

On the Limits of Monetary Policy*

Stefano Eusepi[†]

Marc Giannoni[‡]

Bruce Preston[§]

September 17, 2018

[PRELIMINARY AND INCOMPLETE]

ABSTRACT

This paper provides theory and evidence that distorted long-term interest-rate expectations represent a fundamental constraint on monetary policy design. Permitting beliefs to depart from those consistent with rational expectations equilibrium breaks the tight link between policy rates and long-term expectations, even when long-term interest rates are determined by the expectations hypothesis of the yield curve. Because bond prices are excessively sensitive to short-term interest rates, the central bank faces an intertemporal trade-off which results in optimal policy being less aggressive relative to rational expectations. More aggressive policy leads to sub-optimal volatility in long-term interest rates and aggregate demand through standard intertemporal substitution effects. These effects are quantitatively important over the Great Inflation and Great Moderation periods of US monetary history.

Keywords :Optimal Monetary Policy, Learning Dynamics, Expectations Stabilization, Great Inflation

JEL Codes: E32, D83, D84

*The authors particularly thank Greg Kaplan, Eric Leeper and our discussant Leonardo Melosi. Also thanked are seminar and conference participants at the Reserve Bank of Australia, University of Sydney, Australian National University, Riksbank conference on “Deflation: Causes, Consequences and Policies”, the 2015 “Tsinghua-CAEPR Conference” and the 2018 NBER Summer Institute Dynamic Equilibrium Models Working Group. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of Dallas or the Federal Reserve System. The usual caveat applies.

[†]University of Texas, Austin. E-mail: eusepi73@yahoo.com.

[‡]Federal Reserve Bank of Dallas. E-mail: marc.giannoni@dal.frb.org.

[§]The University of Melbourne. E-mail: bruce.preston@unimelb.edu.au.

I INTRODUCTION

In models frequently used for policy evaluation it is not the current interest rate, but anticipated movements in future interest rates that are central to aggregate demand management. Movements in the term structure of interest rates are linked by arbitrage relationships. Through the appropriate choice of current interest rates, good policy seeks to have expectations evolve in a way that achieves the most desirable short-run trade-off between inflation and the output gap. And in rational expectations analyses, the control of expectations is precise and strikingly effective: in many models it is possible to stabilize the macroeconomy completely in response to movements in the real natural rate of interest. This Divine Coincidence result appears to be quite general, arising in the canonical New Keynesian model, medium-scale models for policy evaluation, and heterogeneous agent models.¹ Obvious questions emerge: to what extent does this policy advice depend on rational expectations? Is monetary policy compromised when current interest-rate movements are not precisely transmitted to various longer-term interest rates relevant to economic decision making?

Loose control of long-term interest rates can arise from two sources: movements in term premia and movements in long-term interest-rate expectations. This paper studies the latter, adopting a modeling approach guided by two types of empirical evidence. First, Crump, Eusepi, and Moench (2015) provide evidence of low-frequency movements in various survey data on long-term expectations. These patterns are evident in both nominal and real variables, and display strong correlation with short-term surprise movements of these same series. Second, Gurkaynak, Sack, and Swanson (2005), Crump, Eusepi, and Moench (2015) and Steinsson and Nakamura (2017) document evidence that monetary disturbances have sizable and significant effects on long-term nominal and real rates of interest. Both types of evidence are difficult to reconcile with standard models used for monetary policy evaluation, which typically impose strong restrictions on the long-run behavior of dynamics: expectations of inflation and the real interest rate are constant. And because these models have limited internal propagation mechanisms, surprise movements in the policy rate tend to have effects on aggregate dynamics that are resolved over the near to medium-term.

¹See, for example, Clarida, Gali, and Gertler (1999), Justiniano, Primiceri, and Tambalotti (2013) and Debortoli and Gali (2017).




Figure 1: **Drift in Expectations.**

The figure shows long-term expectations data from Blue Chip Economics and Financial. The red circles and diamonds are the average 1-10 and 5-10 years inflation expectations. The blue diamonds are the difference between these data and the corresponding expectations for the three-month Treasury Bill rate.

Figure 1 underscores the challenge data pose for modelling, plotting measures of long-term inflation expectations along with long-term real interest-rate expectations. Long-run inflation expectations exhibit a pronounced downward trend over the 1980s and 1990s, the Great Moderation period, with some cyclical variation, before stabilizing around 2000. In contrast, real rate expectations have no evident trend, though display large cyclical variation, rising in booms, and falling in recessions. The fact that both nominal and real interest rates exhibit variation suggests multiple sources of low-frequency movement in these data. Though not shown, strong correlations with short-run forecast errors also reveal dependence on macroeconomic conditions.

We build a New Keynesian model that is consistent with these facts. Following Kozicki and Tinsley (2001) and Eusepi and Preston (2018a) agents have imperfect knowledge about the long run. This reflects fundamental uncertainty about the inflation target and future production possibilities. They forecast long-run outcomes using a model with ‘shifting end-

points'. Econometric filtering attributes part of observed forecast errors to low-frequency developments in the macroeconomy. Long-run beliefs respond to short-run forecast errors as documented in the data. Because agents base decisions on subjective beliefs which differ from objective beliefs, the model has self-referential dynamics — see Marcet and Sargent (1989). For this reason, despite the expectations hypothesis of the term structure being a core building block of the model, with distorted beliefs long-run bond prices are excessively sensitive to surprise movements in short rates. Consistent with Sinha (2016), the model can account for the rejections of the expectations hypothesis documented in Campbell and Shiller (1991).

In a special case of our empirical model, we solve for optimal Bayesian decisions. We then characterize optimal monetary policy assuming the central bank minimizes the welfare-theoretic loss function knowing the true dynamics implied by firm and household behavior. In general the Divine Coincidence does not emerge under optimal policy, even though an otherwise equivalent model with rational expectations would predict joint stabilization of inflation and the output gap in response to demand shocks. The result arises because distorted beliefs about long-run interest rates represent a fundamental constraint on stabilization policy. Optimal policy is less aggressive than under rational expectations, because large movements in short-term interest rates induce low-frequency movements in interest-rate beliefs. Because long-term bond prices display excess sensitivity, such movements translate into large movements in long-term real rates and instability in aggregate demand.

Formally, we show that the aggregate demand equation is a binding constraint on the central bank's optimal policy problem. For beliefs which are sufficiently sensitive to short-run forecast errors, the Lagrange multiplier on this constraint is strictly positive. Under rational expectations, it would be zero. It is because aggregate demand is a constraint that policy has diminished capacity to respond. This model property distinguishes our analysis from other contributions on optimal policy design, and specifically those finding the Divine Coincidence holds even in very general settings — see Justiniano, Primiceri, and Tambalotti (2013) and Debortoli and Gali (2017). That optimal policy is less aggressive in response to aggregate disturbances under imperfect knowledge relative to rational expectations also contrasts with Orphanides and Williams (2005a), Ferrero (2007) and Molnar and Santoro

(2013). The different conclusions arise from assumptions about the transmission mechanism of monetary policy. These earlier papers assume that only current interest rates matter for aggregate demand, rather than the entire path of future expected one-period rates, as in the New Keynesian model. We show that the results are not because optimal control theory exploits the true structure of the economy, or because the central bank is assumed to have rational expectations: the insight holds for simple rules — it is the transmission mechanism of interest-rate policy that is pivotal. Indeed, if the central bank could control aggregate demand directly, or if interest-rate beliefs were not distorted (households had rational interest rate beliefs), then Divine Coincidence would be restored.

The final part of the paper establishes that these effects are quantitatively relevant in US data. An advantage of our modeling assumptions is the model has a state-space representation that is linear. This permits use of Bayesian likelihood-based methods for statistical inference using standard time series, as well data on survey expectations to discipline beliefs. The sample spans the Great Inflation and the Great Moderation periods, periods during which data exhibit substantial low-frequency movement. The model fits the data well and provides a plausible historical narrative. Monetary policy errors in the late 1960s and early 1970s have a substantial and significant effect on macroeconomic outcomes, despite being independent and identically distributed over time. Excessively loose monetary policy generates higher inflation which gets entrenched in inflation and nominal interest-rate expectations. These effects, which cause further increases in inflation, are highly persistent. Importantly, the model-theoretic output gap provides an account of real activity that accords with conventional thinking, clearly identifying recessionary periods, and exhibiting high correlation with other independent measures, such as the CBO’s output gap measure.

Counterfactual experiments demonstrate that while there is evidence monetary policy mistakes contribute to the Great Inflation, complete stabilization of aggregate demand is infeasible. Indeed, the optimal policy under imperfect knowledge still permits substantial fluctuations in real activity, moderating the size of recessions to only a small degree relative to the benchmark policy. In contrast, reflecting the relative importance nominal and real rigidities in goods and labor markets, the optimal policy stabilizes wage inflation to a larger degree than goods price inflation. A rational expectations analysis of the model corrob-

rates the findings of Justiniano, Primiceri, and Tambalotti (2013) — the Divine Coincidence approximately holds. Imperfect knowledge represents a fundamental constraint on policy.

The paper proceeds as follows. Section 2 lays out a medium-scale New Keynesian model with features required for a plausible account of aggregate data. Section 3 develops the theory of beliefs. Section 4 provides theory of optimal policy in a special case of the model. The results and a simple example build intuition for the analysis of optimal policy in the empirical model. Section 5 estimates the model and discusses basic properties. Finally, section 6 performs a number of counterfactual experiments to isolate core mechanisms and the central result on optimal policy.

II A MEDIUM-SCALE NEW KEYNESIAN MODEL

This section develops a version of the New Keynesian model widely used for monetary policy analysis. The principle modeling innovation concerns the treatment of expectations formation. This feature, and wanting a tightly-specified empirically plausible model, dictated assumptions on scale. Further details on the microfoundations can be found in Woodford (2003) and Giannoni and Woodford (2004).

Firms. A continuum of monopolistically competitive firms $f \in [0, 1]$ each produce differentiated goods, $Y_t(f)$, using the linear production technology in composite labor services, $N(f)$,

$$Y_t(f) = A_t [Z_t N_t(f)]$$

where Z_t , labor-augmenting technological progress, evolves deterministically as $Z_t = \gamma Z_{t-1}$, with $\gamma > 1$, and A_t denotes a stationary technology shock

$$\log A_t = \rho_a \log A_{t-1} + \sigma_a \varepsilon_t^a$$

where ε_t^a is IID $N(0, 1)$, $\sigma_a > 0$, and $0 < \rho_a < 1$. Each firm faces a demand curve

$$Y_t(f) = \left(\frac{P_t(f)}{P_t} \right)^{-\theta_{p,t}} Y_t$$

where $\theta_{p,t} > 1$, the elasticity of substitution across differentiated goods, follows an exogenous

process

$$\log \left(\frac{\theta_{p,t}}{\theta_p} \right) = \rho_{\theta_p} \log \left(\frac{\theta_{p,t-1}}{\theta_p} \right) + \sigma_{\theta_p} \varepsilon_t^{\theta_p}$$

where $\varepsilon_t^{\theta_p}$ is IID $N(0, 1)$, $\sigma_{\theta_p} > 0$, $0 < \rho_{\theta_p} < 1$ and $E[\theta_{p,t}] = \theta_p$.

Following Calvo (1983) and Yun (1996) a fraction of firms $0 < \xi_p < 1$ cannot optimally choose their price, but reset it according to the indexation rule

$$P_t(f) = P_{t-1}(f) \pi_{t-1}^{\iota_p}$$

where $\pi_t = P_t/P_{t-1}$ is the inflation rate, and $0 < \iota_p < 1$. The remaining fraction of firms choose a price $P_t(f)$ to maximize the expected discounted value of profits

$$\hat{E}_t^f \sum_{T=t}^{\infty} \xi_p^{T-t} Q_{t,T} \Gamma_T^f(f)$$

where the stochastic discount factor, $Q_{t,T} = \beta^{T-t} \lambda_T / \lambda_t$, values future profits

$$\Gamma_T^f(f) = Y_T(f) \left((1 - \tau_f) \frac{P_t(f)}{P_T} \left(\frac{P_{T-1}}{P_{t-1}} \right)^{\iota_p} - \frac{W_T}{A_T P_T Z_T} \right)$$

for constant sales revenue tax τ_f , and λ_t the marginal value of household wealth. The conditional expectations of firms, \hat{E}_t^f , is discussed below.

Households. A continuum of households $i \in [0, 1]$ maximize intertemporal utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[\frac{C_{H,T}(i)^{1-\sigma}}{1-\sigma} - \varphi_T \int \frac{N_T(i, j)^{1+\phi^{-1}}}{1+\phi^{-1}} dj \right]$$

where

$$C_{H,t}(i) = \frac{C_t(i)}{A_t} - b \frac{C_{t-1}(i)}{A_{t-1}}$$

with $\sigma, \phi > 0$, $0 < b < 1$. Each household comprises a large family, whose members $j \in [0, 1]$ supply specialized labor, $N(i, j)$, to the production of each differentiated good j . The large family assumption insures each household member against labor market risk from nominal wage contracting. The dis-utility of labor supply shock is a stationary exogenous process

$$\log \left(\frac{\varphi_t}{\varphi} \right) = \rho_{\varphi} \log \left(\frac{\varphi_{t-1}}{\varphi} \right) + \sigma_{\varphi} \varepsilon_t^{\varphi}$$

ε_t^{φ} is IID $N(0, 1)$, $\sigma_{\varphi} > 0$, $0 < \rho_{\varphi} < 1$ and $E[\varphi_t] = \varphi$. The conditional expectations of

households, \hat{E}_t^i , is discussed below.

The household's flow budget constraint is

$$C_t(i) + \frac{B_t(i)}{P_t} \leq R_{t-1}\pi_t^{-1} \frac{B_{t-1}(i)}{P_{t-1}} + (1 - \tau_w) \int \frac{W_t(j)}{P_t} N_t(i, j) dj + \Gamma_t^f - T_t + T_t^w + T_t^f$$

where: R_t is the gross one-period nominal interest rate; $B_t(i)$ holdings of one-period nominal government debt; Γ_t^f dividend payments net of sales taxes; τ_w the labor income tax rate whose proceeds are rebated lump-sum to households as T_t^w ; T_t^f the lump-sum rebate of sales revenue taxes; and T_t lump-sum taxes.² Household's optimal consumption and portfolio choice must also satisfy the No-Ponzi condition

$$\lim_{T \rightarrow \infty} \hat{E}_t^i \left(\prod_{s=0}^{T-t} R_{t+s} \pi_{t+s}^{-1} \right)^{-1} \frac{B_T(i)}{P_T} \geq 0.$$

Households have market power in the supply of differentiated labor inputs.³ The demand for labor type j by firm f is

$$N_t(j, f) = \left(\frac{W_t(j)}{W_t} \right)^{-\theta_{w,t}} N_t(f) \quad (1)$$

where

$$N_t(f) = \left[\int_0^1 N_t(j, f)^{\frac{\theta_{w,t}-1}{\theta_{w,t}}} dj \right]^{\frac{\theta_{w,t}}{\theta_{w,t}-1}} \quad \text{and} \quad W_t = \left[\int_0^1 W_t(j)^{1-\theta_{w,t}} dj \right]^{\frac{1}{1-\theta_{w,t}}}$$

define the composite labor input used in production and the associated wage rate. The elasticity of demand across differentiated labor inputs satisfies the exogenous process

$$\log \left(\frac{\theta_{w,t}}{\theta_w} \right) = \rho_{\theta_w} \log \left(\frac{\theta_{w,t-1}}{\theta_w} \right) + \sigma_{\theta_w} \varepsilon_t^{\theta_w}$$

and $\theta_{w,t} > 1$ and $\varepsilon_t^{\theta_w}$ is IID $N(0, 1)$, $\sigma_{\theta_w} > 0$, $0 < \rho_{\theta_w} < 1$ and $E[\theta_{w,t}] = \theta_w$. This shock is set to zero in the current version of the paper. Following Erceg, Henderson, and Levin (2000) a fraction of household members $0 < \xi_w < 1$ cannot optimally reset their wage, but follow the

²The assumptions on tax policy ensure an efficient steady state level of output.

³The assumption is interpreted as follows. For each type of labor, which is sourced from all households, there is an employment agency that has market power. See Giannoni and Woodford (2004) and Justiniano, Primiceri, and Tambalotti (2013).

indexation rule

$$W_t(j) = W_{t-1}(j) \pi_{t-1}^{\iota_w} \gamma \quad (2)$$

for $0 < \iota_w < 1$. For the remaining fraction, ξ_w , each member j of household i , choose optimally their nominal wage, $W_t(j)$, to maximize

$$\hat{E}_t^i \sum_{T=t}^{\infty} (\xi_w \beta)^{T-t} \left[Q_{t,T}(i) \frac{W_t(j)}{P_T} \left(\frac{P_{T-1}}{P_t - 1} \right)^{\iota_w} \frac{Z_{T-1}}{Z_{t-1}} N_T(i) - \varphi_T \frac{N_T(j)^{1+\phi^{-1}}}{1 + \phi^{-1}} \right]$$

subject to (1).⁴

Government Policy. The central bank implements monetary policy using the interest rate rule

$$R_t = (R_{t-1})^{\rho_R} \left[R(P_t/P_{t-1})^{\phi_\pi} X_t^{\phi_x} \right]^{1-\rho_R} m_t$$

where $\phi_\pi, \phi_x \geq 0$, R the steady-state gross interest rate, and $X_t = Y_t/Y_t^n$ denotes the model-theoretic output gap, where Y_t is the level of output and Y_t^n the natural rate of output in a flexible-price equilibrium of the model. Interest-rate policy exhibits inertia and responds to deviations of inflation and output gap from steady-state levels. The steady-state inflation rate is zero; $\log m_t = \sigma_m \varepsilon_t^m$ denotes a mean-zero IID monetary shock.

To give focus to how learning dynamics constrain monetary policy, we assume fiscal policy is Ricardian, and that this is understood by agents. Eusepi and Preston (2018) show that in general learning will imply departures from Ricardian equivalence, with holdings of the public debt perceived as net wealth. The associated wealth effects on aggregate demand can be sizable, which impairs the standard intertemporal substitution channel of monetary policy. We also assume that agents know the tax rules in place, including the rebate of sales and income taxes. Together these assumptions imply agents do not need to forecast various taxes and that debt will not have monetary consequences. This permits focus on how belief distortions affect long-term interest rates and monetary policy design, understanding that imperfect knowledge about fiscal and monetary policy both serve to complicate inflation policy. With this in mind, we consider an economy with zero government debt and balanced

⁴Members supplying labor of type j , being represented by an employment agency, re-optimize at the same time in all households i .

budget policy

$$T_t = G_t$$

where exogenous government purchases satisfy

$$\log \left(\frac{G_t}{G} \right) = \rho_G \log \left(\frac{G_{t-1}}{G} \right) + \sigma_G \varepsilon_t^G$$

where ε_t^G is IID $N(0, 1)$, $\sigma_G > 0$, $0 < \rho_G < 1$ and $E[G_t] = G$. Motivated by empirical fit we follow Smets and Wouters (2007) and permit correlation between government purchases and technology shocks.

Market Clearing and Equilibrium. We consider a symmetric equilibrium in which all households are identical, even though they do not know this to be true. Given that households have the same initial asset holdings, preferences, and beliefs, and face common constraints, they make identical state-contingent decisions. Similarly, all firms having the opportunity to re-optimize choose an identical re-set price. Equilibrium requires all goods, labor and asset markets to clear providing the restrictions

$$\begin{aligned} \int_0^1 C_t(i) di + G_t &= \int_0^1 Y_t(f) df \\ \int_0^1 \int_0^1 N_t(i, j) didj &= \int_0^1 N_t(f) df \end{aligned}$$

and

$$\int_0^1 B_t(i) di = 0$$

with initial condition $B_{-1}(i) = 0$. Given exogenous processes $\{G_t, \theta_{p,t}, \theta_{w,t}, m_t, Z_t, A_t, \varphi_t\}$, equilibrium then is a sequence of prices $\{P_t, W_t, R_t\}$ and allocations $\{C_t, N_t, Y_t, T_t, T_t^w, T_t^f, \Gamma_t^f\}$ satisfying individual optimality — detailed in the appendix — and market clearing conditions.

III BELIEFS

The appendix shows a first-order approximation to optimal decisions and market clearing conditions give aggregate dynamics

$$A_0 z_t = \sum_{s=1}^3 A_s \hat{E}_t \sum_{T=t}^{\infty} \lambda_s^{-(T-t)} z_{T+1} + A_4 z_{t-1} + A_5 \varepsilon_t \quad (3)$$

where the vector z_t collects all model variables, the vector ε_t collects exogenous innovations and the matrices A_i , for $i \in 1, \dots, 5$, collect relevant model coefficients. This representation holds for arbitrary beliefs, including rational expectations. Dynamics depend on a set of projections into the indefinite future, reflecting the intertemporal decision problems solved by households and firms. The projected variables are those macroeconomic objects taken as given and beyond the control of each decision maker. Firms must forecast real wages and goods price inflation; households must forecast goods price inflation, wage inflation, the real wage, nominal interest rates, and aggregate demand. The discount factors λ_s are the model's unstable eigenvalues, so that the infinite sums encode the usual forward recursion to suppress the effects of explosive roots.

An assumption on belief formation closes the model. We make a number of choices to ensure tractability in estimation and optimal policy exercises. Specifically, we analyze a belief structure that delivers a linear state-space representation of the model so standard likelihood methods can be employed. At the same time, these choices ensure a linear-quadratic optimal policy problem.

Subjective beliefs. Consistent with the assumption of a symmetric equilibrium, each agent has a common forecasting model

$$z_t = S\bar{\omega}_t + \Phi z_{t-1} + e_t \quad (4)$$

$$\bar{\omega}_t = \rho \bar{\omega}_{t-1} + u_t \quad (5)$$

where Φ is a matrix to be discussed; $0 \leq \rho \leq 1$ a parameter; e_t and u_t IID with $R = E[e_t e_t']$ and $Q = E[u_t u_t']$. The vector $\bar{\omega}_t$ is an unobserved state, capturing imperfect knowledge about the conditional mean of the process z_t . For example, when forecasting inflation, the unobserved state represents an estimate of the inflation target; when forecasting real vari-

ables it reflects fundamental uncertainty about long-term production possibilities. We refer to these terms as low-frequency drift, drift in beliefs, or distorted beliefs. The matrix S is a selection matrix which determines which low-frequency drift is relevant for each macroeconomic variable z_t . The beliefs nest rational expectations as a special case: $\bar{\omega}_t = \bar{\omega}_{t-1} = 0$ when $Q = 0$ — that is the prior belief about the variance-covariance matrix of the drift terms is zero.

The forecasting model implies conditional expectations satisfy

$$E_t z_{t+n} = \Phi^n z_t + \sum_{j=0}^n \Phi^j S \rho^{n-j} \bar{\omega}_t.$$

Medium to long-term forecasts are determined by two components: the first term is the conventional auto-regressive impact of the current state. The second term captures the effects of drifting beliefs on conditional expectations. The empirical work resolves an identification question: which component is more important for projections? For the model to explain the properties of survey data requires either highly persistent exogenous shocks, or highly persistent low-frequency movements in beliefs. We present evidence in support of the latter. In the special case $\rho = 1$ we have an example of a shifting end-point model in the language of Kozicki and Tinsley (2001). Beliefs then satisfy

$$\lim_{n \rightarrow \infty} E_t z_{t+n} = (I - \Phi)^{-1} S \bar{\omega}_t. \quad (6)$$

Objective Beliefs. Given an estimate of the unobserved state, ω , we can evaluate expectations required for optimal decisions as

$$E_t \sum_{T=t}^{\infty} \lambda_s^{-(T-t)} z_{T+1} = F_0(\lambda_s) S \omega_t + F_1(\lambda_s) z_t$$

where $F_0(\lambda_s)$ and $F_1(\lambda_s)$ are composites of structural parameters and eigenvalues λ_s . The structural equations (3) then provide

$$\begin{aligned} Z_t &= \left(A_0 - \sum_{j=1}^3 A_s F_1(\lambda_s) \right)^{-1} \left[\sum_{j=1}^3 A_s F_0(\lambda_s) S \omega_t + A_4 z_{t-1} + A_5 \varepsilon_t \right] \\ &= T(\Phi^*) S \omega_t + \Phi^* z_{t-1} + \Phi_\varepsilon^* \varepsilon_t \end{aligned}$$

where

$$\begin{aligned}\Phi^* &\equiv \left(A_0 - \sum_{j=1}^3 A_s F_1(\lambda_s) \right)^{-1} A_4 \\ \Phi_\varepsilon^* &\equiv \left(A_0 - \sum_{j=1}^3 A_s F_1(\lambda_s) \right)^{-1} A_5\end{aligned}$$

represent a fixed point of the beliefs (4).⁵ We therefore assume that agents understand the true dynamics of aggregate variables up to the unobserved mean. This preserves linearity of aggregate belief dynamics and gives focus to imperfect knowledge of conditional means.⁶

Drifts in beliefs are encoded into the intercept of the true data-generating process, and represent the only difference between subjective and objective beliefs in the model. That beliefs affect the true data-generating process, which in turn affects beliefs, is an example of what Marcet and Sargent (1989) call self-referentiality. When $T(\Phi^*) = I$ beliefs are perfectly validated by the data, generating a self-confirming equilibrium — see Sargent (1999). If $T(\Phi^*) = 0$ we have rational expectations. For intermediate cases, beliefs are partially self-confirming. Such beliefs present a challenge for stabilization policy. Eusepi and Preston (2018b) reviews relevant theory, showing good policy limits self-referential dynamics.

Subjective belief updating. Beliefs are updated using the recursion

$$\begin{aligned}\omega_{t+1} &= \rho\omega_t + \rho P_t (P_t + R)^{-1} S' F_t \\ P_{t+1} &= \rho^2 P_t - \rho^2 P_t (P_t + R)^{-1} P_t + Q\end{aligned}$$

where the matrix P_t is the mean square error associated with the estimate ω_{t+1} . The vector F_t denotes the current prediction error

$$F_t = (z_t - S\omega_{t-1} - \Phi^* z_{t-1}).$$

Following Sargent and Williams (2005), we make the following simplifying assumptions. Rescale the posterior estimate using $P_t = \Xi_t R$ and use the approximation $(I + \Xi_t)^{-1} \simeq I$ for

⁵Formally an example of the method of undetermined coefficients.

⁶Eusepi and Preston (2011, 2018a, 2018b) adduce theoretical and empirical evidence that together demonstrate learning about intercepts generates empirically relevant variation and creates policy challenges. Learning about the coefficients Φ would make the filtering problem and the state-space representation of the model non-linear.

small Ξ_t to give

$$\begin{aligned}\omega_{t+1} &= \rho\omega_t + \rho\Xi_t S' F_t \\ \Xi_{t+1} &= \rho_t^2 \Xi - \rho_t^2 \Xi_t \Xi_t + Q R^{-1}\end{aligned}$$

Study the steady state of this filter assuming prior beliefs satisfy the restriction $Q = g^2 R$ for scalar g . Under these assumptions the belief updating equation becomes

$$\omega_{t+1} = \rho\omega_t + \rho\alpha S' F_t$$

where $\Xi = \alpha I$ and $0 < \alpha < 1$ is a function of the parameters ρ and g . In the special case $\rho = 1$, $\alpha = g$.

The restriction on prior beliefs about low and high-frequency components of data is important to policy exercises in the sequel. Because we study counterfactuals in which the central bank implements optimal policy conditional on knowing beliefs, we want beliefs to be endogenous to the policy framework (at least to some extent). As the policy regime changes the transmission of exogenous disturbances and therefore R , scaling the prior variance Q in proportion ensures low-frequency effects of prior beliefs don't change in relative importance. The signal-to-noise ratio is policy invariant.

Evaluating the forecast error implies beliefs are updated as

$$\begin{aligned}\omega_{t+1} &= \rho\omega_t + \rho\alpha S' (z_t - S\omega_{t-1} - \Phi^* z_{t-1}) \\ &= [\rho + \alpha S' (T(\Phi^*) - I) S] \omega_t + \alpha S' \Phi_\varepsilon^* \varepsilon_t.\end{aligned}\tag{7}$$

Subsequent estimation and policy evaluation exercises require beliefs to be stationary. This implies a restriction on the matrix

$$\rho + \alpha S' (T(\Phi^*) - I) S$$

whose eigenvalues determine the evolution of the first-order difference equation in beliefs.

State-space representation. Finally, combining aggregate dynamics with beliefs, provides the linear state-space representation of the model

$$Z_t = F(\Theta) Z_{t-1} + Q(\Theta) \varepsilon_t$$

where Θ defines the set of model parameters with

$$F(\Theta) = \begin{bmatrix} \Phi^* & T(\Phi^*) S\rho & T(\Phi^*) S\alpha \\ 0 & \rho I & \alpha I \\ 0 & S'[T(\Phi^*) - I] S\rho & S'[T(\Phi^*) - I] S\alpha \end{bmatrix}$$

and

$$Z_t = \begin{bmatrix} z_t \\ \omega_t \\ S'F_t \end{bmatrix} \text{ and } Q(\Theta) = \begin{bmatrix} \Phi_\varepsilon^* \\ 0 \\ S'\Phi_\varepsilon^* \end{bmatrix}.$$

This permits standard likelihood-based estimation.

IV INTERTEMPORAL TRADE-OFFS UNDER OPTIMAL POLICY

This section frames basic conceptual issues, theory and intuition that is relevant to understand the empirical results. We do this in a special case of our empirical model, in which there is a frictionless labor market, and purely forward-looking optimal pricing and consumption decisions. The analysis of optimal monetary policy shows that in general the aggregate demand curve is a binding constraint on feasible choices of interest-rate paths, even though this is never true of an equivalent model with rational expectations. Drifting interest-rate beliefs confront policy with an intertemporal trade-off that limits what monetary policy can and cannot achieve when responding to any aggregate disturbance.

A THE POLICY PROBLEM

The policymaker minimizes the period loss function

$$L_t = \pi_t^2 + \lambda_x x_t^2$$

where $\lambda_x > 0$ determines the relative weight given to output gap versus inflation stabilization. This welfare-theoretic loss function represents a second-order approximation to household utility under maintained beliefs.⁷ Feasible sequences of inflation and the output gap must

⁷The absence of an output gap target x^* reflects the assumptions on tax policy which deliver an efficient steady-state level of output.

satisfy the aggregate demand and supply equations

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta) x_{T+1} - (R_T - \pi_{T+1} - r_T^n)] \quad (8)$$

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} (\xi_p \beta)^{T-t} [\kappa x_T + (1 - \alpha) \beta \pi_{T+1}] \quad (9)$$

where all variables are interpreted as log-deviations from steady state; x_t is the output gap; r_t^n the natural rate of interest an exogenous process; and $\kappa = (1 - \xi_p \beta)(1 - \xi_p)/\xi_p$ the slope of the short-run trade-off between inflation and the output gap.⁸ Optimal consumption and price-setting requires households and firms to forecast future output, interest rates and inflation. Assume agents have a forecasting model of the form (4) and (5), with

$$z_t = \begin{bmatrix} \pi_t \\ x_t \\ R_t \end{bmatrix} \text{ and } \omega_t = \begin{bmatrix} \omega_t^\pi \\ \omega_t^x \\ \omega_t^R \end{bmatrix}$$

and where $\Phi = 0$ and $\rho = 1$ to give a shifting end-point model.⁹ This belief assumption with decisions (8) and (9) gives the optimal Bayesian solution to the model, and is an example of what Adam and Marcet (2011) call internal rationality.

Subject to aggregate demand and supply, and the evolution of beliefs, the central bank solves the problem

$$\min_{\{x_t, \pi_t, R_t, \omega_t\}} E_t^{RE} \sum_{T=t}^{\infty} \beta^{T-t} L_T \quad (10)$$

taking as given initial beliefs, ω_{-1} . Assume that the central bank has rational expectations and has complete information about the true structural relations describing household and firm behavior. Interpret this as a best-case scenario. To the extent that learning dynamics impose constraints on what the central bank can achieve, these difficulties will only be more acute with limited information. Moreover, the nature of these constraints might also inform the choice of less sophisticated approaches to monetary policy.

⁸Derivation of these expressions assume a unity elasticity of intertemporal substitution, and infinite Frisch elasticity of labor supply.

⁹Under rational expectations, because the model is purely forward looking, the minimum state variables solution is a linear function of aggregate disturbances. We therefore assuming a belief structure consistent with this solution.

The first-order conditions are described in the appendix. Because beliefs are state variables there is no distinction between optimal commitment and discretion. The policy maker can only influence expectations through current and past actions — not through announced commitments to some future course of action. There is a further important difference to a rational expectations analysis of this model: the aggregate demand schedule is generally a binding constraint on feasible state-contingent choices over inflation and the output gap under learning. That is the Lagrange multiplier attached to (8) is positive for beliefs that are sufficiently sensitive to current forecast errors.

This property of optimal control under arbitrary beliefs arises from the structure of aggregate demand. Even if the optimal policy problem determines unique paths for inflation and the output gap, and, therefore, expectations about future values of these variables, current interest-rate policy still depends on beliefs about future interest-rate policy. And for arbitrary beliefs it need not be feasible to choose a bounded interest-rate sequence. Beliefs are a state variable so that subjective beliefs do not in general coincide with the objective probabilities implied by the economic model, in contrast to a rational expectations analysis. This means the aggregate demand equation is necessarily a constraint on what a central bank can achieve, because it takes appropriate account of the effects of interest-rate choices on interest-rate beliefs. A concrete example will be given later in this section.

The first-order conditions constitute a linear rational expectations model.¹⁰ The system can be solved using standard methods. Using results from Giannoni and Woodford (2017), Eusepi, Giannoni, and Preston (2018) establish conditions on beliefs for a unique bounded rational expectations equilibrium. We develop that analysis further here, to provide the following result.

Proposition 1. *Let $\bar{g} = \frac{(1-\alpha\beta)(\lambda_x+\kappa^2)}{\lambda_x(1-\beta)+\kappa^2}$. For beliefs $g \in (0, \bar{g})$ that satisfy either $g < 2(1-\beta)$ or $g > \beta^{-1}-\beta$ the optimal policy problem has a unique bounded solution. When $g < 2(1-\beta)$ the aggregate demand constraint is not binding, and the associated Lagrange multiplier is equal to zero. When $g > \beta^{-1}-\beta$ the aggregate demand constraint is binding, and the associated Lagrange multiplier is strictly positive.*

Proof. See the appendix. □

¹⁰In an innovative study, Molnar and Santoro (2013) explore optimal policy under learning in a model where only one-period-ahead expectations matter to the pricing decisions of firms. Gaspar, Smets, and Vestin (2006) provide a global solution to the same optimal policy problem but under a more general class of beliefs.

This result formalizes the central insight of the paper. When long-term interest rate beliefs are sufficiently sensitive to short-run forecast errors, aggregate demand limits the movements in interest rates. The central bank has imprecise control of long-term interest rates, even though the model satisfies the expectations hypothesis of the term structure. Belief distortions prevent changes in short-term rates being efficiently transmitted to long-term rates relevant for aggregate demand. These effects are shown to be quantitatively relevant in the sequel.

A further implication concerns a special case of beliefs. When the gain coefficient converges to zero the optimal policy coincides with optimal discretion under rational expectations. This result is intuitive: for small gains beliefs are almost never revised. Because policy cannot influence beliefs, which is precisely the assumption of optimal discretion, dynamics will correspond to those predicted by optimal discretion. For sufficiently small gains, policy is well approximated by rational expectations equilibrium analysis, and the central bank will have precise control of long-term inflation expectations.

Corollary 1. *In the special case $g \rightarrow 0$ optimal policy will give the same dynamic responses to disturbances as optimal discretion under rational expectations.*

This type of result has been discussed by Sargent (1999), Molnar and Santoro (2013) and Eusepi, Giannoni, and Preston (2018).

B A SIMPLE EXAMPLE

To appreciate the implications of aggregate demand as a constraint confronting policy, consider a central bank faced only with IID shocks to the natural rate r_t^n , and private agent beliefs initially consistent with rational expectations equilibrium so that $\omega_{t-1} = 0$. Because initial forecasts satisfy $E_t z_T = 0$ for all $T > t$, period t equilibrium is determined by the aggregate demand and supply curves (8) and (9) which simplify to

$$\pi_t = \kappa x_t \quad \text{and} \quad x_t = -(R_t - r_t^n).$$

Given a disturbance to the natural rate of interest, complete stabilization is possible in period t . Nominal interest-rate policy must track the natural rate, $R_t = r_t^n$, giving $\pi_t = x_t = 0$. But

this implies subsequent movements in long-run interest-rate beliefs according to

$$\omega_t^R = \omega_{t-1}^R + g(r_t^n - \omega_{t-1}^R).$$

The next-period's stabilization problem — and every subsequent period — is given by the pair of equations

$$\begin{aligned}\pi_{t+1} &= \kappa x_{t+1} \\ x_{t+1} &= -(R_{t+1} - r_{t+1}^n) - \frac{1}{1-\beta} \beta \omega_t^R\end{aligned}$$

where the final term in the demand equation reflects the restraining effects of long-term interest rates on aggregate demand. These effects operate through the expectations hypothesis of the term structure of interest rates. Because short-term forecast errors are in part perceived to signal a permanent change in interest rates, the model predicts excess sensitivity of bond prices. Low-frequency drift in beliefs therefore shift the entire yield curve which has implications for current monetary policy. Complete stabilization of inflation and the output gap is again possible by having nominal interest rates track long-run expectations and the natural rate of interest.

But is this interplay sustainable? Imposing full stabilization, $x_{t+1} = \pi_{t+1} = 0$, the aggregate demand constraint defines the implicit policy rule

$$R_{t+1} = r_{t+1}^n - \frac{\beta}{1-\beta} \omega_t^R \tag{11}$$

in every period t . Optimal policy not only responds to natural-rate disturbances, but also movements in long-term interest rates, driven by expectations.¹¹ Substituting into the updating rule for beliefs, ω_t^R , gives

$$\omega_{t+1}^R = \left(1 - \frac{\bar{g}}{1-\beta}\right) \omega_t^R + \bar{g} r_{t+1}^n$$

which is a first-order difference equation. Sustainable policy requires the dynamics of beliefs

¹¹The implied interest rates of a bond of any maturity can be shown to be a function of the long-term interest rate belief. This is an example of the expectations hypothesis of the term structure of interest rates.

to be stationary. The following restriction must hold

$$g < 2(1 - \beta).$$

For larger gains, stability is not feasible, implying beliefs and interest rates are explosive. This is not a permissible, or at least desirable, feature of optimal policy if only because the zero lower bound on interest rates obviates such solutions.¹²




Figure 2: **Volatility as a function of the constant gain.**

This figure shows the volatility of output and interest rates as a function of the constant gain \bar{g} . The welfare theoretic loss gives the volatility of the interest rate (red circles) and the output gap (blue triangles); while a policy maker with a concern of interest rate volatility delivers the interest rate shown by the black line, and the output gap given by the grey dashed line.

Figure 2 gives further insight, plotting the standard deviation of the output gap and interest rate as a function of the constant gain g under optimal policy. Assume the discount factor is $\beta = 0.995$; the frequency of price changes determined by $\xi_p = 0.8$; and the weight on output gap stabilization $\lambda_x = 0.05$. Under these assumptions there is relatively small

¹²While some might not object to nominal explosions, if beliefs about real activity depend on nominal interest-rate forecast errors, there would also be unbounded paths for real variables.

variation in inflation, so it matters little whether we plot the sum of the output gap and inflation variation or the output gap alone. Only variations in the natural rate, r_t^n , drive economic fluctuations. The figure describes outcomes under the welfare-theoretic loss (10), and under a loss function

$$L_t = \pi_t^2 + \lambda_x x_t^2 + \lambda_R R_t^2$$

that also penalizes volatility in the interest rate. Recall optimal discretion corresponds to the case $g = 0$. Under the standard loss function a knife-edge result obtains: for $g < 0.01$ the output gap is fully stabilized even if this induces substantial volatility in the interest rate. For large values of g , the policy maker loses the ability to stabilize the output gap. Feasible policy restricts variation in the policy rate, translating into increasing volatility in the output gap. If the policy maker has some preference for interest-rate stabilization, perhaps reflecting zero-lower bound considerations, then the increase in output volatility occurs continuously with the size of the gain. Even relatively small values of the gain lead to output gap volatility.

These exercises point to a fundamental property of optimal policy under long-term drift in expectations: current interest rates move less in response to disturbances when compared to a rational expectations analysis, including movements in the natural rate of interest.

Proposition 2. *In the model given by (8) and (9), Divine Coincidence will in general not hold even in a model with only disturbances to the natural real rate of interest.*

Aggressive changes in the stance of policy engender low-frequency movements in interest-rate beliefs, which can move long rates move too much, creating instability in demand. This constrains aggregate demand management policy. The inability of the central bank to stabilize both output gap and inflation in the face of aggregate demand shocks stems from agents' expectations about the policy rate. For example, suppose as in Molnar and Santoro (2013) the policymaker can directly control the output gap as the instrument of policy, and solves the problem

$$\min_{\{x_t, \pi_t, \omega_t^x, \omega_t^\pi\}} E_t^{RE} \sum_{T=t}^{\infty} \beta^{T-t} L_T$$

subject only to the Phillips curve (9), taking as given initial beliefs ω_{-1}^x and ω_{-1}^π . Equivalently, suppose interest-rate beliefs are anchored so that $\omega_t^i = 0$ for all t , giving households rational

expectation forecasts of interest rates. Then the Divine Coincidence holds, despite long-term drift in expectations about inflation and real activity.

Corollary 2. *Absent low-frequency drift in interest-rate beliefs, the central bank can directly control aggregate demand, and the Divine Coincidence holds.*

Proof. See the appendix. □

The result underscores the importance of central bank communications policy. Because agents face complex forecasting problems, resolving uncertainty about future interest-rate policy can promote stability. Eusepi and Preston (2010) provide a theoretical treatment of central bank communication in a closely related model, when monetary policy is implemented using simple rules.

C FURTHER IMPLICATIONS AND DISCUSSION

These findings contrast to earlier results on optimal monetary policy with non-rational expectations. Molnar and Santoro (2013) analyze optimal monetary in a model with learning in which

$$\begin{aligned} x_t &= \hat{E}_t x_{t+1} - \sigma^{-1} (R_t - E_t \pi_{t+1} - r_t^n) \\ \pi_t &= \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t \end{aligned}$$

are taken as the decision rules describing aggregate demand and supply, and beliefs described in section 4.1. They conclude that the Divine Coincidence holds in response to a disturbance in the natural rate, and that optimal policy should be more aggressive relative to rational expectations. Similar conclusions on the aggressive stance of policy have been documented by Bomfim, Tetlow, von zur Muehlen, and Williams (1997), Orphanides and Williams (2005) and Ferrero (2007).

What is the source of these distinct conclusions? The substantive difference concerns the transmission of monetary policy: aggregate demand in these models does not depend on interest-rate expectations, only the contemporaneous policy rate. The feed back effects of short-run forecast errors on long-term beliefs, and therefore policy choice, does not arise. The aggregate demand equation is not a constraint. This leads to a simpler policy design problem.

One might argue the result arises because the central bank is assumed to exploit knowledge of the true structure of the economy when determining interest rates — a quirk of optimal control theory. However, the consequences of the aggregate demand constraint are not specific to fully optimal policy. Figure 3 describes stability properties under constant-gain learning when monetary policy is implemented according to the Taylor rule

$$i_t = \phi_\pi \left(\pi_t + \tilde{\phi}_x x_t \right)$$

where $\phi_x \equiv \phi_\pi \tilde{\phi}_x$.¹³ Each of the three contours describe the stability frontier in the constant-gain and inflation response coefficient space. Parameter regions above a plotted contour indicate local instability of the equilibrium. Higher contours correspond to progressively weaker responses to the output gap. For many gain coefficients aggressive monetary policy is not desirable.

Indeed, even an arbitrarily large inflation response for some gain coefficients would not deliver stability. Rewriting the above Taylor rule as

$$\phi_\pi^{-1} i_t = \pi_t + \tilde{\phi}_x x_t$$

and taking the limit $\phi_\pi \rightarrow \infty$ gives the target criterion

$$\pi_t + \tilde{\phi}_x x_t = 0.$$

Evans and Honkapohja (2006), Woodford (2007) and Preston (2008) argue the target criterion approach to implementing policy has the advantage of being robust to alternative assumptions about belief formation.¹⁴ The above makes clear that for many gains such policies are infeasible under our theory of the transmission mechanism of interest-rate policy, even though they would deliver stability in the Molnar-Santoro model.

V ESTIMATION AND MODEL IMPLICATIONS

¹³This is an example of ‘robust learning stability’ proposed by Evans and Honkapohja (2009).

¹⁴This example delivers the target criterion under optimal discretion in the model of this section. Optimal commitment target criteria would be subject to similar concerns.




Figure 3: **Stability Frontiers.**

The figure shows stability frontiers corresponding to alternative Taylor rules. In particular (\bar{g}, ϕ_π) above the frontier correspond to locally unstable equilibria under constant-gain learning. The black solid line corresponds to the standard Taylor Rule. The solid (dashed) grey line corresponds to $\phi_x = \phi_x^*/2$ ($\phi_x = \phi_x^*/3$).

A ESTIMATION

The Data. To estimate model parameters we use thirteen US time series. Five are standard macroeconomic variables: the log-difference of the GDP deflator, the output gap (as measured by the Congressional Budget Office), the three-month Treasury-Bill interest rate, and, following Justiniano, Primiceri, and Tambalotti (2013), and two measures of nominal wage growth from NIPA and the BLS Establishment survey.¹⁵ The remaining eight time series are short and long-term professional forecasts of the three-month Treasury-bill rate and inflation. We use these series to discipline beliefs. For each of these two variables, the one-quarter and four-quarter ahead forecasts from the Survey of Professional forecasts measure short-term

¹⁵We use the CBO measure of the output gap to detrend output, not to fit the model-theoretic output gap.

forecasts. The mean one-to-ten years-ahead and the five-to-ten-years ahead forecasts from Blue Chip Economics and Financial measure long-term beliefs. Together these short and long-term data on expectations permit inference on the gain parameter.

The estimation uses quarterly data over the period 1964Q1 to 2007Q3. The end of the sample is chosen to exclude the period when the policy rate is at the zero lower bound on nominal interest rates. Short-term forecasts of inflation are available from 1968Q3; short-term forecasts of nominal interest rates from 1981Q3; long-term forecasts of inflation from 1979Q3 and long-term interest-rate forecasts from 1985Q1.

Observation Equation. Section 3 showed the model has a time-invariant state-space representation

$$Z_t = F(\Theta) Z_{t-1} + Q(\Theta) \varepsilon_t \quad (12)$$

where Θ is a vector of parameters and Z the state vector of variables, which include the perceived drifts. The measurement equation

$$Y_t = \mu_t(\Theta) + H_t(\Theta) Z_t + o_t$$

attaches ten measurement errors, o_t , to the eight survey forecasts and the two measures of the nominal wage growth. The vector μ_t contains the long-run mean of the observables. The matrix H_t and μ_t are time varying because of missing observations. We estimate the model using Bayesian inference.¹⁶

Calibrated Parameters. The quarterly growth rate in technical progress $\gamma = 1.04$ matches the average GDP per-capita growth over the sample. Elasticity of demand across differentiated goods and labor services, θ_p and θ_w , are both set equal to 5. The parameter ρ which determines the persistence in beliefs is 0.995.¹⁷ And the government spending-to-output ratio is $G/Y=0.16$.

Prior Distributions. Tables 1, 2 and 3 provide details on the priors. The priors for the exogenous shock processes are the same across variables. The persistence of the

¹⁶Details are in the appendix.

¹⁷The data suggest a unit root would be appropriate. However, to fit the steady-state real rate of interest requires a household discount factor quite close to unity. This makes model dynamics highly sensitive to shifting expectations. A value of ρ slightly below unity effectively discounts expectations, permitting jointly fitting expectations data and the steady-state real rate.

	Dist.	Prior		Posterior			
		Mean	Std	Mode	Mean	5%	95%
β	Gamma	0.500	0.100	0.205	0.202	0.176	0.228
σ	Gamma	1.00	0.600	9.45	8.80	6.81	11.3
ϕ_n	Gamma	0.500	0.100	0.404	0.365	0.248	0.506
b	Beta	0.350	0.100	0.560	0.590	0.485	0.684
ξ_w	Beta	0.500	0.100	0.912	0.896	0.866	0.923
ι_w	Beta	0.500	0.150	0.583	0.644	0.507	0.782
ξ_p	Beta	0.500	0.100	0.922	0.900	0.848	0.937
ι_p	Beta	0.500	0.150	0.071	0.101	0.040	0.182
ϕ_π	Normal	1.50	0.150	1.01	1.05	1.01	1.09
ρ_i	Beta	0.500	0.100	0.816	0.818	0.792	0.844
ϕ_x	Normal	0.120	0.050	0.162	0.178	0.129	0.233
$\bar{\pi}$	Gamma	0.500	0.100	0.581	0.588	0.463	0.715
g	Gamma	0.035	0.030	0.053	0.050	0.039	0.060

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

Table 1: **Prior and Posterior Distribution of Structural Parameters**

autocorrelated processes have a beta distribution with mean 0.5 and standard deviation 0.2; the standard deviations of the innovations and all measurement errors have an inverse-gamma distribution with mean 0.1 and standard deviation of 2. The priors to the parameters of the monetary policy reaction function are based on the Taylor rule. Given evidence in Hall (1988) and Ravina (2011), the inverse intertemporal elasticity of substitution, σ , has a gamma distribution with mean 1 and fairly large standard deviation of 0.6, while the degree of habit persistent has a beta prior with mean 0.35. Turning to optimal price and wage setting, the Calvo adjustment parameters, ξ_p and ξ_w , have prior means which imply contracts have an average duration of one-half a year. The parameters capturing price and wage indexation, ι_p and ι_w , have means 0.5. Following Slobodyan and Wouters (2012), the constant-gain coefficient has a gamma distribution with mean 0.035 and standard deviation 0.03.

Posterior Distributions. Tables 1, 2 and 3 also show the mean, the mode and the 5 and 95 percentiles of the posterior distribution of parameters. The data are informative. The mean inflation rate is estimated to be 2.4% per annum. The estimated policy parameters are quite different from prior values. In particular, the inflation response coefficient is only slightly above unity. This reflects our choice to capture dynamics of the Great Inflation and

	Prior			Posterior			
	Dist.	Mean	Std	Mode	Mean	5%	95%
ρ_{θ_p}	Beta	0.500	0.100	0.265	0.240	0.153	0.323
ρ_{φ}	Beta	0.500	0.100	0.494	0.506	0.396	0.606
ρ_g	Beta	0.500	0.100	0.885	0.881	0.848	0.911
ρ_a	Beta	0.500	0.100	0.975	0.975	0.966	0.983
σ_{θ_p}	InvGamma	0.100	2.00	0.215	0.220	0.190	0.250
σ_{φ}	InvGamma	0.100	2.00	0.023	0.027	0.016	0.040
σ_g	InvGamma	0.100	2.00	0.859	0.877	0.789	0.976
σ_m	InvGamma	0.100	2.00	0.195	0.198	0.179	0.219
σ_a	InvGamma	0.100	2.00	0.701	0.578	0.348	0.927
$\sigma_{g\gamma}$	Beta	0.500	0.200	0.072	0.245	0.061	0.542

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

Table 2: **Prior and Posterior Distribution of Shock Processes**

	Prior			Posterior			
	Dist.	Mean	Std	Mode	Mean	5%	95%
$\sigma_{o,\pi^{1Q}}$	InvGamma	0.100	2.00	0.115	0.116	0.100	0.132
$\sigma_{o,R^{1Q}}$	InvGamma	0.100	2.00	0.041	0.043	0.035	0.051
$\sigma_{o,\pi^{4Q}}$	InvGamma	0.100	2.00	0.098	0.100	0.087	0.114
$\sigma_{o,R^{4Q}}$	InvGamma	0.100	2.00	0.078	0.078	0.061	0.096
$\sigma_{o,R^{510Y}}$	InvGamma	0.100	2.00	0.078	0.084	0.064	0.109
$\sigma_{o,\pi^{510Y}}$	InvGamma	0.100	2.00	0.031	0.033	0.025	0.042
$\sigma_{o,R^{110Y}}$	InvGamma	0.100	2.00	0.070	0.074	0.061	0.089
$\sigma_{o,\pi^{110Y}}$	InvGamma	0.100	2.00	0.027	0.028	0.020	0.037
σ_{o,w_1}	InvGamma	0.100	2.00	0.563	0.573	0.521	0.628
σ_{o,w_2}	InvGamma	0.100	2.00	0.325	0.337	0.305	0.372
Γ	Normal	1.00	0.500	@ 0.816	0.812	0.781	0.843

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

Table 3: **Prior and Posterior Distribution of Shock Processes**

Great Moderation under a single monetary policy regime. The intertemporal elasticity of substitution is remarkably low, within the range 0.1 to 0.15. The Frisch elasticity of labor supply is estimated to be 0.4, broadly consistent with micro evidence. The price and wage stickiness parameters are estimated to be 0.9, implying a long duration of price contracts, common to most estimated New Keynesian DSGE models. However, the implied slope of the wage Phillips curve is an order of magnitude smaller than the price Phillips curve, with important implications for monetary policy. The learning gain g is estimated to be 0.05 which implies a short-term forecast error of 1 percent leads to a 5 basis point revision in

long-term beliefs. Moreover, the gain implies an observation that is five-years old receives a weight of about 15% percent. The shocks have lower persistence than usually found in DSGE models. This reflects the role of learning in soaking low-frequency variation in data. Measurement error on survey data is small, indicating a tight mapping between short-run forecast errors and long-term beliefs, which is central to our model.

B MODEL FIT

[TO BE ADDED]

C MODEL PREDICTIONS: POLICY MISTAKES AND THE GREAT INFLATION

Figure 4 provides basic model predictions. The top and middle panels show the model accounts for the survey data well, both at low and high frequencies. The tight 95 percent credible interval of the model's predictions for long-term inflation and real rate expectations (computed as the difference between long-term nominal interest rate and inflation expectations) reflects the small measurement error attached to survey data in estimation. The estimate of the welfare-theoretic output gap provides a characterization of business cycles that accords with conventional wisdom. The major recessions are all evident. The estimate has strong correlation with the CBO output gap measure, which was not used to estimate the output gap.¹⁸

An advantage of building and estimating a model which includes a behavioral theory of expectations formation is being able to study how structural shocks shape the evolution of the term structure of expectations. We can assess how disturbances drive short-term forecast errors, and how these errors get mapped into long-run beliefs. Figure 5 shows the historical contribution of each shock to a number of model variables, including short and long-term expectations. We interpret supply shocks as the combined effects of labor supply and technology shocks. The contribution of initial conditions is also shown. To assist interpretation the shock decomposition uses an annual frequency.

The dominant driver of nominal variables and their expectations are monetary policy shocks. During the 1970s much of the variation in inflation, and short and long-term fore-

¹⁸The CBO output gap was only used to de-trend output, not to compute the difference between actual de-trended output and the natural rate of output.

casts of inflation is due to excessively loose monetary policy, with evidence of favorable supply shocks providing some restraining influence on inflation dynamics very early in the sample. The importance of adverse supply shocks is evident in the period before the Volcker disinflation and favorable supply shocks clearly restrain inflation during the 1990s, consistent with the view of substantial productivity improvements over this period.

Interest rates and their long-term expectations tell a similar story, with a somewhat more pronounced role of supply shocks during the Volcker disinflation to explain persistently high interest rates. The role of demand shocks is small in comparison, though more evident than for inflation. This reflects the links forged by interest-rate policy. De-trended output displays a prominent role for demand shocks in both booms and recessions. Monetary policy is clearly stimulatory during the early 1970s, and contractionary during the Volcker disinflation.

The central role of monetary policy shocks is somewhat surprising given they are IID. Subsequent counterfactual analysis shows the interaction of these nominal disturbances with expectations explains their prominent role. However, the shock decompositions provide a clear narrative. In the late 1960s monetary policy was overly stimulatory: there are a sequence of large negative monetary policy shocks. Strong real economic activity and rising inflation engender higher one-year-ahead and one-to-ten year ahead inflation expectations. Because the monetary policy reacts weakly to inflation (recall the policy response coefficient is 1.05), these monetary policy shocks are propagated over time through expectations drift and self-referentiality: rising long-term inflation expectations generate higher inflation (through strategic complementarity in price-setting and higher demand from lower projected real interest rates) which feeds back into expectations. A more aggressive monetary policy would restrain these developments. For this reason monetary policy shocks become an increasingly important determinant of interest rates in the latter part of the 1970s. It is also why these shocks appear to be both a source of stimulus to de-trended output and rising interest rates (the contributions reflect both current and past innovations). Policy mistakes in the early part of the sample and weak adjustment of nominal interest rates in response to inflation entrench low-frequency movements in inflation expectations which generate a significant part of the Great Inflation of the 1970s.

Figure 6 provides further insight and corroborates the role of monetary policy shocks in

generating low-frequency movements in inflation and self-referential dynamics. The figure plots the variance decomposition of inflation and interest-rate data, for both these variables, and expectations at the four-quarter and one-to-ten year horizons. At short-horizons most of the variance of inflation is accounted for by price markup shocks, while for interest rates monetary policy shocks explain most variation. In contrast, for the four-quarter ahead and long-term expectations, supply shocks play the dominate role at short-horizons. Over longer horizons, all expectations data display a similar profile: supply shocks become progressively less important, while the influence of monetary disturbances grows. In the case of one-to-ten-year inflation expectations, monetary shocks ultimately explain some 70% of variation. Because low-frequency movements in beliefs drive actual inflation outcomes, the relative importance of monetary shocks to inflation variance grows with the horizon of the decomposition, from zero to roughly 25%. These patterns are consistent with self-referential dynamics, and an endogenous inflation trend which propagates into the future the effects of past policy error to the broader macroeconomy.

VI OPTIMAL POLICY COUNTERFACTUALS

This section provides a quantitative evaluation of the trade-off confronting policy from belief distortions. Counterfactual analysis shows that optimal policy is unable to jointly stabilize, output, wage and price inflation in contrast to a rational expectations analysis of the model. Importantly, disturbances that result in efficient movements in output are a non-trivial source of variation under optimal policy.

A THE LOSS FUNCTION

Under arbitrary beliefs, the period welfare-theoretic loss is

$$L_t = \lambda_p (\pi_t - \iota_p \pi_{t-1})^2 + \lambda_w (\pi_t^w - \iota_w \pi_{t-1}^w)^2 + \lambda_x (x_t - \bar{b}x_{t-1})^2$$

where the weights

$$\begin{aligned} \lambda_p &= \frac{\theta_p \kappa_p^{-1}}{\theta_p \kappa_p^{-1} + \theta_w \kappa_w^{-1}} \\ \lambda_w &= \frac{\theta_w \kappa_w^{-1}}{\theta_p \kappa_p^{-1} + \theta_w \kappa_w^{-1}} \end{aligned}$$

and

$$\lambda_x = \frac{\phi^{-1}\sigma(1 - \beta b)^{-1}}{\theta_p \kappa_p^{-1} + \theta_w \kappa_w^{-1}}$$

determine the relative priority given to stabilizing prices, wages and output, and are functions of the slopes of the wage and price Phillips curves, κ_p and κ_w . Finally, the parameter $\bar{b} \leq b$ is a function of structural parameters. Details are found in Giannoni and Woodford (2004).

The derivation of the second-order approximation to household utility is valid under both rational expectations and learning. The architecture of the loss function reflects well-understood sources of inefficiency which arise from monopoly power in goods and labor markets. In our model, equilibrium price and wage markups can vary for two reasons. First, exogenous time variation in the elasticity of demand across differentiated goods and labor services shifts firms' and workers' desired markups. Second, staggered price setting in goods and labor markets means prevailing prices depart from the optimal flexible-price levels, which lead to endogenous variation in markups in response to all aggregate disturbances. Optimal policy mitigates this second source of variation due to nominal rigidities. By stabilizing endogenous variation in markups, policy reduces cross sectional dispersion in price and wage setting, and the associated inefficiencies in supply of goods and labor.

B THE POLICY PROBLEM

We consider two versions of the policy problem: one assuming the benchmark model under learning; another assuming a counterfactual economy in which agents have rational expectations. In each case the policy maker takes as given the set of equations characterizing private sector behavior. The policy maker has only short-term interest rates as an instrument of policy.

Regardless of the assumptions on expectations formation, these policy problems are standard, and it is straightforward to solve for the optimal state-contingent path of interest rates that maximizes welfare. However, to assist interpreting the differences in policy across belief structures, we instead look for optimal policy within a class of interest-rate rules. This permits direct comparison of policy rule coefficients, and therefore inference on how drift in long-term interest rates constrain optimal policy. The approach also resolves the question of how to implement optimal policy. It is well known that purely fundamentals-based rules are

prone to indeterminacy of rational expectations equilibrium and expectations instability under learning. While the optimal target criterion can be derived under rational expectations, it is rather complicated, involving a large number of leads and lags of various endogenous variables. To facilitate transparency we work with simple rules.

Under learning the central bank minimizes the expected discounted loss

$$E_t^{RE} \sum_{T=t}^{\infty} \beta^{T-t} L_T \quad (13)$$

subject to

$$\begin{aligned} Z_t &= F(\Theta; \phi) Z_{t-1} + Q(\Theta; \phi) \varepsilon_t \\ R_t &= \rho_R R_{t-1} + \phi_\pi (\pi_t - \iota_p \pi_{t-1}) + \phi_w (\pi_t^w - \iota_w \pi_{t-1}) + \phi_x x_t \end{aligned}$$

by choice of policy parameters $\phi = \{\rho_R, \phi_\pi, \phi_w, \phi_x\}$. The first equation implicitly drops the interest rate from the true data-generating process (12). Under rational expectations the central bank minimizes (13) subject to

$$\begin{aligned} z_t &= \Phi^*(\phi) z_{t-1} + \Phi_\varepsilon^*(\phi) \varepsilon_t \\ 0 &= \pi_t - \iota_p \pi_{t-1} + \phi_w (\pi_t^w - \iota_w \pi_{t-1}) + \phi_x (x_t - x_{t-1}) \end{aligned}$$

by choice of policy parameters $\phi = \{\phi_\pi, \phi_x\}$, where the first equation gives the true data-generating process under the special case of rational expectations: $T(\Phi^*) = 0$. The choice of a slightly modified rule reflects the fact that it approximates the optimal commitment policy very well. Ideally we would use identical rules in each problem, though it ultimately doesn't matter for results. Because we can compute the optimal policy for the rational expectations economy, we know we cannot do better than the optimal commitment policy under the timeless perspective. And the proposed policy gets quite close to this. In the case of learning, we searched a wide variety of rules.

[ADD POLICY COEFFICIENTS]

C THE COUNTERFACTUALS

We make the following assumptions when analyzing optimal policy in the counterfactual economies. The economy experiences the same sequence of shocks and starts in the same initial state as estimated in the benchmark model. While contemplating different policy regimes, we assume that agents know the new transition dynamics associated with the regime. The thought experiment is interpreted as one in which agents have inhabited the regime since the distant past. The question of how to design the optimal transition from one regime to another is left for future research. Furthermore, we assume the gain coefficient is policy invariant. While we assume the perceived signal-to-noise ratio is invariant across policy regime (that is, agents perceive the same volatility of long-term drift relative to short-term disturbances), belief dynamics can be very different.

Figures 7 and 8 provide two counterfactual experiments to elucidate the core mechanisms of the model. Figure 7 shows the evolution of inflation expectations in a counterfactual economy with the same historical policy rule and the same shocks but where agents have rational expectations. The model is unable to account for the Great Inflation, despite having highly stimulatory policy during this period. Figure 8 shows the benchmark model in absence of monetary policy shocks. Once again, the model can't explain the Great Inflation. Together these figures highlight a central mechanism of the paper. IID monetary policy shocks interact with beliefs to generate lasting effects on aggregate dynamics. Moreover, they generate substantial movement in long-term interest-rate expectations, consistent with the excess sensitivity of bond prices in our model.

Figure 9 plots results for two economies under rational expectations. The first, the dashed red line, gives the dynamics predicted by the benchmark model when agents have rational expectations rather than imperfect knowledge. The second, the solid black lines, gives the dynamics of the counterfactual rational expectations model under the optimal policy. The top panel shows paths for the output gap. Under the benchmark policy rule, there are fairly wide variations in the output gap. In contrast the optimal policy delivers almost complete stabilization of the real economy. The middle and bottom panels reveal real stability does not imply nominal instability. In fact the optimal policy almost completely stabilizes the

wage variations, and provides greater stability of the inflation gap.¹⁹ That optimal policy stabilizes nominal wages more than goods prices reflects the relative slopes of the wage and price Phillips curves. Even though the frequency of price and wage adjustment are similar, the slopes are radically different, with the wage Phillips curve flatter by a factor of 10. This implies the weight on wage stabilization in the loss function is close to unity, and the weight on price stabilization close to zero.

These results replicate the findings of Justiniano, Primiceri, and Tambalotti (2013). Viewing historical data through the lens of a medium-scale model of the kind proposed by Smets and Wouters (2007), they show that once appropriate account is taken of low-frequency movements in hours data, as well as measurement error in wage data, there is little evidence of a fundamental trade-off confronting monetary policy. Exogenous variation in desired markups explains little variation in observed data. A striking implication is that observed fluctuations in the output gap, and events such as the Great Inflation, are the result of policy error.

Figure 10 casts a more positive light on the historical performance of the Federal Reserve. It plots dynamics under learning dynamics for both the benchmark model and the counterfactual economy assuming optimal policy. The top panel reveals complete stabilization of the output gap is infeasible. Moreover, output fluctuations are remarkably similar under both the benchmark and optimal policies: while there are some small differences in the early part of the sample, and in the depths of each recession, the basic tenor of the real activity implications of these policies are the same. Differences emerge in the middle and bottom panels showing the predictions for goods price inflation and wage inflation. As under rational expectations, optimal policy largely stabilizes wage inflation. And variations in good price inflation are reduced, but less so, again reflecting the relative weights assigned to these objectives in the welfare-theoretic loss, and the role of exogenous markup shocks in goods pricing.

The imperfect knowledge economy makes clear that stabilization policy may not be as effective as predicted by rational expectations analysis. There are limits to what optimal

¹⁹Note the predictions early in the sample for wages are driven by the estimated initial state — which would never have occurred under the counterfactual policy.

monetary policy can achieve. More subtle insights also emerge. There remains evidence the Great Inflation is due to policy error, with much more stabilization of inflation under the optimal policy in the mid to late 1970s. The optimal policy also reveals the advantages of stabilizing inflation and inflation expectations over this period: during the early 1980s optimal policy delivers falling inflation and rising output, despite the importance of demand shocks at this time ???. Good management of expectations locates the short-run trade-off in a desirable way. But beliefs ultimately represent a fundamental constraint on policy. For example, in the recession of early 2000s, optimal policy achieves higher output with the cost of higher inflation.

VII CONCLUSIONS

[TO BE ADDED]

A APPENDIX

[TO BE ADDED]

REFERENCES

- ADAM, K., AND A. MARCET (2011): “Internal Rationality, Imperfect Knowledge and Asset Prices,” *Journal of Economic Theory*, 146, 1224–1252.
- BOMFIM, A., R. J. TETLOW, P. VON ZUR MUEHLEN, AND J. WILLIAMS (1997): “Expectations, Learning and the Costs of Disinflation: Experiments using the FRB/US Model,” manuscript, Federal Reserve Board of Governors.
- CAMPBELL, J. Y., AND R. SHILLER (1991): “Yield Spreads and Interest Rate Movements: A Bird’s Eye View,” *Review of Economic Studies*, 58, 495–514.
- CLARIDA, R., J. GALI, AND M. GERTLER (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661–1707.
- CRUMP, R., S. EUSEPI, AND E. MOENCH (2015): “The Term Structure of Expectations and Bond Yields,” unpublished, Federal Reserve Bank of New York.
- DEBORTOLI, D., AND J. GALI (2017): “Monetary Policy with Heterogeneous Agents: Insights from TANK Models,” unpublished manuscript, Universitat Pompeu Fabra CREI.
- ERCEG, C. J., D. W. HENDERSON, AND A. T. LEVIN (2000): “Optimal Monetary Policy with Staggered Wage and Price Contracts,” *Journal of Monetary Economics*, 46(281-313).
- EUSEPI, S., M. GIANNONI, AND B. PRESTON (2018): “Some Implications of Learning for Price Stability,” *European Economic Review*, 106, 1–20.
- EUSEPI, S., AND B. PRESTON (2010): “Central Bank Communication and Macroeconomic Stabilization,” *American Economic Journal: Macroeconomics*, 2, 235–271.
- (2011): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, 101(6), 2844–2872.
- (2018a): “Fiscal Foundations of Inflation: Imperfect Knowledge,” *American Economic Review*, 108(9), 2551–89.

- (2018b): “The Science of Monetary Policy: An Imperfect Knowledge Perspective,” *Journal of Economic Literature*, 56(1), 3–61.
- EVANS, G. W., AND S. HONKAPOHJA (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.
- (2009): “Robust Learning Stability with Operational Monetary Policy Rules,” in *Monetary Policy under Uncertainty and Learning*, ed. by K. Schmidt-Hebbel, and C. Walsh, pp. 27–76.
- FERRERO, G. (2007): “Monetary Policy, Learning and the Speed of Convergence,” *Journal of Economic Dynamics and Control*, 39(9), 3006–3041.
- GASPAR, V., F. SMETS, AND D. VESTIN (2006): “Adaptive Learning, Persistence and Optimal Monetary Policy,” *Journal of the European Economics Association*, 4, 376–385.
- GIANNONI, M., AND M. WOODFORD (2004): “Optimal Inflation-Targeting Rules,” in *The Inflation-Targeting Debate*, NBER Chapters, pp. 93–172. National Bureau of Economic Research, Inc.
- GIANNONI, M. P., AND M. WOODFORD (2017): “Optimal target criteria for stabilization policy,” *Journal of Economic Theory*, 168(C), 55–106.
- GURKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, 95(1), 425–436.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2013): “Is There a Trade-Off between Inflation and Output Stabilization?,” *American Economic Journal: Macroeconomics*, 5(2), 1–31.
- KOZICKI, S., AND P. A. TINSLEY (2001): “Shifting endpoints in the term structure of interest rates,” *Journal of Monetary Economics*, 47(3), 613–652.

- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, (48), 337–368.
- MOLNAR, K., AND S. SANTORO (2013): “Optimal Monetary Policy When Agents Are Learning,” *European Economic Review*, 66, 39–62.
- ORPHANIDES, A., AND J. C. WILLIAMS (2005a): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation Targeting Debate*, ed. by B. S. Bernanke, and M. Woodford. University of Chicago Press.
- (2005b): “Inflation scares and forecast-based monetary policy,” *Review of Economic Dynamics*, 8(2), 498–527.
- PRESTON, B. (2008): “Adaptive Learning and the Use of Forecasts in Monetary Policy,” *Journal of Economic Dynamics and Control*, 32(4), 2661–3681.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.
- SARGENT, T. J., AND N. WILLIAMS (2005): “Impacts of Priors on Convergence and Escape Dynamics,” *Review of Economic Dynamics*, 8(2), 360–391.
- SINHA, A. (2016): “Learning and the Yield Curve,” *Journal of Money, Credit and Banking*, 48(2-3), 513–547.
- SLOBODYAN, S., AND R. WOUTERS (2012): “Learning in an estimated medium-scale DSGE model,” *Journal of Economic Dynamics and Control*, 36(1), 26–46.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: a Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- STEINSSON, J., AND E. NAKAMURA (2017): “High Frequency Identification of Monetary Non-Neutrality,” *Quarterly Journal of Economics*, 133(3), 1283–1330.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

——— (2007): “Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice,” unpublished, Columbia University.

YUN, T. (1996): “Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles,” *Journal of Monetary Economics*, 37, 345–370.

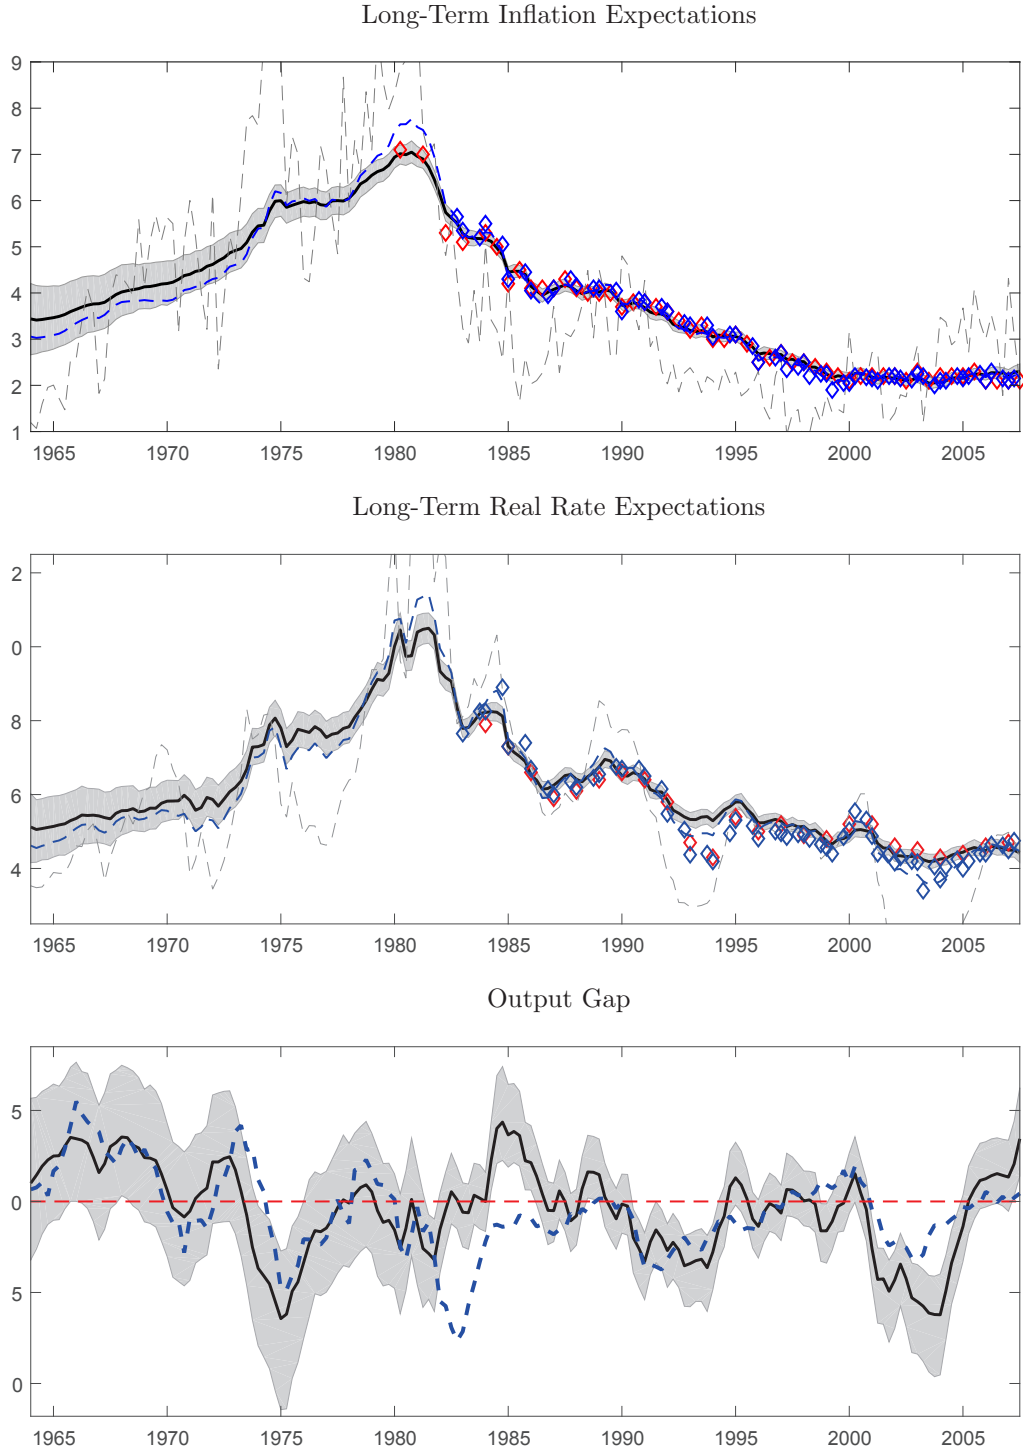


Figure 4: **Baseline Predictions.**

The top and middle panels show the evolution of long-term survey expectations data for inflation and the real rate of interest. Actual variable (dashed black), the two survey expectations measures (red and blue diamonds), the model implied 1-10 year average expectations (the dashed blue line); and the model implied 5-10 year average expectations with 95% posterior probability band (solid black line). The bottom panel shows the estimated output gap (solid black) with the CBO output gap (dashed blue).

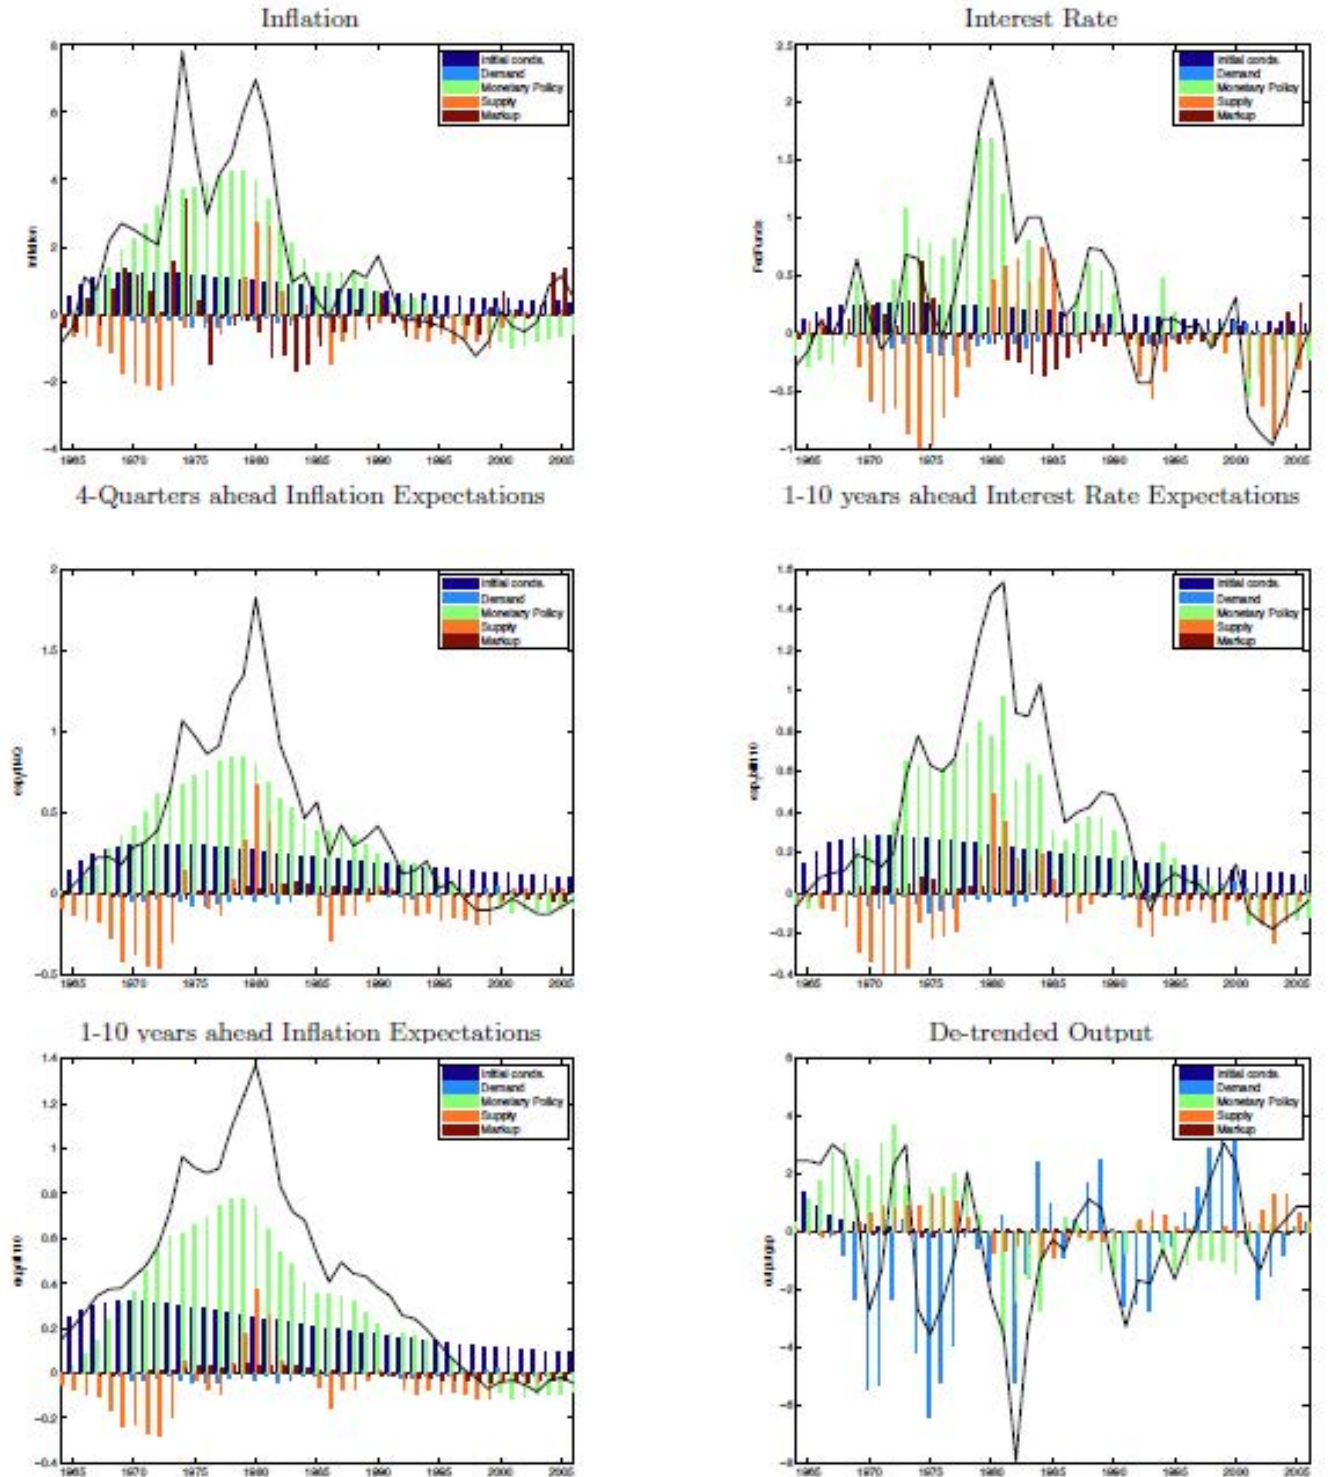


Figure 5: Historical Shock Decomposition

The panels show the decomposition of selected variables calculated at the posterior mode. Data are plotted at an annual frequency; inflation is expressed in term of a four-quarter average; interest rate; output gap and expectations are the fourth-quarter realizations (not annualized).

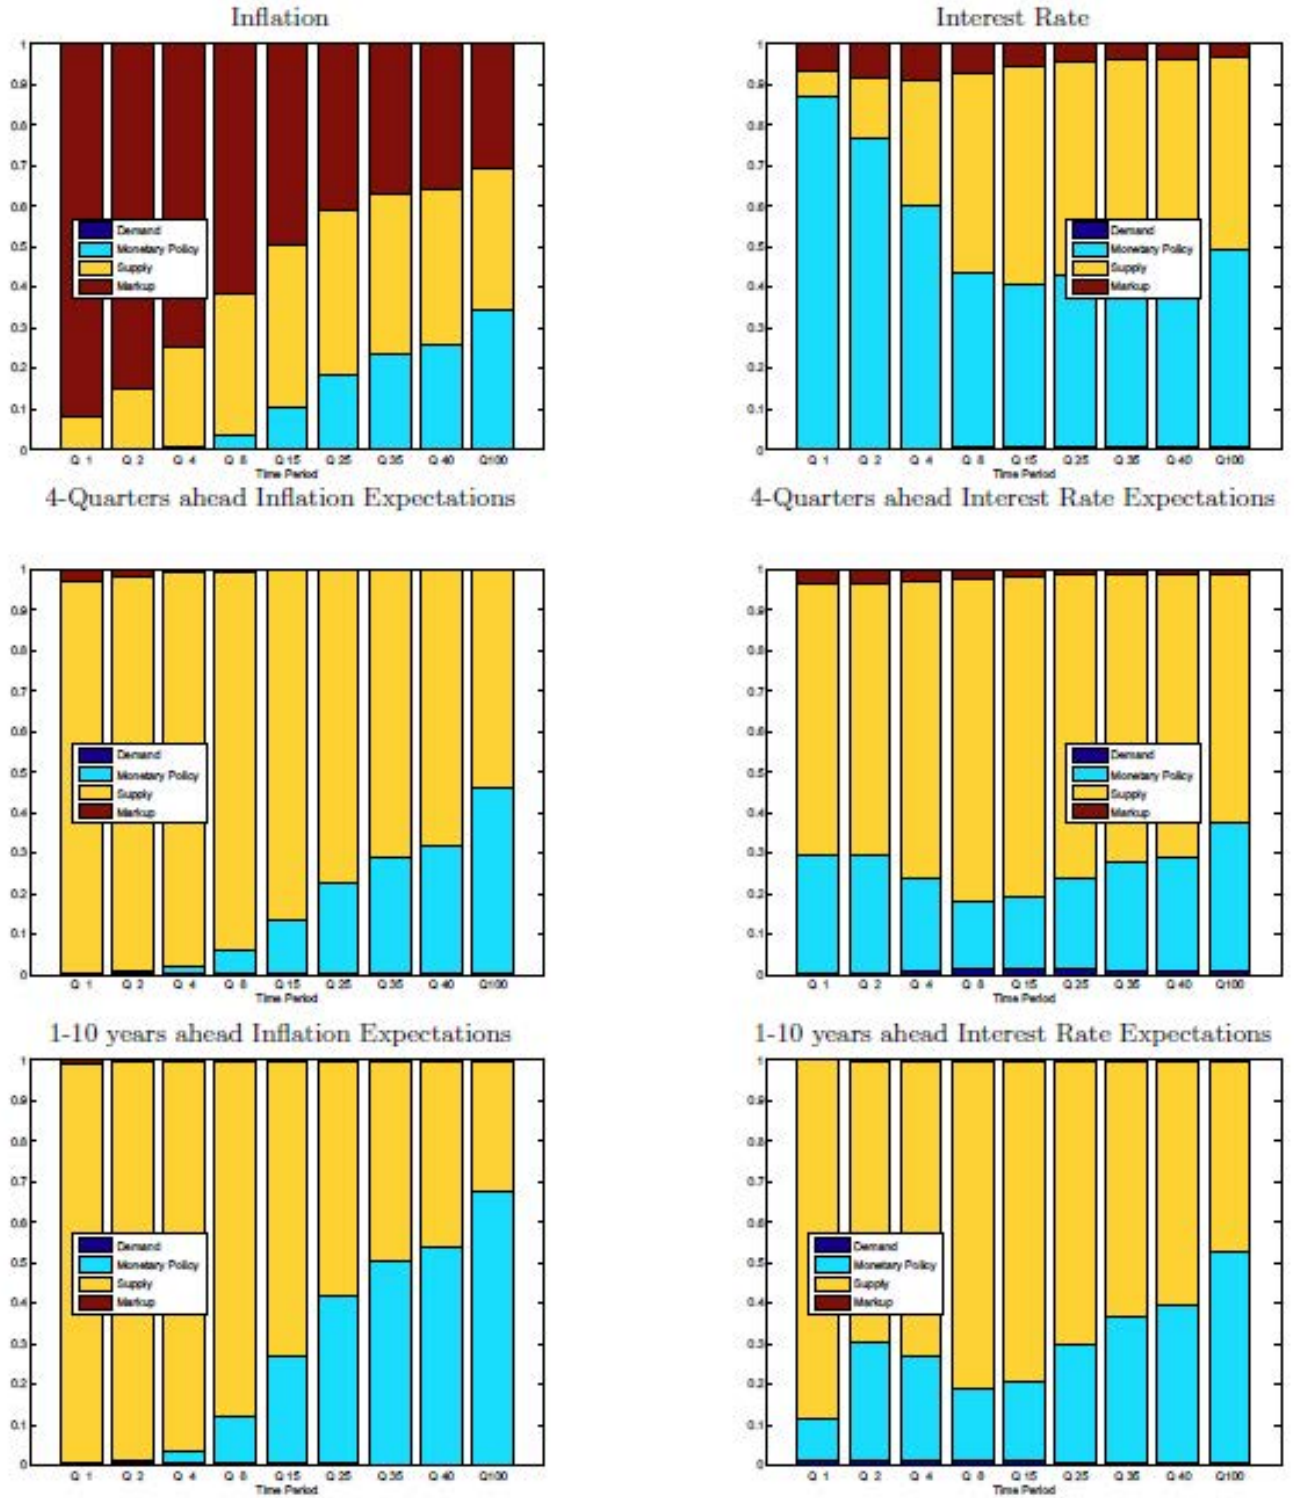


Figure 6: Variance Decomposition

The panels show the variance decomposition of selected variables calculated at the posterior mode.




Figure 7: **Long-term inflation expectations under rational expectations.**

The figure shows the counterfactual evolution of long-term inflation expectations (black solid line) under the assumption of rational expectations, compared with the data.




Figure 8: **Long-term inflation expectations without monetary policy shocks.**

The figure shows the counterfactual evolution of long-term inflation expectations (black solid line) in absence of monetary policy shocks, compared with the data.

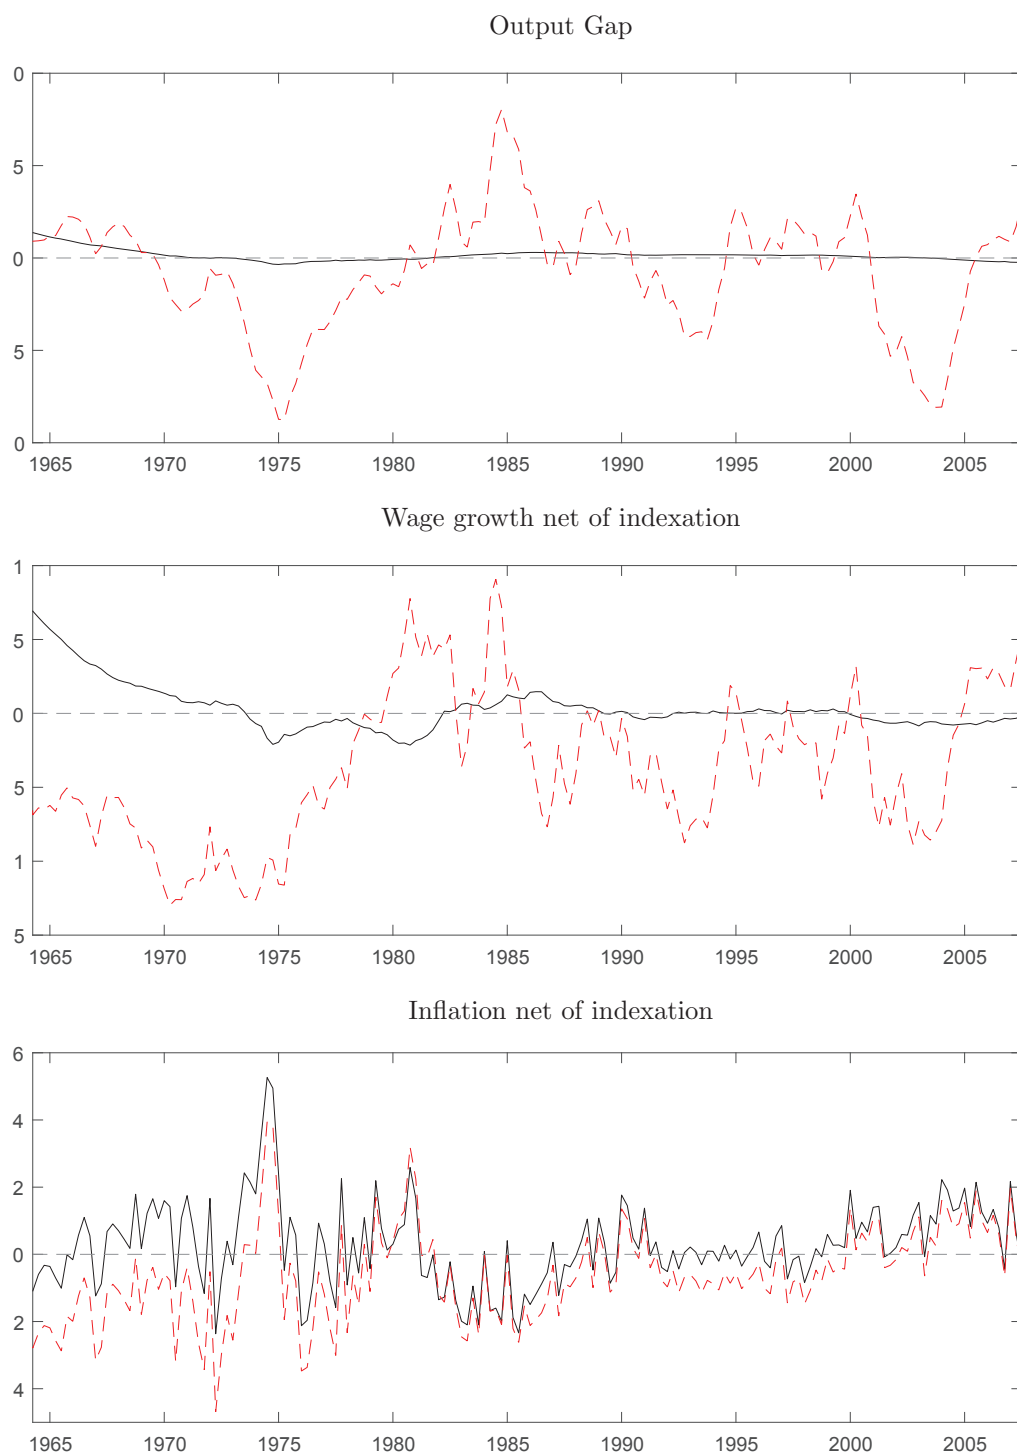


Figure 9: **Optimal Policy Under Rational Expectations**

Each panel shows the model dynamics under rational expectations with either the benchmark Taylor rule (dashed red line) or optimal monetary policy (solid black line).

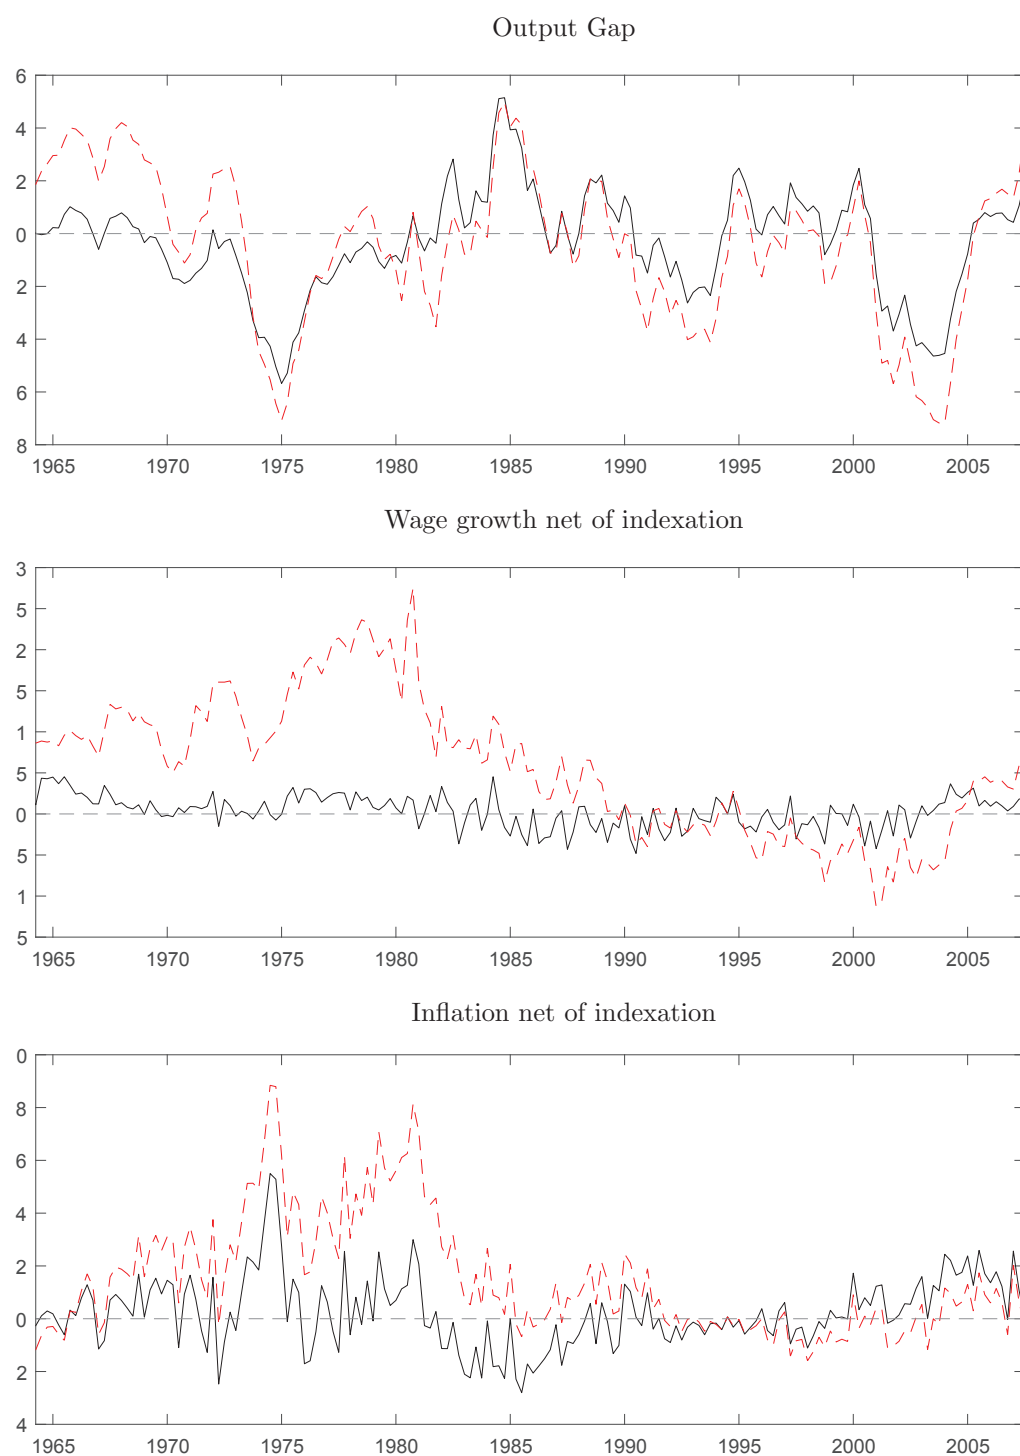


Figure 10: **Optimal Policy Under Learning.**

Each panel shows the model dynamics under learning with either the benchmark Taylor rule (dashed red line) or optimal monetary policy (solid black line).