Advances in Macroeconomics

Volume 3, *Issue* 1 2003 *Article* 1

Where Is the Natural Rate? Rational Policy Mistakes and Persistent Deviations of Inflation from Target

Ricardo Reis*

*Harvard University, rreis@princeton.edu

Where Is the Natural Rate? Rational Policy Mistakes and Persistent Deviations of Inflation from Target

Ricardo Reis

Abstract

Empirical research has shown that there is large uncertainty concerning the value of the natural rate of unemployment at any point in time. I incorporate this feature in a model of monetary policy where the policymaker targets an inflation rate and the natural rate of unemployment and solve for the optimal policy. Two interesting results emerge. First, under a realistic shock profile, the model generates long-lasting deviations of inflation from target, providing an alternative (but also a complement) to the popular Barro-Gordon framework. Second, the economy exhibits large inflation persistence and can have very rich inflation dynamics. The model is able to account for approximately one third of the increase in inflation in the United States in the late 1970s, and suggests an explanation for the low inflation of the late 1990s. Moreover, I present empirical evidence for the United States and other countries that support the model including a new empirical finding: across countries there is a positive statistical relation between the persistence of unemployment and the persistence of inflation.

KEYWORDS: Monetary policy under uncertainty, Natural rate of unemployment, Inflation persistence.

1 Introduction

The study of uncertainty in the context of monetary policy is hardly new. From the initial contributions of Poole (1970) and Brainard (1967) focusing on parameter uncertainty to the more challenging issue of overall model uncertainty, a large number of studies in theoretical and applied monetary policy have studied how different forms of uncertainty affect the choice of optimal policy by the Central Bank. Nevertheless, most models assume the policymaker knows with certainty one key policy variable: the natural rate of unemployment first discussed by Phelps (1968) and Friedman (1968). The natural rate was defined by Phelps as "...the equilibrium unemployment rate - the rate at which the actual and expected price increases (or wage increases) are equal...".[1] Even though most models in the theory of monetary policy treat it as known, being in essence a theoretical construct, its exact value is (and should be treated as) an unknown. Since the introduction of the concept more than thirty years ago, it is impossible to find at any point in time a consensus within economics on what the value of the natural rate was. Staiger, Stock and Watson (1994) confirmed this perception. Examining a series of different methods for estimating the natural rate of unemployment they find that all lead to very imprecise estimates with wide confidence intervals. This imprecision leads the authors to go as far as challenging the use of the concept as a policy tool altogether, in stark contrast with most theory of monetary policy.[2]

Not only is the natural rate uncertain, but it should also be expected to vary considerably over time. Immediately at the introduction of the natural rate concept, Friedman stated: "...by using the term "natural" rate of unemployment, I do not mean to suggest that it is immutable and unchangeable."[3] The empirical studies of Gordon (1997) and Staiger, Stock and Watson (1994) estimate large changes in the natural rate in the postwar period in the United States.

This paper studies the implications for monetary policy and the behavior of inflation of having an uncertain, variable natural rate. I build a simple model of monetary policy in which the natural rate of unemployment is unknown to the policymaker, but she forms optimal forecasts of its value. The optimal policy is derived and the equilibrium path of inflation and unemployment are determined. Some interesting new results emerge.

First, large and persistent deviations of inflation from target arise, as in Barro and Gordon (1983). In the Barro-Gordon model, policymakers wish to lower unemployment from its natural level and so are tempted to surprise agents by inflating the economy. As rational agents foresee this temptation they raise inflation expectations to a level at which the marginal gain from surprise inflation (and bringing unemployment closer to the target rate) is just offset by the marginal loss in terms of extra inflation. In equilibrium, unemployment is stuck at the natural rate and inflation is above target. This "inflation bias" provides an explanation for the high inflation of the 1970s and has led to a large amount of research on ways to eliminate or at least partially offset the problem.[4] Yet, central bankers have always rejected the notion that they are attempting to trick agents by targeting an unemployment rate below the natural rate. As Blinder (1998) put it, in response to such a suggestion most

[1]Phelps (1968), page 682.

[2]More recently, Stock and Watson (1999b) reinforce this result with U.S. data.

[3]Friedman (1968), page 9.

[4]See Rogoff (1989) for a survey.

central bankers would say "Of course that would be inflationary. That's why we don't do it."[5]

The model presented in this paper is able to accommodate this criticism, while still generating large and persistent deviations of inflation from target. Unemployment will move either because of changes in the natural rate or due to short-run supply shocks. The Central Bank wishes to counteract only the latter, so it will target an optimally formed forecast of the natural rate. Persistent deviations of inflation from target can then occur in response to a shock to the natural rate. If the natural rate unexpectedly rises, the policymaker, induced by her previous observations, will underestimate its value. She will then optimally choose to increase inflation in order for unemployment to approach its biased-down forecast of the natural rate, and thus inflation above target will result. As the forecast error is persistent and only disappears asymptotically, inflation will be set above target for a long period of time. Yet, note that this inflation above target does not arise from a desire to trick economic agents, but simply results from the policymakers' incorrect forecast of the natural rate. Even though the policymaker believes she is behaving in an appropriate way by targeting its forecast of the natural rate, she is in fact generating the observed high inflation.

The paper is organized as follows. Section 2 outlines the basic model. I derive the optimal forecast and examine the impact of shocks to the natural rate. I then analyze the impulse response functions to different shocks to develop some intuition on the key features of the model and contrast its general predictions with the historical behavior of inflation in the United States. Using some reasonable parameter values, I show that the model accounts for between one fifth and one third of the increase in inflation in the 1970s. Next, I examine alternative explanations for the movements in U.S. inflation and discuss ways in which these can be distinguished from the explanation given in this paper. Section 3 extends the model by enriching the stochastic structure of shocks and derives new predictions. I compare these new predictions with the evidence from the time-series properties of unemployment and inflation for a sample of countries. A calibration of the model is able to replicate the high serial correlation of inflation observed in the data. Section 4 relates this work to other papers in the literature. Section 5 concludes.

2 A model of monetary policy with an uncertain natural rate of unemployment

The model has two key components: a Phillips curve and an objective function for policy. Below, I discuss two reduced-form relations for these components. Appendix A provides one specific micro-foundation in the form of a fully articulated general equilibrium model with nominal rigidities which generates precisely these reduced-form relations.

[5] Blinder (1998), page 43.

2.1 The short-run behavior of employment

The short-run behavior of the economy is described by a familiar expectations-augmented downward-sloping Phillips curve:

$$u_t = u_t^N - \alpha(\pi_t - \pi_t^e) + \varepsilon_t. \tag{1}$$

Equation (1) states that unemployment at time t (u_t) deviates from its natural rate (u_t^N) if the inflation rate (π_t) is different from its expected rate (π_t^e) or an unexpected supply shock occurs (ε_t). Nominal shocks, in the form of deviations of inflation from their expected value ($\pi_t - \pi_t^e$), affect unemployment either via changes in the supply decisions of imperfectly informed agents who confuse relative and absolute changes in prices as in Lucas (1973), or due to predetermined nominal wages as in Fischer (1977). The parameter $\alpha > 0$ is the inverse of the slope of the Phillips curve. Finally, ε_t is a short-run supply shock that induces deviations of unemployment from its equilibrium natural rate. In the model presented in Appendix A, they are identified as shocks to the markup of prices over marginal costs. More generally, they represent short-run disturbances that the policymaker wants to offset.

The policymaker is assumed to observe unemployment realizations at time t without any lag. Still, both the natural rate of unemployment and the supply shocks are unobservable. The natural rate at t is not known at t or at any period after. Even today, nowhere in the official statistics is there an exact measure of the natural rate 30 years ago. Accordingly, the short-run supply shocks ε_t are also not observed at any point, since otherwise it would be possible to identify the natural rate exactly. Consistent with the underlying assumption of rational expectations by agents, the unemployment rate cannot deviate systematically from its long-run equilibrium natural level, so $E(\varepsilon_t) = 0$. Finally, I assume the ε_t are identically, independently distributed (i.i.d.) draws from a distribution with finite, constant, known variance σ_ε^2.

An alternative assumption on short-run supply shocks is that they contain an observable component to the policymaker and an unobservable component (e.g. $\varepsilon_t = \varepsilon_{k,t} + \varepsilon_{u,t}$, with $\varepsilon_{k,t}$ known but $\varepsilon_{u,t}$ unknown). The policymaker could then have an explicit stabilization role if it has an information advantage over agents in offsetting the $\varepsilon_{k,t}$ shocks. This would not alter the conclusions as long as there is some component of short-run shocks that is not observable and thus prevents identification of the natural rate. By the same argument, I could assume that the Central Bank is unaware of the exact level of the contemporaneous unemployment rate when it sets monetary policy. The model easily accommodates this extension by reinterpreting the ε_t to now also include estimation errors of the unemployment rate.

2.2 The natural rate

Two broad facts are generally taken from the literature regarding the natural rate. First, that it is unknown and imprecisely estimated. Second, that it exhibits a large degree of persistence. Theoretically, this fits the descriptions by Friedman and Phelps and empirically the results of Staiger, Stock and Watson.

For now, I assume the natural rate evolves according to the following AR(1) process:[6]

$$u_t^N = \theta u_{t-1}^N + v_t. \tag{2}$$

The first term in the right-hand side captures the persistence in the natural rate and $\theta \in [0, 1]$ should be high enough to capture this persistence.[7] The second term captures shocks to the natural rate, with v_t being i.i.d. random draws, unobservable by the policymaker, with $E(v_t) = 0$ and $Var(v_t) = \sigma_v^2$. In terms of the model in Appendix A, the v_t represent shocks to marginal costs either from movements in productivity or in the disutility of supplying labor. More generally, they stand for any disturbance in the economy that moves the unemployment rate and which the policymaker does not want to offset.

This specification can be compared with the ones used in the empirical literature on the natural rate. Gordon (1997) assumed that the natural rate would follow a random walk, a hypothesis nested within our model by setting $\theta = 1$. Staiger, Stock and Watson (1994) model the natural rate in many different ways, including: a constant, a constant with two changes in time, a cubic spline, and as coming out of a Phillips curve relation with time-varying coefficients. The behavior of all of these can be reasonably well approximated by an AR(1) process.

2.3 The policymakers' objective function

The policymaker minimizes the present discounted value of expected losses:

$$V_t = \sum_{i=0}^{\infty} \psi^i E_t L_{t+i}, \tag{3}$$

where $\psi \in (0, 1)$ is a discount factor and L_t is the period loss function which is quadratic in deviations of inflation from target and unemployment from the natural rate:

$$L_t = \frac{1}{2}(\pi_t - \pi^*)^2 + \frac{\lambda}{2}(u_t - u_t^N)^2. \tag{4}$$

Appendix A derives this loss function as a second-order approximation to the utility of a representative agent in an economy with nominal rigidities. Intuitively, variability in inflation is costly since it leads to pricing errors by firms and thus to an inefficient allocation of resources. Variations in unemployment are costly since they translate into variability in consumption which is undesirable by risk-averse consumers. The parameter λ measures the relative weight given to employment stabilization relative to inflation stabilization, and u_t^N is the natural rate. The inflation target π^* is exogenously given to the Central Bank, and the Central Bank perfectly sets inflation.[8]

---

[6] Section 3 relaxes this assumption by allowing for any stationary process for the shocks.

[7] As defined, the natural rate converges to zero, but this is simply a normalization. Letting $u_t^N = \theta u_{t-1}^N + (1-\theta)\bar{u} + \nu_t$, where $\bar{u}$ is the value to which the NAIRU converges, leads to the same conclusions.

[8] This assumption could be relaxed by instead assuming that the Central Bank chooses a desired inflation level $\hat{\pi}_t$, while actual realized inflation differs from this by an error term $\varepsilon_{\pi,t}$, i.i.d. with zero mean and constant variance, so $\pi_t = \hat{\pi}_t + \varepsilon_{\pi,t}$. This error could be interpreted as imperfect control over the inflation rate by the policymaker or more generally as capturing shocks to demand. In this model, it would have a similar impact as the short-run shocks ε_t.

Taking expectations of the loss function in equation (4) conditional on the information set of the Central Bank (which includes π_t, u_t and u_t^T), it follows that the Central Bank wishes to minimize:

$$\frac{1}{2}(\pi_t - \pi^*)^2 + \frac{\lambda}{2}(u_t - u_t^T)^2, \tag{5}$$

where $u_t^T \equiv E_t u_t^N$. I now turn to the construction of this forecast of the natural rate of unemployment.

2.4 Forecasts of the natural rate

The policymaker's problem of forecasting the natural rate is an optimal signal extraction problem. Since the policymaker contemporaneously observes unemployment as well as inflation expectations[9] and sets inflation π_t, it can build the observation variable:

$$x_t = u_t + \alpha(\pi_t - \pi_t^e). \tag{6}$$

The economy is then described by the system of equations:

$$x_t = u_t^N + \varepsilon_t, \tag{7}$$
$$u_t^N = \theta u_{t-1}^N + v_t. \tag{8}$$

The optimal forecast[10] is a geometrically weighted average of present and past observations on x_t:

$$u_t^T = \frac{\theta - \beta}{\theta} \sum_{i=0}^{\infty} \beta^i x_{t-i}, \tag{9}$$
$$\beta = \frac{\frac{\sigma_v^2}{\sigma_\varepsilon^2} + 1 + \theta^2 - \sqrt{(\frac{\sigma_v^2}{\sigma_\varepsilon^2} + 1 + \theta^2)^2 - 4\theta^2}}{2\theta} \leq \theta. \tag{10}$$

This result is derived in Part A of Appendix B. There it is also shown that the forecast can be expressed in terms of the updating formula :

$$u_t^T = \beta u_{t-1}^T + \frac{(\theta - \beta)}{\theta} x_t. \tag{11}$$

[9] Treating inflation expectations as observed amounts to the definition of Nash equilibrium in the game between economic agents and the Central Bank where both have rational expectations. In equilibrium, the Central Bank will treat inflation expectations as constant and equal to a value which then turns out to be the equilibrium strategy by economic agents. Alternatively, as we will see later, inflation expectations in equilibrium will equal the inflation target at all points in time so they are easy to identify by the Central Bank.

[10] By optimal, I mean the estimator that minimizes mean squared error loss within the class of linear estimators, using the methods outlined in Whittle (1983). Adding the assumption of joint normality of the shocks would make these forecasts optimal in the class of all (including non-linear) estimators.

If the natural rate follows a random walk ($\theta = 1$), this collapses to the well-known Muth (1960) result:

$$u_t^T = (1-\beta)\sum_{i=0}^{\infty} \beta^i x_{t-i}. \tag{12}$$

Note that the smaller is $\sigma_\nu^2/\sigma_\varepsilon^2$, the smaller is the weight given to the most recent observation of unemployment because β approaches θ. If most of the variation in unemployment is driven by short-run supply shocks, recent observations will likely be dominated by these shocks and thus have little influence on the natural rate forecasts. Consequently, given a shock to the natural rate, there will be little updating and non-negligible forecast errors arise.

2.5 Solving for equilibrium

The optimal policy is to choose inflation every period to minimize the loss function in equation (4), subject to the constraint posed by the economy's short-run behavior in equation (1). Note that even though the objective function involves lagged unemployment (in the forecast of the natural rate), this is not a dynamic programming problem, but rather a succession of one-shot problems. While the choice of the inflation rate has an effect on the unemployment rate, it does not affect the forecasts of the natural rate of unemployment next period. Mathematically, this is captured by the fact that the optimal forecasts use observations of x_t and not u_t.

The first-order condition of the minimization problem determines the optimal rule for setting inflation:

$$\pi_t = \pi^* + \alpha\lambda(u_t - u_t^T). \tag{13}$$

Given a deviation of unemployment from target, the policymaker faces a dilemma. The optimal reaction to a short-run supply disturbance (ε_t) is to fully accommodate it, whereas a change in the natural rate (u_t^N) should have inflation unchanged. Being unable to distinguish between the two, the best the policymaker can do is to establish a forecast of the natural rate and react only to deviations from this target level. If unemployment is above target, inflation will be increased in order to try to lower unemployment back to its target value. If it is below, inflation will be set below target.

Private agents in the economy are unaware of the value of the natural rate of unemployment and form expectations rationally.[11] Their forecast of the natural rate is also u_t^T, because this is announced by the Central Bank or because they solve the same signal extraction problem that the Central Bank solved. Replacing u_t using equation (1) in equation (13) and taking expectations shows that inflation expectations always equal the announced target:

$$E\pi_t = \pi_t^e = \pi^*. \tag{14}$$

[11] I could alternatively assume that private agents know the value of the natural rate of unemployment, but this would lead to very similar qualitative conclusions. Moreover, it is difficult to understand why in such a world, the Central Bank would not be aware of a variable that all price-setters in the economy know.

The equilibrium outcome is then found by combining this best response for private agents with the best response for the Central Bank in (13) to find the equilibrium outcome level of inflation:

$$\pi_t = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda}\varepsilon_t + \frac{\alpha\lambda}{1+\alpha^2\lambda}(u_t^N - u_t^T). \tag{15}$$

Consider what this expression implies for inflation in response to different shocks. If there is a short-term supply shock ($\varepsilon_t > 0$), then, with perfect information on the natural rate, the optimal reaction would be to increase inflation by $\frac{\alpha\lambda}{1+\alpha^2\lambda}\varepsilon_t$, partially offsetting the shock to lower the variability of employment. The shock is not fully offset since there is also a loss in setting inflation away from target. With incomplete information, though, the higher observation of unemployment caused by the shock raises the estimate of the natural rate, so the third term in the right hand side of equation (15) becomes negative. Thus the policymaker will increase inflation by less than with perfect information. Similarly, given an increase in the natural rate ($\nu_t > 0$), the optimal reaction with complete information would be to leave inflation on target. Observing the increase in unemployment, the policymaker will give some weight to the possibility that this is caused by a short-run supply shock. She will therefore underestimate the natural rate, so the third term in equation (15) will be positive, and inflation will be set above target.

Generally, deviations of the natural rate from target ($u_t^T \neq u_t^N$) will lead to deviations of inflation from target. For instance, by pursuing a target that underestimates the natural rate, the policymaker will set higher inflation to lower unemployment below the natural rate, towards the target. Since the process of updating forecasts has a geometric form, this will lead to long-lived deviations of inflation from target, which can easily be confused for a permanent inflation bias.

Yet, this is not a "bias" in the Barro-Gordon sense. First, since the estimator of the natural rate is consistent, the forecast errors disappear asymptotically. Second, unlike in Barro-Gordon, the deviations of inflation from target may be negative. As long as some unexpectedly favorable shock hits the natural rate, the forecast error will be positive, and the precise same mechanism works in reverse, inducing inflation below target. Third, this is not a constant bias, but rather it changes each period with new realizations of the short-run supply shock and the long-run natural rate. Fourth, note that these deviations do not arise out of deliberate attempts by the policymaker to trick agents. The policymaker will believe she is always targeting the natural rate as asserted by the Blinder quote in the introduction. Large and persistent deviations of inflation from target arise simply as the result of imperfect information regarding a key policy variable: the natural rate of unemployment. Finally, note that it is straightforward to incorporate a constant inflation bias along the lines of Barro and Gordon in this framework. This model is therefore best seen as a complement to the Barro and Gordon model.

2.6 Comparative statics

A few experiments will illustrate the properties of this model of inflation dynamics. The proofs of all the results are in Part B of Appendix B.

First, consider the impact of a shock to the natural rate at date 0. The following result establishes the path of inflation and unemployment, shown in Figure 1.

Figure 1: Impulse responses of inflation and unemployment to a natural rate shock

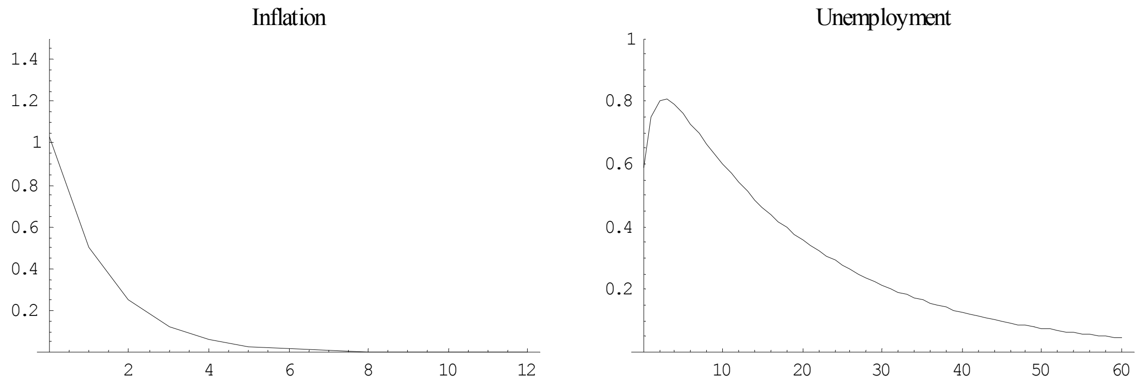

Result 1: *In response to a one-time unit-shock to the natural rate at date 0 (ν_0), inflation and unemployment are given by:*

$$\pi_t = \pi^* + \frac{1}{\theta}\frac{\alpha\lambda}{1+\alpha^2\lambda}\beta^{t+1}, \tag{16}$$

$$u_t = \theta^t - \frac{1}{\theta}\frac{\alpha^2\lambda}{1+\alpha^2\lambda}\beta^{t+1}. \tag{17}$$

We can see in Figure 1 that a long-lived deviation of inflation from target results.[12] Intuitively, given the initial shock, the policymaker observes higher unemployment and is unable to distinguish if this is due to a shock to the natural rate or simply to a short-run supply shock. A rational policymaker gives some weight to both hypotheses, and thus partially offsets the observed increase in unemployment by raising inflation. In the following periods, given the persistence in the shock to the natural rate, unemployment is still higher than the updated estimate of the natural rate, and the policymaker again finds it optimal to increase inflation, although this increase is by less than before because some updating has already occurred.

As for unemployment, two forces are in operation. On the one hand, the natural rate is higher pulling unemployment up. However, on the other hand, surprise inflation is being generated, lowering unemployment. It turns out that the first effect always outweighs the second. The intuition comes from two previous results. First, since there is the possibility that the shock is to short-run unemployment, the updating of the target rate occurs at a slower rate than the evolution of the natural rate itself (i.e., $\theta > \beta$). Moreover, given that the policymaker gives some weight to both unemployment and inflation stabilization (i.e.,

[12] All the Figures are drawn for the parameter configuration $\sigma_\nu^2/\sigma_\varepsilon^2 = 0.5$ and $\theta = 0.95$. The results are not qualitatively sensitive to variations around these values, and I will provide justifications for these choices in Section 3.2. The two other parameters set were $\frac{\alpha\lambda}{1+\alpha^2\lambda} = 2$ and $\frac{\alpha^2\lambda}{1+\alpha^2\lambda} = 0.8$. It is difficult to assess wether these are reasonable; however, since they work solely as scale parameters, they do not affect the overall shape of the figures.

$0 < \lambda < \infty$), the countercyclical policy will aim at not offsetting the shock fully. Thus, the surprise inflation generated will never be so strong as to lead to a fall in unemployment following a rise in the natural rate.

Result 2: *After a shock to the natural rate (v_t), inflation converges back to target and unemployment to the steady state natural rate as long as $\theta < 1$. In the transition, unemployment is always positive and will be increasing for an initial period of time if α and λ are large enough.*

A continued increase in the unemployment rate for many periods following the shock is more likely the flatter the Phillips curve is (larger α) and the larger the weight given to employment stabilization (larger λ). The intuition here comes again from realizing that the two forces described above are in action. As time passes, the natural rate falls progressively, converging back to zero and pushing unemployment down. Yet, the forecast errors of the natural rate are also getting smaller and so is the surprise inflation generated. This pushes unemployment up from the last period. The larger is λ, the larger is the reaction of the policymaker to the observed change in unemployment. Thus, the larger is the inflation generated. The larger is α, the larger is the impact of this surprise inflation on unemployment. Thus, the larger is $\alpha\lambda$, the stronger is the second effect. Thus, in the initial periods after the shock, unemployment may actually be rising before it starts falling back to its steady state.

An interesting special case is when the natural rate follows a random walk ($\theta = 1$):

Result 3: *When the natural rate follows a random walk, inflation initially rises before falling back to target, but unemployment approaches its new steady state from below, rising throughout the entire transition.*

Figure 2 shows the paths of inflation and unemployment. Whereas the impulse response of inflation is similar to before, the impulse response of unemployment now converges to a higher long-run level. Actual unemployment will be increasing throughout the convergence to the new steady state. Of the two forces identified in the previous paragraph only the second remains, since the natural rate no longer falls after the shock.

The response of inflation and unemployment to a short-run supply shock at date 0 (ε_0) is described by the following two results:

Result 4: *In response to a short-run supply unit shock at date 0 (ε_0), inflation and unemployment at date 0 are given by:*

$$\pi_0 = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda}\frac{\beta}{\theta}, \tag{18}$$

$$u_0 = 1 - \frac{\beta}{\theta}\frac{\alpha^2\lambda}{1+\alpha^2\lambda}. \tag{19}$$

At $t \geq 1$:

$$\pi_t = \pi^* - \frac{\theta-\beta}{\theta}\frac{\alpha\lambda}{1+\alpha^2\lambda}\beta^t, \tag{20}$$

$$u_t = \frac{\theta-\beta}{\theta}\frac{\alpha^2\lambda}{1+\alpha^2\lambda}\beta^t. \tag{21}$$

Figure 2: Impulse responses of inflation and unemployment to a natural rate shock if the natural rate follows a random walk

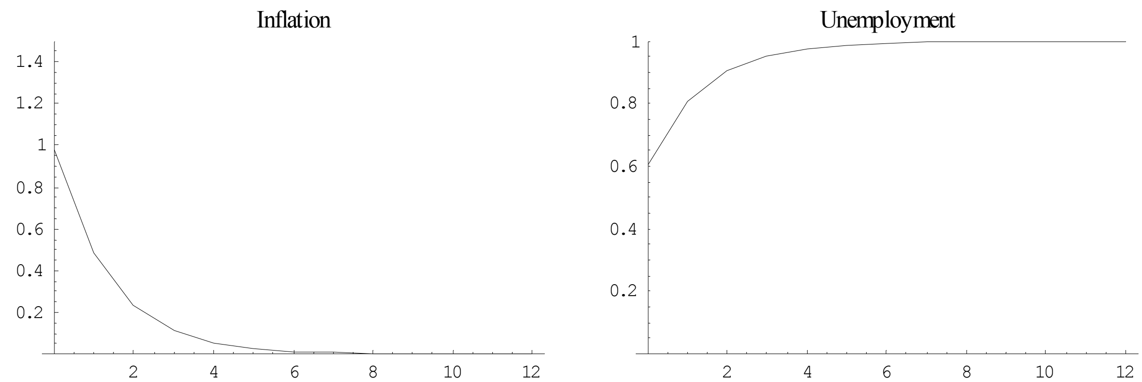

Figure 3: Impulse responses of inflation and unemployment to a short-run supply shock

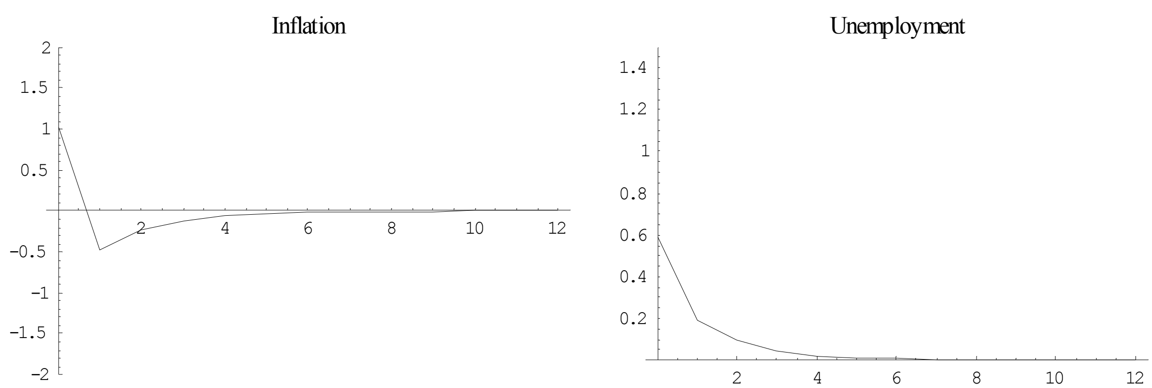

Result 5: *Both inflation deviations from target and unemployment deviations from the natural rate will be "small" for $t \geq 1$ as long as shocks to the natural rate are less volatile than shocks to short-run unemployment. Both disappear asymptotically.*

Figure 3 shows the impulse response functions for inflation and unemployment. Following the shock, inflation is above target as the policymaker observes the increase in unemployment and responds by raising inflation. Moreover, she adjusts the forecast of the natural rate upwards given this observation of higher unemployment. From $t = 1$ onwards, the target for the natural rate will be too high, and this positive forecast error will lead to inflation being set below target. Note nevertheless that this effect for $t \geq 1$ will in general be small as long as long-run shocks are less variable than short-run shocks. If $\sigma_\nu^2/\sigma_\varepsilon^2$ is small, then β is close to θ, so recent individual observations receive a small weight in the optimal forecast.

As for unemployment, initially it rises above the natural rate since the countercyclical policy does not fully offset the shock. This occurs both because (a) it would not be optimal to do so since there is some weight in the objective function on stabilizing inflation and (b) the forecast of the natural rate is updated upwards, reducing the gap the policymaker aims at closing. After the first period, the slight disinflation will lead to unemployment above target, but again this effect should be small.

To sum up, in response to an unfavorable shock to the natural rate of unemployment, inflation will exhibit a prolonged positive deviation from target and unemployment will be higher for a number of periods, possibly increasing during the initial periods. Short-run supply shocks will lead to significant movements in unemployment and inflation only at the time they occur.

The figures above were drawn for the case when $\sigma_\nu^2/\sigma_\varepsilon^2 = 0.5$ and $\theta = 0.95$. Figure 4 plots the response of inflation to a natural rate shock for different parameter combinations. While the main qualitative conclusions remain unchanged, quantitatively we can see that the smaller is $\sigma_\nu^2/\sigma_\varepsilon^2$, the larger and more persistent are the effects of the shock on inflation. Intuitively, if the natural rate shocks are less preponderant, then the forecasts of the natural rate will respond less to variations in unemployment and the forecast errors that generate the movement in inflation will be larger and more persistent. A smaller θ on the other hand will imply that the errors in forecasting the natural rate are less persistent and the impulse response of inflation is more short-lived.

2.7 Predictions and some evidence

The key prediction of the model was highlighted in the previous section. Given an unfavorable shock to the natural rate, the optimally forecasted natural rate will lag behind the actual natural rate and policy will push for too low unemployment by inflating the economy consistently above the target level. Symmetrically, if a favorable long-run employment shock occurs, inflation will be below target for a sustained period of time.

Most economists would agree that the 1970s are a perfect example of decade afflicted by a series of unfavorable shocks. The end of Bretton Woods, the two oil shocks, the demographic changes in the labor force composition described by Perry (1970), and the productivity slowdown are some examples of these shocks, many of which can be interpreted

Figure 4: Reaction of inflation to a natural rate shock for different parameter configurations

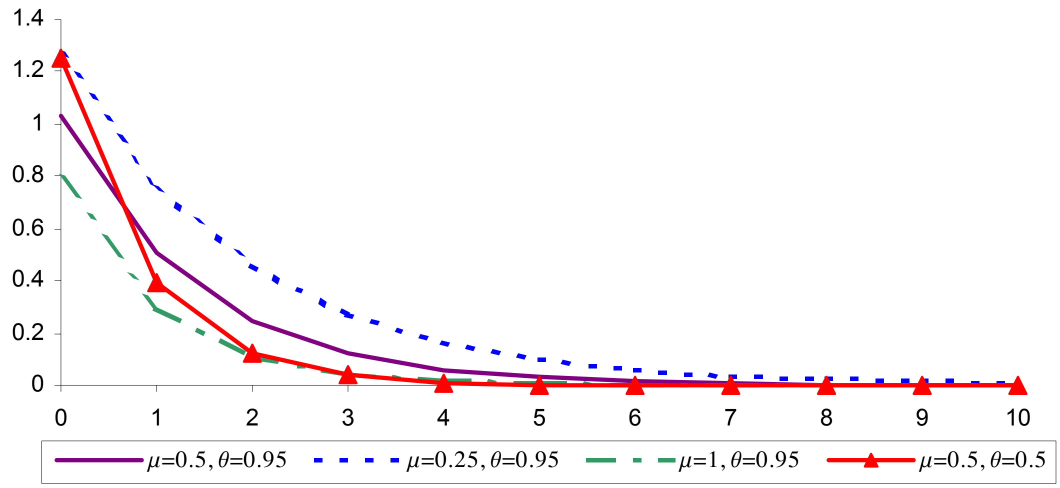

as shocks to the natural rate. Milton Friedman's Nobel lecture (1977) puts forward a series of reasons why one should expect the natural rate to have risen in the United States during this period. The point estimates from Gordon (1997) or Staiger, Stock, Watson (1994) show the natural rate of unemployment increased substantially in the 1970s. The model would then predict high long-lasting inflation, which is precisely what happened (see Figure 5 for annual U.S. inflation). Moreover, the Central Bank at the time justified its policy by stating that it gave some weight to the possibility that these were simply short-run disturbances to be offset by active policy, and its estimates of the natural rate seem too low from today's perspective (Orphanides, 2003). This is exactly what the model predicts. The observed inflation bias of the 1970s can therefore be explained without putting the blame on foolish policymakers deliberately trying to push for an infeasible level of unemployment. This decade was simply one of much uncertainty, with many supply shocks of unknown nature and effect, which ended up leading to a large forecast error of the natural rate, creating persistently high inflation.

Using equation (15), we can get an idea of how much of the difference in inflation between the 1960s and the 1970s can be explained by the inability to forecast the natural rate. The first step in this exercise is to obtain estimates of the Central Bank's forecast errors of the natural rate: $u_t^N - u_t^T$. I obtain u_t^T from reading through the Economic Reports of the President published by the Council of Economic Advisors (CEA). Following the Report of 1962, the CEA started reporting the rate of unemployment consistent with "full employment" and referring to the difference between this and the actual unemployment rate as a gap that policy should aim to eliminate. In all reports until 1981, I was able to find explicit references to this rate; Orphanides (2003) argues that the Federal Reserve Board used these gap estimates to set monetary policy during this period. As for u_t^N, I estimate it using a Hodrick-Prescott filter with an adjustment coefficient of 100 on annual unemployment data from 1954 to 1999. While it is a premise of this paper that we can never obtain an exact

Figure 5: Annual CPI inflation in the United States

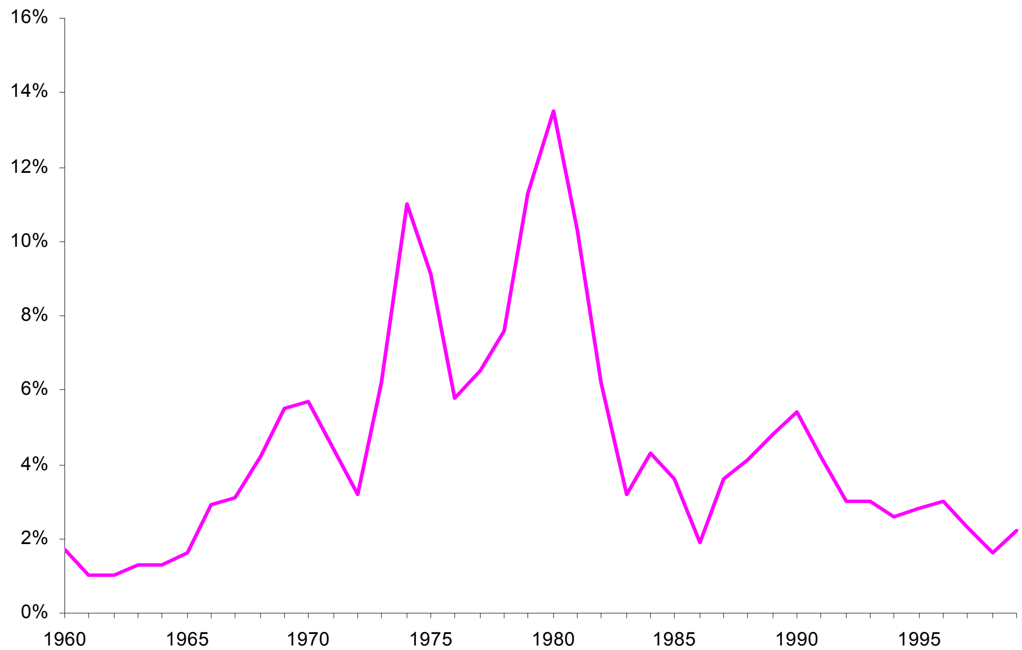

measure of u_t^N, the hope is that using this large sample of data we can obtain some reasonable estimates of the value of u_t^N for the 1962-1981 period. Figure 6 displays these measures.

Implementing equation (15) also requires knowledge of the factor $\alpha\lambda/(1+\alpha^2\lambda)$, and for this I use three different estimates. The first comes from Broadbent and Barro (1997), who estimate an extended Barro-Gordon model on U.S. data and obtain the estimates $\alpha = 0.23$ and $\lambda = 6.30$, which implies that $\alpha\lambda/(1+\alpha^2\lambda) = 1.09$. A second set of estimates comes from using data from 1954-99 on annual inflation and unemployment and using last year's inflation as a measure of expected inflation to obtain $\alpha = 0.64$. This estimate of the slope of the Phillips curve implies that reducing inflation by 1% raises unemployment by 1.56%, which is consistent with conventional estimates of the sacrifice ratio. Then, as will be described in more detail in Section 3.2, matching the model's predicted serial correlation of unemployment with the one observed in the data, I estimate that $\alpha^2\lambda/(1+\alpha^2\lambda) = 0.982$. These two estimates imply that $\alpha\lambda/(1+\alpha^2\lambda) = 1.54$. Finally, a third set of estimates comes from using the microfoundations in Appendix A to realize that $\lambda = \alpha\eta$, where η is the elasticity of product demand. Chari, Kehoe and McGrattan (2000) use information on price-cost margins to calibrate this parameter at $\eta = 10$. I will use this value, which together with my earlier estimate of the slope of the Phillips curve implies that $\alpha\lambda/(1+\alpha^2\lambda) = 1.13$.

I then compare the average inflation in the decades 1961-72 and 1972-81. Taking the average over 10 years approximately eliminates the influence of short-run supply shocks ε_t. Using the estimates described in the previous paragraph, I can compute by how much the model in this paper predicts that inflation should have changed between these two periods. Table 1 performs this exercise. The theory described in this paper can account for between

Figure 6: Unemployment rate and forecast errors of the natural rate 1962-81

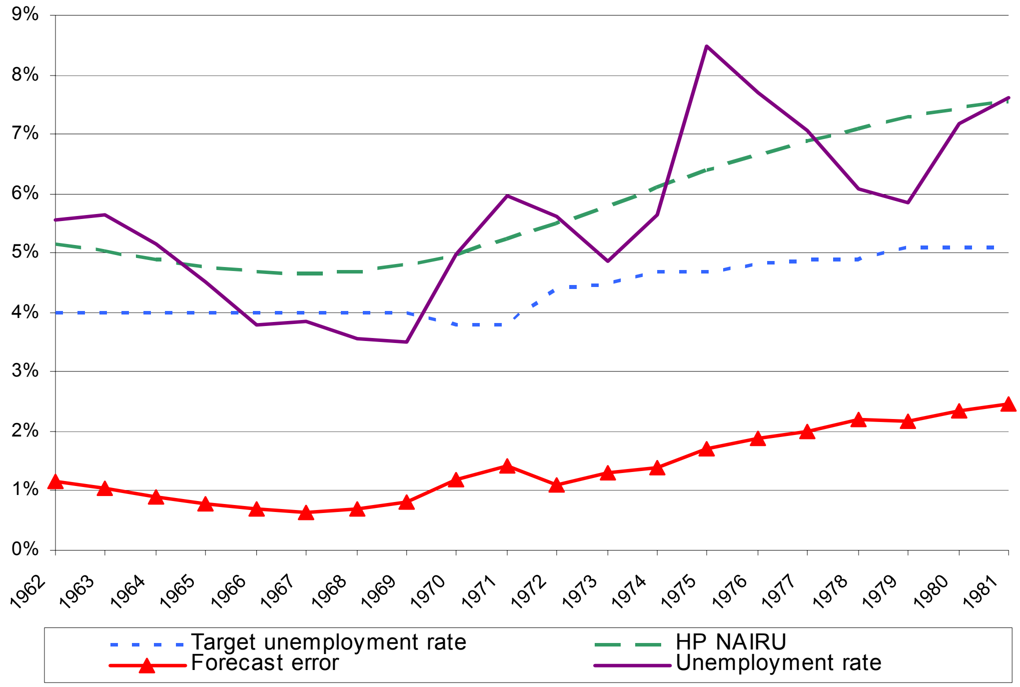

1/5 and 1/3 of the increase in inflation from the 1960s to the 1970s. While forecast errors of the natural rate may not be the full story behind the inflation of the 1970s, these calculations indicate that they may carry a substantial share of the blame.

Table 1

$\frac{\alpha\lambda}{1+\alpha^2\lambda}$	Predicted change in inflation	Actual change in inflation	Share of increase explained
Broadbent-Barro			
1.09	1.01%	4.54%	22.13%
Estimate from section 3.2			
1.54	1.43%		31.41%
Micro-foundations			
1.13	1.05%		23.05%

Moreover, the model gives some insights concerning the recent behavior of U.S. inflation. If the late 1990s have indeed been an age of technological revolution, as the evidence seems to suggest (Jorgenson and Stiroh, 2000), this would probably lower the natural rate of unemployment. Ball and Mankiw (2002) take this view and offer alternative stories to account for the declining natural rate in the 1990s. The model predicts that the Central Bank overestimates the natural rate and thus pushes for disinflationary policies even below target. Throughout the late 1990s, the Federal Reserve repeatedly stated that it believed the economy was "overheated," which should probably be interpreted as seeing the unemployment rate below its forecast of the natural rate: $u_t < u_t^T$.[13] Perhaps it was right in its forecast, $u_t^T = u_t^N$. But perhaps it was wrong, overpredicting the natural rate ($u_t^T > u_t^N$) after a favorable shock in the late 1990s and conducting a too restrictive monetary policy leading to too low inflation or a negative inflation bias.

2.8 Alternative explanations

Other theories have been put forward to explain the long-term movements in U.S. inflation.

Barro and Gordon (1983) assumed that there is perfect knowledge of the natural rate, but the Central Bank's loss function penalizes deviations of unemployment from a level below the natural rate by an amount k. Equilibrium inflation then depends on k, and the different inflation rates observed in the 1960s, 1970s and 1980s could be explained by changes in this desire by the policymaker to target a low unemployment rate. Yet, as De Long (1997) argues, it is difficult to identify in the history of the Federal Reserve any significant institutional changes that would affect this inflation bias and explain the decade-to-decade variations in the rate of inflation.

A slightly different set-up of the Barro-Gordon model offers more promise. If the loss function of the Central Bank involves targeting a constant unemployment rate rather than a constant difference from the natural rate, so the loss function has a term $(u_t - \bar{u})^2$ rather than a term $(u_t - u_t^N + k)^2$, then equilibrium inflation in the Barro-Gordon model depends on the value of the natural rate. In such a model, the high inflation of the 1970s could then also be the result of an unfavorable shock to the natural rate as in the model in this paper. Yet, this would be due to a stronger inflation bias by the policymaker since the discrepancy between the natural rate and the Central Bank's target rate is now higher, rather than due to the forecast errors discussed in this paper. These two stories are empirically distinguishable: the Barro-Gordon story has inflation depending on u_t^N, whereas in this paper inflation depends on $u_t^N - u_t^T$. In the 1962-1981 period displayed in Figure 6 though, these two variables move very closely together (with a correlation coefficient of 0.98), so it is almost impossible to distinguish the two theories. On a larger sample, tests along the lines of those in the classic work by Barro (1977) could be used to distinguish whether only forecast errors of the natural rate or any changes in the natural rate affect the inflation rate. As in that old literature on unanticipated money, obtaining estimates of u_t^T is quite a challenge, and this is even harder in this case since we do not have good measures of u_t^N.

An alternative explanation for the high inflation of the late 1970s argues instead that during this period policymakers believed that there was a long-run trade-off between inflation and unemployment and thus chose higher inflation while trying in vain to reach lower unemployment (Sargent, 1999). Yet, this story has a timing problem. Romer and Romer (2002) use contemporary discussions of the economy by policymakers and find that by the early 1970s, policymakers had already changed their beliefs into accepting the Friedman-Phelps natural rate hypothesis, while the largest increase in inflation occurs in the late 1970s.

[13] One of many possible examples is Governor Meyer's (1998) speech, in which he stated: "...most estimates would put the actual unemployment rate at the end of the year [1997] perhaps 3/4 percentage points below the NAIRU."

The explanation of the behavior of inflation in this paper contained two hypothesis: first, that inflation was driven by a gap between the natural rate and the policymaker's forecast, $u_t^N - u_t^T$, and second that this forecast was formed rationally. The evidence in the previous section favors the first of these hypothesis but does not test the second. Romer and Romer's (2002) study of policy in the 1970s supports the first hypothesis as well, but they argue that the forecasts of the natural rate 1970s were irrationally optimistic, mostly as a result of using incorrect models of the economy. Empirically, to assess the rationality of forecasts of the natural rate one would need a statistical model that took into account the data limitations in real time. In practice, it is hard to powerfully test the hypothesis of rationality. The persistence of the forecast errors displayed in Figure 6 suggests that the Romer and Romer explanation may have some truth to it.

Finally, note that the previous section suggested that the mechanism identified in this paper can reasonably account for only up to 1/3 of the increase in inflation in the 1970s. It is perfectly acceptable that the alternatives discussed in this section may pick up the remaining 2/3. Only a more detailed quantitative exploration that leaves room for all of these hypotheses could say for sure, but for now I leave this for future research to determine.

3 Extending the model

In this section, I allow for a richer stochastic structure of the shocks. In the previous section, I assumed that the short-run supply shocks (ε_t) were independently identically distributed and the natural rate followed an AR(1) process. In this section, I start by allowing for serial correlation in the short-run supply shocks and derive two new results, which are then contrasted with some data. Finally, I analyze the general case in which the shocks are arbitrary ARMA processes.

3.1 A model with first-order autoregressive short-run supply shocks

The natural rate still follows the same AR(1) process in equation (2). The short-run Phillips curve is now given by:

$$u_t = u_t^N - \alpha.(\pi_t - \pi_t^e) + s_t, \tag{22}$$

similar to before. Yet now the short-run supply shocks follow an AR(1) process:

$$s_t = \rho s_{t-1} + \varepsilon_t, \tag{23}$$

with ε_t i.i.d. $(0, \sigma_\varepsilon^2)$.

In many aspects, a specification that allows for serial correlation in short-run supply shocks is more natural. Within the context of the model above, short-run supply shocks refer to disturbances that push the unemployment rate away from its "natural rate" and which the Central Bank wishes to neutralize. Examples are changes in sales taxes or in the degree of competition through entry of new firms and periodic collusion. Also, as explained in Section 2, measurement errors of contemporaneous unemployment and inflation would enter the model as s_t shocks. Even if all these shocks are short-lived, it is still reasonable to suppose that they exhibit some degree of persistence. Moreover, as an AR(1) specification for the shocks nests the i.i.d. case, it generalizes the analysis, which is desirable insofar as there is disagreement as to what the exact properties of the shocks are.

The new optimal forecast of the natural rate can be expressed as the updating rule:

$$u_t^T = \hat{\beta} u_{t-1}^T + A(x_t - \rho x_{t-1}), \tag{24}$$

where $\hat{\beta}$ and A are known functions in the parameters of the model, derived in Part A of Appendix B. Given this forecast, the problem facing the policymaker is the same as before, so the first-order condition is still given by equation (13) and inflation by equation (15).

Two new results now emerge. First, inflation is positively serially correlated. Second, unemployment and inflation persistence are generally positively correlated. The intuition can again be seen from examining equation (15), repeated here for convenience:

$$\pi_t = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda} s_t + \frac{\alpha\lambda}{1+\alpha^2\lambda}(u_t^N - u_t^T).$$

Consider first the case where there is no uncertainty about the value of the natural rate so the third term on the right-hand side is zero. Then, given that the short-run supply shocks s_t are now positively correlated then so is inflation π_t. Indeed both series would exhibit the exact same first-order serial correlation. Intuitively, given a short-run shock the Central Bank, wishing to keep unemployment at the natural rate, will change inflation to push the unemployment rate closer to the natural rate. If these shocks are positively correlated, a shock today also implies a deviation of unemployment from the natural rate tomorrow in the same direction and thus also a deviation of inflation from target in that same direction.

Introducing uncertainty about the natural rate brings in an additional effect since the forecast errors are also serially correlated. In fact it can be shown that the forecast errors have a first-order serial correlation of exactly θ. This further reinforces the previous effect and leads to potentially considerable serial correlation of inflation well above ρ.

3.2 Empirical evidence on persistence

The persistence of unemployment in this framework follows directly from the assumed AR(1) process for both shocks. The novel result is that this is associated with inflation persistence. The U.S. post-war inflation rate is very persistent (Pivetta and Reis, 2001), and this fact is usually presented as evidence against the Lucas-type Phillips curve used in this paper (Taylor, 1999). The model in this paper shows that serially correlated forecast errors in estimating the natural rate of unemployment can generate this inflation persistence. Currently there is some debate over how to theoretically explain inflation persistence (see Taylor, 1999 section 6) and the model in this paper offers one particular mechanism generating inertia.

Overall, four key predictions regarding observed inflation and unemployment are spelled out by the theory. First, the unemployment rate should be serially correlated and, second, so should inflation. Third, the more persistent the unemployment rate, the more persistent the inflation rate should be. The serial correlation coefficient on short-run supply shocks ρ is a structural parameter of the economy that should vary from country to country according to labor market regulations and the existing market structure, among others. The model predicts that a larger ρ will lead to higher serial correlation of both the unemployment rate and the inflation rate. A positive association between the two can therefore be seen as support for the model. Finally, a fourth implication of the theory is that inflation is only serially correlated due to the effect of serially correlated unemployment. A necessary condition for the serial correlation of unemployment to be zero is that the natural rate and short-run supply shocks are both serially uncorrelated, $\theta = \rho = 0$, and this is a sufficient condition for the serial correlation of inflation to be zero.

I test these by looking at a sample comprising 33 countries with quarterly observations on the unemployment rate and CPI inflation. The time period varies from country to country according to data availability, but generally lies in the 1982-1999 range. Twenty-two countries are OECD members, and the data were obtained from the OECD, whilst for the other eleven, data is from the IMF.[14]

Table 2 contains the first set of results. As a first test I compute the persistence of unemployment using the regression $u_t = c + r_u u_{t-1}$. The t-statistics for the null $r_u = 0$ are shown and we can see that for all but two countries the null is rejected. Second, by the same process I compute the serial correlation of the inflation rate. Again we reject the null $r_\pi = 0$ at the 5% level for all series but one. Third, Figure 7 shows the scatter plot of the 33 coefficients on inflation on the coefficients obtained from the unemployment relation. There is a clear positive relation. Regressing the serial correlation of inflation on the serial correlation of unemployment gives:

$$\begin{aligned} r_\pi \;&=\; \underset{\substack{(0.0862)\\ [0.1097]}}{0.3566} + \underset{\substack{(0.0939)\\ [0.1217]}}{0.5531} r_u \qquad R^2 = 0.52 \end{aligned} \tag{25}$$

The OLS standard errors are in parentheses, while in square brackets are bootstrap-generated standard errors.[15] The coefficient on r_u is positive and statistically significant at the 1% significance level. In Figure 7 two points in the sample (Cyprus and Brazil) seem to be outliers. Refitting the equation excluding these two observations, I obtain a very similar result:

$$\begin{aligned} r_\pi \;&=\; \underset{\substack{(0.1611)\\ [0.0288]}}{0.1064} + \underset{\substack{(0.1705)\\ [0.0311]}}{0.8141} r_u \qquad R^2 = 0.44 \end{aligned} \tag{26}$$

The t-statistic on the r_u coefficient is still very high and statistically significant.

[14] The only criteria used in selecting the countries used was that there would be at least 3 years of consecutive data, i.e., 12 observations, for each of the series.

[15] Since this is a regression between generated regressors, the OLS standard errors are not correct. Yet, none of the standard solutions (Pagan, 1984) apply to this particular problem. I generate the standard errors in square brackets by fitting a VAR to inflation and unemployment for each country, and drawing from the residuals to obtain alternative histories of inflation and unemployment, each of the same size as the original sample for that country. For each of these histories and for all countries I calculate the persistence of inflation and unemployment and then regress the former on the latter across countries, as I did with the original data. I repeat this procedure 100,000 times and report in square brackets the standard deviation of the estimates of the parameters in the regression.

Table 2: First-order serial correlation of inflation and unemployment in 33 countries

Country	Sample size	Unemployment *Serial correlation*	*t-statistic*	*95% C.I.*	Inflation (CPI) *Serial correlation*	*t-statistic*	*95% C.I.*
Australia	70	0.944	26.005	0.871 , 1.016	0.955	29.499	0.890 , 1.020
Austria	26	0.920	14.019	0.785 , 1.056	0.920	13.388	0.778 , 1.063
Belgium	70	0.986	49.333	0.946 , 1.026	0.953	34.172	0.897 , 1.008
Canada	70	0.961	25.688	0.886 , 1.035	0.895	28.713	0.833 , 0.957
Denmark	46	1.006	31.914	0.942 , 1.069	0.896	18.611	0.799 , 0.993
Finland	62	0.988	57.342	0.953 , 1.022	0.942	32.078	0.883 , 1.001
France	70	0.961	57.053	0.927 , 0.994	0.913	49.527	0.876 , 0.950
Germany	70	0.980	38.367	0.929 , 1.031	0.923	23.691	0.845 , 1.000
Hungary	30	1.024	16.689	0.899 , 1.150	0.971	11.482	0.797 , 1.144
Ireland	70	1.030	59.437	0.995 , 1.064	0.883	33.741	0.831 , 0.936
Italy	70	0.966	68.263	0.938 , 0.995	0.950	58.275	0.917 , 0.983
Japan	70	1.065	49.888	1.023 , 1.108	0.864	14.048	0.742 , 0.987
Korea	34	0.915	11.071	0.745 , 1.085	0.786	6.933	0.555 , 1.012
Luxembourg	70	0.966	30.560	0.903 , 1.030	0.942	32.116	0.883 , 1.000
Netherlands	70	1.031	55.763	0.994 , 1.068	0.869	24.169	0.797 , 0.941
New Zealand	55	0.956	35.425	0.901 , 1.010	0.911	20.122	0.820 , 1.002
Norway	70	0.956	38.419	0.895 , 1.017	0.938	37.469	0.888 , 0.988
Portugal	66	0.985	41.323	0.938 , 1.033	0.972	65.653	0.917 , 1.026
Spain	70	0.957	37.258	0.905 , 1.008	0.952	38.424	0.902 , 1.001
Sweden	70	0.992	67.420	0.963 , 1.022	0.958	26.165	0.885 , 1.031
Switzerland	70	0.991	65.192	0.960 , 1.021	0.944	25.192	0.869 , 1.018
U. K.	70	1.010	50.574	0.970 , 1.050	0.889	20.483	0.804 , 0.975
U.S.A.	70	0.981	44.279	0.937 , 1.026	0.844	17.582	0.748 , 0.940
Brazil	14	0.349	1.155	-0.315 , 1.013	0.844	4.535	0.434 , 1.253
Chile	28	0.769	6.132	0.511 , 1.028	0.884	19.766	0.792 , 0.976
Colombia	25	0.899	5.971	0.587 , 1.212	0.855	9.970	0.678 , 1.033
Cyprus	16	0.160	0.570	-0.445 , 0.765	0.293	1.348	-0.177 , 0.764
Israel	32	0.824	9.058	0.638 , 1.011	0.792	6.821	0.554 , 1.029
Malta	28	0.707	4.837	0.406 , 1.009	0.650	4.485	0.353 , 0.949
Romania	23	0.850	6.288	0.568 , 1.132	0.746	9.425	0.581 , 0.911
Russia	22	0.945	18.368	0.838 , 1.053	0.669	15.159	0.576 , 0.761
Slovak Rep.	16	0.850	6.492	0.567 , 1.132	0.857	17.804	0.753 , 0.961
South Africa	16	0.760	4.869	0.423 , 1.097	0.527	2.169	0.002 , 1.051

Data sources: From Australia to USA, the OECD Economic Outlook. From Brazil to South Africa, IMF International Financial Statistics.

Figure 7: First-order serial correlation of inflation and unemployment in 33 countries

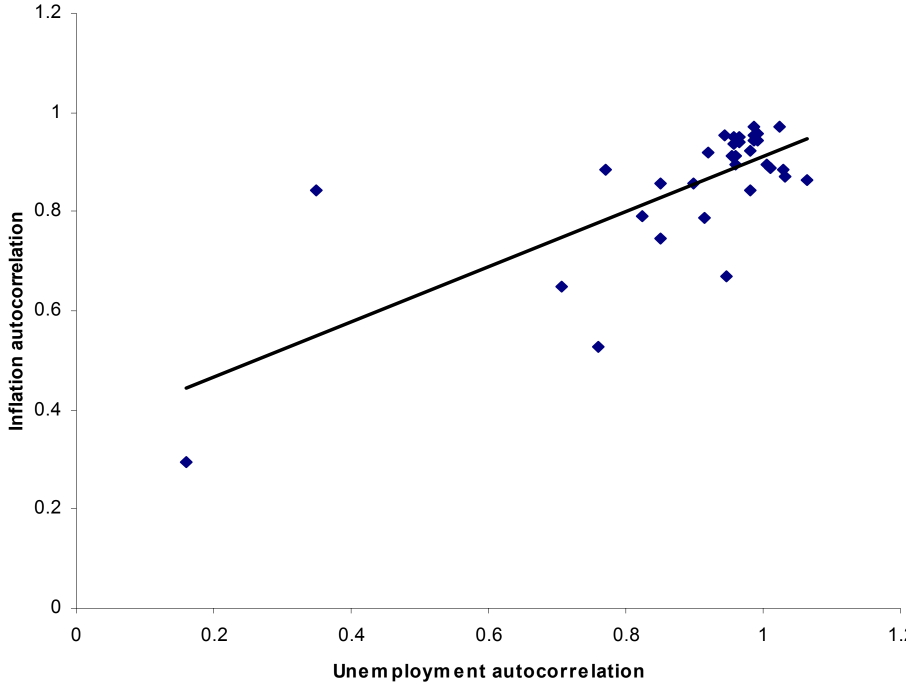

A more formal testing procedure consists of estimating the system of equations:

$$u_{i,t} = a_i + b_i u_{i,t-1} + w_{i,t}, \qquad (27)$$

$$\pi_{i,t} = c_i + (\delta + \gamma b_i)\pi_{i,t-1} + z_{i,t}, \qquad (28)$$

where i indexes countries and t time and $w_{i,t}$ and $z_{i,t}$ are the residuals. The four predictions of the theory spelled out before can be written in terms of this system as:

$(H1)$: $b_i > 0$ - unemployment is persistent.

$(H2)$: $\delta + \gamma b_i > 0$ - inflation is persistent.

$(H3)$: $\gamma > 0$ - higher unemployment persistence raises inflation persistence.

$(H4)$: $\delta = 0$ - inflation is only persistent insofar as unemployment is persistent.

The system is identified by the assumption that δ is the same for all countries, so that there is a common serial correlation of inflation across countries once the effect of the unemployment persistence is taken into account. Estimating the system by iterative least squares yields the results in Table 3.[16] Regarding the first and second hypothesis, the data clearly reject the possibility that the serial correlation of inflation and unemployment is zero. This is rejected

[16] In the panel of estimation, only countries with 30 quarters of consecutive, overlapping unemployment and inflation data are included. This leaves 21 countries, for the sample period 1992 Q1 to 1999 Q2: Australia, Belgium, Canada, Denmark, France, Finland, Germany, Hungary, Italy, Ireland, Israel, Japan, Luxembourg, Netherlands, Norway, New Zealand, Portugal, Spain, Sweden, the United Kingdom and the United States.

not only for the joint test, but also for all the individual coefficients at the 1% significance level.[17] The third hypothesis is also accepted by the data, since the hypothesis $\gamma = 0$ is rejected against the one-sided alternative $\gamma > 0$ at the 1% significance level. Finally, $(H4)$ is tested against a two-sided alternative and the p-value obtained is 42%, so that common significance levels fail to reject the hypothesis that all of the persistence in inflation is due to persistence in the unemployment rate. Thus, the model passes all the tests $(H1)$ to $(H4)$ in the data.[18]

Table 3

	Average of coefficients	F-statistic
b_i	0.952	3066
$\delta + \gamma b_i$	0.862	3328
	Coefficient	t-statistic
γ	1.306	2.595
δ	-0.381	-0.798
	Unemployment regression	Inflation regression
Adjusted R^2	0.992	0.967
Durbin-Watson	1.124	1.459

A further test of the model is to see whether it can generate, for reasonable parameter values, the amount of persistence we observe in the data. Using the results in the previous section, I can calculate the exact predicted serial correlation for inflation, given three parameters: $\sigma_\nu^2/\sigma_\varepsilon^2$, ρ, θ. Staiger, Stock and Watson (1994) provide estimates of the natural rate of unemployment in the United States from 1953 to 1994. Their time series does not correspond exactly to either u_t^N nor u_t^T. The actual natural rate (u_T^N) is always unknown and the Central Bank's forecasts (u_t^T) are based on information available up until t, whereas the Staiger, Stock and Watson estimates use information on the full sample to estimate the natural rate at any point in time. Still, u_t^T has a serial correlation exactly equal to θ, which is the serial correlation of u_t^N, as should any good forecast of the natural rate including those of Staiger, Stock and Watson. From their figures, I calculate that θ is 0.98.[19] The other two parameters, ρ and $\sigma_\nu^2/\sigma_\varepsilon^2$, are then picked to maximize the fit between the serial correlation of unemployment predicted by the model and the value of 0.981 that we observe in U.S. data.[20] This procedure estimates that $\rho = 0.513$, so that the effect of a short-run supply shock is approximately only half its initial impact on the following quarter, and around 1/15

[17] Individual b_i coefficients and significance levels are not reported for brevity but are available from the author.

[18] Interestingly, Fuhrer and Moore (1995) have shown that U.S. inflation is persistent after controlling for detrended output, whereas I find that controlling for the unemployment rate in a panel of 21 countries, there is no persistence in inflation beyond that in unemployment. I leave it to future research to explore the underlying source of these different results.

[19] I am referring to the time-varying estimates by Staiger, Stock, Watson, or TVP, shown in Figure 5.6 in their paper. I use the quarterly estimates.

[20] The serial correlation of unemployment predicted by the model depends on a further parameter: $\alpha^2\lambda/(1+\alpha^2\lambda)$. This is picked jointly with ρ and $\sigma_v^2/\sigma_\varepsilon^2$ to fit the data. The estimated value of this parameter is 0.982.

by one year after the shock; and that $\sigma_\nu^2/\sigma_\varepsilon^2 = 0.501$, so that short-run supply shocks are twice as variable as natural rate shocks.

Given these parameters, the model then predicts that the serial correlation of inflation is 0.948. Looking back at Table 2 we see that the serial correlation for inflation in the United States is 0.844. If anything, the model is therefore generating too much inflation persistence. To check on the robustness of this result, Table 4 presents the serial correlation of inflation obtained for values of $\sigma_\nu^2/\sigma_\varepsilon^2$ in the range [0.3, 0.7], and of ρ in the range [0.3, 0.7]. From the table we see that the model is well capable of producing values for the serial correlation of inflation consistent with those observed in the data.

Table 4: Autocorrelation of inflation

ρ—autocorrelation coefficient of the short-run supply shocks
μ—ratio of the variances of natural rate shocks to short-run shocks

μ \ ρ	0.3	0.4	0.5	0.6	0.7
0.3	0.539	0.761	0.901	0.968	0.993
0.4	0.588	0.804	0.925	0.978	0.996
0.5	0.626	0.832	0.940	0.984	0.997
0.6	0.655	0.852	0.949	0.987	0.997
0.7	0.678	0.868	0.956	0.989	0.997

Note: The autocorrelation of inflation in the 33-country sample lies in (0.84, 0.97).

An alternative explanation for the persistent movements in inflation and unemployment is that both are driven by persistent changes in policy. For instance, in the model in Ball (1995), policy can change stochastically between an inflation-biased regime and a zero-inflation regime and the public cannot observe which type is currently in power. While in tranquil times the inflation-prone policymaker has an incentive to disguise herself as a zero-inflation type to lower inflation expectations, following an adverse shock she reveals her type and raises inflation, which then remains high until a zero-inflation type takes over. While this is a reasonable explanation of persistent bouts of high inflation as in the late 1970s, it cannot account for persistent periods of low inflation, as in the late 1990s in the United States. More generally, whether it is persistent changes in policy or in the natural rate that drive persistent movements in inflation could be empirically tested, subject to being able to obtain good measures of the two. It seems though that while in the late 1970s we can find evidence of both changes in the natural rate and in policy regime, it is very hard to see any policy changes explaining the low inflation in the United States in the late 1990s, whereas there are convincing reasons to expect there has been a favorable shock to the natural rate.

3.3 Further generalizations

The model presented in this paper can be generalized to almost fully general stochastic structures of the exogenous disturbances. Indeed, I can specify the shocks so that $s_t = \frac{R(L)}{S(L)}\varepsilon_t$ and $u_t^N = \frac{N(L)}{D(L)}\nu_t$, where $D(L)$, $N(L)$, $R(L)$, and $S(L)$ are all finite lag polynomials of respective orders d, n, r and s, and ε_t and ν_t are independent, mean zero, finite variance disturbances. By Wold's theorem, any stationary process can be approximated in this way. Given the short-run Phillips curve, equation (1), I can derive the optimal forecast of the natural rate just as before. Given the first-order condition from optimal inflation-setting by the policymaker, equation (13), I obtain the path for inflation. This leads to the following general result:

Result 6: *Given the general representation of the shocks $s_t = \frac{R(L)}{S(L)}\varepsilon_t$ and $u_t^N = \frac{N(L)}{D(L)}\nu_t$:*

a) The optimal forecast of the natural rate is $u_t^T = \Omega(L)x_t$ where $\Omega(L) = \frac{S(L)N_\nu(L)}{T(L)}$ and $N_\nu(L)$ is a lag polynomial of order $max(n, d-1, 0)$ and $T(L)$ a lag polynomial of order $max(n+s, d+r)$.

b) The autocovariance generating function of inflation is given by:

$$g_{\pi\pi}(z) = \frac{\alpha\lambda}{1+\alpha^2\lambda}(1-\Omega(z))(1-\Omega(z^{-1}))g_{xx}(z),$$

where $g_{xx}(z)$ is the autocovariance generating function of the observation variable x_t.

Part a) shows that the optimal forecast of the natural rate still has the form of a distributed lag of past observations of unemployment. Part b) infers the associated stochastic properties of the inflation rate. Note that, even though the policymaker is freely setting inflation at every period and can only affect unemployment by inducing unsystematic deviations from expected inflation, inflation can still exhibit very rich dynamics and follow a very general stochastic structure, as seems to be the case in reality.[21] Moreover, again we see the tight relation between inflation and unemployment autocovariance structures (recall that $x_t = u_t^N + s_t$). Thus, the key lessons taken from the analysis of the special cases above still hold for more general stochastic structures.

Another interesting extension is to allow for the two shocks in the model to have an observable component. Letting g_t and h_t stand for the observable components of the natural rate of unemployment and the short-run supply shocks respectively, the economy is now described by:

$$u_t = u_t^N - \alpha(\pi_t - \pi_t^e) + h_t + \varepsilon_t, \quad (29)$$

$$u_t^N = \theta u_{t-1}^N + g_t(1-\theta) + v_t. \quad (30)$$

The policymaker can now observe

$$x_t = u_t + \alpha(\pi_t - \pi_t^e) + h_t - g_t. \quad (31)$$

and given the two equations above, the economy is described by:

$$x_t = (u_t^N - g_t) + \varepsilon_t, \quad (32)$$

$$(u_t^N - g_t) = \theta(u_{t-1}^N - g_t) + v_t. \quad (33)$$

[21] See Stock and Watson (1999a) on the complex issue of forecasting inflation.

Comparing these two equations with (7) and (8), we see that the forecasting problem is precisely the same as before. Following the same steps as in Section 2.5, it is easy to find the equilibrium inflation rate in this case:

$$\pi_t = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda}\varepsilon_t + \frac{\alpha\lambda}{1+\alpha^2\lambda}(u_t^N - u_t^T) + \alpha\lambda h_t. \tag{34}$$

The last term in this expression is the only difference from (15). Inflation optimally varies with changes in the short-run supply shocks in order to lower the variability of unemployment at the expense of some extra variability in inflation. On the other hand, observable changes in the natural rate have no effect on inflation. They raise both unemployment and the natural rate by the same amount. Since the policymaker is affected only by the difference between the two, they have no effect on the inflation policy.

4 Relationship to the literature

Few papers have examined optimal monetary policy with an uncertain natural rate. Two exceptions are Wieland (1998) and Meyer, Swanson and Wieland (2001). The first examines how uncertainty about the natural rate may affect the balance between caution and experimentation in policy, while the second studies the optimality of a non-linear interest rate policy rule. The model in this paper instead focuses on the implications of uncertainty concerning the natural rate for the behavior of inflation.

Orphanides and van Norden (1999) document the large uncertainty concerning real time estimates of the output gap (which is closely related to the natural rate of unemployment) by the Federal Reserve Board, thus supporting the point of departure for this paper. Orphanides (2003) discusses optimal rules for setting interest rates given this uncertainty. As with the model in this paper, he also explains the great inflation of the 1970s via under-estimates of the natural rate. Finally, after the first draft of this paper was written, Lansing (2001) studied the impact of unforeseen changes in trend productivity on monetary policy. His paper is the closest in the literature to this one. Still, his model of the economy and of monetary policymaking is different from the one in this paper, and his empirical implementation focuses on a distinct set of issues.

This paper is also related to the literature on learning and monetary policy. It differs from the line of work in Cukierman (1992), in which the Central Bank has better information about the economy than the uninformed public. The focus is on the consequences of having the public learning and the policymaker taking account of this learning process to trick the public with surprise inflations to lower the unemployment rate from the natural rate. In the model in this paper, there is no incentive to trick the public and the focus is on the learning process of the Central Bank. Sargent (1999) also has policymakers learning about the state of the economy.[22] Yet, the uncertainty regards what model better describes the economy with each period the policymaker using historical data to estimate the model of the economy and setting optimal policy. This gives rise to the possibility of multiple self-confirming equilibria, which can explain regime changes and large persistent changes in inflation. The model in

[22] Sargent models learning by the use of least squares as in Friedman (1979).

this paper shares with Sargent (1999) the emphasis on rational learning by the policymaker, but focuses the learning on the exact value of the natural rate of unemployment.

Finally, Ireland (1999) structurally estimates a version of the Barro-Gordon model that is close to the one in this paper. While he finds that the long-run restrictions of this model are accepted in the data, the Barro-Gordon model has difficulty capturing the short-term dynamics of inflation and unemployment mainly due to its inability to generate persistent inflation. The results in this paper suggest that incorporating uncertainty about the natural rate of unemployment may remedy this deficiency.

5 Conclusion

Policymakers do not know what the exact value of the natural rate of unemployment is at any point in time. Given uncertainty about the natural rate together with uncertain short-run supply shocks, movements in actual unemployment could be caused by either of the two. If the aim is to target the (forecast of the) natural rate, short-run shocks should be offset, whereas long-run shocks should not. Being unable to distinguish between the two, the policymaker will allow inflation to deviate from its target level and unemployment will fluctuate. Even if the policymaker is behaving optimally in forming his forecast, given shocks to the unknown natural rate, inflation may deviate from target in a persistent way for many periods.

As in Barro and Gordon's (1983) seminal contribution, inflation in this model will deviate from target, but these deviations will vary with time and can either be positive or negative. The model can therefore account for not only the high inflation of the 1970s, but also the low inflation of the 1990s. Moreover, it can generate inflation persistence of the magnitude that we typically observe in the data. This persistence is very related to the persistence of the unemployment rate, and the predicted link between the two is confirmed in the cross-country data.

Colophon

Acknowledgments: I am grateful to N. Gregory Mankiw for many discussions on this topic and to Alberto Alesina, Robert Barro, Mafalda Cardim, Christopher Foote, Benjamin Friedman, Kenneth Rogoff, David Romer, Bryce Ward, and anonymous referees for valuable comments and suggestions. Financial support from the Fundacao Ciencia e Tecnologia, Praxis XXI, is gratefully acknowledged.

E-mail: reis@fas.harvard.edu.

Address: Littauer 200, Department of Economics, Harvard University, Littauer Center, Cambridge, MA 02138.

Appendix A

This Appendix presents a general equilibrium model with nominal rigidities, very close to the one in Woodford (2003), which generates the reduced form model used in the text. The model is as simple as possible, but could be made much more general.

The economy is populated by a continuum of yeoman farmers indexed by i and distributed uniformly in the interval $[0,1]$. All agents are identical in all respects except that each is the monopolist producer of good i. There are perfect financial markets that allow individuals to diversify specific production risk. Thus, equilibrium aggregates can be found using a representative consumer who maximizes $\sum \psi^t U_t$, with period utility function:

$$U_t = \frac{C_t^{1-\kappa} - 1}{1-\kappa} - \int \frac{(\chi L_t(i))^{1+\varphi}}{1+\varphi} di. \tag{A1}$$

Utility of consumption is of the constant relative risk aversion form with $\kappa > 0$. The consumption bundle C_t is a CES aggregator of the consumption of the different goods:

$$C_t = \left[\int_0^1 C_{tj}^{\frac{\eta-1}{\eta}} dj \right]^{\frac{\eta}{\eta-1}}. \tag{A2}$$

It is then easy to derive, as in Dixit and Stiglitz (1977) the demand for each good is

$$C_{tj} = \left(\frac{P_{tj}}{P_t} \right)^{-\eta} C_t, \tag{A3}$$

where P_t refers to the ideal price index, given by:

$$P_t = \left[\int P_{tj}^{1-\eta} dj \right]^{\frac{1}{1-\eta}}. \tag{A4}$$

The disutility of labor supplied by agent i is of the constant elasticity form, and $\varphi > 0$ in order to ensure concavity of the utility function with respect to both of its arguments. The term χ determines the relative disutility of labor supply. It is stochastic, capturing variations in preferences. The agent has access to a linear production technology $Y_t(i) = A_t L_t(i)$, where

A_t is also stochastic capturing variations in productivity. Substituting for labor in the utility function using the production function:

$$U_t = \frac{C_t^{1-\kappa}}{1-\kappa} - \int e^{\nu_t} \frac{Y_t(i)^{1+\varphi}}{1+\varphi} di, \tag{A5}$$

where I define $\nu_t \equiv (1+\varphi)\log(\chi_t/A_t)$. As a normalization, I assume $E(\nu_t) = 0$.

Each period, U_t is maximized subject to the demand for C_{ti} above and to the budget constraint:[23]

$$P_t C_t = \int (1-\tau_t) P_t(i) Y_t(i) di + T_t. \tag{A6}$$

The first-order condition yields the optimal production plan for each variety i:

$$\frac{P_t(i)}{P_t} = \frac{\eta}{\eta-1}\frac{1}{1-\tau_t} e^{\nu_t} \frac{Y_t(i)^{\varphi}}{C_t^{-\kappa}}. \tag{A7}$$

The term $\eta/(\eta-1)(1-\tau_t)$ is the markup over marginal cost charged by the producer. I introduce the term $(1-\tau_t)$ to capture possible sales taxes on the producers (in which case T_t are the corresponding lump-sum transfers), but also more generally variations in the markup due to changes in the extent of competition or collusion, as surveyed in Rotemberg and Woodford (1999). I define the composite stochastic disturbance: $\mu_t = \log(\eta/(\eta-1)(1-\tau_t))$ so that shocks to this should be interpreted as markup shocks, and assume that $E(\mu_t) = 0$, which implies that on average the economy is perfectly competitive. This assumption is necessary to ensure the validity of the welfare approximations and is discussed at length in Woodford (1999). Using the market clearing condition $C_{ti} = y_t(i)$ in the demand funtion (A3) to substitute for $y_t(i)$ in the pricing equation (A7), and using the overall equilibrium condition $C_t = Y_t$ gives the condition that describes equilibrium in the model:

$$\left(\frac{P_t(i)}{P_t}\right)^{1+\eta\varphi} = e^{\mu_t+\nu_t} Y^{\kappa+\varphi}. \tag{A8}$$

Taking logs and letting lower case letters represent the log of a variable this reduces to:

$$p_t(i) = p_t + \frac{(\kappa+\varphi)}{1+\eta\varphi} y_t + \frac{\mu_t+\nu_t}{1+\eta\varphi}. \tag{A9}$$

I can now start deriving the key equations used in the reduced form model (1)-(4).

First, I define the efficient level of output as the output level that would prevail if there were no markup shocks ($\mu_t = E(\mu_t)$) and prices were perfectly flexible. Since the right-hand side of equation (A2) is identical for all firms in these circumstances $p_t(i) = p_t$ for all i. The efficient level is then: $y_t^N = -\nu_t/(\kappa+\varphi)$. In logs, unemployment is given from the production function by: $u_t = -l_t = -y_t - \log(A_t)$. The natural rate of unemployment used in the text is then just defined as $u_t^N = \nu_t/(\kappa+\varphi) + \log(A_t)$. It fluctuates over time in response to marginal cost shocks, either driven by changes in preferences (χ_t via ν_t) or changes in

[23] I abstract from intertemporal trade, again solely for the sake of simplicity, but this is not important for the results.

productivity (A_t directly and via ν_t). Equation (2) in the text follows from modelling these marginal cost disturbances by an AR(1) process.

Second, I specify the nature of the nominal rigidity in the economy. The easiest way to do so is to, following Fischer (1977), assume that while some firms set prices based on contemporaneous information, others must set predetermined prices prior to observing the shocks. Each period, a share $1-\xi$ sets its price with incomplete information at $Ep(i)$, after which the random variables are realized, and then a share ξ of firms sets its price based on perfect information at $p(i)$. The price level is then given by $p = \xi p(i) + (1-\xi)Ep(i)$. Since inflation is defined as $\pi = p_t - p_{t-1}$, it then follows that:

$$\begin{aligned} \pi_t - E\pi_t &= p_t - Ep_t \\ &= \xi(p_t(i) - Ep_t(i)) \\ &= \frac{\xi}{1-\xi}(p_t(i) - p_t) \\ &= \frac{\xi}{1-\xi}\frac{1}{1+\eta\varphi}\left[(\kappa+\varphi)\, y_t + \mu_t + \nu_t\right] \\ &= \frac{\xi}{1-\xi}\frac{1}{1+\eta\varphi}\left[(\kappa+\varphi)\,(y_t - y_t^N) + \mu_t\right] \\ &= \frac{\xi}{1-\xi}\frac{1}{1+\eta\varphi}\left[-(\kappa+\varphi)\,(u_t - u_t^N) + \mu_t\right], \end{aligned} \tag{A10}$$

which rearranging and defining $\alpha \equiv \frac{1-\xi}{\xi}\frac{1+\eta\varphi}{(\kappa+\varphi)}$, and $\varepsilon_t \equiv \frac{\mu_t}{(\kappa+\varphi)}$, yields the short-run Phillips Curve in the text in equation (1). The short-run supply shocks (ε) in the text then correspond to markup shocks in terms of this general equilibrium model.

The final reduced-form equation in the model is the loss function for the Central Bank in equation (4). I can obtain this from our model, by following the approach in Woodford (2002) and log-linearizing the objective function of the representative consumer around the efficient rate of output without shocks: $\bar{Y} = 1$.[24] A second-order approximation of the first component is:

$$\frac{e^{(1-\kappa)y_t} - 1}{1-\kappa} \simeq y_t + \frac{1}{2}(1-\kappa)y_t^2. \tag{A11}$$

I linearize each component of the integral in the second term in turn. Note that the point of approximation involves both $Y_t = \bar{Y}$ and $\nu_t = 0$. Moreover, note that any term that involves only ν_t is beyond the control of the central bank, and will therefore contribute only as a constant to the objective function. These terms can therefore be ignored as can other constants. As a second-order approximation:

$$\frac{e^{\nu_t + (1+\varphi)y_t(i)}}{1+\varphi} \simeq y_t(i) + \frac{1}{2}(1+\varphi)y_t(i)^2 + y_t(i)\nu_t. \tag{A12}$$

[24] This application differs from Woodford by allowing for markup shocks. Moreover, he works with a general utility function whereas in this paper I use specific functional forms.

The next step is to integrate this over the population of households. Using the notation $E_i(y_t(i)) \equiv \int y_t(i)di$ and $VAR_i(y_t(i)) \equiv (E_i(y_t(i)))^2 - E_i(y_t(i)^2)$:

$$E_i(y_t(i)) + \frac{1}{2}(1+\varphi)\left[(E_i(y_t(i)))^2 + VAR_i(y_t(i))\right] + E_i(y_t(i))\nu_t. \tag{A13}$$

Using then the fact that total output is defined by the CES aggregate in equation (A2) to take a second-order approximation of this around the steady state $\bar{Y} = 1$ gives:

$$y_t \simeq E_i(y_t(i)) + \frac{1}{2}(1-\eta^{-1})VAR_i(y_t(i)), \tag{A14}$$

which I will use to replace for $E_i(y_t(i))$ above. Note that the terms in $y_tVAR_i(y_t(i))$, $\nu_tVAR_i(y_t(i))$ or in $VAR_i(y_t(i))^2$ are of third, third and fourth stochastic order respectively. Since I am interested in a second-order approximation, I can drop them to obtain:

$$\int e^{\nu_t}\frac{Y_t(i)^{1+\varphi}}{1+\varphi}di \simeq y_t + \frac{1}{2}(1+\varphi)y_t^2 + y_t\nu_t + \frac{1}{2}(\varphi+\eta^{-1})VAR_i(y_t(i)). \tag{A15}$$

Combining the two components gives the approximation of the welfare function:

$$\begin{aligned} U_t &\simeq -\frac{1}{2}(\varphi+\kappa)y_t^2 - y_t\nu_t - \frac{1}{2}(\varphi+\eta^{-1})VAR_i(y_t(i)) = \\ &= -\frac{1}{2}(\varphi+\kappa)\left\{y_t^2 + y_t\frac{2\nu_t}{\varphi+\kappa} + \frac{\varphi+\eta^{-1}}{\varphi+\kappa}VAR_i(y_t(i))\right\}. \end{aligned} \tag{A16}$$

Finally, recall that $y_t^N = -\nu_t/(\kappa+\varphi)$ and note that this term squared is beyond the control of the policymaker and can therefore be added to the objective function without changing the problem. For the same reason the term $(\varphi+\kappa)$ multiplying the brackets can be dropped to obtain:

$$U_t \simeq -\frac{1}{2}(u_t - u_t^N)^2 - \frac{1}{2}\frac{\varphi+\eta^{-1}}{\varphi+\kappa}VAR_i(y_t(i)). \tag{A17}$$

Finally, realize that from the CES demand function:

$$y_t(i) = -\eta(p_t(i) - p_t) + Y_t, \tag{A18}$$

so the only term varying across the price setters is $p_t(i)$, and so $VAR_i(y_t(i)) = \eta^2 VAR_i(p_t(i))$. But since there are only two groups in the population this takes the simple form:

$$\begin{aligned} VAR_i(p_t(i)) &= (1-\xi)\xi(p_t(i) - Ep_t(i))^2 = \\ &= \frac{(1-\xi)}{\xi}(\pi_t - E\pi_t)^2. \end{aligned} \tag{A19}$$

From this and defining $\lambda \equiv \frac{\eta^2\varphi+\eta}{\varphi+\kappa}\frac{(1-\xi)}{\xi}$ gives equation (4) in the text as the loss function for the policymaker. Expected inflation is replaced by the target π^*.

Appendix B

A. The optimal forecasts of the natural rate

The structure of the economy is described by the two equations (7) and (8). To solve for the optimal forecast, first I must obtain the canonical factorization of x_t, i.e. the lag polynomial $B(L)$ and the i.i.d. random variable ϵ_t with $E(\epsilon_t) = 0$, $Var(\epsilon_t) = \sigma^2$, such that $x_t = B(L)\epsilon_t$. The lag operator L is defined such that for any variable a_t, then $La_t = a_{t-1}$. From (7):

$$(1 - \rho L)(1 - \theta L)x_t = (1 - \rho L)\nu_t + (1 - \theta L)\varepsilon_t. \tag{B1}$$

Defining $\omega_t = \nu_t - \rho\nu_{t-1} + \varepsilon_t - \theta\varepsilon_{t-1}$, then since this is the sum of two MA(1) processes it will be also an MA(1) process: $\omega_t = \epsilon_t - \hat{\beta}\epsilon_{t-1}$ for some $\hat{\beta}$.

I can then find the two unknown parameters $\hat{\beta}$ and σ^2 from the set of equalities arising from the two representations of the ω_t process:

$$\begin{aligned} E(\omega_t^2) &= (1+\rho^2)\sigma_\nu^2 + (1+\theta^2)\sigma_\varepsilon^2 \\ &= (1+\hat{\beta}^2)\sigma^2. && \text{(B2)} \\ E(\omega_t\omega_{t-1}) &= -\rho\sigma_\nu^2 - \theta\sigma_\varepsilon^2 \\ &= -\hat{\beta}\sigma^2. && \text{(B3)} \\ E(\omega_t\omega_{t-i}) &= 0, i \geq 2. && \text{(B4)} \end{aligned}$$

Define $\mu = \sigma_\nu^2/\sigma_\varepsilon^2$, the ratio of the variance of shocks to the natural rate to the variance of real short-run supply shocks. Solving the two equations given in (B2) and (B3) gives the solution:[25]

$$\begin{aligned} \sigma^2 &= \frac{\sigma_\varepsilon^2(\rho\mu + \theta)}{\hat{\beta}}, && \text{(B5)} \\ \hat{\beta} &= \frac{\mu(1+\rho^2) + 1 + \theta^2 - \sqrt{[\mu(1+\rho^2) + 1 + \theta^2]^2 - 4(\rho\mu+\theta)^2}}{2(\rho\mu+\theta)}. && \text{(B6)} \end{aligned}$$

Finally, using the definition of ω_t and (B1) leads to the lag polynomial characterizing the canonical factorization of x_t:

$$B(L) = \frac{1 - \hat{\beta}L}{(1 - \rho L)(1 - \theta L)}. \tag{B7}$$

In the second step, I use the Wiener-Kolmogorov prediction formula, shown by Whittle (1983), Chapters 3 and 6, to minimize the mean squared error within the class of linear estimators $u_t^T = \Omega(L)x_t$:

$$\Omega(z) = \frac{1}{\sigma^2 B(z)} \left[\frac{g_{u^N x}}{B(z^{-1})}\right]_+, \tag{B8}$$

[25] Taking the smallest root of the quadratic.

where z is a complex variable, $g_{u^N x}$ the cross-covariance generating transform[26] between u_t^N and x_t, $B(z)$ and $B(z^{-1})$ come from $g_{xx}(z) = \sigma^2 B(z)B(z^{-1})$ and $[...]_+$ refers, for a sum with both positive and negative powers of z, to the sum containing only elements on non-negative powers of z. I have already found $B(z)$ and σ^2. Using (7), (8) and the assumption that ε_t and ν_t are independent:

$$g_{u^N x} = g_{u^N u^N} = \frac{\sigma_\nu^2}{(1-\theta z)(1-\theta z^{-1})}. \tag{B9}$$

I possess all the elements to apply (B8) and obtain:[27]

$$\Omega(z) = \frac{\mu\hat{\beta}}{\mu+\theta}\frac{1-\rho z}{1-\hat{\beta}z}(1-\theta z)\left[\frac{1-\rho z^{-1}}{(1-\theta z)(1-\hat{\beta}z^{-1})}\right]_+. \tag{B10}$$

Then, applying Whittle's Theorem[28] (Whittle, 1983, page 93):

$$(1-\theta z)\left[\frac{1-\rho z^{-1}}{(1-\theta z)(1-\hat{\beta}z^{-1})}\right]_+ = 1+\theta\frac{\hat{\beta}-\rho}{(1-\hat{\beta}\theta)(1-\rho\theta)}. \tag{B11}$$

Therefore, the optimal forecast of the natural rate is:

$$u_t^T = A\frac{1-\rho L}{1-\hat{\beta}L}x_t, \tag{B12}$$

$$A = \frac{\hat{\beta}(\theta-\hat{\beta})[(1-\hat{\beta}\theta)(1-\rho\theta)+\theta(\hat{\beta}-\theta)]}{(1-\theta\rho)[\rho(1-\theta\hat{\beta})(\theta-\hat{\beta})+\theta(1-\rho\hat{\beta})(\hat{\beta}-\rho)}. \tag{B13}$$

For $\rho \neq 0$, this is the formula in equation (24) in the main text. With i.i.d. short-run supply shocks ($\rho = 0$), then A reduces to $(\theta-\beta)/\beta$ and $\hat{\beta}$ to the β in equation (10) in the main text. Combining equations (B12) and (B13) in this case yields equation (9).

B. Results in Section 2.6

Result 1: From equation (15), inflation at t is:

$$\pi_t = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda}(u_t^N - u_t^T). \tag{B14}$$

[26] The autocovariance generating function of a zero mean random variable π_t is defined as $g_{\pi\pi}(z) = \sum_{i=-\infty}^{+\infty} E(\pi_t \pi_{t-i})z^i$, where z is a complex variable and the series converges for z contained in an annulus about the unit circle (see Nerlove, Grether and Carvalho, 1979; or Hamilton, 1994, for a discussion).

[27] Using the fact that $\mu = \frac{(1-\theta\beta)(\theta-\beta)}{(1-\rho\beta)(\beta-\rho)}$ to simplify the expression.

[28] The version of Whittle's Theorem used here is:

$$R(z) = (1-\theta z)\left[\frac{Q(z)}{1-\theta z}\right]_+ = \Pi_0(z) + [Q(z)]_+$$

where $\Pi_0(z)$ is a polynomial of degree zero in z, such that the differential coefficients of order 0 of $R(z)$ are equal to those of $Q(z)$ at $z = \theta^{-1}$, i.e. $\Pi_0(z) = [Q(z)]_-$ evaluated at $z = \theta^{-1}$.

For all $t < 0$, $u_t = u_t^N = 0$. For $t > 0$, replacing the path of inflation for $t \geq 0$: $u_t^N = \theta^t$ and the definition of the target rate from the optimal forecast equation (9) and simplifying gives the desired expression in (16). As for unemployment, since there are no supply shocks:

$$u_t = u_t^N - \alpha(\pi_t - \pi^*). \tag{B15}$$

Since I have already found the two components in the right-hand side of this expression, just substituting and simplifying gives equation (17).■

Result 2: From the formula for the forecast updating coefficient β, equation (6), as long as $\sigma_\nu^2/\sigma_\varepsilon^2 > 0$ then $\beta < \theta$ and in turn, since $\theta \leq 1$, then $\beta < 1$. Taking limits as $t \to \infty$ of (16) and (17) then shows that π_t converges back to π^* and u_t goes to 0. To prove that $u_t \geq 0, \forall t \geq 0$, note that given the expression (17), this is equivalent to having $\frac{\alpha^2\lambda}{1+\alpha^2\lambda} \leq (\frac{\theta}{\beta})^{t+1}$. As the left-hand side is at most one and the right-hand side at least one as β is bounded above by θ, this inequality holds and so u_t is always non-negative. For the final claim, note that

$$\frac{\partial u_t}{\partial t} \mid_{t=0} = \ln(\theta) - \frac{\beta}{\theta}\frac{\alpha^2\lambda}{1+\alpha^2\lambda}\ln(\beta). \tag{B16}$$

This is positive as long as:

$$\frac{\alpha^2\lambda}{1+\alpha^2\lambda} > \frac{\ln(\theta)}{\ln(\beta)}\frac{\theta}{\beta}. \tag{B17}$$

This inequality holds as long as α and λ are large enough and μ is sufficiently large so β is sufficiently small relative to θ.■

Result 3: Follows by the same steps used to prove results 1 and 2. ■

Result 4: At date 0: the shock occurs $\varepsilon_0 = 1$, the natural rate is unchanged, and the target rises to $(\theta - \beta)/\theta$. Combining all these in the expression for inflation (15), then:

$$\pi_0 = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda} - \frac{\alpha\lambda}{1+\alpha^2\lambda}\frac{\theta-\beta}{\theta}, \tag{B18}$$

which after simplifying gives the desires expression. For $t \geq 1$ then there is no longer a term in ε_t, and the forecast natural rate is just $u_t^T = \frac{\theta-\beta}{\theta}\beta^t$. Again replacing in equation (15) and rearranging gives the result. Unemployment at $t = 0$ is given by $u_0 = -\alpha(\pi_0 - \pi^*) + 1$. Replacing the term in inflation gives the result. For $t \geq 1$, since the natural rate is at zero, and there are no more ε_t shocks, unemployment is just given by α times the expression for the deviation of inflation from target.■

Result 5: The smaller $\sigma_\nu^2/\sigma_\varepsilon^2$, the closer β is to θ and from the expressions for the path of $\pi_t - \pi^*$ and u_t, for $t \geq 1$ the smaller these are. In the limit, as $\sigma_\nu^2/\sigma_\varepsilon^2 \to 0$, then as $\beta \to \theta$, $\pi_t = \pi^*$ and $u_t = 0$, for $t \geq 1$. The proof of asymptotic convergence comes from just taking limits and noting that $\beta < 1$. ■

C. Correlation of inflation and unemployment

First I need to find the first-order serial correlation of inflation. Combining the first-order condition from loss minimization in (13), the equation for x_t in (7) and the optimal forecast of the natural rate in (9) gives the reduced form of inflation:

$$\pi_t = \pi^* + \frac{\alpha\lambda}{1+\alpha^2\lambda}\left[1 - A\frac{1-\rho L}{1-\beta L}\right]\left[\frac{\nu_t}{1-\theta L} + \frac{\varepsilon_t}{1-\rho L}\right]. \tag{B19}$$

Using the fact that ν_t and ε_t are independent and mean zero:

$$COV(\pi_t, \pi_{t-1}) = \left(\frac{\alpha\lambda}{1+\alpha^2\lambda}\right)^2 \left\{ \begin{array}{c} COV\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_t}{1-\rho L}, \left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_{t-1}}{1-\rho L}\right] + \\ COV\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{v_t}{1-\rho L}, \left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{v_{t-1}}{1-\rho L}\right] \end{array} \right\}. \tag{B20}$$

Let me focus on each of the two terms in brackets in turn. On the first:

$$\begin{aligned} & COV\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_t}{1-\rho L}, \left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_{t-1}}{1-\rho L}\right] \\ = \; & COV\left((B_\pi - C_\pi L)\, z_t, (B_\pi - C_\pi L)\, z_{t-1}\right) \\ = \; & (B_\pi^2 + C_\pi^2) COV\left(z_t, z_{t-1}\right) - B_\pi C_\pi VAR\left(z_t\right) - B_\pi C_\pi COV\left(z_t, z_{t-2}\right), \end{aligned} \tag{B21}$$

where $B_\pi = 1 - A$, $C_\pi = \beta - \rho A$, $\phi_1 = \rho + \beta$, $\phi_2 = -\rho\beta$ and $z_t = \varepsilon_t/(1 - \phi_1 L - \phi_2 L^2)$. For an AR(2) process like z_t, it is easy to derive the autocovariance function from which:

$$VAR\left(z_t\right) = \frac{(1-\phi_2)\,\sigma_\varepsilon^2}{(1+\phi_2)[(1-\phi_2)^2 - \phi_1^2]}, \tag{B22}$$

$$COV\left(z_t, z_{t-1}\right) = \frac{\phi_1\sigma_\varepsilon^2}{(1+\phi_2)[(1-\phi_2)^2 - \phi_1^2]}, \tag{B23}$$

$$COV\left(z_t, z_{t-2}\right) = \frac{[\phi_1^2 + \phi_2\,(1-\phi_2)]\sigma^2}{(1+\phi_2)[(1-\phi_2)^2 - \phi_1^2]}. \tag{B24}$$

Using these in (B21):

$$\begin{aligned} & COV\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_t}{1-\rho L}, \left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_{t-1}}{1-\rho L}\right] \\ = \; & \frac{\sigma_\varepsilon^2\left[(B_\pi^2 + C_\pi^2)(\rho+\beta) - B_\pi C_\pi(1 - \rho^2\beta^2 + (\rho+\beta)^2)\right]}{(1-\rho\beta)[(1+\rho\beta)^2 - (\rho+\beta)^2]}. \end{aligned} \tag{B25}$$

Consider now the second term in (B20). Following exactly the same steps, the only difference is now that $\phi_1 = \theta + \beta$ and $\phi_2 = -\theta\beta$, so:

$$\begin{aligned} & COV\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{v_t}{1-\rho L}, \left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{v_{t-1}}{1-\rho L}\right] \\ = \; & \frac{\sigma_\nu^2\left[(B_\pi^2 + C_\pi^2)(\theta+\beta) - B_\pi C_\pi(1 - \theta^2\beta^2 + (\theta+\beta)^2)\right]}{(1-\theta\beta)[(1+\theta\beta)^2 - (\theta+\beta)^2]}. \end{aligned} \tag{B26}$$

So, I have found a formula for the first-order autocovariance of inflation by combining the terms in (B25) and (B26) into (B20).

Turn now attention to the variance. Again using the independence of the errors:

$$VAR(\pi_t) = \left(\frac{\alpha\lambda}{1+\alpha^2\lambda}\right)^2 \left\{ \begin{array}{c} VAR\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\varepsilon_t}{1-\rho L}\right] + \\ VAR\left[\left(1 - A\frac{1-\rho L}{1-\beta L}\right)\frac{\nu_t}{1-\theta L}\right] \end{array} \right\}. \tag{B27}$$

By very similar steps as above:

$$VAR\left(\left[1 - A\frac{1-\rho L}{1-\beta L}\right]\frac{\varepsilon_t}{1-\rho L}\right) = \frac{\sigma_\varepsilon^2\left[(B_\pi^2 + C_\pi^2)(1+\rho\beta) - 2B_\pi C_\pi(\rho+\beta)\right]}{(1-\rho\beta)[(1+\rho\beta)^2 - (\rho+\beta)^2]}, \tag{B28}$$

$$VAR\left(\left[1 - A\frac{1-\rho L}{1-\beta L}\right]\frac{\nu_t}{1-\theta L}\right) = \frac{\sigma_\nu^2\left[(B_\pi^2 + C_\pi^2)(1+\theta\beta) - 2B_\pi C_\pi(\theta+\beta)\right]}{(1-\theta\beta)[(1+\theta\beta)^2 - (\theta+\beta)^2]}. \tag{B29}$$

So, combining (B27), (B28) and (B29) gives the variance of inflation. I therefore have derived the first order serial correlation of inflation as a (complicated but known) function $f(.)$ of the relative variance of the two shocks and their autoregressive coefficients:

$$CORR(\pi_t, \pi_{t-1}) = \frac{COV(\pi_t, \pi_{t-1})}{VAR(\pi_t)} = f(\mu, \rho, \theta). \tag{B30}$$

Now, I would like to show that this is always positive. An algebraic proof would be exhausting. Instead, I numerically evaluated (B30) in the range: $0 \leq \rho < \theta < 1$ and $0 < \mu \leq 1$. In this entire range, the correlation of inflation is positive.

For the serial correlation of unemployment, plugging (B19) into the Phillips curve in equation (1) yields the law of motion:

$$u_t = \left(1 - \frac{\alpha^2\lambda}{1+\alpha^2\lambda} + \frac{\alpha^2\lambda}{1+\alpha^2\lambda}A\frac{1-\rho L}{1-\beta L}\right)\left[\frac{\nu_t}{1-\theta L} + \frac{\varepsilon_t}{1-\rho L}\right]. \tag{B31}$$

But note then that this can be rewritten as:

$$u_t = \frac{B_u - C_u L}{1-\beta L}\left[\frac{\nu_t}{1-\theta L} + \frac{\varepsilon_t}{1-\rho L}\right], \tag{B32}$$

with $B_u = 1 - \frac{\alpha^2\lambda}{1+\alpha^2\lambda} + \frac{\alpha^2\lambda}{1+\alpha^2\lambda}A$ and $C_u = \beta(1 - \frac{\alpha^2\lambda}{1+\alpha^2\lambda}) + \frac{\alpha^2\lambda}{1+\alpha^2\lambda}A\rho$. Comparing this with (B19), it becomes obvious that this is precisely the same problem as before, only with B_u and C_u rather than B_π and C_π. An analogous formula to (B30) for unemployment can then be found by just replacing B_π and C_π by B_u and C_u :

$$CORR(u_t, u_{t-1}) = g(\mu, \rho, \theta, \frac{\alpha^2\lambda}{1+\alpha^2\lambda}) \tag{B33}$$

D. Proof of Result 6

a) is Theorem 8.3 in Nerlove, Grether and Carvalho (1979), pages 308-9. Note only that $T(z)$ is found from the canonical factorization of x_t:

$$\sigma^2 T(z)T(z^{-1}) = \sigma_\varepsilon^2\left[\mu N(z)S(z)N(z^{-1})S(z^{-1}) + R(z)R(z^{-1})D(z)D(z^{-1})\right]. \tag{B34}$$

The $N_\nu(z)$ come from using Whittle's theorem to evaluate:

$$N_\nu(z) = D(z)\left[\sigma_\nu^2 \frac{N(z)N(z^{-1})S(z^{-1})}{T(z^{-1})D(z)}\right]. \tag{B35}$$

For b), from the first-order condition, equation (13), using the Phillips Curve, equation (1) and the optimal forecasting rule $u_t^T = \Omega(L)x_t$:

$$\pi_t = \frac{\alpha\lambda}{1+\alpha^2\lambda}(1-\Omega(L))x_t, \tag{B36}$$

from which b) follows immediately from the definition of the autocovariance generating function. ■

References

Ball, Laurence (1995). "Time-consistent Policy and Persistent Changes in Inflation," *Journal of Monetary Economics*, vol. 36, pp. 329-350.

Ball, Laurence and N. Gregory Mankiw (2002). "The NAIRU in Theory and Practice," *Journal of Economic Perspectives*, vol. 16, pp. 115-136.

Barro, Robert (1977). "Unanticipated Money Growth and Unemployment in the United States," *American Economic Review*, vol. 67, pp. 101-115.

Barro, Robert and David Gordon (1983). "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, vol. 91, pp. 589-610.

Blinder, Alan (1998). *Central Banking in Theory and Practice*, The MIT Press, Cambridge.

Brainard, William (1967). "Uncertainty and the Effectiveness of Policy," *American Economic Review*, vol. 57, pp. 411-425.

Broadbent, Ben and Robert Barro (1997). "Central Bank Preferences and Macroeconomic Equilibrium," *Journal of Monetary Economics*, vol. 39, pp. 17-43.

Cukierman, Alex (1992). *Central Bank Strategy, Credibility and Independence*, The MIT Press, Cambridge.

Chari, V. V., Patrick Kehoe and Ellen McGrattan (2000). "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, vol. 68, pp. 1151-1179.

De Long, J. Bradford (1997). "America's Only Peacetime Inflation: The 1970s," in C. Romer and D. Romer, eds., *Reducing Inflation: Motivation and Strategy*, Chicago University Press, Chicago.

Dixit, Avinash and Joseph Stiglitz (1977). "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, vol. 67 (3), pp. 297-308.

Fischer, Stanley (1977). "Long-term Contracts, Rational Expectations, and the Optimal Money Supply Rule," *Journal of Political Economy*, vol. 85 (1), pp. 191-205.

Friedman, Benjamin (1979). "Optimal Expectations and the Extreme Information Assumptions of "Rational Expectations" Macromodels," *Journal of Monetary Economics*, vol. 5 (1), pp. 23-41.

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

Friedman, Milton (1977). "Nobel Lecture: Inflation and Unemployment," *Journal of Political Economy*, vol. 85 (3), pp. 451-472.

Fuhrer, Jeff and George Moore (1995). "Inflation Persistence," *Quarterly Journal of Economics*, vol. 110 (1), pp. 127-160.

Gordon, Robert (1997). "The Time-Varying NAIRU and Its Implications for Monetary Policy," *Journal of Economic Perspectives*, Winter, pp. 11-32.

Hamilton, James (1994). *Time Series Analysis*, Princeton University Press, Princeton.

Ireland, Peter (1999). "Does the Time-consistency Problem Explain the Behavior of Inflation in the United States?," *Journal of Monetary Economics*, vol. 44, pp. 279-291.

Jorgenson, Dale and Kevin Stiroh (2000). "Raising the Speed Limit: US Economic Growth in the Information Age," *Brookings Papers on Economic Activity*, vol. 2, pp. 125-211.

Lansing, Kevin (2001). "Learning About a Shift in Trend Output," Federal Reserve Bank of San Francisco Working Paper 00-16.

Lucas, Robert (1973). "Some International Evidence on Output-Inflation Trade-offs," *American Economic Review*, vol. 63, pp. 326-334.

Meyer, Laurence (1998). "The Economic Outlook and Challenges Facing Monetary Policy," Speech before the Economic Strategy Institute, Washington, D.C, available at http://www.federalreserve.gov/boarddocs/speeches/1998/19980108.htm.

Meyer, Laurence, Eric Swanson and Volker Wieland (2001). "NAIRU Uncertainty and Nonlinear Policy Rules," *American Economic Review Papers and Proceedings*, vol. 91 (2), pp. 226-231.

Muth, John (1960). "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association*, vol. 55, pp. 290-306.

Nerlove, Marc, David Grether and Jose Carvalho (1979). *Analysis of Economic Time Series: A synthesis*, Academic Press, New York.

Orphanides, Athanasios (2003). "The Quest for Prosperity Without Inflation," *Journal of Monetary Economics*, vol. 50, pp. 633-663.

Orphanides, Athanasios and Simon van Norden (1999). "The reliability of output gap estimates in real-time," Finance and Economics Discussion Series 99-38, Board of Governors of the Federal Reserve System.

Perry, George (1970). "Changing Labor Markets and Inflation," *Brookings Papers on Economic Activity*, no. 3, pp. 411-448.

Phelps, Edmund (1968). "Money-Wage Dynamics and Labor Market Equilibrium," *Journal of Political Economy*, vol. 76, pp. 678-711.

Pivetta, Frederic and Ricardo Reis (2001). "The Persistence of Inflation in the United States," Harvard University unpublished paper.

Poole, William (1970). "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model," *Quarterly Journal of Economics*, vol. 84, pp. 197-216.

Rogoff, Kenneth (1989). "Reputation, Coordination and Monetary Policy," in Robert Barro, ed., *Modern Business Cycle Theory*, Harvard University Press, Cambridge.

Romer, Christina and David Romer (2002). "The Evolution of Economic Understanding and Postwar Stabilization Policy," in Federal Reserve Bank of Kansas City, *Rethinking Stabilization Policy*, Jackson Hole, Wyoming.

Rotemberg, Julio and Michael Woodford (1999). "The Cyclical Behavior of Prices and Costs," in Michael Woodford and John Taylor, eds., *Handbook of Macroeconomics*, North-Holland, Elsevier.

Sargent, Thomas (1998). *The Conquest of American Inflation*, Princeton University Press, Princeton.

Staiger Douglas, James Stock and Mark Watson (1994). "How Precise are Estimates of the Natural Rate of Unemployment?," in Christina Romer and David Romer, eds., *Reducing Inflation: Motivation and Strategy*, University of Chicago Press, Chicago.

Stock, James and Mark Watson (1999a). "Forecasting Inflation," *Journal of Monetary Economics*, vol. 44, pp. 293-335.

Stock, James and Mark Watson (1999b). "Business Cycle Fluctuations in Macroeconomic Time Series," in Michael Woodford and John Taylor, eds., *Handbook of Macroeconomics*, North-Holland, Elsevier.

Taylor, John (1999). "Staggered Price and Wage Setting in Macroeconomics," in Michael Woodford and John Taylor, eds., *Handbook of Macroeconomics*, North-Holland, Elsevier.

Whittle, Peter (1983). *Prediction and Regulation by Linear Least Square Methods*, University of Minnesota Press, Minneapolis.

Wieland, Volker (1998). "Monetary Policy under Uncertainty about the Natural Unemployment Rate," Finance and Economics Discussion Series 98-22, Board of Governors of the Federal Reserve System.

Woodford, Michael (2002). "Inflation Stabilization and Welfare," *The B.E. Journals in Macroeconomics, Contributions to Macroeconomics*, vol. 2 (1), article1.
http://www.bepress.com/bejm/contributions/vol2/iss1/art1.

Woodford, Michael (2003). *Interest and Prices*, in preparation for Princeton University Press.