



ELSEVIER

Journal of Monetary Economics 44 (1999) 279–291

Journal of
MONETARY
ECONOMICS

www.elsevier.nl/locate/econbase

Does the time-consistency problem explain the behavior of inflation in the United States?[☆]

Peter N. Ireland*

Department of Economics, Boston College, Chestnut Hill, MA 02467-3806, USA

Received 31 August 1998; received in revised form 17 March 1999; accepted 20 May 1999

Abstract

This paper derives the restrictions imposed by Barro and Gordon's theory of time-consistent monetary policy on a bivariate time-series model for inflation and unemployment and tests those restrictions using quarterly US data from 1960 through 1997. The results show that the data are consistent with the implications of the theory for the long-run behavior of the two variables, indicating that the theory can explain the initial rise and subsequent fall of inflation over the past four decades. The results also suggest that the theory must be extended to account more fully for the short-run dynamics that appear in the data. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: E31; E52; E61

Keywords: Inflation; Time-consistency; Monetary policy

'We kid ourselves into thinking, "... the Seventies will never happen again"'.
Billy Corgan, *Rolling Stone* Interview, 24 December 1998

[☆]I would like to thank Laurence Ball, Marianne Baxter, Lawrence Christiano, Jon Faust, Michael Kiley, Robert King, Serena Ng, Robert Rich, and conference participants at the Federal Reserve Bank of San Francisco, the National Bureau of Economic Research, and the Studienzentrum Gerzensee for many helpful comments and suggestions.

*Tel.: 617-552-3687; fax: 617-552-2308.

E-mail address: peter.ireland@bc.edu (P.N. Ireland)

1. Introduction

The behavior of the inflation rate in the United States, as measured by quarter-to-quarter percentage changes in the GDP implicit price deflator, is shown in Fig. 1. The graph serves to identify two major episodes in post-war US monetary history: the first, a period of rising inflation that extends from the early 1960s through the early 1980s, and the second, a period of falling inflation that begins in the early 1980s and continues to the present day. These two long-run trends appear more clearly in the 10-year centered moving average that is also displayed in Fig. 1.¹ Why did US policymakers allow inflation to drift higher throughout the 1960s, 1970s and early 1980s? And, conversely, what factors have contributed to their ability to bring inflation down more recently?

Barro and Gordon's (1983) celebrated model of time-consistent monetary policy offers answers to these questions. In Barro and Gordon's model, a policymaker desires to reduce unemployment and, lacking the ability to commit in advance to a monetary policy rule, is tempted to do so in each period by creating surprise inflation in an effort to exploit an expectational Phillips curve. Private agents in the model have rational expectations, however; they recognize that the government faces this temptation to inflate and adjust their decisions accordingly. In equilibrium, therefore, unemployment is no lower than it would otherwise be, and yet the rate of inflation is inefficiently high. Moreover, given the convex costs assigned to unemployment in the model, the policymaker's temptation to inflate – and hence the magnitude of the inflationary bias itself

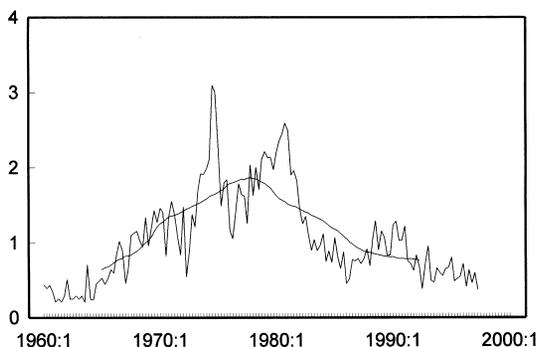


Fig. 1. Quarterly inflation rate and 10-year centered moving average, United States.

¹ King and Watson (1994) find that a similar picture emerges when the data are passed through a filter that is optimally designed to isolate long-run trends.

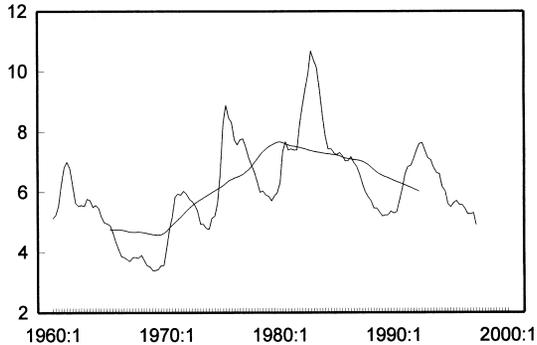


Fig. 2. Unemployment rate and 10-year centered moving average, United States.

– becomes increasingly significant as the natural rate of unemployment rises. Interpreting their model as a positive theory of monetary policy, Barro and Gordon argue that an upward trend in the natural rate of unemployment during the 1960s and 1970s allows their theory to account for the coincident upward trend in inflation.

In fact, estimates of the time-varying natural rate of unemployment constructed, for example, by Gordon (1997) and Staiger et al. (1997) typically indicate not only that the natural rate followed an upward trend during the 1960s, 1970s and early 1980s but also that the natural rate has reversed course and fallen since then. Indeed, Fig. 2 shows that trends appear in the actual civilian unemployment rate that parallel those found in inflation.

Figs. 1 and 2, therefore, suggest that Barro and Gordon's model can potentially explain inflation's initial rise and subsequent fall over the past four decades. They call out for a more detailed analysis that uses the Barro–Gordon model as a basis for statistical tests of the hypothesis that the time-consistency problem underlies the behavior of inflation in the United States. This paper performs such an analysis.² Accordingly, the following sections outline a version of Barro and Gordon's model of time-consistent monetary policy, derive the restrictions that the theory imposes on a bivariate time-series model for inflation and unemployment, and test those restrictions using quarterly US data from 1960 through 1997.

² Parkin (1993) also suggests that the time-consistency problem can explain the behavior of inflation in the United States, but does not test this hypothesis. Baxter (1988) emphasizes the need for statistical tests that establish the empirical relevance of the time-consistency problem, but stops short of actually performing such tests.

2. The Barro–Gordon model and its implications

This section modifies Barro and Gordon's (1983) model by allowing the natural rate of unemployment to follow a more general autoregressive process that contains a unit root and by introducing control errors for inflation. The first extension permits the model to account for the fact, documented below, that the actual unemployment rate in the United States is nonstationary. The second extension permits the model to account for transitory deviations between the actual unemployment rate and the natural rate, which do not arise under Barro and Gordon's original specification.

The actual unemployment rate U_t fluctuates around the natural rate U_t^n in response to deviations of the actual inflation rate π_t from the expected rate π_t^e according to the expectational Phillips curve

$$U_t = U_t^n - \alpha(\pi_t - \pi_t^e), \quad (1)$$

where $\alpha > 0$. The natural rate, in turn, fluctuates over time in response to a real shock ε_t according to the autoregressive process

$$U_t^n - U_{t-1}^n = \lambda(U_{t-1}^n - U_{t-2}^n) + \varepsilon_t, \quad (2)$$

where $1 > \lambda > -1$ and ε_t is serially uncorrelated and normally distributed with mean zero and standard deviation σ_ε .

The monetary authority cannot commit to a policy rule. Instead, at the beginning of each period $t = 0, 1, 2, \dots$, after private agents have formed their expectation π_t^e but before the realization of the real shock ε_t , the policymaker chooses a planned rate of inflation π_t^p . Actual inflation for the period is then determined as the sum of π_t^p and a control error η_t , so that

$$\pi_t = \pi_t^p + \eta_t, \quad (3)$$

where η_t is serially uncorrelated and normally distributed with mean zero, standard deviation σ_η , and covariance $\sigma_{\varepsilon\eta}$ with ε_t .

The policymaker selects π_t^p in an effort to minimize a loss function that penalizes variations of unemployment and inflation around target values kU_t^n and zero:

$$(1/2)(U_t - kU_t^n)^2 + (b/2)\pi_t^2,$$

with $1 > k > 0$ and $b > 0$ so that, in particular, the policymaker wishes to push the actual unemployment rate below the natural rate. Using Eqs. (1) and (3), the policymaker's problem becomes

$$\min_{\pi_t^p} E_{t-1} \{ (1/2)[(1-k)U_t^n - \alpha(\pi_t^p + \eta_t - \pi_t^e)]^2 + (b/2)(\pi_t^p + \eta_t)^2 \},$$

where E_{t-1} denotes the expectation at the beginning of period t or, equivalently, t the end of period $t-1$. The first-order condition for this problem is

$$\alpha E_{t-1} [(1-k)U_t^n - \alpha(\pi_t^p + \eta_t - \pi_t^e)] = b E_{t-1} (\pi_t^p + \eta_t). \quad (4)$$

Private agents know the structure of the economy and understand the policymaker's time-consistency problem. In equilibrium, therefore, they correctly anticipate the policymaker's actions, so that $\pi_t^e = \pi_t^p$. Using this equilibrium condition, along with the fact that $E_{t-1}\eta_t = 0$, Eq. (4) simplifies to

$$\pi_t^p = \pi_t^e = \alpha A E_{t-1} U_t^n, \quad (5)$$

where

$$A = (1 - k)/b > 0.$$

If the policymaker could commit to a policy rule at the beginning of time, he would choose a planned rate of inflation equal to zero in every period. Thus, Eq. (5) reveals that the inflationary bias resulting from the policymaker's inability to commit depends positively on the expected natural rate $E_{t-1} U_t^n$. Here, as in Barro and Gordon's original model, the equilibrium inflation rate moves together with the natural rate of unemployment.

Eqs. (1), (3) and (5) imply that

$$U_t = U_t^n - \alpha \eta_t, \quad (6)$$

which shows how the control error for inflation allows the actual unemployment rate to fluctuate, in equilibrium, around the natural rate. Combining (6) with (2) yields

$$U_t = U_{t-1}^n + \lambda \Delta U_{t-1}^n + \varepsilon_t - \alpha \eta_t, \quad (7)$$

where $\Delta U_{t-1}^n = U_{t-1}^n - U_{t-2}^n$ denotes the change in the natural rate during period $t - 1$. Eqs. (2), (3) and (5), meanwhile, imply that

$$\pi_t = \alpha A U_{t-1}^n + \alpha A \lambda \Delta U_{t-1}^n + \eta_t. \quad (8)$$

Separately, Eqs. (7) and (8) indicate that both inflation and unemployment are nonstationary, inheriting unit roots from the underlying process for the natural rate. Together, however, they imply that

$$\pi_t - \alpha A U_t = -\alpha A \varepsilon_t + (1 + \alpha^2 A) \eta_t, \quad (9)$$

which shows that a linear combination of inflation and unemployment is stationary. Eq. (9), therefore, summarizes the constraint that Barro and Gordon's theory imposes on the long-run behavior of inflation and unemployment; according to the model, these variables should be nonstationary but cointegrated. A statistical test of this cointegration constraint will determine whether the Barro–Gordon model can successfully explain inflation's initial rise and subsequent fall over the past four decades.

Taking first differences of (6), solving for $\Delta U_t^n = U_t^n - U_{t-1}^n$, and substituting the result into (2) yields

$$\Delta U_t = \lambda \Delta U_{t-1} + \varepsilon_t - \alpha \eta_t + \alpha(1 + \lambda) \eta_{t-1} - \alpha \lambda \eta_{t-2}, \quad (10)$$

where $\Delta U_t = U_t - U_{t-1}$ denotes the change in the actual unemployment rate during period t . Together, Eqs. (9) and (10) can be written in the form of a vector ARMA(1,2) for the stationary linear combination of inflation and unemployment and the stationary change in unemployment:

$$\begin{bmatrix} \pi_t - \alpha A U_t \\ \Delta U_t \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & \lambda \end{bmatrix} \begin{bmatrix} \pi_{t-1} - \alpha A U_{t-1} \\ \Delta U_{t-1} \end{bmatrix} + \begin{bmatrix} -\alpha A & 1 + \alpha^2 A \\ 1 & -\alpha \end{bmatrix} \begin{bmatrix} \varepsilon_t \\ \eta_t \end{bmatrix} \\ + \begin{bmatrix} 0 & 0 \\ 0 & \alpha(1 + \lambda) \end{bmatrix} \begin{bmatrix} \varepsilon_{t-1} \\ \eta_{t-1} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & -\alpha\lambda \end{bmatrix} \begin{bmatrix} \varepsilon_{t-2} \\ \eta_{t-2} \end{bmatrix}. \quad (11)$$

The within-equation and cross-equation restrictions appearing in (11) summarize the constraints that Barro and Gordon's theory imposes on the short-run behavior of inflation and unemployment. A statistical test of these restrictions will determine how well the Barro–Gordon model explains the dynamic relationships between these two variables that can be found in the data.

3. Testing the model's implications

This section tests the cointegration constraint shown in (9) to see whether the time-consistency problem can explain the long-run behavior of inflation and unemployment in the United States. In addition, it tests the restrictions shown in (11) to assess the model's ability to account for the short-run dynamics that appear in the data. Both sets of statistical tests use quarterly, seasonally-adjusted data that are drawn from the Federal Reserve Bank of St. Louis' FRED database. The inflation rate is measured, as in Fig. 1, using quarter-to-quarter changes in the GDP implicit price deflator; the unemployment rate is measured, as in Fig. 2, using the civilian unemployment rate. King and Watson (1997) find evidence of a significant break in the time-series behavior of inflation and unemployment occurring around 1970. Therefore, all tests are performed with data from both the full sample period, 1960:1–1997:2, and the subsample beginning in 1970:1.

3.1. Testing the long-run restriction

Eqs. (7) and (8) show that according to the model, both the inflation rate and the unemployment rate ought to be nonstationary. Thus, Table 1 reports results from the Phillips and Perron (1988) test described by Hamilton (1994, Chapter 17), applied to check for unit roots in the two series. The table shows estimates of ρ , the slope coefficient from a regression of each variable on a constant and its own lagged value, as well as τ , the conventional t -statistic for testing the hypothesis that $\rho = 1$. The Phillips–Perron statistic, denoted Z_τ , adjusts the

Table 1
Unit root tests

	ρ	τ	q	Z_τ
1960:1–1997:2				
Unemployment rate	0.9743	– 1.3627	5	– 2.2161
Inflation rate	0.8928	– 2.8739	0	– 2.8739*
1970:1–1997:2				
Unemployment rate	0.9519	– 1.7880	4	– 2.5186
Inflation rate	0.8788	– 2.5362	0	– 2.5362

Notes: The table shows estimates of ρ , the slope coefficient from a regression of each variable on a constant and its own lagged value, as well as τ , the conventional t -statistic for testing the hypothesis that $\rho = 1$. The Phillips and Perron (1988) statistic, Z_τ , adjusts the conventional t -statistic to allow for serial correlation in the regression error. This adjustment uses Newey and West's (1987) method to estimate the variance of the regression error and Andrews' (1991) method to select a value for the lag truncation parameter q required to form the Newey–West estimator, assuming that the process for the regression error is well-approximated by a first-order autoregression. Critical values for Z_τ are reported under the heading case 2 in Hamilton's (1994, p. 763) Table B.6.

*Denotes significance at the 10% level.

conventional t -statistic to allow for serial correlation in the regression error. This adjustment uses Newey and West's (1987) method to estimate the variance of the regression error; Andrews' (Andrews, 1991) method is used to select a value for the lag truncation parameter q required to form the Newey–West estimator, assuming that the process for the regression error is well-approximated by a first-order autoregression.

The results show that the null hypothesis that the process for unemployment contains a unit root cannot be rejected in either sample period. The results for inflation are a little less clear. In the full sample, the unit root hypothesis for inflation can be rejected, but only at the 0.10 significance level; in the post-1970 sample, the unit root hypothesis cannot be rejected.³ Based on these results, along with the others presented below, it seems appropriate to regard both inflation and unemployment as nonstationary.

Table 2 reports results from the Phillips and Ouliaris (1990) test described by Hamilton (1994, Chapter 19), applied to check for the cointegrating relationship between inflation and unemployment implied by (9). The table shows estimates

³ Perron and Ng (1996) show that the Phillips–Perron test tends to reject the null hypothesis of a nonstationary US inflation rate too often in finite samples; this difficulty may account for the conflicting results obtained here.

Table 2
Cointegration tests

Phillips–Ouliaris test					
Sample period	γ	ρ	τ	q	Z_τ
1960:1–1997:2	0.1680	0.8865	– 2.9805	0	– 2.9805**
1970:1–1997:2	0.1791	0.8709	– 2.7603	0	– 2.7603**

Johansen test				
Sample period	λ_1	λ_2	Cointegrating vector	LR = $-T \ln(1 - \lambda_1)$
1960:1–1997:2	0.1010	0.0022	$1.7059\pi_t - 0.3173U_t$	15.7516***
1970:1–1997:2	0.1189	0.0050	$1.6985\pi_t - 0.3303U_t$	13.6701**

Notes: For the Phillips and Ouliaris (1990) test, the table shows estimates of γ , the coefficient from a regression of inflation on unemployment, as well as ρ , the coefficient from a regression of the residual from the inflation-unemployment regression on its own lagged value, and τ , the conventional t -statistic for testing the hypothesis that $\rho = 1$. The statistics q and Z_τ are computed as described in the notes to Table 1. Critical values for the test statistic Z_τ are reported under the heading “case 1” in Hamilton’s (1994, p. 766) Table B.9. For the Johansen (1988) test, the table shows the two eigenvalues λ_1 and λ_2 used in evaluating Johansen’s likelihood function, the estimated cointegrating vectors, and the likelihood ratio statistic LR for testing the null hypothesis of no cointegration. Critical values for LR are reported under the heading “case 1” in Hamilton’s (1994, p. 768) Table B.11.

**Denote significance at the 5% level.

***Denote significance at the 1% level.

of γ , the coefficient from a regression of π_t on U_t , along with the statistics needed to test for a unit root in the residual from this regression. For both sample periods, the hypothesis that inflation and unemployment are not cointegrated can be rejected at the 0.05 significance level.⁴

One potential drawback of the residual-based, Phillips–Ouliaris test, also discussed by Hamilton (1994, pp. 589–590), concerns the fact that the results may hinge sensitively on which variable, inflation or unemployment, appears as the dependent variable in the initial regression. Here, however, Eq. (9) indicates that the hypothesized cointegrating relationship is of the form $\pi_t - \gamma U_t$, making inflation the obvious choice of dependent variable. Nevertheless, the robustness of the Phillips–Ouliaris test results is easily established using Johansen’s (1988) maximum likelihood approach, described by Hamilton (1994, Chapter 20), which does not require a choice of normalization.

⁴ The strong evidence that unemployment is nonstationary, together with the strong evidence that inflation and unemployment are cointegrated, supports the hypothesis, accepted above, that inflation is nonstationary.

Thus, Table 2 also reports statistics associated with the Johansen test for cointegration: the two eigenvalues λ_1 and λ_2 used in evaluating Johansen's likelihood function, the estimates of the cointegrating vector for π_t and U_t , and the likelihood ratio statistic $LR = -T \ln(1 - \lambda_1)$ used to test the null hypothesis of no cointegration. The results of the Johansen tests confirm those of the Phillips–Ouliaris tests, rejecting the null hypothesis of no cointegration between inflation and unemployment at the 0.01 significance level for the full sample and at the 0.05 significance level for the post-1970 sample.

Thus, the results in Table 2 strongly support the Barro–Gordon model's implications for the long-run behavior of inflation and unemployment; as predicted by (9), the two variables are cointegrated. These results are, in particular, consistent with the view that inflation's initial rise and subsequent fall over the past four decades can be explained by the Federal Reserve's inability to commit to a monetary policy that reduces its temptation to exploit the Phillips curve. When coupled with the upward trend in unemployment during the 1960s, 1970s and early 1980s, this inability to commit yields a coincident rise in inflation; when coupled with the downward trend in unemployment since the early 1980s, this inability to commit yields a coincident fall in inflation.

3.2. Testing the short-run restrictions

Focusing now on the theory's implications for the short-run behavior of inflation and unemployment, Table 3 presents maximum likelihood estimates of the model's parameters, obtained by mapping (11), the constrained AMRA(1,2) for $\pi_t - \alpha AU_t$ and ΔU_t , into state-space form and using the Kalman filter to evaluate the likelihood function, as suggested by Hamilton (1994, Chapter 13). The parameter estimates are stable across the two sample periods. Moreover, the standard errors, computed by taking the square roots of the diagonal elements of the inverted matrix of the maximized log likelihood function's second derivatives, are small, implying that each parameter is estimated with a fair degree of accuracy.

The estimates of α suggest that the Phillips curve is quite steep, with each percentage-point of surprise inflation generating a fall in the unemployment rate of only 0.15 percentage points. Although the parameters k and b are not individually identified, the estimates of $A = (1 - k)/b$ exceed unity; the restriction $1 > k > 0$ then implies that b must be less than one. Evidently, the Federal Reserve placed more weight on its goals for unemployment than on its goals for inflation over both sample periods.⁵ The estimates of λ are positive, as expected. Finally, the positive estimates of the covariance σ_{en} indicate that unfavorable

⁵ This conclusion must be accepted with some caution, however: when confidence intervals are placed around the estimates of A , the implied values for b are not necessarily less than one.

Table 3
Maximum likelihood estimates

1960:1–1997:2

	Estimate	Standard error
α	0.1474	0.0579
A	1.1550	0.4546
λ	0.5686	0.0682
σ_ε	0.2675	0.0159
σ_η	0.6299	0.0370
$\sigma_{\varepsilon\eta}$	0.0573	0.0148

L^c = maximized value of constrained log-likelihood = - 148.6597

L^u = maximized value of unconstrained log-likelihood = - 11.2222

LR = $2(L^u - L^c)$ = 274.8750

1970:1–1997:2

	Estimate	Standard error
α	0.1537	0.0641
A	1.1744	0.4894
λ	0.5505	0.0810
σ_ε	0.2905	0.0202
σ_η	0.6530	0.0446
$\sigma_{\varepsilon\eta}$	0.0725	0.0200

L^c = maximized value of constrained log-likelihood = - 119.8246

L^u = maximized value of unconstrained log-likelihood = - 21.7904

LR = $2(L^u - L^c)$ = 196.0684

shocks to the natural rate tend to coincide with unfavorable shocks to inflation; these estimates support the idea that ε_t represents a real, or supply-side, disturbance.

The within-equation and cross-equation restrictions that appear in (11) may be tested by comparing the fit of the constrained ARMA(1,2) to the fit of an unconstrained ARMA(1,2) of the same form. The constrained model has 6 parameters, while an unconstrained ARMA(1,2) for a stationary linear combination of inflation and unemployment of the form $\pi_t - \gamma U_t$ and the stationary change in unemployment ΔU_t has 16 parameters.⁶ Thus, the theory places 10 restrictions on a bivariate time-series model for these stationary variables: if L^c denotes

⁶ The 16 parameters of the unconstrained model include γ , the parameter in the cointegrating relationship, the four autoregressive parameters, the eight moving-average parameters, and the three distinct parameters in the covariance matrix for the innovations.

the maximized value of the log likelihood function for the constrained model, and if L^u denotes the maximized value of the log likelihood function for the unconstrained model, then the likelihood ratio statistic $LR = 2(L^c - L^u)$ has a chi-square distribution with 10 degrees of freedom under the null hypothesis that the constraints hold.

Table 3 shows that $LR = 274.8750$ for the full sample and $LR = 196.0684$ for the post-1970 subsample. The 0.001 critical value for a chi-square random variable with 10 degrees of freedom is 29.6. Thus, the likelihood ratio tests overwhelmingly reject the model's short-run restrictions in (11).

What is the source of this massive rejection of the model's dynamic restrictions? According to (9) – which is also the first equation in (11) – the stationary linear combination of inflation and unemployment identified by the model ought to be serially uncorrelated. The values of ρ reported in the top panel of Table 2, however, indicate that $\pi_t - \alpha AU_t$, while stationary, is also highly persistent. Clearly, Barro and Gordon's model must be extended to explain this persistence if it is to provide a more accurate characterization of the dynamic behavior of inflation and unemployment.

Given the large number of restrictions that appear in (11), however, it comes as no surprise that statistical tests reject the constrained model. After all, the underlying theory is described by three simple equations, listed above as (1)–(3); it would be unreasonable to expect this simple structure to account for all of the dynamics that can be found in the data. More impressive, therefore, is the fact that the time-series model (11) that is implied by this simple theory can be used to obtain reasonable, stable, and precise estimates of the parameters listed in Table 3.

4. Conclusion

Does the time-consistency problem explain the behavior of inflation in the United States? Barro and Gordon's (1983) model of time-consistent monetary policy implies that long-run trends in the natural rate of unemployment will introduce similar trends into the inflation rate when the central bank cannot commit to a monetary policy rule. The tests results presented above, which indicate that inflation and unemployment are cointegrated, are consistent with this implication. Thus, at a minimum, the results suggest that the answer to the question posed in this paper is 'yes', that the time-consistency problem may underlie inflation's initial rise and subsequent fall over the past four decades.

Tests of the model's short-run restrictions, also presented above, indicate that the model is less successful at accounting for the dynamic, quarter-to-quarter comovement of inflation and unemployment in the United States. The version of Barro and Gordon's model used here, however, is extremely simple; consisting of just three linear equations, it makes no attempt, for instance, to explain the

lagged but highly persistent effects on both real and nominal variables that are frequently associated with monetary policy in the US economy. As surveys by Blackburn and Christensen (1989) and Rogoff (1989) make clear, Barro and Gordon's model has already been extended in a variety of ways in order to produce additional insights into the optimal design of monetary policymaking procedures and institutions. Perhaps, similarly, the model can be extended to more accurately capture the short-run behavior of inflation and unemployment.⁷

Looking forward, what implications do the results obtained here have for monetary policy in the United States? DeLong (1997) and Taylor (1997) argue that the behavior of inflation in the US economy is attributable not to the time-consistency problem, but to the gradual process through which policymakers and academic economists learn about the economy.⁸ They suggest that during the 1960s and 1970s, as academic economists came to believe in an exploitable Phillips curve trade-off, policymakers were encouraged to accept a higher rate of inflation in exchange for a lower rate of unemployment; the failure of this experiment then paved the way for the rational expectations revolution and the more recent era of lower inflation. According to this view, the fundamental source of high inflation – economists' and policymakers' ignorance about the structure of the economy – has been removed; hence, the inflationary experiences of the 1960s and 1970s will not be repeated.

The theory and evidence presented here do not invite such optimism. According to the Barro–Gordon model, the inflation of the 1960s and 1970s has as its proximate cause a string of bad luck, in the form of a series of negative and persistent supply-side shocks that worked to increase the natural rate of unemployment. Similarly, from the viewpoint of this model, the disinflation of the 1980s and 1990s represents the product of good luck, in the form of a series of positive and persistent supply-side shocks that have worked to decrease the natural rate. According to this alternative view, therefore, nothing fundamental has changed. The Fed must receive a mandate that allows it to place less weight on developments in the real economy and to commit more credibly to a low-inflation policy. Otherwise, sooner or later, bad luck – and high inflation – will return to haunt us.

⁷ Broadbent and Barro (1997) make some progress in this direction by extending the original Barro–Gordon model to provide a more detailed description of the economy and of the exogenous shocks that hit the economy. In subsequent work, Broadbent (1998) tests a version of this extended model that assigns a key role to regime changes – modeled as discrete shifts in the parameters of the Federal Reserve's objective function – in explaining the behavior of US inflation since 1955. Clearly, more work needs to be done along these lines.

⁸ Sargent (1999) also explores this idea.

References

- Andrews, D.W.K., 1991. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica* 59, 817–858.
- Barro, R.J., Gordon, D.B., 1983. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* 91, 589–610.
- Baxter, M., 1988. Toward an empirical assessment of game-theoretic models of policymaking: A comment. *Carnegie-Rochester Conference Series on Public Policy* 28, 141–151.
- Blackburn, K., Christensen, M., 1989. Monetary policy and policy credibility: Theories and evidence. *Journal of Economic Literature* 27, 1–45.
- Broadbent, B., 1998. Monetary policy regimes and the costs of discretion. Columbia University, Department of Economics (manuscript).
- Broadbent, B., Barro, R.J., 1997. Central bank preferences and macroeconomic equilibrium. *Journal of Monetary Economics* 39, 17–43.
- DeLong, J.B., 1997. America's peacetime inflation: The 1970s. In: Romer, C.D., Romer, D.H. (Eds.), *Reducing inflation: Motivation and strategy*, University of Chicago Press, Chicago, pp. 247–276.
- Gordon, R.J., 1997. The time-varying NAIRU and its implications for economic policy. *Journal of Economic Perspectives* 11, 11–32.
- Hamilton, J.D., 1994. *Time Series Analysis*. Princeton University Press, Princeton.
- Johansen, S., 1988. Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control* 12, 231–254.
- King, R.G., Watson, M.W., 1994. The post-war U.S. Phillips curve: A revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41, 157–219.
- Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Parkin, M., 1993. Inflation in North America. In: Shigehara, K. (Ed.), *Price stabilization in the 1990s: Domestic and international policy requirements*. Macmillan Press, London, pp. 47–83.
- Perron, P., Ng, S., 1996. Useful modifications to some unit root tests with dependent errors and their local asymptotic properties. *Review of Economic Studies* 63, 435–463.
- Phillips, P.C.B., Ouliaris, S., 1990. Asymptotic properties of residual based tests for cointegration. *Econometrica* 58, 165–193.
- Phillips, P.C.B., Perron, P., 1988. Testing for a unit root in time series regression. *Biometrika* 75, 335–346.
- Rogoff, K., 1989. Reputation, coordination, and monetary policy. In: Barro, R.J. (Ed.), *Modern business cycle theory*. Harvard University Press, Cambridge, pp. 236–264.
- Sargent, T.J., 1999. *The Conquest of American Inflation*. Princeton University Press, Princeton.
- Staiger, D., Stock, J.H., Watson, M.W., 1997. How precise are estimates of the natural rate of unemployment? In: Romer, C.D., Romer, D.H. (Eds.), *Reducing Inflation: Motivation and Strategy*. University of Chicago Press, Chicago, pp. 195–242.
- Taylor, J.B., 1997. Comment. In: Romer, C.D., Romer, D.H. (Eds.), *Reducing Inflation: Motivation and Strategy*. University of Chicago Press, Chicago, pp. 276–280.