the NKPC, and show that the leading indictor argument works better in this case. Specifically, when they use IV to estimate

$$\pi_t = \beta\pi_{t+1} + \gamma s_t - \beta\epsilon_t \tag{36}$$

where s_t is the labor share, the coefficient on s_t is positive.

However, not everyone is a big fan of this approach. The Rudd-Whelan papers have critiqued the labor share approach in a couple of ways:

- Average and marginal cost are not the same thing and are quite likely to have different cyclical properties. Indeed, it may often be the case that the two are moving in the opposite direction. For example, employed labour is often underutilized in recessions. This can lead to an increase in average costs of production: This shows up in the data as spikes in the labour share of income during recessions (see the attached Figure 1 from Rudd-Whelan for US evidence on this). However, real marginal cost almost

surely falls in recessions because elements such as overtime payments decline.

- This gives another explaination for they "success" of the positive estimated γ coefficient in the NKPC. Output gaps (and probably true real marginal cost) are procyclical and the model estimates a coefficient on them that is negative. So, the way to get a positive coefficient is to replace the output gap with something that is countercyclical. Since the labor share has a countercyclical element (rising in recessions) this does the trick.

- Econometric model used to construct present discounted values of expected future labour shares usually deliver a positive correlated with inflation, but the fits—which vary with the specific forecasting VAR—are generally very poor. For example, see the lower panel of Figure 2: It shows that the present value of labour shares generated from a VAR that fits the data well has an R^2 of only one percent.

Figure 1: Output Gap and Labor Share

NBER Recessions Shaded

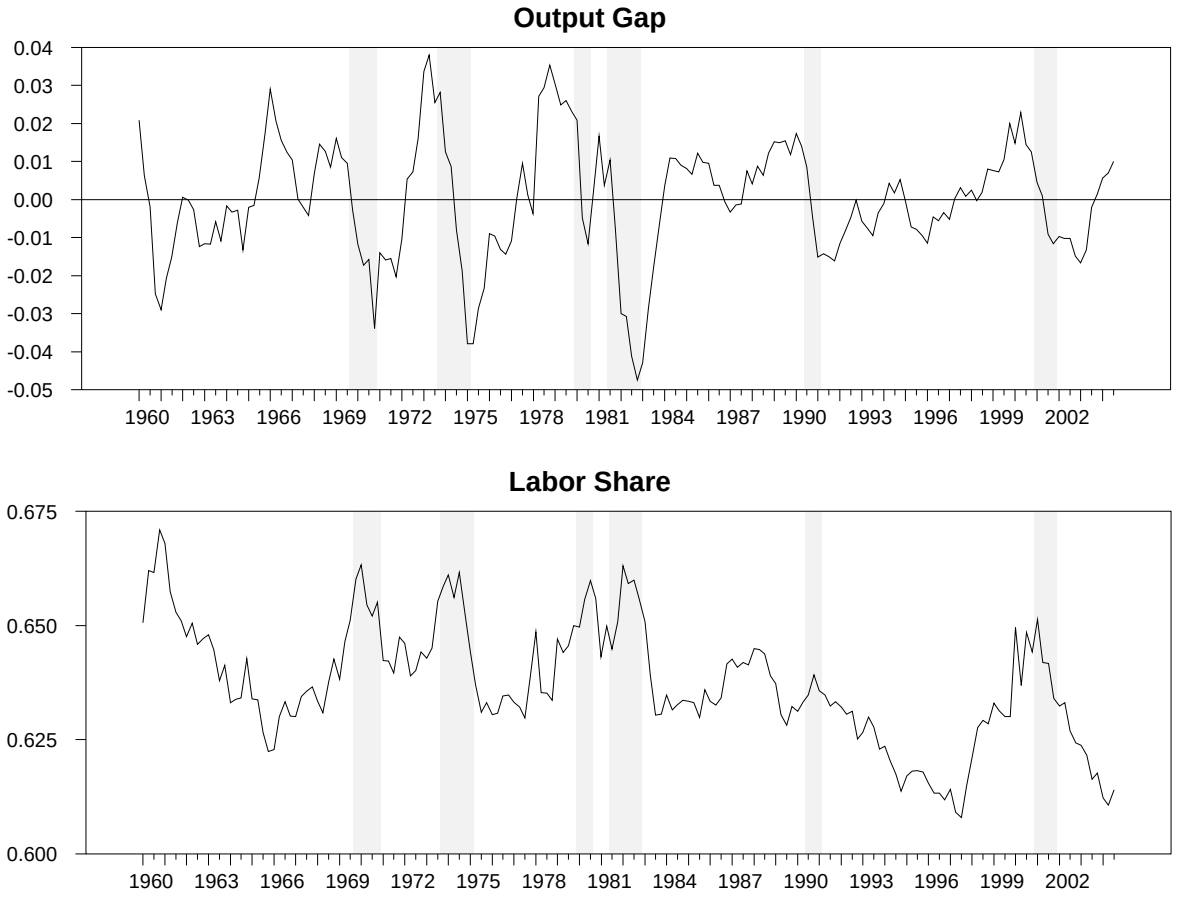

Figure 2
Actual and Predicted Inflation--Present-Value Method
(VAR models include GDP gap, labor's share, and unit labor cost growth)

A. Present Value of GDP Gaps from VAR System

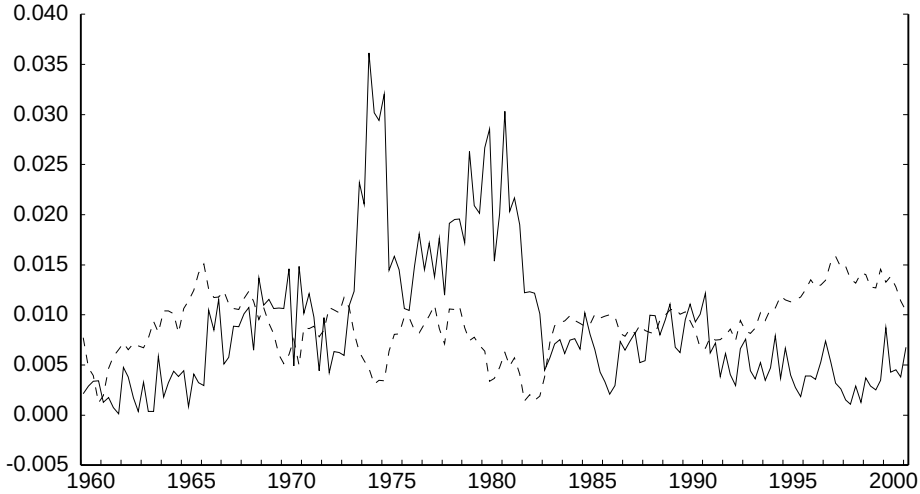

B. Present Value of Labor Income Shares from VAR System

