Monetary Policy Mistakes and the Evolution of Inflation Expectations

Athanasios Orphanides
Central Bank of Cyprus
and
John C. Williams*
Federal Reserve Bank of San Francisco

March 2009

Abstract

What monetary policy framework, if adopted by the Federal Reserve, would have avoided the Great Inflation of the 1960s and 1970s? We use counterfactual simulations of an estimated model of the U.S. economy to evaluate alternative monetary policy strategies over the past 40 years. We show that policies constructed using optimal control techniques aimed at stabilizing inflation, economic activity, and interest rates would have succeeded in achieving a high degree of economic as well as price stability assuming that the Fed had excellent information regarding the structure of the economy. However, in the presence of realistic informational imperfections, such a policy approach would have failed to keep inflation expectations well-anchored, with the result that inflation would have been highly volatile in the 1970s. Optimal control policies would have succeeded only if the weight placed on stabilizing the real economy was relatively modest—with the best results achieved if virtually all the weight was placed on stabilizing prices. We also show that a strategy of following a robust first-difference policy rule would have been more successful than optimal control policies in the presence of informational imperfections. Finally, this robust monetary policy rule yields simulated outcomes that are close to those seen during the period of the Great Moderation starting in the mid-1980s.

Keywords: Great Inflation, rational expectations, robust control, model uncertainty, natural rate of unemployment.

JEL Classification System: E52

*We thank Nick Bloom, Bob Hall, Seppo Honkapohja, Andy Levin, John Taylor, Bharat Trehan, and participants at the NBER Conference on the "Great Inflation" for helpful comments and suggestions. The opinions expressed are those of the authors and do not necessarily reflect views of the Central Bank of Cyprus, the Governing Council of the European Central Bank, the management of the Federal Reserve Bank of San Francisco, or the Board of Governors of the Federal Reserve System. Correspondence: Orphanides: Central Bank of Cyprus, 80 Kennedy Avenue, 1076 Nicosia, Cyprus, Tel: +357-22714471, email: Athanasios.Orphanides@centralbank.gov.cy. Williams: Federal Reserve Bank of San Francisco, 101 Market Street, San Francisco, CA 94105, Tel.: (415) 974-2240, e-mail: John.C.Williams@sf.frb.org.

> "The promise of modern economic policy, managed with an eye to maintaining prosperity, subduing inflation, and raising the quality of life, is indeed great. And although we have made no startling conceptual breakthroughs in economics in recent years, we *have*, more effectively than ever before, harnessed the existing economics—the economics that has been taught in the nation's college classrooms for some twenty years—to the purposes of prosperity, stability, and growth." (Heller, 1966, p. 116, emphasis in the original.)

1 Introduction

Numerous explanations have been put forward for the macroeconomic policy failures that led to the U.S. Great Inflation of the late 1960s and 1970s. Of these, one that is perhaps most disconcerting to current and future policymakers is that these policy mistakes were made by well-informed benevolent policymakers acting without institutional and political constraints. In this paper, we provide a historical account of the Great Inflation and subsequent evolution of the economy in the United States based on an estimated model with such a benevolent and sophisticated policymaker. We assume that the central bank uses state-of-the-art optimal control techniques in deriving its policy decisions. Such methods gained popularity in the 1960s, as evidenced by the epigraph quoted above, and have returned to the forefront of research in recent years (see Svensson and Woodford, 2003; Woodford, 2003; Giannoni and Woodford, 2005; and Svensson and Tetlow, 2005). One concern is that such a policy approach is not robust to realistic degrees of model uncertainty (Levin and Williams, 2003; and Orphanides and Williams, 2008a,b).

This paper "stress-tests" optimal control policies and other policies to see how they would have fared in times of particular macroeconomic turmoil and where the central bank faced informational imperfections. The 1970s provides a perfect laboratory for such an experiment. The U.S. economy was buffeted by large shocks, providing stress to the test. Realized macroeconomic performance during this period was dismal. Our analysis is related to that of Orphanides (2002) and Orphanides and Williams (2005a) that show that a strong response to measures of economic slack contributed to the poor performance during the 1960s and 1970s. The contribution of the current paper is to analyze the performance of optimal control policies and alternative policy approaches. In so doing, we aim to use the experiences of the past to glean lessons for the design of robust monetary policy for the future.

Our model respects the natural rate hypothesis and shares key features with modern models used for monetary policy analysis. We investigate what would have happened over history had policymakers implemented state-of-the-art optimal control methods under the assumption of rational expectations. We focus on the difficulties associated with anchoring inflation expectations when policymakers attempt to maintain a high degree of employment stability relative to price stability in an environment with informational imperfections. The estimated model confirms the presence of adverse supply shocks and natural rate misperceptions during the 1970s, which caused policy to become overly expansionary. However, we find that these errors alone cannot account for the Great Inflation experience.

Using counterfactual simulations, we show that in the absence of informational imperfections, following optimal control policy during the 1960s and 1970s would have maintained reasonably well-anchored inflation expectations and would have succeeded in achieving a favorable degree of economic stability without compromising price stability. However, our model simulations also suggest that informational imperfections, such as policymakers' misperceptions of the natural rate unemployment, would make this approach to policy considerably less successful. Imperfect knowledge amplifies the effects of the underlying shocks and optimal control monetary policies designed assuming complete information would have failed to keep inflation expectations well-anchored. Indeed, optimal control policies would have avoided the Great Inflation only if the weight given to stabilizing the real economy was relatively modest—with the best results achieved if virtually all the weight was placed on stabilizing prices.

We also examine an alternative policy that could have been more robust and avoided this experience, even in the presence of supply shocks and natural rate misperceptions.

We show that such a strategy would have been very successful at stabilizing inflation and economic activity, despite the large shocks of the 1970s. A striking result is that this policy rule yields simulated outcomes close to the realized behavior of the economy during the Great Moderation starting in the mid-1980s, suggesting that the actual practice of monetary policy during this period changed in ways that incorporated the key properties of the robust monetary policy rule.

The remainder of the paper is organized as follows. Section 2 describes the data on natural rate estimates. Section 3 describes the model and its estimation. Section 4 describes the optimal control policy. Section 5 describes the models of expectations formation. Section 6 discusses the simulation methods. Section 7 examines the performance of the optimal control policy using counterfactual simulations. Section 8 analyzes the performance of a simple robust monetary policy rule. Section 9 concludes.

2 The Great Inflation, Inflation Expectations, and Natural Rate Misperceptions

Figure 1 shows the four-quarter average of the U.S. inflation rate, measured by the GDP price deflator, from 1966 to 2003. (Note that throughout this paper, unless otherwise indicated, the figures show the four-quarter moving average of the inflation rate so that the visual clutter caused by quarterly volatility is reduced.) For comparison, the horizontal line shows the 2 percent inflation target that we assume reflects the policymaker's price stability objective for our counterfactual simulations. The inflation rate was around this level before the Great Inflation and returned once again to this level in the last decade of our sample.

Inflation expectations became unmoored during the Great Inflation. Figure 2 compares the current-quarter forecast of inflation from the Survey of Professional Forecasters (SPF) to the corresponding four-quarter-ahead forecast of inflation. In each case, the annualized quarterly inflation rate is shown. The SPF data is first available for the fourth quarter of 1968. (Note that there are missing data for the four-quarter-ahead forecasts in a number of quarters early in the sample.) At the start of the 1970s, the four-quarter-ahead inflation forecast was relatively stable, despite large fluctuations in the current-quarter inflation forecasts, suggesting that inflation expectations were still reasonably well anchored. By the late 1970s, however, four-quarter-ahead inflation expectations were moving almost one-for-one with current-quarter inflation, indicating that inflation expectations had become completely untethered from their previously low level. In the final decade of our sample, there is again evidence that expectations were reasonably well anchored, with four-quarter-ahead inflation expectations more stable than current-quarter inflation. The stability of the long-run inflation expectations according to the SPF (not shown in the figure) supports the conclusions that inflation expectations have been reasonably well anchored during that latter period.

As noted above, a key ingredient in our analysis is the natural rate of unemployment. Figure 3 plots real-time estimates of the natural rate of unemployment and a retrospective measure of the natural rate equal to the current Congressional Budget Office (CBO) estimates. The real-time series for the natural rate is taken from Orphanides and Williams (2005a), extended to include more recent data. These real-time estimates were constructed drawing on a number of sources. During the 1960s, four percent was widely accepted as a reasonable working definition of the full employment rate of unemployment. We do not have precise information for the evolution of real-time perceptions of the natural rate of unemployment in the early 1970s, but many estimates rose during that period, as reflected in various policy-related studies. (See Orphanides and Williams, 2002, for additional information on such studies.) We therefore assume that perceptions of the natural rate rose to 4.5 percent in 1970. Published accounts of Federal Reserve Board model exercises and estimates by the Council of Economic Advisers reported in the *Economic Report of the President* indicate that natural rate estimates continued to rise during the 1970s. From the late 1970s to the present, the CBO has regularly reported, explicitly or implicitly, its estimates of the natural rate in its publications regarding the economic and budget outlook so for these years we use the contemporaneous values for our real-time estimates from these CBO publications.

Differences between real-time estimates of the natural rate of unemployment and current retrospective estimates were especially large and persistent during the second half of the 1960s and the 1970s. The mean absolute difference between the real-time and current estimates was 1.2 percentage points over this period. But such natural rate "misperceptions" are not merely a historical curiosity, as seen clearly in the figure. The mean absolute difference between the two measures is 0.6 percentage points over the period 1980-2003.[1]

3 An Estimated Model of the U.S. Economy

Our analysis uses simulations of the simple estimated quarterly model of the U.S. economy described in Orphanides and Williams (2008a). The specification of the model is motivated by the recent literature on micro-founded models incorporating some inertia in inflation and output (see Woodford, 2003, for a fuller discussion). The main difference from other monetary policy models is that the unemployment gap is substituted for the output gap in the model to facilitate estimation using real-time data. The two concepts are closely related in practice by Okun's law, and the key properties of the model are largely unaffected by this choice.

3.1 The Model

The structural model consists of two equations that describe the behavior of the unemployment rate and the inflation rate and equations describing the time series properties of the exogenous shocks. To close the model, the short-term interest rate is set by the central bank, as described in the next section.

The "IS curve" equation is motivated by the Euler equation for consumption with adjustment costs or habit:

$$u_t = \phi_u u^e_{t+1} + (1 - \phi_u) u_{t-1} + \alpha_u \left(i^e_t - \pi^e_{t+1} - r^* \right) + v_t, \tag{1}$$

$$v_t = \rho_v v_{t-1} + e_{v,t}, \quad e_v \sim \mathrm{N}(0, \sigma^2_{e_v}). \tag{2}$$

[1]Note that this measure of natural rate misperceptions does not take into account the uncertainty regarding the CBO's estimates of the natural rate. Instead, it merely measures changes in the estimates that reflect changes in methodology and the effects of new data.

Equation (1) relates the unemployment rate, u_t, to the unemployment rate expected in the next period, one lag of the unemployment rate, and the difference between the expected *ex ante* real interest rate, r_t^e—equal to the difference between the nominal short-term interest rate, i_t, and the expected inflation rate in the following period, π_{t+1}—and the natural rate of interest, r^*. The unemployment rate is subject to a shock, v_t, that is assumed to follow an AR(1) process with innovation variance $\sigma_{e_v}^2$. The AR(1) specification for the shocks is based on the evidence of serial correlation in the residuals of the estimated unemployment equation, as discussed below.

The "Phillips curve" equation is motivated by the New Keynesian Phillips curve with indexation:

$$\pi_t = \phi_\pi \pi_{t+1}^e + (1 - \phi_\pi)\pi_{t-1} + \alpha_\pi (u_t - u_t^*) + e_{\pi,t}, \quad e_\pi \sim \mathrm{N}(0, \sigma_{e_\pi}^2). \tag{3}$$

It relates inflation, π_t, (measured as the annualized percent change in the GNP or GDP price index, depending on the period) during quarter t to lagged inflation, expected future inflation (denoted by π_{t+1}^e), and the difference between the unemployment rate, u_t, and the natural rate of unemployment, u_t^*, during the current quarter. The parameter ϕ_π measures the importance of expected inflation on the determination of inflation, while $(1 - \phi_\pi)$ captures the effects of inflation indexation. The "mark up" shock, $e_{\pi,t}$, is assumed to be a white noise disturbance with variance $\sigma_{e_\pi}^2$.

We model the low frequency behavior of the natural rate of unemployment as an exogenous AR(1) process independent of all other variables:

$$u_t^* = (1 - \rho_{u^*})\bar{u}^* + \rho_{u^*} u_{t-1}^* + e_{u^*,t}, \quad e_{u^*} \sim \mathrm{N}(0, \sigma_{e_u^*}^2). \tag{4}$$

We assume this process is stationary based on the finding using the standard ADF test that one can reject the null of nonstationarity of the unemployment rate over 1950–2003 at the 5 percent level.

3.2 Model Estimation and Calibration

The IS curve and Phillips curve equations are estimated using forecasts from the Survey of Professional Forecasters (SPF) as proxies for the expectations that appear in the equations.[2] Expectations are assumed to be formed in the previous quarter; that is, the expectations affecting inflation and unemployment in period t are those collected in quarter $t-1$. This matches the informational structure in many theoretical models (see Woodford, 2003, and Giannoni and Woodford, 2005). To match the inflation and unemployment data as closely as possible with these forecasts, the first announced estimates of these series are used. These are obtained from the Real-Time Dataset for Macroeconomists maintained by the Federal Reserve Bank of Philadelphia. In estimating the inflation equation, the Congressional Budget Office (2001) estimates of the natural rate of unemployment are used as proxies for the true values over time. The data sample used for estimating of the model runs from 1968:4 to 2004:2, where the starting date is the first sample point in the SPF.[3]

Estimation results are reported below, with standard errors indicated in parentheses. The IS curve equation is estimated using least squares with AR(1) residuals. Unrestricted estimation of the IS curve equation yields a point estimate for ϕ_u of 0.39, with a standard error of 0.15. This estimate is below the lower bound of 0.5 implied by theory; however, the null hypothesis of a value of 0.5 is not rejected by the data.[4] Thus the restriction $\phi_u = 0.5$ is imposed in estimating the remaining parameters of the equation. Note that the estimated equation also includes a constant term (not shown) that provides an estimate of the natural real interest rate, which is assumed to be constant:

$$u_t = 0.5\, u^e_{t+1} + \; 0.5\, u_{t-1} \; + \underset{(0.022)}{0.056} \; (r^e_t - r^*) + v_t, \tag{5}$$

[2]Specifically, the mean forecasts of the unemployment rate and the three-month Treasury bill rate are used. The inflation forecasts are constructed using the annualized log difference of the GNP or GDP price deflator, taken from the reported forecasts of real and nominal GNP or GDP. The Survey is currently maintained by the Federal Reserve Bank of Philadelphia. See Croushore (1993) and Croushore and Stark (2001) for details on the survey methodology.

[3]Expectations for the Treasury bill rate were not collected in the first few years of the sample. When these are not available, the expectations of the three-month rate implied by the slope of the term structure under the expectations hypothesis are used.

[4]This finding is consistent with the results reported in Giannoni and Woodford (2005) who, in a similar model, find that the corresponding coefficient is constrained to be at its theoretical lower bound.

$$v_t = \underset{(0.085)}{0.513}\ v_{t-1} + e_{v,t}, \quad \hat{\sigma}_{e_v} = 0.30, \tag{6}$$

$$\pi_t = 0.5\,\pi^e_{t+1} + 0.5\,\pi_{t-1} \ \underset{(0.087)}{-0.294}\ (u^e_t - u^*_t) + e_{\pi,t}, \quad \hat{\sigma}_{e_\pi} = 1.35, \tag{7}$$

Unrestricted estimation of the Phillips curve equation yields a point estimate for ϕ_π of 0.51, just barely above the lower bound implied by theory.[5] For symmetry with the treatment of the IS curve, the restriction $\phi_\pi = 0.5$ is imposed and the remaining parameters are estimated using OLS. The estimated residuals for this equation show no signs of serial correlation in the price equation (DW = 2.09), consistent with the assumption of the model.

We do not estimate the model of the natural rate of unemployment; instead, we set the autocorrelation parameter, ρ_{r^*}, to 0.99 and set the unconditional mean to the sample average of the unemployment rate.

4 Monetary Policy

We focus on two alternative approaches to monetary policy. The first is the optimal control approach. The second is a simple monetary policy that is closely related to nominal income growth targeting. In both cases, the policy instrument is the nominal short-term interest rate. We assume that the central bank observes all variables from previous periods when making the current-period policy decision. We further assume that policy is conducted under commitment.

4.1 Optimal Control Monetary Policy

The optimal control approach stipulates that the policy instrument is chosen to minimize the central bank's loss function given the constraints imposed by the central bank's model. We construct the optimal control policy, as is typical in the literature, assuming that the policymaker knows the true parameters of the structural model and assumes all agents use rational expectations and the central bank knows the natural rates of unemployment.[6] Note

[5] For comparison, Giannoni and Woodford (2005) find that the corresponding coefficient is constrained to be at its theoretical lower bound of 0.5.

[6] See, for example, Sargent's (2007) description of the optimal policy approach.

that for optimal control policy, as well as the simple monetary policy rules described below, we use lagged information in the determination of the interest rate, reflecting the lag in data releases.

We assume that the central bank's objective is to minimize a loss equal to the weighted sum of the unconditional variances of the inflation rate, the difference between the unemployment rate and the natural rate of unemployment, and the first-difference of the nominal federal funds rate:

$$\mathcal{L} = Var(\pi - 2) + \lambda Var(u - u^*) + \nu Var(\Delta(i)), \tag{8}$$

where $Var(x)$ denotes the unconditional variance of variable x. We assume an inflation target of 2 percent. In the following, we consider different values of the parameters of the loss function.[7]

The optimal control policy is described by a set of equations that describe the first-order optimality condition for policy and the behavior of the Lagrange multipliers associated with the constraints on the optimization problem implied by the structural equations of the model economy. Because we are interested in describing the setting of interest rates in a potentially misspecified model, it is useful to represent the optimal control policy in an equation that relates the policy instrument to macroeconomic variables, rather than in terms of Lagrange multipliers that depend on the model. There are infinitely many such representations. In the following, we focus on one representation of the optimal control policy, denote by "OC." In the OC policy, the current interest rate depends on three lags of: the inflation rate, the difference between the unemployment rate and the central bank's estimate of the natural rate of unemployment, and the difference between the nominal interest rate and the natural rate of interest. The OC representation yields a determinate rational expectations equilibrium. We find that including three lags of these variables is sufficient to very closely mimic the optimal control outcome assuming the central bank observes the natural rate of unemployment.[8]

[7]Based on an Okun's Law type relationship, the variance of the unemployment gap is about 1/4 that of the output gap, so a choice of λ of 4 corresponds to equal weights on inflation and output gap variability.

[8]In deriving the OC policy, we use the innovation processes from the estimated model and set the

The resulting OC policy, derived for $\lambda = 16$ and $\nu = 1$, is given by the following equation:

$$i_t = 1.16 i_{t-1} - 0.05 i_{t-2} - 0.21 i_{t-3}) \tag{9}$$

$$+0.13\pi_{t-1} - 0.07\pi_{t-2} + 0.05\pi_{t-3}$$

$$-3.70(u_{t-1} - \hat{u}^*_{t-1}) + 2.81(u_{t-2} - \hat{u}^*_{t-1}) - 0.15(u_{t-3} - \hat{u}^*_{t-1}),$$

plus a constant reflecting the constant natural rate of interest and inflation target, where $\hat{u}*_t$ denotes the central bank's estimate of the natural rate of unemployment. The magnitude of the coefficients on inflation and unemployment in the OC policy depend on the value of λ in the loss, with larger values of λ (greater weight on stabilizing unemployment) implying more muted responses to inflation and stronger responses to unemployment.

In the following sections, we also examine the performance of the OC policy derived for alternative values of λ. The resulting OC policy for $\lambda = 4$ is given by the following equation:

$$i_t = 1.17 i_{t-1} + 0.02 i_{t-2} - 0.28 i_{t-3}) \tag{10}$$

$$+0.18\pi_{t-1} + 0.03\pi_{t-2} + 0.01\pi_{t-3}$$

$$-2.47(u_{t-1} - \hat{u}^*_{t-1}) + 2.11(u_{t-2} - \hat{u}^*_{t-1}) - 0.33(u_{t-3} - \hat{u}^*_{t-1}),$$

Compared to the OC policy derived with $\lambda = 16$, this policy is characterized by a stronger response to inflation and a much smaller response to the unemployment rate. Finally, the OC policy, derived for $\lambda = 0$ and $\nu = 1$, is given by:

$$i_t = 1.12 i_{t-1} + 0.13 i_{t-2} - 0.34 i_{t-3}) \tag{11}$$

$$+0.17\pi_{t-1} + 0.09\pi_{t-2} - 0.01\pi_{t-3}$$

$$-1.63(u_{t-1} - \hat{u}^*_{t-1}) + 1.53(u_{t-2} - \hat{u}^*_{t-1}) - 0.38(u_{t-3} - \hat{u}^*_{t-1}).$$

As expected, this policy is characterized by a stronger response to inflation and a much smaller response to the unemployment rate than the OC policy derived for $\lambda = 4$.

innovation standard deviation of the natural rate of unemployment to 0.07. See Orphanides and Williams (2008b) for details.

4.2 Central Bank Natural Rate Estimates

A key input for the OC policy is the central bank's estimate of the natural rate of unemployment. In deriving the OC policy, we assumed that the central bank knows the true structure of the economy, including the value of the natural rate of unemployment. In the model simulations, however, we also entertain alternative assumptions regarding the central bank's knowledge of the natural rate. One alternative is that the central bank's estimates of the natural rate follow the historical pattern of the real-time estimates reported in Figure 3. We refer to this case as "historical natural rate misperceptions." A second alternative is that the central bank estimates the natural rate using the Kalman filter. We refer to this case as "Kalman Filter estimates." In each case, we assume that the true values of the natural rate of unemployment follows the current CBO estimates shown in Figure 3.

In the case of Kalman filter estimation of the natural rate of unemployment, we assume that central bank uses an appropriate Kalman filter consistent with the data. In particular, the central bank's real-time Kalman filter estimate of the natural rate of unemployment, $\hat{u}_t^*$, is given by:

$$\hat{u}_t^* = a_1 \hat{u}_{t-1}^* + a_2 (u_t^* - e_{\pi,t}/\alpha_\pi) \tag{12}$$

where a_1 and a_2 are the Kalman gain parameters. The term within the parenthesis is the current-period "shock" to inflation that incorporates the effects of the transitory inflation disturbance and the deviation of the natural rate of unemployment rate from its unconditional mean, scaled in units of the unemployment rate. Note that the central bank only observes this "surprise" and not the decomposition into its two components.

The optimal values of the gain parameters depend on the variances of the various shocks in the model. Based on a calibrated Kalman filter model, we assume that the central bank uses the following values: $a_1 = 0.982$ and $a_2 = 0.008$ (see Orphanides and Williams, 2008b, for the derivation of these values). We assume that the central bank starts the simulation with the value of 4 percent, consistent with the evidence reported above.

5 Expectations and Simulation Methods

We assume that private agents and, in some cases, the central bank form expectations using an estimated reduced-form forecasting model. Specifically, following Orphanides and Williams (2005b), we posit that private agents engage in perpetual learning, that is, they reestimate their forecasting model using a constant-gain least squares algorithm that weights recent data more heavily than past data. (See Sargent, 1999; Cogley and Sargent, 2001; and Evans and Honkapohja, 2001 for related treatments of learning.) This approach to modeling learning allows for the possible presence of time variation in the economy, including the natural rates of interest and unemployment. It also implies that agents' estimates are always subject to sampling variation, that is, the estimates do not eventually converge to fixed values.

Private agents forecast inflation, the unemployment rate, and the short-term interest rate using a unrestricted vector autoregression model (VAR) containing three lags of these three variables and a constant. Note that we assume that private agents do not observe or estimate the natural rates of unemployment and interest directly in forming expectations. The effects of time variation in natural rates on forecasts are reflected in the forecasting VAR by the lags of the interest rate, inflation rate, and unemployment rate. First, variants of VARs are commonly used in real-world macroeconomic forecasting, making this a reasonable choice on realism grounds. Second, the rational expectations equilibrium of our model with known natural rates is very well approximated by a VAR of this form. As discussed in Orphanides and Williams (2008a), this VAR forecasting model provides accurate forecasts in model simulations.

At the end of each period, agents update their estimates of their forecasting model using data through the current period. Let Y_t denote the 1×3 vector consisting of the inflation rate, the unemployment rate, and the interest rate, each measured at time t: $Y_t = (\pi_t, u_t, i_t)$. Further, let X_t be the 10×1 vector of regressors in the forecast model: $X_t = (1, \pi_{t-1}, u_{t-1}, i_{t-1}, \ldots, \pi_{t-3}, u_{t-3}, i_{t-3})$. Also, let c_t be the 10×3 vector of coefficients of the forecasting model. Using data through period t, the coefficients of the forecasting model can be written in recursive form as follows:

$$c_t = \quad c_{t-1} + \kappa R_t^{-1} X_t (Y_t - X_t' c_{t-1}), \tag{13}$$

$$R_t = \quad R_{t-1} + \kappa (X_t X_t' - R_{t-1}), \tag{14}$$

where κ is the gain. Agents construct the multi-period forecasts that appear in the inflation and unemployment equations in the model using the estimated VAR.

The matrix R_t may not be full rank at times. To circumvent this problem, in each period of the model simulations, we check the rank of R_t. If it is less than full rank, we assume that agents apply a standard Ridge regression (Hoerl and Kennard, 1970), where R_t is replaced by $R_t + 0.00001 * I(10)$ and $I(10)$ is a 10×10 identity matrix.

5.1 Calibrating the Learning Rate

A key parameter in the learning model is the private agent updating parameter, κ. Estimates of this parameter tend to be imprecise and sensitive to model specification, but tend to lie between 0 and 0.04.[9] We take 0.02 to be a reasonable benchmark value for κ, a value that implies that the mean age of the weighted sample is about the same as for standard least squares with a sample of 25 years.

6 Model Simulations

We examine a set of alternative counterfactual simulations to investigate the implications of alternative monetary policy frameworks on macroeconomic developments over the past 40 years. We start our simulations at the beginning of 1966, which corresponds to what many observers consider to be the beginning of the Great Inflation in the United States.

6.1 Initial Conditions

The state variables of the model economy with learning are: the current and lagged values of the inflation rate, the federal funds rate, the unemployment rate, the true natural rate of unemployment, the real-time estimate of the natural rate, the shocks to the structural

[9]See Sheridan (2003), Orphanides and Williams (2005a), Branch and Evans (2006), and Milani (2007).

equations, and the matrices C and R for the forecasting model. We initialize the C and R matrices using the values implied by the reduced-form solution of the model under rational expectations for the stipulated monetary policy rule.

To compute the history of equation residuals, we first compute the implied forecasts from our forecasting model of inflation, the unemployment rate, and the federal funds rate over the period 1966 – 2003. We initialize the coefficients of the learning model by estimating the VAR using data from 1949-1965. We treat the forecasts generated by the learning model as the true data for agents' expectations and then compute tracking residuals, that is, the values of the historical residuals for the equations for the unemployment rate, the inflation rate, and the natural rate of unemployment. Thus, given these residuals and the historical path for the nominal interest rate, the model's predictions will exactly match the historical paths for all endogenous variables. We then conduct counterfactual experiments in which we modify assumptions regarding policy or the learning process, but do not change the paths for the equation residuals for unemployment, inflation, and the natural rate of interest, which we assume are exogenous. Each counterfactual simulation starts in the first quarter of 1966 and ends in the fourth quarter of 2003.

Note that we assume that the initial conditions for the agents' learning model are consistent with the policy rule in place. That is, we assume that at the start of the simulation expectations are well aligned with the monetary policy regime under consideration. Over time, expectations then evolve as described above. This assumption implies that the initial conditions for these state variables are different across the counterfactual simulations. As a result, the simulated paths will often differ significantly from the historical patterns.

7 Performance of Optimal Control Policies

If the Fed had correct estimates of the natural rate of unemployment, then the the OC policy derived assuming a moderate weight for unemployment stabilization would have avoided the Great Inflation. Figure 4 shows the simulated paths for the four-quarter inflation rate, inflation expectations, and the unemployment rate, assuming that the Fed follows the OC policy derived under $\lambda = 16$ and $\nu = 1$. The left column of the figure show the outcomes assuming the Fed knew the most recent value of the natural rate of unemployment. Inflation would have been somewhat volatile during the 1970s, reflecting the effects of the large shocks hitting the economy at the time, but the deviation of the four-quarter inflation rate from target would not have exceeded three percentage points during that period.

In the absence of natural rate misperceptions, inflation expectations would have remained reasonably contained during the 1970s. The middle left panel shows the simulated four-quarter-ahead inflation expectations under the OC policy For comparison, the figure also shows the corresponding SPF inflation forecasts, which rose dramatically in the 1970s.[10] As seen in the left panel of Figure 5, the OC policy acts to raise the unemployment rate up to the natural rate by 1967 and holds the unemployment rate moderately above the natural rate through most of the 1970s, offsetting the inflationary effects of the supply shocks of that period. These policy actions helps stabilize inflation and inflation expectations and avoids the need of a disinflationary policy at the end of the decade.

This same OC policy performs dismally in the face of the historical natural rate misperceptions and leads to a Great Inflation outcome in the 1970s. The panels on the right column of Figure 4 show the outcomes when the Fed uses the historical real-time natural rate estimates. The simulated path of inflation during the 1970s is similar to that seen in the actual data. But, unlike in the data, the high volatility of inflation continues through to the end of the sample. Owing to the low real-time estimate of the natural rate of unemployment, in this simulation the Fed does not act to raise unemployment during the latter part of the 1960s and early 1970s, as seen in the right panel of Figure 5. This extended period of easy policy leads to a sustained rise in inflation and inflation expectations. By the time that the supply shocks of the 1970s strike, inflation expectations are completely untethered from the assumed 2 percent target.

Could the high inflation of the late 1970s have been mitigated by following an optimal control policy predicated on placing a much lower weight on unemployment stabilization?

[10]For this figure and those that follow, the SPF three-quarter-ahead inflation forecast is substituted for the four-quarter-ahead forecast in the periods when the latter is missing from the survey.

Orphanides and Williams (2008a,b) show that "robust optimal control" policies derived assuming downward biased values of λ and ν are robust to imperfect knowledge of the type studied in the present paper. We examine the effectiveness of such an approach by evaluating the performance of the OC policies derived assuming alternative weights on unemployment in the central bank loss of 4 and 0.

The OC policy derived assuming $\lambda = 4$ avoids the worst of the Great Inflation during the 1970s, even with natural rate misperceptions. The left column of panels in Figure 6 show the simulation results when the natural rate is known by the Fed. The results are similar to the case of the OC policy derived assuming $\lambda = 16$. In the case of natural rate misperceptions, monetary policy is too easy during the late 1960s and early 1970s and, as a result, inflation and inflation expectations trend upwards. The rise in inflation during this period is not as extreme as seen in the actual data.

In the absence of natural rate misperceptions, the OC policy that places no weight on unemployment stabilization performs very well during the 1970s (and indeed for the entire sample period). The left column panels in Figure 7 show the simulated paths of inflation, inflation expectations, and unemployment under this policy in the case of no natural rate misperceptions. Under this policy, fluctuations in inflation and inflation expectations are far more muted than under the OC policy derived assuming $\lambda = 4$ or 16. This greater stabilization of inflation comes at the cost of only somewhat greater variability in the unemployment rate.

With the historical natural rate misperceptions, the OC policy derived with a zero weight placed on the stabilization of unemployment in the loss function avoids the Great Inflation, but still allows considerable inflation volatility to develop. The panels on the right column of Figure 7 show the simulated paths of inflation, inflation expectations, and unemployment under this policy in the case of historical natural rate misperceptions. Given the incorrect low estimate of the natural rate of unemployment at the start of the simulation, this policy keeps unemployment too low for too long. This mistake allows inflation to rise and inflation expectations to become untethered. Note that the policy error does not stem from a concern of stabilizing unemployment for its own good, but instead reflects the importance of deviations of unemployment from its natural rate for the future path of inflation. With inflation reaching 6 percent by mid-decade, policy acts aggressively to bring inflation back down to target by the end of the 1970s and a major stagflation is averted.

Interestingly, given the presence of natural rate misperceptions, the OC policies derived with a nonnegligible weight on stabilizing unemployment yield much greater inflation variability in the final 20 years of our sample than is seen in the data. Although these policies describe the Great Inflation period reasonably well, they do not match the experience since the disinflation of the early 1980s. In contrast, the OC policy derived assuming no weight on unemployment stabilization does a much better job of describing inflation during the latter part of the sample.

8 Performance of a Simple Policy Rule

We now examine the performance of an alternative monetary policy rule that has proven to be robust to various forms of model uncertainty in other contexts (see Tetlow, 2006, and Orphanides and Williams, 2008a,b). The rule was proposed by Orphanides and Williams (2007) and takes the form:

$$i_t = i_{t-1} + \theta_\pi(\bar{\pi}^e_{t+3} - \pi^*) + \theta_{\Delta u}(u_{t-1} - u_{t-2}). \tag{15}$$

A key feature of this policy is the absence of any measures of natural rates in the determination of policy. This policy rule is related to the elastic price standard proposed by Hall (1984), whereby the central bank aims to maintain a stipulated relationship between the forecast of the unemployment rate and the price level. It is also closely related to the first difference of a modified Taylor-type policy rule in which the forecast of the price level is substituted for the forecast of the inflation rate.

We choose the parameters of these simple rules to minimize the central bank loss for $\lambda = 4$ and $\nu = 1$, under the assumptions of rational expectations and constant natural rates.[11] The resulting optimized simple rule is given by:

$$i_t = i_{t-1} + 1.74\,(\bar{\pi}^e_{t+3} - \pi^*) - 1.19\,(u_{t-1} - u_{t-2}). \tag{16}$$

This is the same rule as analyzed in Orphanides and Williams (2008a,b), where it was shown to perform well under imperfect knowledge.

According to the model simulation, if the Fed had followed this simple rule over the past 40 years, inflation would have been relatively stable and the Great Inflation would never have occurred. Figure 8 compares the simulated path of inflation, inflation expectations, and unemployment under this simple robust policy rule to the actual data. Because this simple policy rule does not respond to the natural rate of unemployment, the simulations are invariant to the assumed path of central bank natural rate estimates. Inflation does fluctuate a bit during the 1970s, reflecting the large shocks of that period, but the deviations from target are short lived. The simulated path for inflation is very stable since the mid-1980s.

This simple policy rule is extremely effective at keeping inflation expectations well anchored. Although the inflation rate itself fluctuates under the simple policy rule, inflation is expected to return to near its target rate of 2 percent within one year. As discussed in Orphanides and Williams (2008a), the anchoring of inflation expectations is key to the success of this rule in stabilizing inflation and unemployment. A striking result is that this simple rule does better at stabilizing inflation and inflation expectations than the OC policy derived for $\lambda = 0$. The anchoring of inflation expectations implies that the gap between the unemployment rate and the natural rate is considerably smaller throughout the sample than in the actual data.

Interestingly, the simulated behavior of inflation, inflation expectations, and unemployment over the latter part of our sample is very close to that of the actual data. This finding suggests that the actual policy framework during this period may not have been very different from that prescribed by this robust simple rule.

[11] If we allow for time-varying natural rates that are *known by all agents*, the optimized parameters of this simple rule under rational expectations are nearly unchanged. The relative performance of this policy is also unaffected.

The anchoring of long-run inflation expectations under the simple robust policy rule is illustrated by the small variance in the simulated path for inflation expectations over the next 10 years. The thick solid line in Figure 9 shows the simulated path for 10-year inflation expectations when monetary policy follows the simple robust policy rule. This line fluctuates very little over the entire sample.[12] For comparison, the thin dashed line in the chart shows the path of 10-year inflation expectations under the OC policy optimized for $\lambda = 4$ and assuming historical natural rate misperceptions. This line fluctuates considerably over the sample, reflecting the relatively poor of anchoring of inflation expectations under this regime. The dashed line shows the corresponding outcomes under the OC policy optimized for $\lambda = 4$ and assuming no natural rate misperceptions. Not surprisingly, long-run inflation expectations are generally reasonably well-anchored in this case. However, there are extended episodes during the 1970s and early 2000s when long-ruun inflation expectations become unmoored.

9 Conclusion

The model simulations in this paper suggest that modern optimal control policies would have done a poor job of stabilizing inflation during the 1970s in the presence of realistic informational imperfections such as misperceptions of the natural rate of unemployment. Such a policy would have failed to keep inflation expectations well-anchored, resulting in highly volatile inflation. An optimal control policy would have succeeded only if the weight placed on stabilizing the real economy was relatively modest—with the best results achieved if virtually all the weight was placed on stabilizing prices. Finally, we show that a strategy of following a simple first-difference policy rule would have been more successful than optimal control policies in the presence of realistic informational imperfections. In addition, this policy rule yields simulated outcomes close to the realized behavior of the economy during

[12]No consistent data series for 10-year inflation expectations exists for the entire sample against which we can compare these simulated data. Data from surveys of long-run inflation expectations start at the end of the 1970s. These show 10-year consumer price index inflation expectations reaching around 8 percent at the start of the 1980s, and then gradually falling to around 2-1/2 percent in the late 1990s. Since the late 1990s, these estimates fluctuate very little.

the Great Moderation starting in the mid-1980s, suggesting that the actual practice of monetary policy during this period changed in ways that incorporated the key properties of the robust monetary policy rule.

References

Branch, William A. and George W. Evans (2006), "A Simple Recursive Forecasting Model," *Economics Letters* 91, 158-166.

Cogley, Timothy and Sargent, Thomas (2001), "Evolving Post-World War II U.S. Inflation Dynamics," In: Bernanke, B.S., Rogoff, K.S. (Eds.), NBER Macroeconomics Annual 2001. Cambridge, MA: The MIT Press, 331–373.

Croushore, Dean (1993), "Introducing: The Survey of Professional Forecasters," Federal Reserve Bank of Philadelphia *Business Review*, November/December, 3–13.

Croushore, Dean and Tom Stark (2001), "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105, 111–130, November.

Evans, George and Honkapohja, Seppo (2001), *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press.

Gaspar, Vitor, Frank Smets, and David Vestin (2006). "Adaptive Learning, Persistence, and Optimal Monetary Policy" ECB Working paper, June.

Giannoni, Marc P. and Michael Woodford (2005). "Optimal Inflation Targeting Rules," In: Bernanke, B.S., Woodford, M. (Eds.), *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 93–162.

Hall, Robert E. (1984), "Monetary strategy with an elastic price standard," *Price Stability and Public Policy: A Symposium Sponsered by the Federal Reserve Bank of Kansas City*, Kansas City: Federal Reserve Bank of Kansas City, 137–159.

Heller, Walter W. (1966), *New Dimensions of Political Economy*, Cambridge, MA: Harvard University.

Hoerl, A.E. and R.W. Kennard (1970), "Ridge regression: Biased Estimation of Nonorthogonal Problems," *Tecnometrics*, 12, 69–82.

Levin, Andrew T. and John C. Williams, 2003. Robust Monetary Policy with Competing Reference Models. *Journal of Monetary Economics*, 50, 945–975.

Milani, Fabio (2007), "Expectations, Learning, and Macroeconomic Persitence," *Journal of Monetary Economics*, 54, 2065–2082.

Orphanides, Athanasios (2002), "Monetary Policy Rules and the Great Inflation," *American Economic Review,* 92(2), 115-120, May.

Orphanides, Athanasios and John C. Williams (2002), "Robust Monetary Policy Rules with Unknown Natural Rates", *Brookings Papers on Economic Activity*, 2:2002, 63-118.

Orphanides, Athanasios and John C. Williams (2005a). "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations." *Journal of Economic Dynamics and Control*, 29(11), 1927–1950, November.

Orphanides, Athanasios and John C. Williams (2005b) "Imperfect Knowledge, Inflation Expectations and Monetary Policy," in *The Inflation Targeting Debate*, Ben Bernanke and Michael Woodford (eds.), Chicago: University of Chicago Press, 201–234.

Orphanides, Athanasios and John C. Williams (2007) "Inflation Targeting Under Imperfect Knowledge," in *Monetary Policy Under Inflation Targeting*, Frederic Mishkin and Klaus Schmidt-Hebbel (eds.), Santiago: Central Bank of Chile, 77–123

Orphanides, Athanasios and John C. Williams (2008a) "Learning, Expectations Formation, and the Pitfalls of Optimal Control Monetary Policy," in *Journal of Monetary Economics,* forthcoming.

Orphanides, Athanasios and John C. Williams (2008b) "Imperfect Knowledge and the Pitfalls of Optimal Control Monetary Policy," Federal Reserve Bank of San Francisco Working Paper 2008-09, July.

Sargent, Thomas J. (1999), *The Conquest of American Inflation*, Princeton: Princeton University Press.

Sargent, Thomas J. (2007), "Evolution and Intelligent Design," mimeo, New York University, September.

Sheridan, Niamh (2003), "Forming Inflation Expectations," Johns Hopkins University, mimeo, April.

Svensson, Lars E. O. and Robert Tetlow (2005,) "Optimum Policy Projections," *International Journal of Central Banking*, 1, 177–207.

Svensson, Lars E. O. and Michael Woodford (2003), "Optimal Indicators for Monetary Policy," *Journal of Monetary Economics*, 46, 229-256.

Tetlow, Robert J. (2006), "Real-time Model Uncertainty in the United States: 'Robust' policies put to the test," mimeo,Federal Reserve Board, May 22.

Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press.

Figure 1: The Great Inflation

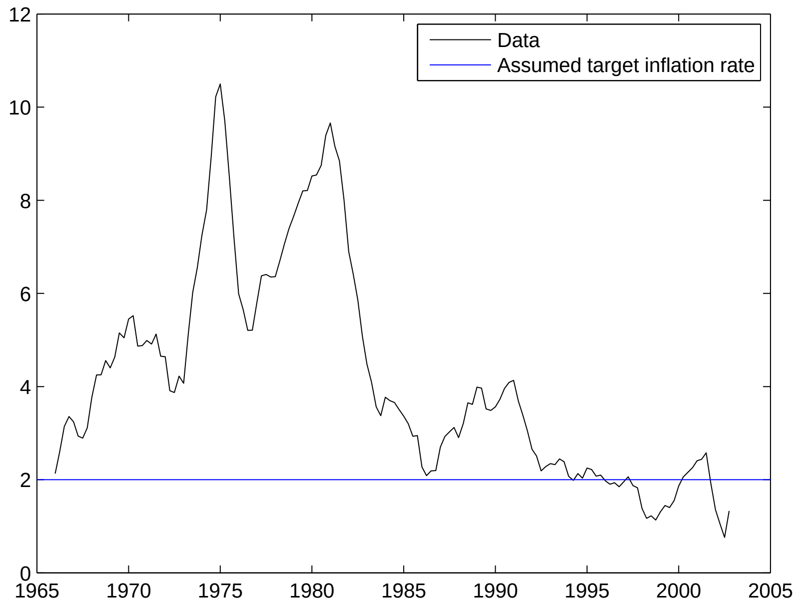

Figure 2: Inflation Expectations Become Untethered

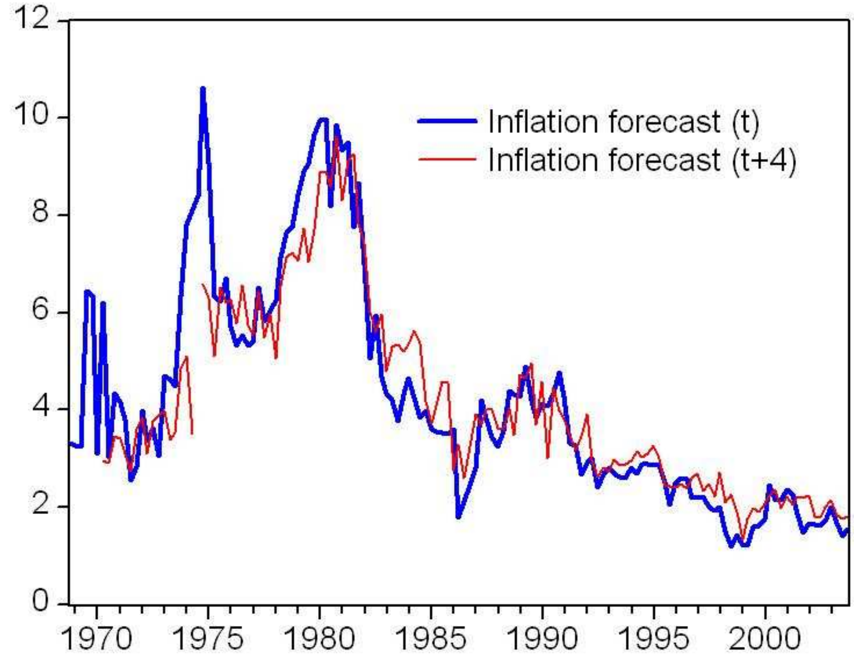

Figure 3: Estimates of the Natural Rate of Unemployment

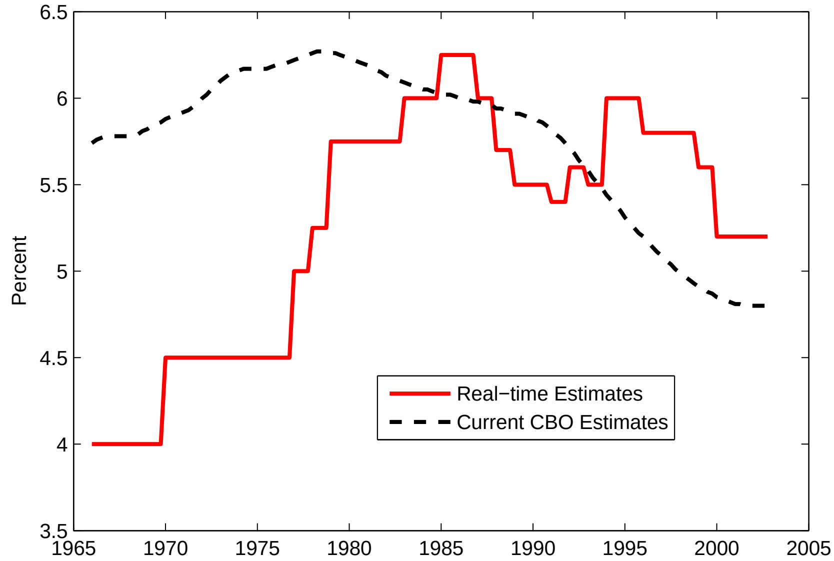

Figure 4: Counterfactual Simulations under OC Policy with $\lambda = 16$

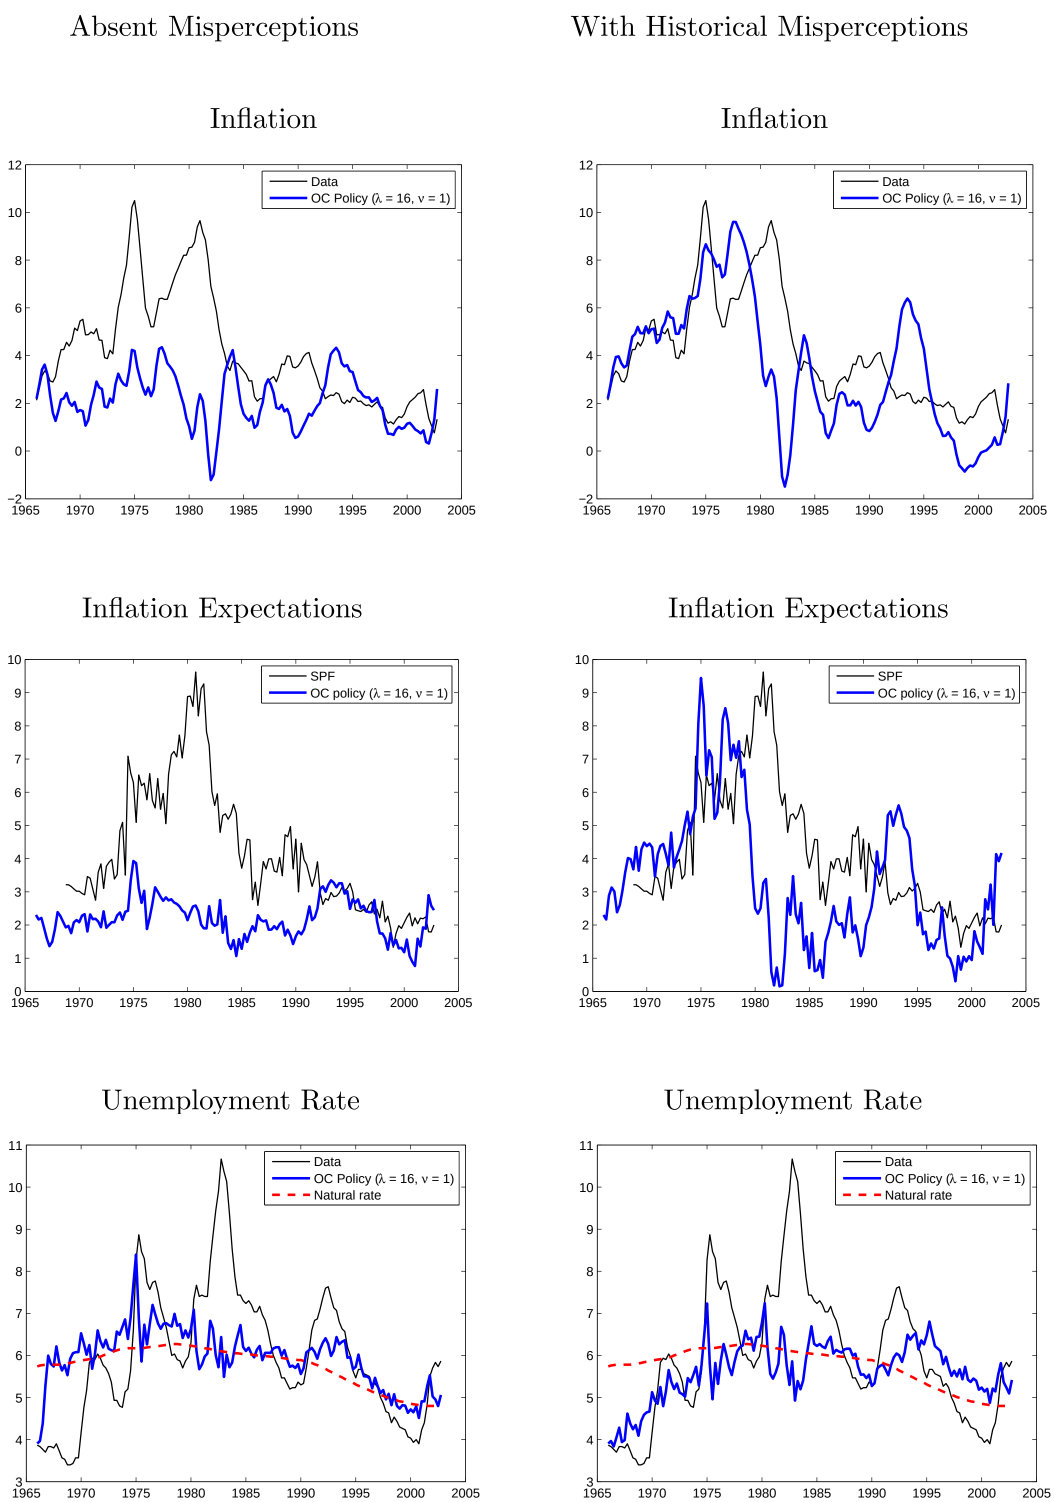

Figure 5: Counterfactual Simulations under OC Policy with $\lambda = 16$

Absent Misperceptions

With Historical Misperceptions

Real Interest Rate

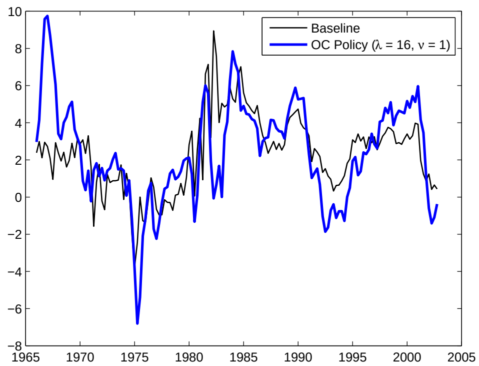

Real Interest Rate

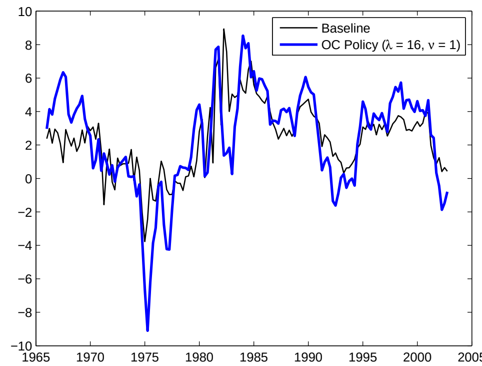

Figure 6: Counterfactual Simulations under OC Policy with $\lambda = 4$

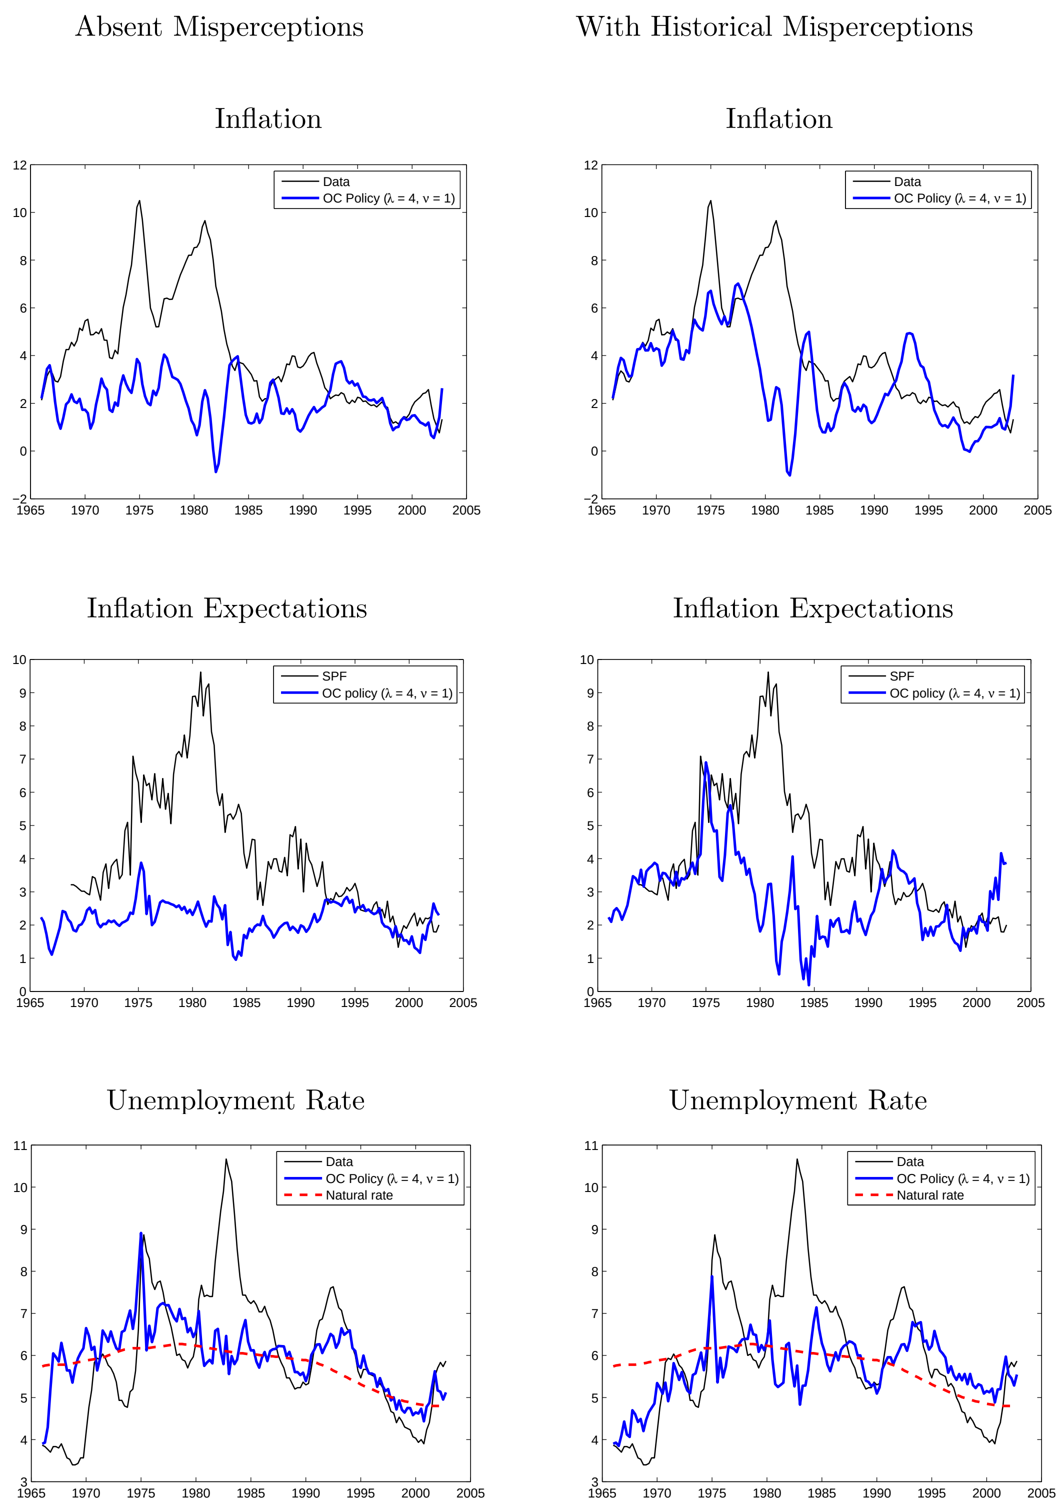

Figure 7: Counterfactual Simulations under OC Policy with $\lambda = 0$

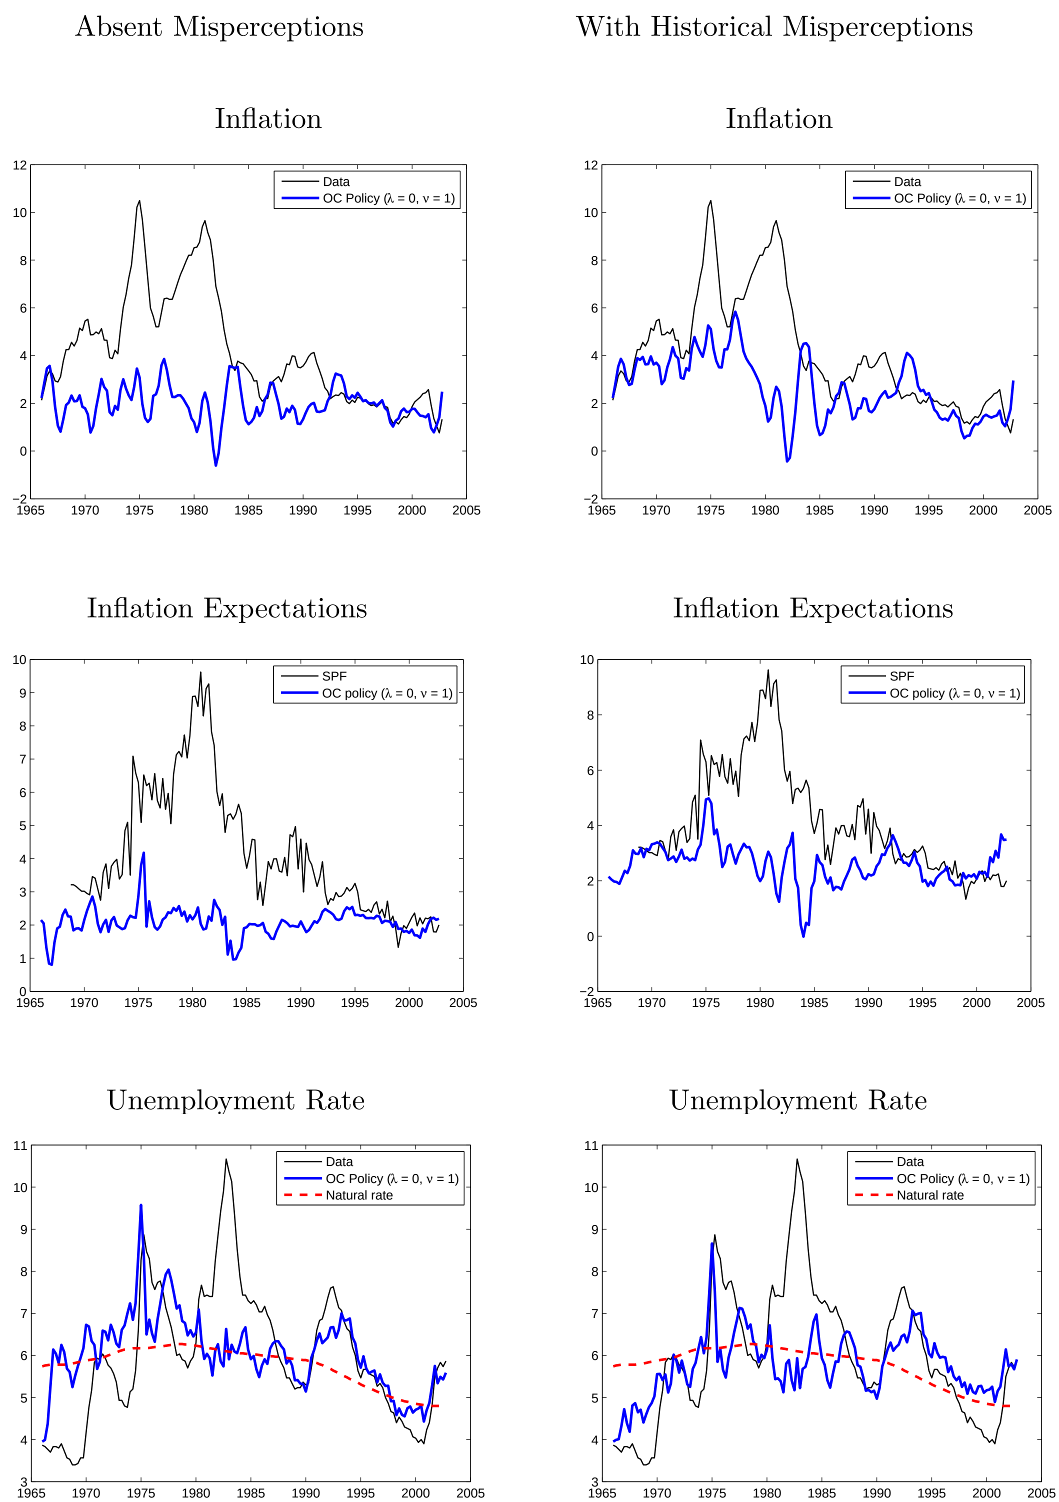

Figure 8: Counterfactual Simulations under Robust Simple Rule ($\lambda = 4$)

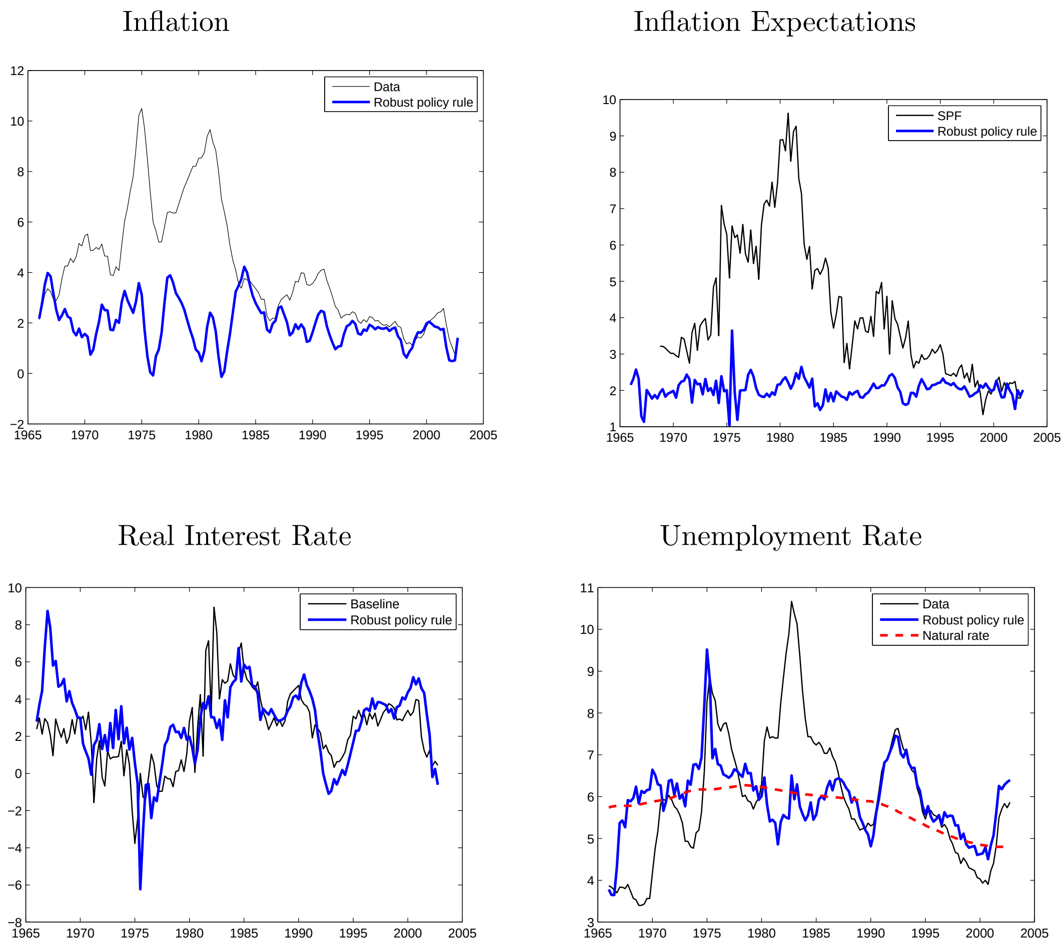

$$i_t = i_{t-1} + 1.74\,(\bar{\pi}^e_{t+3} - \pi^*) - 1.19\,(u_{t-1} - u_{t-2}).$$

Figure 9: Simulations of 10-Year Inflation Expectations ($\lambda = 4$)