

Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps

N. GREGORY MANKIW AND RICARDO REIS

1. INTRODUCTION

How do employment and inflation respond to real and monetary forces? What frictions cause these macroeconomic variables to deviate from the ideals that would arise in a fully classical world? These questions are at the center of macroeconomic research, as well as at the center of much of Ned Phelps's formidable research career. Early in his career Phelps (1967, 1968), together with Milton Friedman (1968), gave us the natural rate hypothesis, which remains the benchmark for understanding monetary nonneutrality. More recently, Phelps's (1994) work on structural slumps examined the real forces that can cause the natural rate of unemployment to change over time.

This chapter offers a model that weaves these two threads of Phelps's work together. In this model, information is assumed to disseminate slowly throughout the population of wage-setters, and as a result, wages respond slowly to news about changing economic conditions. Our model includes two kinds of relevant information: news about aggregate demand, as determined by monetary policy, and news about equilibrium real wages, as determined by productivity.

We introduce this sticky-information model in Section 2. The model generalizes the one in Mankiw and Reis (2001). In the earlier paper, we applied the assumption of sticky information to the price-setting process in order to understand the dynamic response of the economy to monetary policy. The resulting model has three properties that are consistent with the conventional wisdom of central bankers and the empirical evidence of macroeconomists. First, disinflations in the model are always contractionary (although announced disinflations are less contractionary than surprise ones). Second, the model predicts that monetary

policy shocks have their maximum impact on inflation with a substantial delay. Third, given a realistic stochastic process for money growth, the model implies that the change in inflation is positively correlated with the level of economic activity, which is the well-known "accelerationist Phillips curve." The model presented here encompasses the previous model—and thus these results—as a special case.

In Section 3 we offer some new results that come from applying the sticky-information assumption to the labor market and introducing productivity as an additional driving force. We show that disinflations and productivity slowdowns have parallel effects: Both cause employment and output to fall below their classical levels until the information about the new regime works its way through the economy. Thus, the model may help explain why unemployment rose during the productivity slowdown of the 1970s and yet again during the Volcker disinflation of the early 1980s. It may also help explain why the productivity resurgence associated with the "new economy" of the 1990s coincided with falling unemployment. Ball and Moffitt (2001) and Staiger et al. (2001) have recently documented the empirical importance of productivity trends for shifts in the Phillips curve. The sticky-information model presented here offers one interpretation of those findings.

In Section 4 we examine how the economy reacts to a change in trend productivity if the monetary authority responds either by stabilizing employment or by stabilizing inflation. We find that a strict policy of stabilizing inflation causes large fluctuations in employment. We also find that the employment-stabilizing policy keeps both nominal income and the nominal wage on its predetermined path. These results are a warning for the many central banks around the world that have adopted inflation targeting as a policy framework. They suggest that employment would be more stable if the monetary authority used nominal income or the price of labor, rather than the price of goods and services, as its nominal anchor.

In Section 5 we use the model as a lens through which to examine U.S. quarterly time series. According to the model, employment is a function of unexpected movements in inflation and productivity, where expectations are formed at various points in the past. The key parameter that governs this relationship is the one that measures the stickiness of information. When we apply the model and estimate this parameter, we find that the average wage-setter updates his information about once a year. Moreover, the correlation between predicted and actual unemployment is as high as 0.60. Thus, the sticky-information model can be used to explain much of the observed variation in aggregate unemployment as a response to inflation and productivity surprises. One anomaly, however, is that productivity surprises appear to influence unemployment with a lag of several quarters.

This chapter brings together many elements from the literature on inflation-employment dynamics that Phelps helped to spawn. The natural rate hypothesis, imperfect information, infrequent adjustment, and the wage curve all play central roles in the model. Our goal is not to offer a radically new theory but to combine familiar elements in a new way in an attempt to better explain monetary nonneutrality and structural slumps.

We thank Laurence Ball and Andrew Caplin for comments. Ricardo Reis is grateful to the Fundacao Ciencia e Tecnologia, Praxis XXI, for financial support and to Jonathan Leape and the Center for Economic Performance at the London School of Economics for their hospitality during part of this research.

2. THE WAGE CURVE MEETS STICKY INFORMATION

The model we offer is designed to explain the dynamics of wages, prices, employment, and output. There are two key elements: the wage curve and the assumption of sticky information.

2.1. The Wage Curve

The starting point is an equation describing the target nominal wage w^* :

$$w_t^* = p_t + \theta_t + \alpha e_t,$$

where p is the price level, θ is labor productivity, and e is employment. All the variables are in logs, and employment is normalized so that it averages zero. This equation resembles the wage curve, as in Phelps (1994) and Blanchflower and Oswald (1995). Some might call it a “pseudo labor supply curve.”

The simplest way to motivate this equation is from the standpoint of a union that sets wages. The union has a target for the real wage that depends on its workers’ productivity. In addition, high employment makes the union more aggressive, raising the wages they demand.

Another way to view this wage equation is from the standpoint of a firm that sets wages in an efficiency-wage environment. The wage curve describes a “no-shirking condition.” Firms pay productive workers more because they have better outside alternatives. In addition, high employment increases shirking among workers because unemployment is a less potent threat. As a result, high employment induces firms to offer higher wages.

2.2. The Sticky Information Assumption

To this standard wage curve, we add the assumption that wage-setters make their decisions based on imperfect information. Every wage-setter sets a new wage every period, but they collect information about the economy and update their decisions slowly over time.

Our assumption about information arrival is analogous to the adjustment assumption in the Calvo (1983) model of price adjustment. That is, adjustment occurs as a Poisson process. In each period, a fraction λ of the wage-setters obtains new information about the state of the economy and recomputes optimal plans based on that new information. Other wage-setters continue to set wages based on old plans and outdated information. Each wage-setter has the same probability of being one of those updating its information, regardless of how long it has been since its last update.¹

1. A natural question is *whose* expectations suffer from the stickiness of information. Under the union interpretation of the wage curve, the expectations are clearly those of the union workers. The answer is less obvious in the efficiency-wage interpretation, according to which the wage curve represents the no-shirking condition. In this case, firms set the wages, but because the firms are trying to induce workers not to shirk, the expectations of the workers are again relevant.

A wage-setter that last updated its information j periods ago sets the wage

$$w_t^j = E_{t-j} w_t^*.$$

The aggregate wage is the average of wages in the economy:

$$w_t = \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j w_t^j.$$

These two equations, together with the wage-curve equation above, fully describe the process of wage setting.

The rest of the model is conventional. The overall price level is determined by the level of wages and labor productivity:

$$p_t = w_t - \theta_t.$$

This equation is both the supply curve in the output market (prices depend on costs) and the demand curve in the labor market (the real wage depends on productivity). The level of output is determined by employment and labor productivity:

$$y_t = e_t + \theta_t.$$

This is the production function. The above five equations together make up the aggregate-supply side of the model.

It is worth noting what happens in the limiting case when information is perfect ($\lambda = 1$). It is straightforward to show that in this case, $e = 0$ and $y = \theta$. That is, under perfect information, employment is constant and output mirrors productivity. In this model, all employment dynamics follow from the assumed stickiness of information.

2.3. The Sticky-Information Phillips Curve

We can solve for the price level as a function of output by combining the previous five equations. The resulting aggregate supply equation is

$$p_t = \left\{ \lambda \sum_{j=0}^{\infty} (1 - \lambda)^j E_{t-j} [p_t + \alpha y_t + (1 - \alpha)\theta_t] \right\} - \theta_t.$$

With this equation in hand, another special case is apparent. If there are no productivity shocks ($\theta = 0$), the model simplifies to the price adjustment model in Mankiw and Reis (2001). Thus, the results in that paper concerning the dynamic effects of aggregate demand apply here as well.

With some tedious algebra, which we leave to the appendix, this equation for the price level yields the following equation for the inflation rate:

$$\pi_t = [\alpha\lambda/(1-\lambda)]e_t - \Delta\theta_t + \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \Delta\theta_t + \alpha\Delta e_t),$$

where $\Delta e_t = e_t - e_{t-1}$. We call this the *sticky-information Phillips curve*.

Let us examine each of the determinants of inflation in this model: (1) High employment means higher inflation: This is the conventional short-run Phillips curve. Here it arises because high employment raises target wages, which in turn raises costs and thus the prices charged by firms. (2) Higher productivity growth lowers inflation because it lowers firms' costs of production. (3) Higher expected inflation raises inflation because, by virtue of the sticky-information assumption, past expectations are still affecting current wage increases. (4) Higher expected productivity growth raises inflation because it influences the wages that wage-setters thought were appropriate. (5) Higher expected employment growth also raises inflation because it also influences expected target wages.

2.4. Comparison to a Leading Competitor

This sticky-information Phillips curve is very different from the "new Keynesian Phillips curve," which has become popular in recent years. The latter model, based on the work of Taylor (1980), Rotemberg (1982), and Calvo (1983), implies an expectation-augmented Phillips curve, where current inflation depends on the expectation of next period's inflation. Ignoring productivity shocks for the moment, we can write the new Keynesian Phillips curve as

$$\pi_t = \gamma e_t + E_t \pi_{t+1}.$$

For a notable example of a paper applying this model of inflation dynamics, see Clarida et al. (1999). McCallum (1997) has called this model "the closest thing there is to a standard specification," while Goodfriend and King (1997) said that it is part of a "new neoclassical synthesis" in macroeconomics.

Yet this model has some significant empirical flaws. Mankiw (2001) argues that it is inconsistent with conventional views about the effects of monetary policy. Mankiw and Reis (2001) compare it with a simple version of the sticky-information model and show that the latter yields more plausible dynamics. In particular, according to the new Keynesian Phillips curve, inflation responds immediately and substantially to monetary shocks, whereas according to the sticky-information Phillips curve, the effect of monetary shocks on inflation is delayed and gradual.

The key difference between the models comes from the role of expectations. The new Keynesian Phillips curve gives a prominent role to *current* expectations of *future* inflation. These expectations can adjust quickly in response to changes in monetary policy. By contrast, the sticky-information Phillips curve gives a prominent role to *past* expectations of *current* inflation. These expectations are predetermined and thus cannot change quickly. In this way, the sticky-information model more closely resembles an earlier generation of price-adjustment models proposed by Fischer (1977) and Phelps and Taylor (1977).

3. DISINFLATIONS AND PRODUCTIVITY SLOWDOWNS

We now consider the dynamics of employment and inflation in this sticky-information model. To do this, we have to close the model with a specification for aggregate demand. One approach to doing this would be to add a goods market equilibrium condition (i.e., an IS curve) and a monetary policy reaction function for the interest rate (such as a Taylor rule). This approach has several advantages and might be pursued in future work. However, in this chapter, we adopt an approach that, although perhaps less realistic, is simpler and more transparent. We use a simple money market equilibrium condition:

$$m_t - p_t = \beta y_t.$$

The parameter β is the income elasticity of money demand. More important for our purposes, it determines the elasticity of the aggregate demand curve. The smaller β is, the flatter the aggregate demand curve. The variable m can be viewed narrowly as the money supply or broadly as a summary measure of all variables that shift the aggregate demand for goods and services. Thus, we can refer to m as either money or aggregate demand.

With this equation, we have now fully described the model. For purposes of numerical illustration, we have to choose values for the parameters. We pick $\beta = \frac{1}{2}$. For some of the qualitative results reported below, it is important that $\beta < 1$. We also set $\alpha = 0.1$ and $\lambda = 0.25$. The value of α measures the sensitivity of the target real wage to labor market conditions. In the terminology of Ball and Romer (1990), it measures the degree of "real rigidity." If the period is taken to be a quarter, then $\lambda = 0.25$ means that wage-setters update their information on average once a year. These parameter values strike us as plausible, and we use them for the purposes of illustration. However, in a later section, we use the time-series data to estimate λ , the key parameter measuring the stickiness of information, and find that $\lambda = 0.25$ fits the data well.

Here we look at how this economy responds to a sudden change in regime. In the first case, we assume that the annual rate of growth of money m falls by 2 percentage points, or 0.5 percentage points per quarter. In the second case, we assume that the growth rate of productivity θ falls by this amount. In each case, the change in the exogenous growth rate is assumed to be sudden, unexpected, and permanent.

The appendix explains in detail how we solve for the resulting equilibrium path for the economy, whereas here we present and discuss the solution in more intuitive terms. Figure 1 shows the response to each of these disturbances. The top panel shows the path of employment e , and the bottom shows the response of inflation π . Recall that the natural rate of employment in this model—the level of employment that would prevail if information were perfect—is normalized to zero. Thus, the employment variable measures how far the economy is operating from what might be viewed as its "full employment" level.

Figure 1 shows that employment follows a hump-shaped pattern in response to both disinflations and productivity slowdowns. Employment falls for a while as the impact of the change in regime builds over time. But then the effect on




Figure 1. Employment and inflation after a 2 percent fall in the annual growth rate of money or productivity at date 0.

employment dissipates as the information about the change in regime spreads its way throughout the population. For these parameters, the maximum impact on employment occurs at six quarters, although of course this result is sensitive to the parameter values we have chosen.

The two inflation paths in the bottom panel of Figure 1 are quite different from each other. In the case of a slowdown in money growth, inflation declines only gradually. It overshoots the new lower level briefly as agents learn of the new regime and correct their previous mistakes. In the case of a productivity slowdown, inflation immediately spikes up, as firms pass on higher costs to consumers. Inflation gradually declines as wages start reflecting lower productivity growth. After a brief period of overshooting, inflation eventually settles at a new higher level.

The reader might have noticed that the dynamic paths for employment are similar in the case of disinflations and productivity slowdowns. This is not a coincidence. In fact, the appendix proves the following:

PROPOSITION: Productivity has an impact on employment that is exactly $1 - \beta$ times that of aggregate demand.

This naturally leads to:

COROLLARY: If the aggregate demand curve is unit elastic ($\beta = 1$), then productivity has no effect on employment.

The intuition behind the corollary is the following: Adverse news about productivity influences target nominal wages in two ways. They fall because target real wages are lower, and they rise because the price level is higher. If $\beta = 1$ and the path of m is held constant, these two effects exactly cancel. In this special case, wages stay exactly on track to maintain full employment, despite the slow diffusion of information.

4. STABILIZATION POLICY IN THE FACE OF PRODUCTIVITY CHANGE

The outcomes shown in Figure 1 assume that when the economy experiences a productivity slowdown, the monetary authority holds aggregate demand on the same path that it otherwise would have followed. This benchmark is natural, but it need not hold. As employment and inflation fluctuate, the monetary authority may well respond by altering aggregate demand. Here we consider two other natural benchmarks.

One possibility is that the monetary authority may choose a policy to stabilize employment. This means it would adjust aggregate demand m to keep $e = 0$ at all times. Because fluctuations in employment in this model arise because of imperfections in private information and because the monetary authority is assumed to have full information, it might well try to achieve the level of employment and output that private decision-makers would choose on their own if only they were better informed.

Alternatively, the central bank may choose to stabilize inflation. This means it would adjust aggregate demand to maintain a constant inflation rate. This benchmark is interesting in part because many central banks are in fact now committed to inflation targeting as a policy framework (although many such regimes allow temporary deviations from a strict target).

There is also another reason to focus on the employment path that follows under an inflation-stabilizing policy: This path corresponds to what econometricians might measure as the NAIRU (the nonaccelerating inflation rate of unemployment). In this model, the natural rate of employment—that is, the level that would prevail with full information—is assumed constant. Yet the rate of employment associated with nonaccelerating inflation is not constant. To find the level of employment that corresponds to the measured NAIRU, we can solve the model assuming that monetary policy stabilizes inflation.

Figure 2 shows employment and inflation in response to a productivity slowdown under three policies: constant aggregate demand, employment stabilization, and inflation stabilization. The appendix presents the details of solution.

The figure shows that inflation stabilization leads to very large fluctuations in employment. To extinguish the impact on inflation, monetary policy responds to the adverse shift in aggregate supply by contracting aggregate demand. This instability in employment under inflation stabilization can be viewed as a warning to the many central banks around the world that have adopted inflation targeting. In essence, a productivity slowdown in this model looks like a rise in the NAIRU, which is consistent with the empirical results in Ball and Moffitt (2001) and Staiger et al. (2001). Yet the natural rate—the level of employment that would prevail under full information—has not changed. The fall in employment that results from a productivity slowdown is no more desirable in this model than any other downturn. A policy of inflation stabilization not only tolerates the downturn in employment but exacerbates it to keep inflation under control.

Figure 2 also shows that a policy of employment stabilization leads to a permanent rise in the inflation rate. That is, for the central bank to maintain full employment during a productivity slowdown, it has to accommodate the adverse shift in aggregate supply by raising aggregate demand, permitting a higher inflation rate. The rise in inflation needed to stabilize employment equals the magnitude of the slowdown in productivity growth. That is, a slowdown of 2 percent per year requires an increase in the inflation rate of 2 percentage points.

There are other ways to describe the policy of stabilizing employment. It turns out that along the path with stable employment, both nominal income $p + y$ and the nominal wage w equal their previously expected values. Turning this observation around leads to a policy: If the central bank commits to keeping either nominal income or the nominal wage at a constant level (or growing at some constant rate), aggregate demand and inflation in the price of goods and services will automatically respond to a productivity shift by exactly the right amount to maintain full employment.

Overall, this model suggests that inflation targeting may not be the best framework for monetary policy in the face of changing productivity. A better target variable than the price of goods and services is nominal income or the price of labor. Like inflation targeting, targeting nominal income or the nominal wage gives monetary policy a nominal anchor, but it does so in a way that permits greater stability in employment. There is a long tradition suggesting the desirability of nominal income targeting; see Tobin (1980) and Hall and Mankiw (1994) for two




Figure 2. Employment and inflation after a 2 percent fall in the annual growth rate of productivity at date 0, under different policies.

examples. But, as far as we know, the possibility of targeting the price of labor has not received much attention.

One exception is the plan discussed by Phelps (1978). He concludes, “the program envisioned here aims to stabilize wages on a level or rising path, leaving the price level to be buffeted by supply shocks and exchange rate disturbances.”

This is precisely the policy suggested by the sticky-information model of wage setting. We agree with Phelps that nominal wage targeting seems like a promising alternative to inflation targeting as a monetary policy rule.²

5. AN EMPIRICAL APPLICATION

In the model presented here, inflation and productivity surprises drive employment fluctuations. This section examines this prediction using U.S. time-series data to get a sense of how well the model works and to learn where it has problems matching the world.

5.1. The Approach

We start with the sticky-information Phillips curve presented earlier:

$$\pi_t = [\alpha\lambda/(1-\lambda)]e_t - \Delta\theta_t + \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \Delta\theta_t + \alpha\Delta e_t).$$

Now we turn this equation around to use it as a theory of employment:

$$e_t = [(1-\lambda)/\alpha\lambda] \left[\pi_t + \Delta\theta_t - \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \Delta\theta_t + \alpha\Delta e_t) \right],$$

which implicitly expresses employment as a function of inflation, productivity, and the past expectations of these variables. Our goal is to see whether this equation can help explain observed U.S. fluctuations. To do this, we make two auxiliary assumptions.

First, we assume that inflation and productivity growth follow univariate stochastic processes. Their moving average representations can be written as

$$\pi_t = \sum_{j=0}^{\infty} \rho_j \varepsilon_{t-j} \quad \text{and} \quad \Delta\theta_t = \sum_{j=0}^{\infty} \eta_j v_{t-j}.$$

In this case, the expectations of inflation and productivity growth are a function of their past univariate innovations.³ As the appendix shows, the model can now be solved as

2. Aoki (2001) presents a related result: In a world with a sticky-price sector and a flexible-price sector, optimal monetary policy targets inflation in the sticky-price sector. In our model, the labor market is analogous to the sticky-price sector.

3. In reality, expectations may also depend on other information. For example, news about monetary policy may affect expected inflation. But experience suggests that the improvement from multivariate over univariate forecasting is often slight. Future work could improve upon this assumption, but we probably do not go too far wrong in assuming univariate processes for inflation and productivity growth.

$$\alpha e_t = \sum_{j=0}^{\infty} \gamma_j \varepsilon_{t-j} + \sum_{j=0}^{\infty} \Psi_j v_{t-j},$$

where the parameters γ_j and Ψ_j are functions of ρ_j , η_j , and λ . This equation can be used to predict fluctuations in employment, taking as inputs the univariate innovations in inflation and productivity.

Second, we assume that the observed unemployment rate is linearly related to a moving average of e , the employment variable in our model. In particular, the unemployment rate u is assumed to be

$$u_t = \delta_0 - \delta_1 e_t^s,$$

where e_t^s is defined as

$$e_t^s = (e_{t-3} + 2e_{t-2} + 3e_{t-1} + 4e_t + 3e_{t+1} + 2e_{t+2} + e_{t+3})/16.$$

We found that without any smoothing of the predicted employment variable, the model produces too much high-frequency variation in the predicted unemployment rate. This triangular filter smooths out the rapid quarter-to-quarter fluctuations in the predicted series, leaving the business cycle variation and the longer-term trends.⁴

Given these two auxiliary assumptions, we can use the model to produce a predicted path of unemployment for any given set of parameters. The predicted unemployment rate in any period depends on the history of shocks to inflation and productivity.

5.2. Parameter Estimation

The next issue is how to choose the parameters. We estimate the parameters of the univariate processes for inflation and productivity growth as autoregressions using ordinary least squares. We then invert these autoregressions to obtain ρ_j and η_j , the parameters of the moving average representations. The residuals from the autoregressions are the shocks ε_t and v_t , which we feed into the model to explain fluctuations in unemployment.

The remaining three parameters are λ , δ_0 , and δ_1/α . (The parameters α and δ_1 are not identified separately.) We estimate these parameters using a least-squares procedure; that is, we choose these parameters to minimize the sum of squared deviations between predicted and actual unemployment.

The results from this procedure have a simple interpretation. The estimated δ_0 equals the mean unemployment rate. The estimated δ_1/α ensures that the prediction

4. There are two natural hypotheses to explain the high-frequency variation predicted by the model without the filter. (1) Some high-frequency variation in productivity growth and inflation is measurement error. (2) The high-frequency variation in productivity growth and inflation in the data is real, but for some reason, employment responds more sluggishly in the world than it does in our model. Very possibly each explanation has an element of truth.

error is not correlated with the predicted value. These two parameters do not affect the autocorrelations of predicted unemployment or the cross-correlations between predicted and actual. In some sense, these parameters are not interesting: Changing δ_0 and δ_1/α alters the mean and variance of predicted unemployment without altering the dynamics in any other way.

The interesting parameter is λ , which measures the rate of information arrival. Our least-squares procedure picks this parameter to make the model fit the data. In particular, it picks the value for λ that maximizes the correlation between predicted and actual unemployment.

5.3. Data and Results

We use quarterly, seasonally adjusted U.S. data from the first quarter of 1959 to the first quarter of 2001. Inflation is measured by the consumer price index. Productivity growth is the growth of output per hour in the business sector. Unemployment is the civilian unemployment rate.

The univariate autoregressions indicate that inflation and productivity growth follow very different stochastic processes. Productivity growth is indistinguishable from white noise, indicating no persistence. By contrast, inflation is well modeled as an AR(3) with autoregressive parameters 0.37, 0.22, and 0.29. These three coefficients on lagged inflation sum to 0.88, indicating that inflation is borderline nonstationary. We use the white noise specification for productivity growth and the AR(3) parameters for inflation to calculate the moving average parameters and the estimated innovations.⁵

Next we fit the model using the unemployment data. The estimated value of λ , the rate of information arrival, is 0.25. This means that wage setters in the economy update their information on average every four quarters. Figure 3 shows the actual and predicted unemployment rate using this estimate. The correlation between these two series is 0.47.

Figure 4 decomposes the predicted unemployment rate from Figure 3 into two pieces. The top panel shows the piece of predicted unemployment that is driven by inflation innovations. The bottom panel shows the piece of predicted unemployment that is driven by productivity innovations. Actual unemployment is shown in both panels, and all the parameters are held constant at the same values as in Figure 3.

The figure shows that inflation and productivity innovations are both important for explaining unemployment. The correlation of unemployment with the inflation

5. When we tried using an estimated AR(3) process for productivity growth, the results were almost the same as those reported here because the estimated autoregressive coefficients were so small. Note that the case we use, according to which the level of productivity follows a random walk, is similar to the specifications commonly used in the real business cycle literature. Exploring alternative processes may be fruitful in future work. If we had assumed a stochastic process with a slowly changing conditional mean (such as an IMA(1,1) process with a large, negative MA parameter), then employment would be related to productivity growth minus a weighted average of past productivity growth. This specification would more closely resemble the one in Ball and Moffitt (2001). See also Grubb et al. (1982, sec. 5) for a related earlier work.

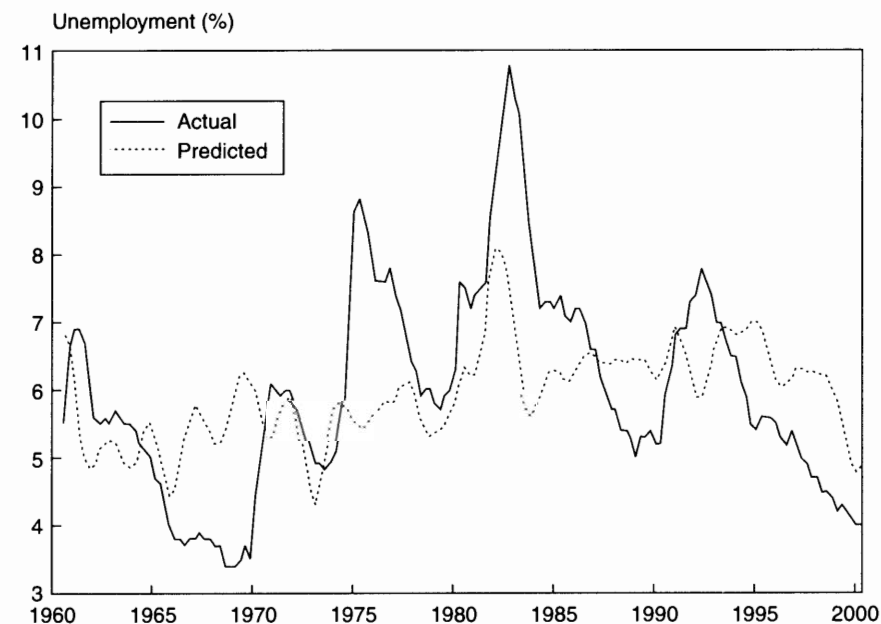


Figure 3. Actual and predicted unemployment.

piece of the predicted series is 0.33. The correlation of unemployment with the productivity piece is 0.17.

Inspection of Figure 4 reveals an anomaly: Unemployment as predicted by productivity innovations moves in advance of actual unemployment. Although this piece of predicted unemployment has a correlation with contemporaneous unemployment of only 0.17, its correlation with unemployment 1 year ahead is 0.48. Productivity appears to take more time to influence unemployment in the world than it does in our model.

As a rough (and, we admit, ad hoc) fix for this lagged response, we add a four-quarter delay to the estimated model. The employment equation then becomes

$$\alpha e_t = \sum_{j=0}^{\infty} \gamma_j \varepsilon_{t-j} + \sum_{j=0}^{\infty} \Psi_j v_{t-j-4}.$$

Otherwise, the model and estimation procedure are the same. The estimate of λ is now 0.26, almost identical to our earlier estimate. The predicted and actual unemployment rates, shown in Figure 5, now have a correlation of 0.60. Adding this lag in the effect of productivity significantly improves the model's fit.

In the end, the empirical verdict on this model is mixed. On the one hand, the model fails to capture some important dynamics in the data. The delay between productivity innovations and unemployment is a puzzle.⁶ On the other hand,

6. This delay in the response to productivity shocks may be related to the mechanism highlighted in Basu et al. (1998). If prices are sticky in the short run, then a good productivity shock can reduce

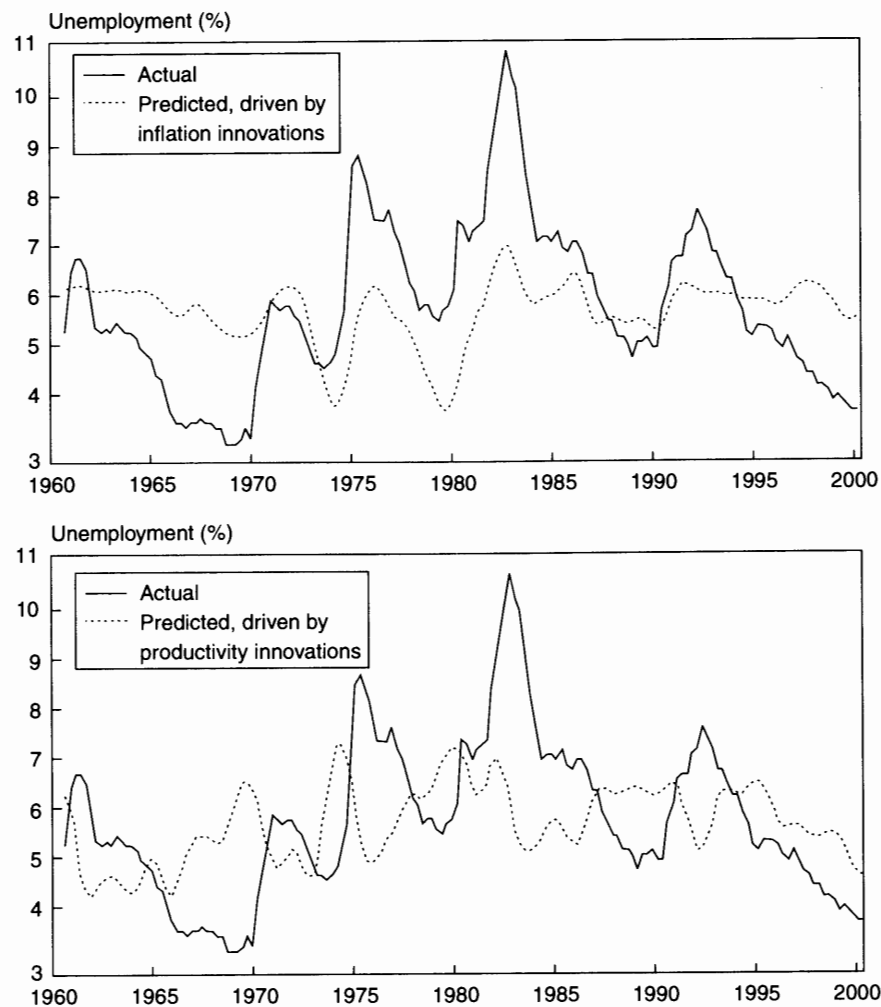


Figure 4. Decomposition of predicted unemployment into inflation- and productivity-driven innovations.

inflation and productivity surprises can account for much of the observed evolution in U.S. unemployment. The simulation in Figure 5 captures many of the main features of the unemployment data, including the big recessions in 1975 and 1982, the rising trend level of unemployment rate during the 1970s and early 1980s, and the declining unemployment rate in the late 1990s.

There are many ways in which this empirical model could be improved. The model excludes many determinants of unemployment, such as changing

employment, because firms need fewer workers to produce the same level of output. Another possibility is that the short-run changes in measured labor productivity are a reflection of labor hoarding over the business cycle rather than exogenous movements in true productivity.




Figure 5. Actual and predicted unemployment with a four-quarter delay to productivity innovations.

demographics, unionization, and minimum wage laws. It also excludes many shocks that might shift the inflation-unemployment relationship, such as exchange rates, food and energy prices, and wage-price controls. The only inputs into these simulations are inflation and productivity surprises. Absent changes in measured productivity, the model interprets all movements in inflation as driven by aggregate demand. Incorporating other sources of unemployment and inflation fluctuations into the model and adding a richer dynamic structure to better explain the link between productivity and employment are tasks we leave for future work.

6. CONCLUSION

Since Friedman (1968) and Phelps (1967, 1968) introduced the natural rate hypothesis in the 1960s, expectations have played a central role in understanding the dynamics of inflation and employment. Early empirical work testing the hypothesis was based on the assumption of adaptive expectations, according to which expected inflation is a weighted average of past inflation. Yet Lucas (1972) and Sargent (1971) forcefully criticized that approach. Since then, economists have most often relied on the assumption of rational expectations.

Over time, however, economists have increasingly realized that while agents may be too smart to form expectations adaptively, they may not be smart enough to form them rationally. Unfortunately, finding a safe haven between the stupidity of adaptive expectations and the hyperintelligence of rational expectations is

not easy. Friedman (1979), Summers (1986), and Sargent (1999) have explored the possibility that least-squares learning can help explain how expectations evolve over time. Sims (2001) has suggested modeling agents' limited capacity for processing information. Ball (2000) has proposed that agents form optimal univariate forecasts, but ignore the possible gains from multivariate forecasts. Each of these approaches has some appeal, but each lacks the parsimony that makes rational expectations so compelling.

Like all these efforts, the sticky-information model we have explored here and in our previous paper attempts to model agents that are smart, but not too smart. Our agents form expectations rationally, but they do not do so often. Because of either costs of acquiring information or costs of recomputing optimal plans, information diffuses slowly throughout the population. As a result, expectations conditional on old information continue to influence current behavior.

Recently, many economists have been pursuing "behavioral economics," a research program that tries to incorporate the insights of psychology into economics. The starting point of this work is the observation that people are not quite as rational as the *homo economicus* assumed in standard models. The sticky-information model of inflation-employment dynamics can be viewed as a modest contribution to this literature. The closest antecedent is the work of Gabaix and Laibson (2001), which tries to explain consumption behavior and the equity premium puzzle by assuming that consumers are slow to recognize changes in the values of their portfolios. The sticky-information model we have presented here applies the analogous assumption to the process of wage setting.

From a theoretical standpoint, assuming sticky information is an attempt to have your cake and eat it too. Like many previous efforts, our model tries to describe agents who do not fully understand their environment and gradually learn about it over time. This lack of full understanding motivates our application of Calvo's assumption of Poisson adjustment to the acquisition of information. Whether this approach is the best way to model the imperfections of human behavior is hard to say. But without doubt, it has a major advantage: Once this leap of faith is taken, the powerful tools of rational expectations become available to solve for the resulting equilibrium.

Of course, the validity of this model of inflation-unemployment dynamics is ultimately an empirical question. The model fits some broad stylized facts. In our previous paper, we showed that the model can explain the dynamic response of the economy to monetary policy: It can explain why disinflations are costly, why monetary shocks have a delayed and gradual effect on inflation, and why changes in inflation are positively correlated with the level of economic activity. This chapter has moved the analysis to the labor market and introduced a role for productivity as a driving force. The model can now explain why productivity slowdowns, such as the one that occurred in the 1970s, are associated with rising unemployment (and a rising NAIRU) and why productivity accelerations, such as the one that occurred in the United States in the 1990s, are associated with declining unemployment (and a declining NAIRU). When we applied the model to U.S. time-series data, we found that inflation and productivity surprises account for a sizable fraction of observed unemployment fluctuations.

The model proposed here suggests a possible problem with the policy of inflation targeting that many of the world's central banks adopted during the 1990s. In this model, stabilizing employment in the face of a productivity slowdown requires a rise in the inflation rate of an equal magnitude. The model suggests that a better nominal anchor than the price of goods and services is nominal income or the nominal wage. Targeting inflation in the price of labor, rather than in the price of goods and services, is a policy rule that deserves a closer look.

APPENDIX

This appendix formally justifies some of the claims asserted in the body of the chapter.

The Sticky-Information Phillips Curve

The three equations describing wage setting, together with the demand for labor, lead to

$$p_t = \left[\lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}(p_t + \alpha e_t + \theta_t) \right] - \theta_t. \quad (\text{A1})$$

Take out the first term to obtain

$$p_t = \lambda[p_t + \alpha e_t + \theta_t] + \left[\lambda \sum_{j=0}^{\infty} (1-\lambda)^{j+1} E_{t-1-j}(p_t + \alpha e_t + \theta_t) \right] - \theta_t \quad (\text{A2})$$

and subtract the p_{t-1} version of (A1) from (A2) to obtain, after some rearranging,

$$\begin{aligned} \pi_t = \lambda(p_t + \alpha e_t + \theta_t) + \left[\lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \alpha \Delta e_t + \Delta \theta_t) \right] \\ - \left[\lambda^2 \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(p_t + \alpha e_t + \theta_t) \right] - \Delta \theta_t, \end{aligned} \quad (\text{A3})$$

where Δ is the first difference operator and the inflation rate is $\pi_t = \Delta p_t$. Multiplying through equation (A2) by λ and rearranging, we get

$$\begin{aligned} (1-\lambda)\lambda^2 \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(p_t + \alpha e_t + \theta_t) \\ = \lambda[(1-\lambda)p_t - \lambda\alpha e_t + (1-\lambda)\theta_t]. \end{aligned} \quad (\text{A4})$$

Replacing (A4) for the last term in (A3) leads to, after rearranging, the sticky-information Phillips curve:

$$\pi_t = [\alpha\lambda/(1-\lambda)]e_t - \Delta\theta_t + \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \Delta\theta_t + \alpha\Delta e_t). \quad (\text{A5})$$

The Response of Employment and Inflation to a Fall in the Growth Rate of Aggregate Demand and Productivity

Replacing for y_t using the aggregate demand equation into the aggregate supply equation, we obtain the law of motion for the price level:

$$p_t = \left\{ \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}[(1-\alpha/\beta)p_t + (\alpha/\beta)m_t + (1-\alpha)\theta_t] \right\} - \theta_t. \quad (\text{A6})$$

First, say that m_t was growing at the rate 0.005, up to date -1 , when it reaches the level 0. Then, unexpectedly, the growth rate falls to 0, so m_t stays at level 0 forever. Formally: $m_t = 0.005(1+t)$ for $t \leq -1$ and $m_t = 0$ for $t \geq 0$. Also, set $\theta_t = 0$ for all t . Then, we can break equation (A6) into two components:

$$p_t = \lambda \sum_{j=0}^t (1-\lambda)^j E_{t-j}[(1-\alpha/\beta)p_t + (\alpha/\beta)m_t] + \lambda \sum_{j=t+1}^{\infty} (1-\lambda)^j E_{t-j}[(1-\alpha/\beta)p_t + (\alpha/\beta)m_t]. \quad (\text{A7})$$

In the first summation term are included the agents that set their expectations after the change in m_t . Thus, they set expectations with full information: $E_{t-j}(p_t) = p_t$ and $E_{t-j}(m_t) = 0$. In the second summation term are agents who set expectations prior to the change: $E_{t-j}(p_t) = E_{t-j}(m_t) = 0.005(1+t)$. Thus, (A7) becomes

$$p_t = (1-\alpha/\beta)[1 - (1-\lambda)^{t+1}]p_t + 0.005(1+t)(1-\lambda)^{t+1}, \quad (\text{A8})$$

which, after rearranging, gives the solution:

$$p_t = \frac{0.005(1+t)(1-\lambda)^{t+1}}{1 - (1-\alpha/\beta)[1 - (1-\lambda)^{t+1}]}. \quad (\text{A9})$$

Employment is given by $e_t = -p_t/\beta$ and inflation is $\pi_t = \Delta p_t$.

The fall in productivity is formalized as: $\theta_t = 0.005(1+t)$ for $t \leq -1$ and $\theta_t = 0$ for $t \geq 0$, and $m_t = 0$ for all t . Starting again from (A6) we obtain, by very similar steps, the response of prices:

$$p_t = \frac{0.005(1-\beta)(1+t)(1-\lambda)^{t+1}}{1 - (1-\alpha/\beta)[1 - (1-\lambda)^{t+1}]}. \quad (\text{A10})$$

Inflation and employment follow immediately.

Proof of the Proposition and Corollary

Starting from (A1), replacing for p_t from the aggregate demand equation and for y_t from the production function, we obtain the expectational difference equation for employment as a function of the two exogenous processes, m_t and θ_t :

$$\beta e_t = m_t + (1-\beta)\theta_t - \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}[m_t + (1-\beta)\theta_t + (\alpha-\beta)e_t]. \quad (\text{A11})$$

If we define the composite disturbance as $m_t + (1-\beta)\theta_t$, employment is driven solely by this term. Shocks to m_t or $(1-\beta)\theta_t$ have exactly the same effect on the equation above and thus on employment. This shows the proposition. The corollary follows from setting $\beta = 1$ in (A11) and noting that θ_t does not enter the stochastic equation determining employment.

Stabilization Policy and Productivity Change

We start with the stable employment policy. From (A11), it follows that if $m_t = -(1-\beta)\theta_t$ then we obtain an expectational difference equation involving only e_t , which has solution $e_t = 0$. This is the policy that stabilizes employment. Given the path for θ_t above, we then have the policy path m_t , and since we know that this ensures $e_t = 0$, then $y_t = \theta_t$ from the production function. The path of prices p_t comes from the aggregate demand equation $p_t = -\theta_t$. A permanent drop in the growth rate of productivity therefore leads to a permanent rise in inflation.

Alternatively, note that combining the wage curve and the labor demand (output supply) equations into the equation for the aggregate wage we obtain

$$w_t = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}(w_t + \alpha e_t). \quad (\text{A12})$$

From (A12), having the wage rate equal to some forever known constant ensures that $e_t = 0$. Without loss of generality, we set that constant to zero. The associated policy is $m_t = -(1-\beta)\theta_t$, from using the aggregate demand equation and the production function. Thus, an employment stabilization policy is equivalent to a policy that stabilizes the wage rate in the sense of having w_t deterministic or known to all.

Similarly, we use the labor demand and the production function equations to replace w_t and θ_t in (A12) above and obtain

$$p_t + y_t - e_t = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}[p_t + y_t + (1-\alpha)e_t]. \quad (\text{A13})$$

Thus, by the same argument as above, stabilizing nominal income $p_t + y_t$ is equivalent to stabilizing employment.

We now turn our attention to inflation targeting. Without loss of generality we set the target to zero, and the level to zero as well, so the policy aim is to have $p_t = 0$ at all periods. From the law of motion for the price level in (A6) and the set path for prices:

$$\theta_t = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-j}[(\alpha/\beta)m_t + (1-\alpha)\theta_t]. \quad (\text{A14})$$

Given a path for θ_t , the necessary path of policy is given by the solution to this equation. For $t \leq -1$, all agents have full information and $\theta_t = 0.005(1+t)$ so the policy is $m_t = 0.005\beta(1+t)$. For $t \geq 0$, breaking the sum in two components, the first referring to informed agents that set their expectations at or after date 0 and the second to uninformed agents that set their expectations before date 0, we obtain

$$0 = (\alpha/\beta) [1 - (1-\lambda)^{t+1}] m_t + 0.005(1+t)(1-\lambda)^{t+1}, \quad (\text{A15})$$

which, after rearranging, becomes

$$m_t = -\frac{0.005\beta(1+t)(1-\lambda)^{t+1}}{\alpha - \alpha(1-\lambda)^{t+1}}. \quad (\text{A16})$$

The path for employment for $t \geq 0$ is $e_t = y_t = m_t/\beta$.

Empirical Implementation

Since the only forces driving employment are inflation and productivity growth, the general solution for employment will be a linear function (with undetermined coefficients) of the shocks to these variables:

$$\alpha e_t = \sum_{j=0}^{\infty} (\gamma_j \varepsilon_{t-j} + \Psi_j v_{t-j}). \quad (\text{A17})$$

Using the MA representations for inflation, productivity growth, and employment; plugging these into the equation for employment in the main text; and taking the relevant expectations, we obtain

$$\sum_{j=0}^{\infty} (\gamma_j \varepsilon_{t-j} + \Psi_j v_{t-j}) = [(1-\lambda)/\lambda] \left\{ \sum_{j=0}^{\infty} (\rho_j \varepsilon_{t-j} + \eta_j v_{t-j}) - \left[\lambda \sum_{j=0}^{\infty} (1-\lambda)^j \sum_{i=j+1}^{\infty} (\rho_i \varepsilon_{t-i} + \eta_i v_{t-i} + \gamma_i \varepsilon_{t-i} + \Psi_i v_{t-i}) \right] + \left[\lambda \sum_{j=0}^{\infty} (1-\lambda)^j \sum_{i=j}^{\infty} (\gamma_i \varepsilon_{t-i-1} + \Psi_i v_{t-i-1}) \right] \right\}. \quad (\text{A18})$$

Matching coefficients gives the solution

$$\gamma_0 = [(1-\lambda)/\lambda]\rho_0, \quad (\text{A19})$$

$$\gamma_j = \{(1-\lambda)^{j+1}/[1-(1-\lambda)^{j+1}]\} \rho_j + \{1 - \{\lambda/[1-(1-\lambda)^{j+1}]\}\} \gamma_{j-1}, \quad (\text{A20})$$

$$\Psi_0 = [(1-\lambda)/\lambda]\eta_0, \quad (\text{A21})$$

$$\Psi_j = \{(1-\lambda)^{j+1}/[1-(1-\lambda)^{j+1}]\} \eta_j + \{1 - \{\lambda/[1-(1-\lambda)^{j+1}]\}\} \Psi_{j-1}. \quad (\text{A22})$$

Given the MA parameters ρ_j and η_j and a series for the shocks ε_t and v_t (all of which we obtain by OLS regressions on the data) we can generate a predicted employment series from (A17), which, after applying the triangular filter in the text, produces a series for smoothed employment: e_t^s . Note then first that, from (A17) and the additive nature of the smoother, $1/\alpha$ factors out of the expressions for e_t^s . Second, that since ρ_j , η_j , ε_t , and v_t are determined by first-stage OLS regressions, the only parameter left determining predicted smooth employment is λ , so that we can write $\alpha e_t^s = g(\lambda)$, where $g(\cdot)$ is the function defined by (A17) and (A19) to (A22). The equation for our predicted unemployment series is

$$\hat{u}_t = \delta_0 - (\delta_1/\alpha)g(\lambda). \quad (\text{A23})$$

We pick parameters by nonlinear least-squares to minimize the sum of squares residuals:

$$\min_{\delta_0, \delta_1/\alpha, \lambda} \sum_{t=1}^T (u_t - \hat{u}_t)^2. \quad (\text{A24})$$

Finally, the decomposition of predicted unemployment in Figure 4 follows directly from the two components of (A17).

REFERENCES

- Aoki, K. (2001), "Optimal Monetary Policy Responses to Relative-Price Changes," *Journal of Monetary Economics* 48:55–80.
- Ball, L. (2000), "Near Rationality and Inflation in Two Monetary Regimes," NBER Working Paper No. 7988.
- Ball, L. and R. Moffitt (2001), "Productivity Growth and the Phillips Curve," mimeo, Johns Hopkins University, June.
- Ball, L. and D. Romer (1990), "Real Rigidity and the Nonneutrality of Money," *Review of Economic Studies* 57:539–52.
- Basu, S., J. G. Fernald, and M. S. Kimball (1998), "Are Productivity Improvements Contractionary?," mimeo, University of Michigan.
- Blanchflower, D. G. and A. J. Oswald (1995), "An Introduction to the Wage Curve," *Journal of Economic Perspectives* 9(3):153–67.
- Calvo, G. A. (1983), "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics* 12:383–98.
- Clarida, R., M. Gertler, and J. Gali (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37(4):1661–1707.

- Fischer, S. (1977), "Long-term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy* 85:191–205.
- Friedman, B. M. (1979), "Optimal Expectations and the Extreme Assumptions of Rational Expectations Macromodels," *Journal of Monetary Economics* 5(1):23–41.
- Friedman, M. (1968), "The Role of Monetary Policy," *American Economic Review* 58:1–17.
- Gabaix, X. and D. Laibson (2001), "The 6D Bias and the Equity Premium Puzzle," mimeo, MIT and Harvard University.
- Goodfriend, M. and R. G. King (1997), "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual*, pp. 231–83.
- Grubb, D., R. Jackman, and R. Layard (1982), "Causes of the Current Stagflation," *Review of Economic Studies* 49:707–30.
- Hall, R. E. and N. G. Mankiw (1994), "Nominal Income Targeting," in N. G. Mankiw, ed., *Monetary Policy*, Chicago: University of Chicago Press.
- Lucas, R. E., Jr. (1972), "Econometric Testing of the Natural Rate Hypothesis," in O. Eckstein, ed., *The Econometrics of Price Determination*, Washington D.C.: Board of Governors of the Federal Reserve System.
- Mankiw, N. G. (2001), "The Inexorable and Mysterious Tradeoff Between Inflation and Unemployment," *Economic Journal* 111(471):C45–61.
- Mankiw, N. G. and R. Reis (2001), "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," NBER Working Paper No. 8290 (forthcoming in the *Quarterly Journal of Economics*).
- McCallum, B. (1997), "Comment," *NBER Macroeconomics Annual*, pp. 355–59.
- Phelps, E. S. (1967), "Phillips Curves, Expectations of Inflation, and Optimal Unemployment over Time," *Economica* 2(3):22–44.
- (1968), "Money-Wage Dynamics and Labor Market Equilibrium," *Journal of Political Economy* 76(2):678–711.
- (1978), "Disinflation without Recession: Adaptive Guideposts and Monetary Policy," *Weltwirtschaftliches Archiv* 114(4):783–809.
- (1994), *Structural Slumps: The Modern Equilibrium Theory of Unemployment, Interest, and Assets*, Cambridge: Harvard University Press.
- Phelps, E. S. and J. B. Taylor (1977), "Stabilizing Powers of Monetary Policy Under Rational Expectations," *The Journal of Political Economy* 85:163–90.
- Rotemberg, J. (1982), "Monopolistic Price Adjustment and Aggregate Output," *Review of Economic Studies* 44:517–31.
- Sargent, T. J. (1971), "A Note on the Accelerationist Controversy," *Journal of Money, Credit, and Banking* 8(3):721–25.
- (1999), *The Conquest of American Inflation*, Princeton: Princeton University Press.
- Sims, C. (2001), "Implications of Rational Inattention," mimeo, Princeton University.
- Staiger, D., J. H. Stock, and M. W. Watson (2001), "Prices, Wages, and the U.S. NAIRU in the 1990s," NBER Working Paper No. 8320.
- Summers, L. H. (1986), "Estimating the Long-Run Relationship Between Interest Rates and Inflation: A Response," *Journal of Monetary Economics* 18(1):77–86.
- Taylor, J. B. (1980), "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy* 88:1–22.
- Tobin, J. (1980), "Stabilization Policy Ten Years After," *Brookings Papers on Economic Activity* 1:19–72.

A Theory of Rational Inflationary Inertia

GUILLERMO A. CALVO, OYA CELASUN, AND MICHAEL KUMHOF

If there is such a thing as an economy with a rock-solid inflation rate of 40 per cent, plus or minus 2 per cent, per year, institutions would surely adapt, so that prices would be announced in catalogs and wage contracts with smooth growth paths paralleling the smooth aggregate price path. Nominal rigidity would set in about this price path in much the same form as we see around the zero inflation rate in low-inflation economies. (Sims, 1988, p. 77)

1. INTRODUCTION

Monetary theory today is dominated by fully microfounded dynamic general equilibrium models incorporating, in one form or another, the assumption of sticky prices. A comprehensive survey of this literature is contained in Clarida et al. (1999) for closed economies and Lane (2001) for open economies. The renewed popularity of sticky-price monetary economics was motivated by empirical findings that demonstrated, at least for the U.S. case, that monetary policy has significant real effects, contrary to the premise of the real business cycle literature. Examples of that empirical literature include Christiano et al. (1996, 1998) and Leeper et al. (1996). As surveyed in Taylor (1998), the assumption of sticky prices does a good job in explaining most, but not all, features of those data. That paper also documents the micro- and macroeconomic evidence supporting the sticky-price assumption itself.

The 1990s resurgence of this model class would have been hard to predict from a 1970s vantage point. At that time, in the wake of the rational expectations revolution, the sticky-price assumption was closely associated with old-style Keynesian models featuring adaptive expectations, and many economists discarded these two concepts together. But at that same time a group of researchers started to build the theoretical foundations for the incorporation of sticky prices into rational expectations models, and it is to their work that the new generation of models

The authors thank the discussant, Mark Gertler, for helpful comments. They also thank Richard Clarida, Timothy Kehoe, Lars Svensson, John Taylor, and seminar participants at Chicago GSB, Wisconsin, Stanford, the 2001 NBER Summer Institute, the 2001 Latin American Meetings of the Econometric Society, and the 2001 Winter Camp in International Finance in Santiago, Chile.