

The Limits of Monetary Policy with Long-term Drift in Expectations*

Stefano Eusepi[†]

Marc Giannoni[‡]

Bruce Preston[§]

February 14, 2018

PRELIMINARY AND INCOMPLETE

Abstract

Survey data on long-horizon forecasts of inflation and interest rates exhibit low frequency drift. This drift, which is correlated with short-term forecast errors, is well-captured by models having time variation in limiting conditional forecasts. Standard structural macroeconomic models are inconsistent with these properties. This paper proposes a structural model for policy evaluation consistent with these facts and argues long-term drift in expectations impose constraints on policy design. In general a Central Bank cannot fully stabilize aggregate demand disturbances, in contrast to a rational expectations analysis of the model. Drifting beliefs are a distortion confronting policy, limiting the degree to which monetary policy optimally responds to evolving macroeconomic conditions. Empirical evidence for long-run drift is adduced and revealed to represent a quantitatively important trade-off for policy.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, or the Federal Reserve System.

[†]Federal Reserve Bank of New York. E-mail: stefano.eusepi@ny.frb.org.

[‡]Federal Reserve Bank of New York. E-mail: marc.giannoni@ny.frb.org.

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

1 Introduction

The modern theory of monetary policy emphasizes the management of expectations. In New Keynesian models frequently used for policy evaluation it is not so much the current interest rate, but anticipated movements in future interest rates that are central to aggregate demand management. Movements in current and future expected interest rates are linked through arbitrage relationships. Through the appropriate choice of current interest rates, good policy seeks to have these expectations evolve in a way that achieves the most desirable short-run trade-off between inflation and the output gap. An important question then is whether the efficacy of monetary policy is compromised when current interest-rate movements are not efficiently transmitted to various longer-term interest rates relevant to spending and pricing plans of agents in the economy. Is the potency of monetary policy diminished when there is imprecise control of interest-rate expectations? That this is a relevant practical concern is well-captured by the following quote from Bernanke (2004):

“[...] most private-sector borrowing and investment decisions depend not on the funds rate but on longer-term yields, such as mortgage rates and corporate bond rates, and on the prices of long-lived assets, such as housing and equities. Moreover, the link between these longer-term yields and asset prices and the current setting of the federal funds rate can be quite loose at times.”

This paper proposes a theory of imprecise control of long-term interest rates motivated by various properties of aggregate time series data. Figure 1 plots time series of expectations data. The top panel shows the evolution of expectations of the 3-month Treasury Bill rate at the one-quarter and four-quarter horizons, taken from the Survey of Professional Forecasters, and the 1-2 year-ahead and 5-10 year-ahead forecasts from the Blue Chip Economic Indicators. In addition to the general decline in the level of interest-rate expectations, all time series show a degree of cyclical sensitivity, with short-horizon expectations exhibiting significantly wider swings. Importantly, movements in long-term expectations are correlated with movements in short-term expectations — there appears to be drift in long-term expectations that depends on current macroeconomic conditions. The bottom panel juxtaposes the 5-10 year nominal interest-rate forecasts with a corresponding forecast of inflation, computed using the GDP deflator, adjusted for an estimate of the real rate of interest. The drift in the latter has

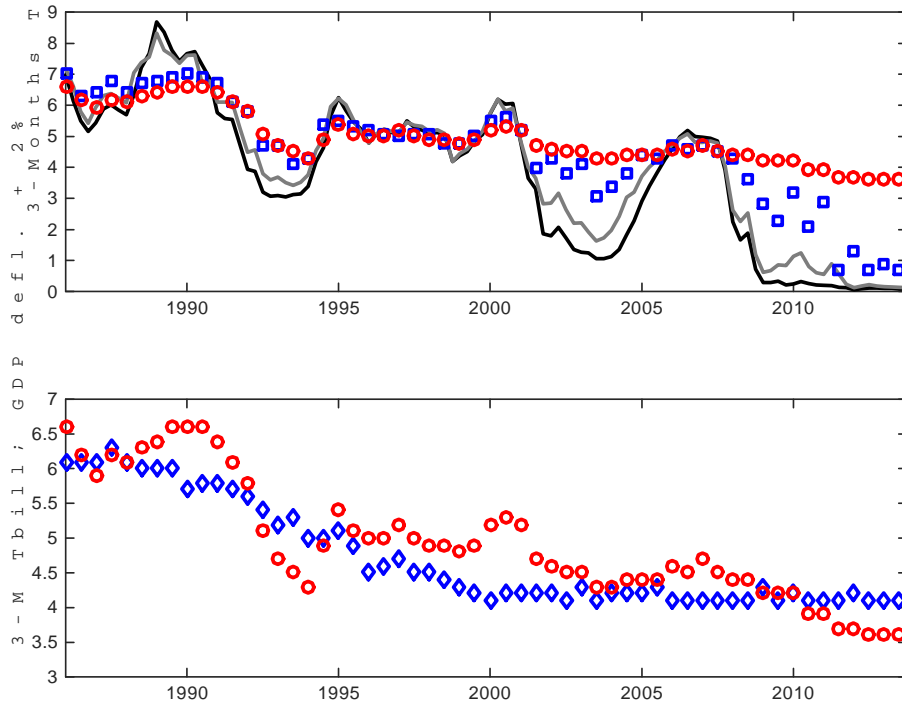


Figure 1: Survey data on 3-month Treasury Bill expectations. Panel 1: the one-quarter (solid black) and fourth-quarter ahead (solid grey) expectations from the Survey of Professional Forecasters, and the 1-2yr (blue square) and 5-10yr (red circle) forecast from BCEI. Panel 2: the 5-10yr Treasury bill forecast (red circle) and the GDP deflator forecast (blue diamond).

much less cyclical variation, though is quite distinct from that in nominal interest rates. This suggests that drift in interest-rate beliefs likely reflect both long-run nominal and real considerations.

Standard models used in macroeconomics are inconsistent with these observations. For example, the baseline New Keynesian model implies near-constant medium-term forecasts even when augmented with habit formation in consumption demand and inflation indexation price setting — see, for example, Gürkaynak, Sack, and Swanson (2005). Long-term outcomes for any macroeconomic variable are assumed to be known and fixed. In contrast, the survey data on expectations suggest beliefs about long-run outcomes exhibit drift over time. Evidence for such drift is also found in empirical models of inflation (see Stock and Watson, 2007; Cogley and Sbordone, 2008; Cogley, Primiceri and Sargent, 2010), the output gap (see Stock and Watson, 1989; Cogley and Sargent, 2005; Laubach and Williams, 2003) and nominal interest rates (see Kozicki and Tinsley, 2001; Gurkaynak, Sack and Swansson, 2005). The

question is whether this matters for monetary policy design?

This paper builds a canonical New Keynesian model consistent with these facts by letting households and firms have prior subjective beliefs that macroeconomic data contain low-frequency drifts. This belief structure implies a true data-generating process under which expectations and macroeconomic time series exhibit drift: the model is self-referential with beliefs affecting the true data-generating process which in turn affects beliefs. Beliefs are to some degree self-fulfilling. The priors, a primitive of the analysis, have clear economic interpretation: drift in nominal variables reflect uncertainty, or imperfect credibility, about the inflation target, while drift in real variables, such as wages, reflect uncertainty about long-run technological possibilities. These are example of “shifting end-points” discussed by Kozicki and Tinsley (2001). Estimates of the drift are obtained from a standard filtering problem, and revised as new data become available. Conditional on this belief structure the remaining model features are standard.

In this environment optimal monetary policy is characterized. The policy maker is assumed to know the true structural relations of the economy and also private agent beliefs. This is a best-case scenario. An important leitmotif of this analysis is that even when a policy maker is unrealistically assumed to possess such knowledge there are important limitations on what monetary policy can achieve. Unlike a rational expectations analysis of the model, the aggregate demand equation is a constraint when choosing interest rates to maximize private sector utility. As a consequence, the transmission mechanism of monetary policy is itself a constraint on policy, and in contrast to optimal policy under rational expectations, a policy maker cannot completely stabilize aggregate demand shocks. The Divine Coincidence does not hold in this model — see Blanchard and Galí (2007).

The mechanism underpinning this result is easily inferred. Households’ consumption demand depends on long-term interest rates which mediates changes in the policy rate. In response to a positive disturbance to the natural rate of interest, a policy maker would like to raise nominal interest rates by precisely the same magnitude. However, the resulting short-term forecast error of leads agents to revise upwards their estimate of long-term interest rates in the subsequent period, which further restrains aggregate demand. Aggressive adjustment of current interest rates presage excessive movements in long rates and macroeconomic volatility. The fact that long-term interest rates are not anchored presents a new intertemporal trade-off

confronting policy design. Optimal policy contends with this intertemporal distortion operating through beliefs by limiting the degree to which policy responds to current macroeconomic conditions.

The sequel adduces evidence on the empirical relevance of long-term drift in expectations and quantifies the magnitude of trade-off induced by such beliefs on the choice of policy. Using data on survey expectations to estimate prior beliefs identifies the existence of one nominal and one real drift. Importantly, long-term beliefs about inflation, interest rates and the output gap are sensitive to short-run forecast errors in each of these same three variables. A model with a constant gain is revealed to be substantially inferior. Consistently with the theoretical results, optimal policy in the empirical model responds less aggressively to both natural-rate and cost-push disturbances. As a result, the volatility of inflation and the output gap under optimal policy is substantial relative to a rational expectations analysis of the model. This underscores the intertemporal distortion imposed by beliefs on policy design is quantitatively relevant, shifting adversely the short-run trade-off between inflation and the output gap.

2 The Basic New Keynesian Model

This section develops a version of the canonical New Keynesian model widely used for monetary policy analysis. Further details on the microfoundations can be found in Woodford (2003) and Galí (2008).

A continuum of households i on the unit interval maximize utility

$$\hat{E}_t^i \sum_{T=t}^{\infty} \bar{C}_T \beta^{T-t} \left[(1 - \sigma)^{-1} c_T(i)^{1-\sigma} - \chi n_T(i) \right],$$

where $0 < \beta < 1$; $\sigma > 0$ and $\chi > 0$, by choice of sequences for consumption, $c_t(i)$, and labor supply, $n_t(i)$, subject to the flow budget constraint

$$c_t(i) + b_t(i) \leq (1 + i_t) \pi_t^{-1} b_{t-1}(i) + W_t n_t(i) / P_t + \Gamma_t(i)$$

and the No-Ponzi condition

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

The variable $b_t(i) \equiv B_t(i) / P_t$ denotes real bond holdings (which in equilibrium are in zero net supply), i_t the nominal interest rate, $\pi_t \equiv P_t / P_{t-1}$ the inflation rate, W_t is the hourly nominal

wage, $\Gamma_t(i)$ real dividends from equity holdings of firms and \bar{C}_T exogenous preference shifter. The operator \hat{E}_t^i denotes agents' subjective expectations, which might differ from rational expectations. The latter is defined by the operator \mathbb{E}_t .

A continuum of monopolistically competitive firms maximize profits

$$\hat{E}_t^j \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} [p_t(j) y_T(j) - W_T n_T(j)]$$

by choice of $p_t(j)$ subject to the production technology and demand function

$$y_T(j) = n_T(j) = (p_t(j) / P_T)^{-\theta_t} Y_T$$

for all $T \geq t$, with the elasticity of demand across differentiated goods an exogenous process satisfying $\theta_t > 1$; and exogenous probability $0 < \alpha < 1$ of not being able to reset their price in any subsequent period. When setting prices in period t , firms are assumed to value future streams of income at the marginal value of aggregate income in terms of the marginal value of an additional unit of aggregate income today giving the stochastic discount factor

$$Q_{t,T} = \beta^{T-t} (P_t Y_t^\sigma) / (P_T Y_T^\sigma).$$

In a symmetric equilibrium $c_t(i) = c_t = \sigma^{-1} w_t \equiv W_t / P_t = n_t = Y_t$ for all i , $p_t(j) = p_t(j)$ and $b_t(i) = b_t(j)$ for all i, j . To a first-order log-linear approximation, in the neighborhood of a zero-inflation steady state, individual consumption and pricing can be expressed as

$$\hat{c}_t(i) = \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta) \hat{w}_{T+1} - \sigma^{-1} (\hat{i}_T - \hat{\pi}_{T+1} - \beta (\bar{c}_T - \bar{c}_{T+1}))] \quad (1)$$

$$\hat{p}_t(j) = E_t^j \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(1 - \alpha\beta) (\hat{w}_T + u_T) + \alpha\beta \hat{\pi}_{T+1}] \quad (2)$$

where for any variable z_t , $\hat{z}_t = \ln(z_t / \bar{z})$ the log-deviation from steady state \bar{z} , with the exceptions $\hat{p}_t(j) = \ln(p_t(j) / P_t)$, $\hat{i}_t = \ln[(1 + i_t) / (1 + \bar{i})]$, $u_t = \ln(\theta_t / \bar{\theta})$ and $\bar{c}_t = \ln(\bar{C}_t / \bar{C})$. With a slight abuse of notation, the caret denoting log deviation from steady state is dropped for the remainder, so long as no confusion results.

Aggregating across the continuum of households and firms, and imposing market-clearing conditions, the economy is described by the aggregate demand and supply equations

$$x_t = \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} [(1 - \beta) x_{T+1} - \sigma^{-1} (i_T - \pi_{T+1} - r_T^n)] \quad (3)$$

$$\pi_t = \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\kappa(x_T + u_T) + (1 - \alpha)\beta\pi_{T+1}] \quad (4)$$

where the output gap is defined as

$$x_t = y_t - y_t^n = \sigma^{-1}w_t$$

the difference between output and the natural rate of output, the level of output determined by a flexible price economy: here $y_t^n = 0$. The associated natural rate of interest $r_t^n = \beta(\bar{c}_t - \hat{E}_t\bar{c}_{T+1})$ is determined by fluctuations in the propensity to consume, which along with the cost-push shock are first-order autoregressive processes. The aggregate demand equation determines the output gap as the discounted expected value of future wages, with the second term capturing variations in the real interest rate, applied in future periods, due to changes in the nominal interest rates and goods price inflation. That expected future dividends are irrelevant to consumption plans, to the first-order, reflects the assumption of an infinite Frisch elasticity of labor supply. The aggregate supply curve determines inflation as the discounted future sequence of marginal costs and the inflation rate. The slope of the Phillips curve is measured by $\kappa = (1 - \alpha\beta)(1 - \alpha)/\alpha$, while the exogenous process u_t captures variations in firms' desired markup reflecting variations in elasticity of demand θ_t .

3 Beliefs, Forecasting and Equilibrium Dynamics

Consistent with the basic data facts detailed above, households and firms have the forecasting model

$$\begin{aligned} z_t &= H'_m \bar{a}_{t-1} + \Omega_s s_{t-1} + e_t \\ \bar{a}_t &= \bar{a}_{t-1} + \nu_t. \end{aligned} \quad (5)$$

The vector $z_t = [\pi_t \ w_t \ i_t]'$ includes endogenous variables that private agents need to forecast, while \bar{a}_t denotes a vector of n_a unobserved random-walk drifts. Agents' priors on the volatility of the innovations ν_t are defined by the variance-covariance matrix $Q_m = \hat{E}(\nu_t\nu_t')$. Subsequent analysis allows for the possibility that the number of the underlying drifts driving the long-term behavior of the economy, n_a , can vary between one and three. Accordingly the matrix H_m has dimension $n_a \times 3$. Finally s_t denotes exogenous disturbances and e_t is a vector

of i.i.d shocks. Given an estimate of the drift, forecasts are computed as

$$\hat{E}_t z_T = H'_m a_{t-1} + \Omega_s \Xi^{T-t} s_t \text{ for } T \geq t \quad (6)$$

where a_{t-1} denotes the current estimate of \bar{a}_{t-1} and where

$$s_t = \Xi s_{t-1} + \varepsilon_t$$

and Ξ a diagonal matrix with eigenvalues inside the unit circle.

A few comments are required. First, beliefs of this kind can be interpreted as a first-order approximation to more general belief structures — see Eusepi and Preston (2013) for further discussion. Second, as shown by Kozicki and Tinsley (2012) for inflation expectations and Crump, Eusepi and Moench (in progress) for the joint behavior of inflation, interest rate and output growth expectations, a simple state-space model of the form (5) is consistent with the term-structure of professional forecasts from a range of surveys. Third, the model implies shifting end-points

$$\lim_{T \rightarrow \infty} \hat{E}_t z_T = H'_m a_{t-1}.$$

The existence of uncertainty about long-term outcomes is a primitive of the analysis and well-motivated by practical considerations. In the case of inflation, long-run drift reflects uncertainty about a central bank's inflation target. And even if the inflation targeting regime has a degree of credibility, it is reasonable to suppose financial market participants continuously evaluate a central bank's resolve to achieve its objectives; in the case of wages, drift embodies fundamental uncertainty about long-run technological advance. As shown by Eusepi and Preston (2011), end-point uncertainty captures the most important effect of uncertainty on equilibrium dynamics; instead uncertainty about autoregressive components of dynamics have limited impact on dynamics because of mean reversion. Fourth, a constant Ω_s retains linearity in the model, permitting application of standard estimation techniques and a linear-quadratic optimal policy problem.

The estimate of the drift, a_t , is updated using the steady-state Kalman filter recursion

$$a_t = a_{t-1} + K_m(z_t - H'_m a_{t-1} - \Omega_s s_{t-1}). \quad (7)$$

The time-invariant Kalman matrix is

$$K_m = P_m H_m (H_m' P_m H_m + R_m)^{-1}$$

$$K_m H_m' P_m = Q_m$$

where $P_m = E [(\bar{a}_t - a_t)(\bar{a}_t - a_t)']$ and R_m denotes agents' priors about the variance-covariance of the shock e_t .

Different prior assumptions about the source of drift generate different Kalman gain matrices. In the learning literature, updating in (7) is typically specified by the assumptions $H_m = I_3$ and $K_m = \bar{g}I_3$. Sargent and Williams (2005) have shown that such a Kalman matrix can be obtained from the Kalman filter recursions provided agents hold a specific prior about the drift volatility

$$Q_m = \bar{g}^2 R_m.$$

While a thorough discussion is deferred until later empirical evaluation of the model, where prior beliefs will be constrained to be consistent with observed data on expectations, it is immediate that this benchmark is quite restrictive. It asserts that long-run beliefs about any given variable depend only on short-run forecast errors about the same variable. For example, it excludes long-run inflation beliefs being revised in the light of short-run forecast errors about nominal interest rates. A contribution of this paper is to demonstrate more general structures have important implications, both for our understanding of belief formation, and the role of beliefs in macroeconomic dynamics and policy design.

The model is closed with an equation describing the behavior of the nominal interest rate, discussed below. It is assumed that the time-invariant matrix Ω_s , which along with the fundamental disturbances govern short-run dynamics, are given by the predictions of the rational expectations model associated with any given monetary policy. Conditional on beliefs and monetary policy these coefficient matrices are determined by a standard fixed point problem.

Evaluating expectations in the aggregate demand and supply curve using (5) permits these relations, along with the monetary policy rule, to be written as

$$z_t = C_0 H_m' a_{t-1} + \Omega_s^* s_{t-1} + \Omega_\varepsilon^* \varepsilon_t \tag{8}$$

which denotes the true data-generating process — the appendix provides details. The matrix C_0 measures the impact that beliefs about the drift have on inflation, the interest rate and the output gap. The matrices Ω_s^* and Ω_ϵ^* denote rational expectations coefficients. As common in this approach to learning, dynamics are self-referential: beliefs feed back into macroeconomic variables which in turn affect beliefs. To see this, combine (7) with (8) to give

$$a_t = (I_{n_a} + K_m C_0 H'_m - K_m H'_m) a_{t-1} + K_m \Omega_\epsilon^* \epsilon_t.$$

Short-term disturbances lead to revisions in the estimated drifts, which depend both on the impact of beliefs on the true data-generating process, C_0 , and the agents' priors about the volatility of drifts and perceived short-term shocks, K_m . The self-referentiality of beliefs can lead to instability. Technically, this occurs if any eigenvalue of the matrix

$$I_{n_a} + K_m C_0 H'_m - K_m H'_m$$

lies outside the unit circle. This plays an important role in policy design, as discussed in the next section.

4 The Optimal Policy Problem

This section characterizes the optimal monetary policy under long-term drift in expectations. 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.

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. Feasible sequences of inflation and the output gap must satisfy the aggregate demand and supply equations, (3) and (4), and the evolution of beliefs

(5). 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.

A more subtle issue warrants remark. The aggregate demand schedule is generally a binding constraint on feasible state-contingent choices over inflation and the output gap under learning. Recall the rational expectations analysis of this model. The central bank minimizes

$$E_t \sum_{T=t}^{\infty} \beta^{T-t} (\pi_T^2 + \lambda_x x_T^2)$$

subject to

$$x_t = E_t x_{t+1} - \sigma^{-1} (i_t - E_t \pi_{t+1} - r_t^n) \quad (9)$$

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t. \quad (10)$$

It is well understood that the aggregate Euler equation is not a binding constraint. For optimal bounded sequences of the output gap and inflation, and their conditional expectations, one can always determine an interest-rate consistent with those sequences. By construction of rational expectations equilibrium interest-rate expectations are also uniquely determined under this optimal policy.

That this logic no longer applies under arbitrary beliefs follows immediately 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. This means the aggregate demand equation is necessarily a constraint on what a central bank can achieve, since 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.

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

$$\min_{\{x_t, \pi_t, i_t, a_t\}} \bar{E}_t^{RE} \sum_{T=t}^{\infty} \beta^{T-t} L_T \quad (11)$$

taking as given initial beliefs, a_{-1} . The first-order conditions are described in the appendix. As first pointed out by Molnar and Santoro (2013), an interesting feature of this decision problem is that the first-order conditions constitute a linear rational expectations model.¹ The system can be solved using standard methods. Using results from Giannoni and Woodford (2010), the following proposition can be stated.

Proposition 1 *The model comprised of (i) the aggregate demand and supply equations (3) and (4); (ii) the law of motion for the beliefs a_t (5); and (iii) the first-order conditions resulting from the minimization of (11) subject the restrictions listed in (i) and (ii) admits a unique bounded rational expectations solution for all parameter values. In particular, model dynamics under optimal monetary policy are unique and bounded for all possible gains.*

Proof. See Appendix. ■

Before delineating the basic properties of optimal policy with long-run drift in expectations, a special case of this result is worth mentioning. When the gain matrix 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.

Corollary 2 *In the special case $K_m \rightarrow 0$ optimal policy will give the same dynamic responses to disturbances as optimal discretion under rational expectations.*

This type of result has been discussed earlier by Sargent (1999) and Molnar and Santoro (2013).

5 Intertemporal Trade-offs

This section gives content to proposition 1 exploring a number of implications. The following maintains two assumptions. First, retain the standard assumption $H_m = I_3$ and $K_m = \bar{g}I_3$. Second, assume that disturbances are i.i.d.. The most important result is that complete stabilization of disturbances to the natural rate of interest is generally not feasible: the model does not satisfy the “Divine Coincidence” — see Blanchard and Galí (2007) and Galí (2008). Additional trade-offs confront policy makers when there is long-term drift in expectations, and most centrally, policy expectations.

¹In 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.

5.1 Basic Properties

Figure 2 plots model impulse response functions to a one percent decline in the natural rate of interest. The following parametric assumptions are made: the discount factor is $\beta = 0.994$; the frequency of price changes determined by $\alpha = 0.8$; the weight on output gap stabilization $\lambda_x = 0.05$; and the gain coefficient $\bar{g} = 0.04$.

Several observations are immediate. First, the optimal policy does not fully stabilize inflation and the output gap. In fact, optimal policy dictates nominal interest rates decline by a relatively small amount when compared to the size of natural-rate disturbances, requiring a degree of nominal and real adjustment. This is a fundamentally different prediction to a rational expectations analysis. Optimal policy under rational expectations, under both commitment and discretion, stipulates nominal interest rates move precisely to offset changes in the natural rate of interest — aggregate demand and inflation are completely stabilized. Hence optimal policy under long-run drift in expectations is less aggressive. Second, optimal policy is inertial: an i.i.d. disturbance requires interest rates to converge gradually back to steady state after the initial impact. This inertia arises solely because of dynamics in beliefs. In this way, policy has some similarities to the optimal commitment policy under rational expectations in the presence of cost-push shocks. Third, there appears to be a negative eigenvalue driving dynamics. The source of this dynamic is discussed in the next section.

By way of comparison, Figure 3 shows corresponding impulse response functions for a standard Taylor rule, with $\phi_\pi = 1.5$ and $\phi_x = 0.5/4$. The principle difference resides in the impact effects of the natural-rate disturbance. Nominal interest rates fall further to accommodate the negative demand shock. This ameliorates the recession in real activity to some degree, though ultimately delivers higher volatility in the output gap. The remaining dynamics are qualitatively similar, a point to which we return.

Further insight is yielded by plotting volatility frontiers as a function of the gain coefficient. Figure 4 plots the standard deviation of the output gap and interest rate as a function of the constant gain \bar{g} under optimal policy. Under maintained parameter assumptions there is relatively small 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

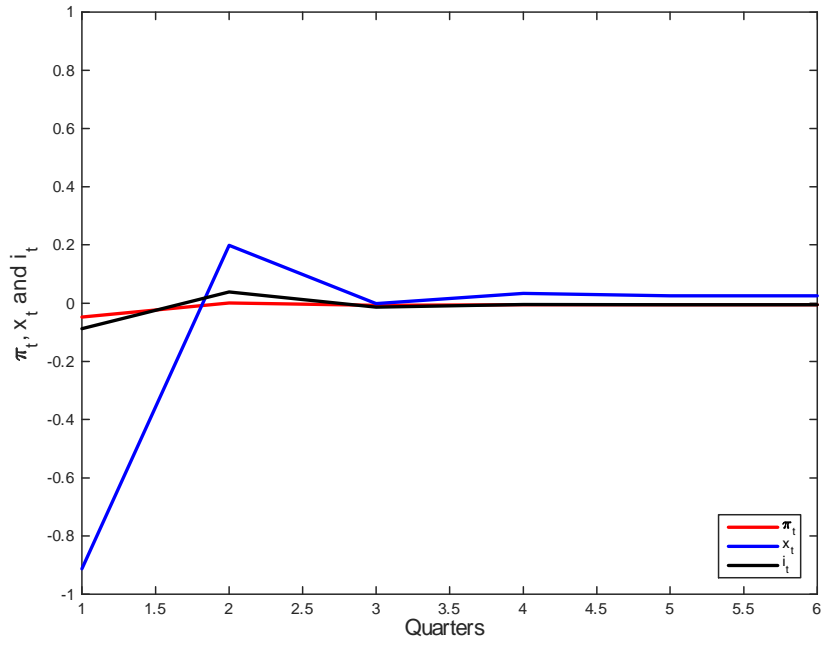


Figure 2: Impulse responses to a one percent natural-rate disturbance under optimal policy.

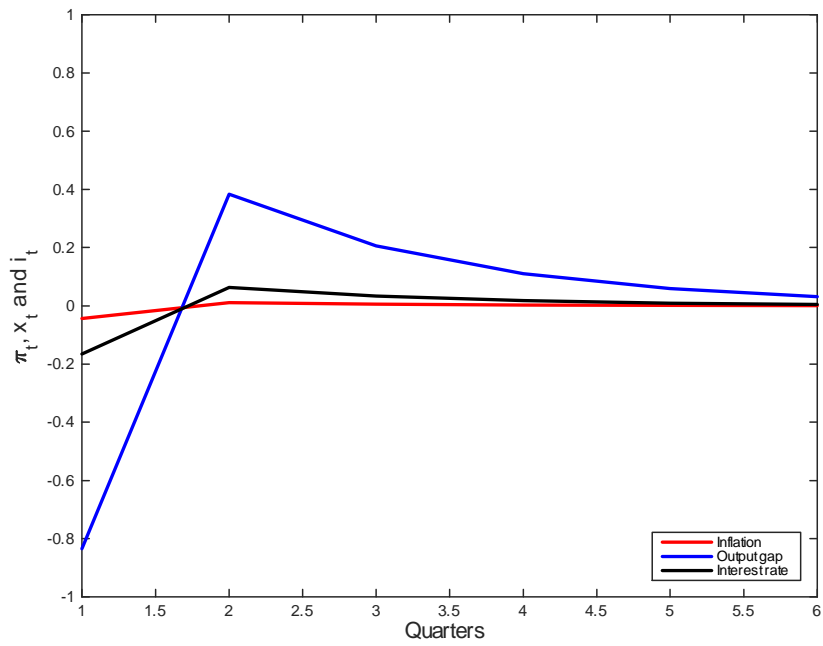


Figure 3: Impulse responses to a one percent natural-rate disturbance under a standard Taylor Rule.

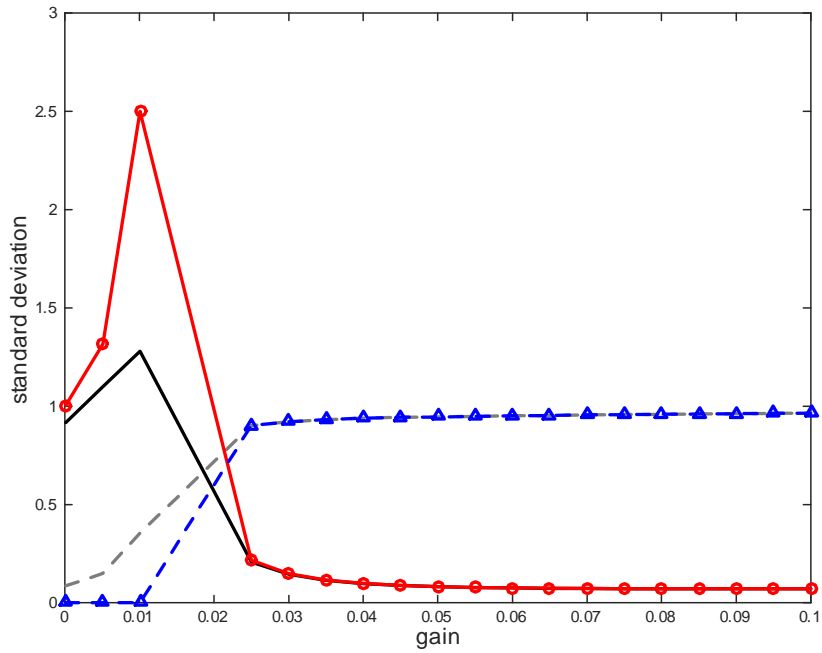


Figure 4: This figure show the volatility of output and interest rates as a function of the constant gain \bar{g} . The welfare theoretic loss gives the volatility of the interest rate (red circles) and the output gap (blue triangles); while a policy maker with a concern of interest rate volatility delivers the interest rate shown by the black line, and the output gap given by the grey dashed line.

(11), and under a loss function

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

that also penalizes volatility in the interest rate. Recall optimal discretion corresponds to the case $\bar{g} = 0$. Under the standard loss function a knife-edge result obtains: for $\bar{g} < 0.02$ the output gap is fully stabilized even if this induces substantial volatility in the interest rate. For large values of \bar{g} , the policy maker loses the ability to stabilize the output gap. Feasible policy permits limited 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 considerable output gap volatility.

These exercises point to a fundamental property of optimal policy under long-term drift in expectations: current interest rates move relatively less in response to demand disturbances when compared to a rational expectations analysis, including movements in the natural rate of interest. This feature can be sourced to the interplay between interest-rate policy and long-term interest-rate beliefs. The following section provides more precise insight.

5.2 A Simple Example

To appreciate the constraint confronting policy consider a central bank faced only with i.i.d. shocks to the natural rate r_t^n ; private agent beliefs initially consistent with rational expectations equilibrium so that $a_{t-1} = 0$; and for simplicity $\sigma = 1$. Because forecasts satisfy $E_t z_T = 0$ for all $T > t$, period t equilibrium is determined by the aggregate demand and supply curves (3) and (4) which simplify to

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

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, $i_t = r_t^n$, giving $\pi_t = x_t = 0$. But this implies subsequent movements in long-run interest-rate beliefs according to

$$a_t^i = a_{t-1}^i + \bar{g} (r_t^n - a_{t-1}^i).$$

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

$$\begin{aligned}\pi_{t+1} &= \kappa x_{t+1} \\ x_{t+1} &= -(i_{t+1} - r_{t+1}^n) - \frac{1}{1-\beta} \beta a_t^i.\end{aligned}$$

Complete stabilization of inflation and the output gap is again possible by having nominal interest rates tracking 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

$$i_{t+1} = r_{t+1}^n - \frac{\beta}{1-\beta} a_t^i \tag{12}$$

in every period t . This expression makes clear the source of the negative eigenvalue manifest in Figure 2. Optimal policy not only responds natural-rate disturbances, but also attempts to offset movements in long-term interest rates, driven by expectations.² Substituting into the updating rule for beliefs, ω_t^i , gives

$$a_{t+1}^i = \left(1 - \frac{\bar{g}}{1-\beta}\right) a_t^i + \bar{g} r_{t+1}^n$$

which is a first-order difference equation. Sustainable policy requires the dynamics of beliefs to be stationary. The following restriction must hold

$$\bar{g} < 2(1-\beta).$$

For larger gains, stability is not feasible, implying beliefs and, concomitantly, 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.

Proposition 3 *In the model given by (3) and (4), and shocks to the natural level of output, Divine Coincidence will in general not hold in absence of cost-push shocks.*

The fundamental issue here is that drifting, or unanchored, long-term interest rate expectations limits the degree to which short-term interest rates can adjust at the time of a natural

²The implied interest rates of a bond of any maturity can be shown to be a function of the long-term interest rate belief. Note also that the negative eigenvalue is prominent because of the absence of other sources of persistence in the model. This kind of dynamics is also likely responsible for some of the observations made in Cogley, Sbordone and Matthius.

rate disturbance — else long rates fall too much, overly stimulating demand. This constrains the central bank’s ability to combat declining demand. It is important to emphasize that 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

$$\min_{\{x_t, \pi_t, a_t^x, a_t^\pi\}} \bar{E}_t^{RE} \sum_{T=t}^{\infty} \beta^{T-t} L_T$$

subject only to the Phillips curve (4), taking as given initial beliefs a_{-1}^x and a_{-1}^π . Then it is straight forward to show the Divine Coincidence holds, despite long-term drift in expectations about inflation and real activity. For further discussion see Eusepi and Preston (2015).

Evidently, the challenge to stabilization policy originates in the interplay between interest-rate beliefs and policy. If agents had correct knowledge of long-term interest rates, giving beliefs $a_t^i = 0$ for all periods t , it would be feasible to stabilize demand shocks, restoring the Divine Coincidence. Anchored beliefs enhance the efficacy of monetary policy. The optimal policy is given by the reaction function (12) which is precisely the optimal policy under commitment and discretion. However, this conclusion depends on the maintained assumption about agents’ priors. For example, suppose agents have a prior belief that there is one source of nominal drift, and employ the forecasting model

$$a_t^\pi = a_{t-1}^\pi + Q_m^{1/2} \varepsilon_t^\pi$$

and

$$z_t = H_m' a_{t-1}^\pi + u_t$$

where $H_m = [1 \ 0 \ 1]$. This model would be interpreted as a situation in which agents are only unsure about the long-run objectives of inflation policy. This single nominal drift affects beliefs about both long-run inflation and also nominal interest rates through the Fisher relation. Beliefs are revised according to

$$a_t^\pi = a_{t-1}^\pi + K_m (z_t - H_m' a_{t-1}^\pi).$$

Because $K_m(R_m, Q_m)$ will in general weight all forecast errors, the logic of the simple example continues to apply. Long-run beliefs about inflation depend on short-run interest-rate

forecast errors, tying current policy decisions to long-term interest-rate movements. The same stabilization challenge emerges.

5.3 Further Implications and Discussion

Two principal related insights emerge from the theory of optimal monetary policy: i) the transmission mechanism, which operates through the expectations hypothesis of the yield curve to aggregate demand, is itself a constraint on policy design; and ii) because of this constraint it is optimal to limit movement in the policy rate when responding to evolving macroeconomic conditions. Optimal policy should not be aggressive relative to the predictions of a rational expectations analysis of the model. This means Divine Coincidence does not hold even in the absence of cost-push shocks. These findings are in direct contrast to earlier work on optimal monetary policy with non-rational expectations.

As discussed, Molnar and Santoro (2013) analyze optimal monetary in a model with learning in which (9) and (10) are taken as the primitive decision rules, and beliefs are given by (5). 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 inconsistent 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.

The consequences of the aggregate demand constraint are not specific to fully optimal policy. Figure 5 describes the model’s 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 $\phi_x \equiv \phi_\pi \tilde{\phi}_x$. This is an example of ‘robust learning stability’ proposed by Evans and Honkapohja (2009). Each of the three contours describe the stability frontier in the constant-gain and inflation response coefficient space. Parameter regions above a plotted contour

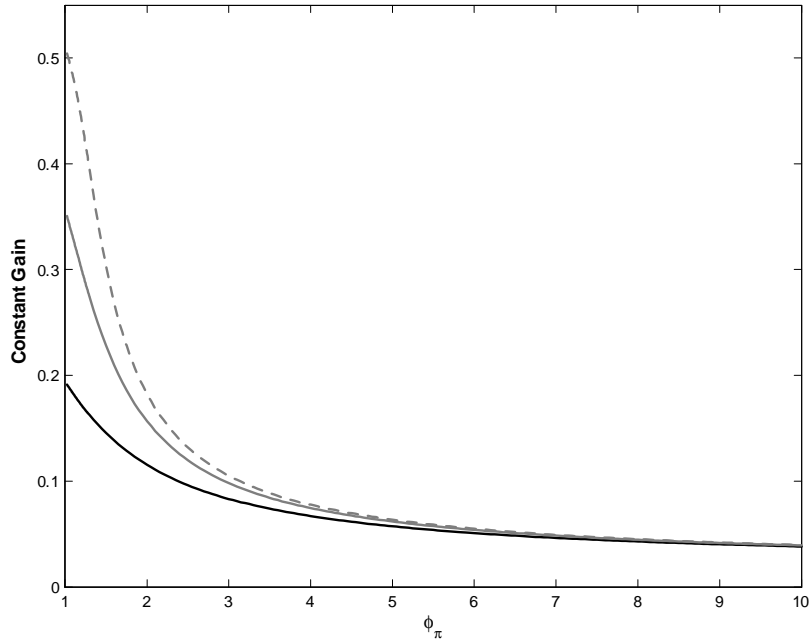


Figure 5: The figure shows stability frontiers corresponding to alternative Taylor rules. In particular (\bar{g}, ϕ_π) above the frontier correspond to locally unstable equilibria under constant-gain learning. The black solid line corresponds to the standard Taylor Rule. The solid (dashed) grey line corresponds to $\phi_x = \phi_x^*/2$ ($\phi_x = \phi_x^*/3$).

indicate local instability of the equilibrium. Higher contours correspond to progressively weaker responses to the output gap. For many gain coefficients aggressive monetary policy may not be desirable.

Interestingly, for some gain coefficients even an “infinite” inflation response would not be sufficient for 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

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

Such an approach to implementing policy under imperfect knowledge and learning has been argued to be desirable by Evans and Honkapohja (2006) and Preston (2008). The above makes clear that for many gains such policies are infeasible.

Finally, a limitation of the analysis is that the proposed belief structure is subject to the Lucas Critique. The agents’ Kalman gain matrix is invariant to the policy regime in place. On the one hand, this assumption prevents agents’ learning rules being revised in response to the policy regime which allows the policymaker to systematically exploit a given belief structure. Nonetheless, the central implication of this property is consistent with the spirit of our analysis: even the most sophisticated policymaker cannot fully stabilize the economy. On the other hand, this assumption prevents the policymaker from inducing a learning rule that might improve the stabilization trade-off. For example, if a policy existed that could induce $K_m \rightarrow 0$ then complete stabilization would be feasible — recall that this special case delivers dynamics isomorphic to optimal discretion. Of course, in practice there are limits on the degree to which a sophisticated policymaker can mould household and firm beliefs. One might reasonably ask whether a policy maker can ever achieve full credibility in the sense that large forecast errors would not influence long-term beliefs. Moreover, some of the perceived drift, such as movements in the real interest rate, might not be directly under the influence of policy makers and should be taken as a constraint invariant to policy.

With these remarks in mind, we acknowledge taking a model of beliefs in which agents have an invariant Kalman matrix is subject to this criticism. The deeper issue of providing microfoundations for the existence of drift in long-term expectations is left to future work.

However, Marcet and Nicolini (2003), and more recently Carvalho, Eusepi, Moench, and Preston (2015), provide one promising approach.

6 Quantitative Implications

The theory of optimal monetary policy under long-term drift in expectations identifies a new trade-off confronting policy design. Potential instability in long-term interest rates limits the degree to which the central bank optimally responds to demand disturbances contemporaneously. Obvious questions present themselves: is there evidence supporting the long-term expectations formation presented here — are short-term forecast errors important in determining long-term expectations? And if so, do such beliefs generate an intertemporal trade-off that is quantitatively important? The remainder of the paper addresses these questions.

6.1 Models of Drift

The constant-gain algorithm underpinning earlier analysis is restrictive, imposing that long-term beliefs about any given variable are revised only in response to short-term forecast errors attached to that variable. It rules out plausible behavior such as revising long-term interest rate forecasts in response to recent inflation forecast errors. Subsequent empirical work seeks to identify the number of drifts and their dependence on short-term forecast errors. The following two models of belief formation are estimated.

Baseline model. The model features two drifts, $n_a = 2$; one nominal, a_t^π , and one real, a_t^y . The nominal interest-rate drift is determined by a linear combination for the nominal and real drift:

$$a_t^i = a_t^\pi + \lambda_{iy} a_t^y.$$

In the notation of section 3:

$$H'_m = \begin{bmatrix} I_2 \\ 1 \quad \lambda_{iy} \end{bmatrix}.$$

To limit the number of estimated parameters, assume agents' priors on the innovation e_t correspond to the short-term forecast errors under rational expectations. This gives

$$R_m = \hat{E}(e_t e_t') = \Omega_\epsilon^* E(\epsilon_t \epsilon_t') \Omega_\epsilon^{*'}.$$

To simplify computations, estimate P_m directly instead of Q_m . Since we have two drifts P_m is a two-dimensional matrix, requiring three parameters to be estimated. Interest-rate beliefs introduce the additional parameter λ_{iy} giving a total of four belief parameters to be estimated, in addition to various structural parameters relevant to household and firm decisions.³ The Kalman matrix and the implied priors for the two drifts are then determined according to

$$K_m = P_m H_m (H_m' P_m H_m + R_m)^{-1}$$

$$Q_m = K_m H_m' P_m.$$

Diagonal gain. We consider an independent constant-gain model, specified by the Kalman matrix

$$K_m = \begin{bmatrix} \bar{g}_\pi & 0 & 0 \\ 0 & \bar{g}_y & 0 \\ 0 & 0 & \bar{g}_i \end{bmatrix}$$

which permits long-term beliefs about each variable to be updated at a different rate. This model requires three belief parameters to be estimated.

6.2 Estimation

To estimate the model using full-information Bayesian inference requires an assumption about the conduct of monetary policy. To begin assume that monetary policy was historically implemented according to a Taylor-type rule of the form

$$\begin{aligned} i_t &= i_t^* + \phi_\pi \pi_t + \phi_x x_t + \phi_\mu \mu_t \\ i_t^* &= \rho_i i_{t-1}^* + \epsilon_t^i. \end{aligned} \tag{13}$$

There is no particular reason to assume monetary policy was conducted optimally in response to drifting beliefs over the sample being considered. Moreover, a substantial literature argues that simple rules of this kind provide a reasonable description of the evolution of US interest rates. Taking this agnostic approach permits adducing evidence on the empirical relevance of low-frequency drift in beliefs. With this evidence in hand, the quantitative implications for

³A model with two nominal drifts and one real drift was also explored. However, the model was poorly identified with the two nominal drifts being virtually identical. This mean the Hessian was not invertible. For this reason, emphasis is given to the two drift model.

optimal policy are evaluated. The only non-standard feature is the inclusion of a response to the cost-push shock in the policy rule. This greatly assists the fit of the simple model, particularly in regards to matching comovements properties of observed data.

When estimating the benchmark model use the following quarterly data, expressed in annual terms, over the sample period 1968Q3-2009Q1, which spans the periods in which various survey data are available: inflation, the output gap and the nominal interest rate are measured by the log-difference of the GDP deflator; the Congressional Budget Office (CBO) measure of the output gap; and the three-month treasury bill rate. To identify information about beliefs, six survey data series are employed: one- and four-quarters-ahead forecasts of inflation and 3-month Treasury Bill rates from the Survey of Professional Forecasters (available from 1968Q3 and 1981Q3), and five-to-ten year inflation and interest-rate forecasts (both available at a biannual frequency from 1986Q2) constructed using the Blue Chip Economic Indicators Survey. This missing data, along with the biannual frequency of the BCEI data, are easily handled by Bayesian methods.

The model consists of the data generating process, (8), the learning rule, (7), and the evolution of the exogenous shocks $s_t = (r_t^n, \mu_t, i_t^*)'$. Some correlation between the cost-push and the demand shocks is permitted to improve the fit of the model. Defining the state vector $\xi_t = (z_t' \ a_t' \ s_t' \ a_{t-1}')$, these relations imply the state-space form

$$\xi_t = \mathbb{F}\xi_{t-1} + \mathbb{Q}^{1/2}\epsilon_t.$$

Details are in the appendix. The observation equation is

$$Y_t = W_t Y_t^*$$

where W_t is a matrix selecting the observables in period t and

$$Y_t^* = \begin{bmatrix} 4\pi_t \\ 4i_t \\ x_t \\ 4\hat{E}_t\pi_{t+1} \\ 4\hat{E}_ti_{t+1} \\ 4\hat{E}_t\pi_{t+4} \\ 4\hat{E}_ti_{t+4} \\ \frac{4}{20}\hat{E}_t\sum_{i=21}^{40}\pi_{t+i} \\ \frac{4}{20}\hat{E}_t\sum_{i=21}^{40}i_{t+i} \end{bmatrix} = \mu + \mathbb{H}'\xi_t + Go_t$$

where observation errors o_t are attached to the survey data. Each of the expectations terms are easily computed recalling

$$\hat{E}_tz_T = a_{t-1}H'_m + \Omega_s\Xi^{T-t}s_t$$

for any period $T > t$.

Table 1 reports the prior and posterior distribution of the parameters. In addition two parametric assumptions are made: the steady-state inflation rate is taken to be 2 percent; and the households' discount factor 0.994, which implies a real interest rate of about 2 percent. The model's estimates are reasonably consistent with values typically found in the empirical New Keynesian literature. The inverse intertemporal elasticity of substitution σ , which affects the sensitivity of consumption to the real interest rate is not too different from one, while the Calvo parameter, α , implies a price duration of about three quarters. Regarding policy, the policy response to inflation, ϕ_π , is about 2, while the response to output gap ϕ_x is 0.85, substantially higher than a standard Taylor rule. The policy rule also responds positively to cost-push shocks with a coefficient $\phi_\mu = 0.65$.

Rudimentary evidence in support of the baseline model is provided by comparing the value of the posterior distribution at the estimated mode. The benchmark model has posterior value of -87.6 . The diagonal model, with three distinct constant gains, has posterior value -229.9 . The independent gain assumption provides a substantially poorer fit of the data. Formal model comparison tests, while not yet completed, would strongly favor the benchmark model. To get a better sense of the model's performance, the final three figures of the paper show

the predictions of the model with three independent gains for the survey forecasts. The model performance is worse compared to our baseline which is discussed further below. Interestingly, the estimated gain are close to zero, so that in this simple model the usual independent gain assumption appears to be poorly supported by the data.

The remaining discussion in the paper takes the baseline model as the preferred model. The baseline model implies the following estimated gain matrix for a_t^π and a_t^x :

$$K_m = \begin{bmatrix} 0.0784 & -0.0099 & 0.0952 \\ -0.0095 & 0.0018 & -0.0123 \end{bmatrix}.$$

Combining with H_m provides the gain on the interest rate drift a_t^i (a linear combination of the first two rows of K_m)

$$H'_m K = \begin{bmatrix} 0.0784 & -0.0099 & 0.0952 \\ -0.0095 & 0.0018 & -0.0123 \\ 0.0789 & -0.0100 & 0.0957 \end{bmatrix}$$

Recalling the observed data in the filtering problem is

$$z_t = \begin{bmatrix} \pi_t \\ x_t \\ i_t \end{bmatrix}$$

several insights emerge. First, the loadings for both inflation and interest rates are very similar, with forecast errors on the interest rate receiving a bit more weight. This reflects the relative volatility of short-term forecast errors both in the data and the estimated model. Over the sample, the variance of inflation forecast errors (annualized) is roughly twice that of the interest-rate forecast errors. The prior on the volatility of e_t implies, however, posterior volatilities that are closer to each other. This is reflected in the similar coefficients of the Kalman matrix.

Second, the Kalman matrix reflects agents' prior of a negative correlation between the two drifts. This is broadly consistent with the negative correlation between inflation and the measured output gap in the sample. HP-filtered trends for the two variables display a -0.20 percent correlation. Negative inflation and interest-rate surprises are associated with higher expected output in the long term.

Figures 6, 7 and 8 provide assurances on the quality of fit. Figure 6 shows the one-quarter-ahead forecast errors for the GDP deflator and the 3 month Treasury Bill computed using the observed data and the model estimates. The model does a reasonable job at fitting these data, though clearly there are greater discrepancies for the inflation forecast data, than for the interest-rate data. This is consistent with the higher standard deviation estimated for the one-quarter-ahead inflation expectations. Figure 7 shows the four-quarter-ahead forecast errors. For inflation the model is consistent with the general pattern of forecast error, though there are periods where the model fails to track high-frequency movements in the data. For interest rates, the model picks up the cyclical movements in expectations nicely, with some periods exhibiting systematic under prediction of the level.

Figure 8 plots the 5-10 year forecasts from the BCEI and the corresponding model predictions for 5-10 years as well as the 4-quarter-ahead forecast. The model picks up the broad trajectory of the long-term forecasts, though generates slightly greater variation in long-term inflation expectations than observed in the data. Model implied long-term expectations for both series exhibit sensitivity to short-term movements in expectations.

Finally, to get a sense of the dynamic properties of the model, Figures 9 and 10 plot the impulse responses to a natural-rate and cost-push shock respectively. The dynamics in each case correspond with standard thinking. The natural-rate disturbance reduces both inflation and the output gap, requiring a stimulatory response from the monetary authority. The cost-push shock embeds a non-trivial trade-off as inflation rises and real activity declines. The monetary authority restrains the degree to which inflation increases by raising nominal interest rates.

6.3 Implications for Optimal Policy

The previous section provides some evidence in support of the two drift model with a non-diagonal Kalman gain matrix. We now study the consequences of this belief structure for optimal policy. Of particular interest is the quantitative relevance of drifting beliefs for policy design. Consider a central bank with a mandate to implement optimal policy. This mandate represents a clear break from historical policy which is assumed to be given by the empirical model. Suppose the central bank takes as given the Kalman gain matrix consistent with historical beliefs attached to the Taylor rule regime. (Note that the primitive of the analysis

is the prior matrix Q_m , not the Kalman gain matrix K_m . A new regime implies a new Kalman gain matrix for fixed prior beliefs.) Interpret the exercise as representing a short-run constraint on what can be achieved by policy. After all it seems plausible that agents would take some time to adjust beliefs in response to the new regime. What are the implications for optimal policy?

6.3.1 A Natural Rate Disturbance

Figure 11 plots the impulse response to a one percent decline in the natural rate for both optimal policy under learning and optimal discretion under rational expectations. There are several interesting properties. First, the only source of difference between policy outcomes under learning and rational expectations are due to the time-varying drifts. Hence comparing interest-rate paths permits immediate inference on the consequence of drifting beliefs for policy. Second, the impact effect on interest rates differs across learning and rational models, with the slightly smaller adjustment in the former delivering a contraction in aggregate demand of about one-third of a percent. Optimal policy does not fully stabilize the natural-rate disturbance — the Divine Coincidence does not hold. Third, in subsequent periods interest-rate paths across each policy are identical. An implication is policy does not respond to beliefs under learning. However, inflation falls by about two-thirds of a percent, because of weaker expected demand conditions, which lower anticipated marginal costs into the indefinite future, and an expectation of lower future inflation which reduces incentives to raise prices today. The dynamics are hump-shaped. Output recovers quickly from lower anticipated real interest rates, leading to standard substitution of future for current consumption. Fourth, the optimal path for inflation exhibits substantial persistence, remaining below the inflation target for several years.

Figure 12 shows analogous impulse response functions for a policymaker with a weaker preference for output gap stabilization. The figure assumes $\lambda_x = 0.005$. The broader contours of adjustment are similar. Reflecting the more aggressive stance towards inflation, the impact effects on interest rates and the output gap are larger than in the benchmark case. The other notable difference is the monotonic convergence of inflation to the steady state.

6.3.2 A Cost-push Disturbance

Figure 13 plots the impulse response to a one percent serially uncorrelated cost-push shock. The dynamics under optimal discretion are familiar: inefficient movements in marginal costs represent a fundamental trade-off for stabilization policy. Inflation rises; the output gap falls. In the presence of long-term drift in expectations, optimal policy leads to a less aggressive adjustment of nominal interest rates, leading to a smaller contraction in real activity, but at the cost of higher inflation. This result extends earlier results on natural-rate disturbances to cost-push shocks. Because of the intertemporal trade-off embodied in distorted long-run interest-rate beliefs, it is optimal to limit adjustment in current interest rates in response to contemporaneous shocks. In the period after the shock, interest rates are lowered below steady state. This leads to a small boom in real activity before gradual convergence to steady state. Despite the rise in long-term inflation expectations, inflation falls in the period after shock due to a decline in anticipated marginal costs. For completeness, Figure 14 shows the case of a lower weight on output gap stabilization. Very similar remarks apply, with the central difference being the magnitude of impact effects.

6.3.3 Efficiency Frontiers

Figure 15 plots efficiency frontiers for the economy under both rational expectations and long-term drift in expectations. The figure plots the standard deviation of inflation and the output gap as the weight on output gap stabilization is systemically varied. Under rational expectations the figure delineates outcomes for both optimal commitment and discretion. As well-known, with commitment or discretion greater preference for output gap stabilization results in rising volatility in inflation. The efficiency frontier for the optimal commitment policy lies strictly inside that for optimal discretion.

The frontier is strikingly different in an economy with long-term drift in expectations. For any given level of output gap volatility inflation is considerably more variable. Again, this reflects the constraint on contemporaneous interest-rate policy: potential instability in long-term interest rates limits the central bank's ability to respond to evolving macroeconomic conditions. The frontier for drift is substantially different relative to the comparison of commitment and discretion. The importance of this result should be underscored. The welfare loss in moving from optimal commitment to optimal discretion is trivial in comparison

to the welfare loss in moving from rational expectations to a model with long-term drift in expectations. Long-term drift severely limits the efficacy of optimal policy, even when the policy maker has complete knowledge of household behavior including beliefs. This suggests a degree of caution is appropriate when evaluating standard policy prescriptions under rational expectations.

7 Additional Exercises

7.1 Optimal Policy under Robust Control

[TO BE ADDED]

7.2 Optimal Policy within a Class of Simple Rules

[TO BE ADDED]

8 Conclusions

[TO BE ADDED]

References

- BLANCHARD, O., AND J. GALÍ (2007): “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 39(s1), 35–65.
- 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.
- CARVALHO, C., S. EUSEPI, E. MOENCH, AND B. PRESTON (2015): “What Drives Long-run Inflation Expectations?,” unpublished manuscript, University of Melbourne.
- COGLEY, T., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2(1), 43–69.
- COGLEY, T., AND T. J. SARGENT (2005): “The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 8(2), 528–563.
- COGLEY, T., AND A. M. SBORDONE (2008): “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98(5), 2101–26.
- EUSEPI, S., AND B. PRESTON (2011): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, 101(6), 2844–2872.
- (2013): “Fiscal Foundations of Inflation: Imperfect Knowledge,” unpublished, Monash University.
- (2015): “The Science of Monetary Policy: An Imperfect Knowledge Perspective,” unpublished, The University of Melbourne.
- 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.

- GALI, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction of the New Keynesian Framework*. Princeton University Press.
- 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. P., AND M. WOODFORD (2010): “Optimal Target Criteria for Stabilization Policy,” unpublished, Columbia University.
- GÜRKAYNAK, 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.
- KOZICKI, S., AND P. A. TINSLEY (2001): “Shifting endpoints in the term structure of interest rates,” *Journal of Monetary Economics*, 47(3), 613–652.
- (2012): “Effective Use of Survey Information in Estimating the Evolution of Expected Inflation,” *Journal of Money, Credit and Banking*, 44(1), 145–169.
- LAUBACH, T., AND J. C. WILLIAMS (2003): “Measuring the Natural Rate of Interest,” *The Review of Economics and Statistics*, 85(4), 1063–1070.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93(5), 1476–1498.
- 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 (2005): “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.
- STOCK, J. H., AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

9 Appendix

9.1 The Data Generating Process

Consider the class of models where the data generating process takes the form:

$$A_0 z_t = A_{11} \hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} z_{T+1} + A_{12} \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} z_{T+1} + A_2 \Xi s_{t-1} + A_2 \epsilon_t$$

$$s_t = \Xi s_{t-1} + \epsilon_t.$$

where agents Perceived Law of Motion (PLM) of the economy takes the form:

$$z_t = H'_m a_{t-1} + \Omega_s s_{t-1} + u_t$$

$$a_t = a_{t-1} + Q_m^{1/2} \nu_t$$

where H'_m is of dimension $n_z \times n_a$.

Agents' forecasts are then obtained as

$$\hat{E}_t z_{T+1} = H'_m a_{t-1} + \Omega_s \Xi^{T-t} s_t; \quad T \geq t$$

$$\hat{E}_t \sum_{T=t}^{\infty} \beta^{T-t} z_{T+1} = \frac{1}{1-\beta} H'_m a_{t-1} + \Omega_s (I - \beta\Xi)^{-1} s_t$$

Substituting for the forecasts we get

$$\begin{aligned} z_t &= C_0 H'_m a_{t-1} + (C_{11} \Omega_s B_{11} + C_{12} \Omega_s B_{12} + C_2) \Xi s_{t-1} \\ &\quad + [C_{11} \Omega_s B_{11} + C_{12} \Omega_s B_{12} + C_2] \epsilon_t \end{aligned}$$

where

$$C_0 = A_0^{-1} \left(\frac{1}{1-\beta} A_{11} + \frac{1}{1-\alpha\beta} A_{12} \right)$$

$$B_{11} = (I_{n_s} - \beta\Xi)^{-1}; \quad B_{12} = (I_{n_s} - \alpha\beta\Xi)^{-1}.$$

$$C_{11} = A_0^{-1} A_{11}; \quad C_{12} = A_0^{-1} A_{12}; \quad C_2 = A_0^{-1} A_2.$$

In order to find the matrices Ω_s and Ω_ϵ consistent with rational expectations, indeterminate coefficients implies

$$\Omega_s^* = [(C_{11}\Omega_s^*B_{11}\Xi) + C_{12}\Omega_s^*B_{12}\Xi + C_2\Xi]$$

and

$$\Omega_\epsilon^* = C_{11}\Omega_s^*B_{11} + C_{12}\Omega_s^*B_{12} + C_2.$$

Vectorizing:

$$vec(\Omega_s^*) = [(B_{11}\Xi)' \otimes C_{11} + (B_{12}\Xi)' \otimes C_{12}] vec(\Omega_s^*) + vec(C_2\Xi)$$

$$vec(\Omega_s^*) = \{I_{n_s \times n_s} - [(B_{11}\Xi)' \otimes C_{11} + (B_{12}\Xi)' \otimes C_{12}]\}^{-1} \cdot vec(C_2\Xi).$$

The true data generating process is

$$z_t = C_0 H_m' a_{t-1} + \Omega_s^* s_{t-1} + \Omega_\epsilon^* \epsilon_t.$$

9.2 State-Space Model

Recall, the true data generating process is

$$z_t = C_0 H_m' a_{t-1} + \Omega_s^* s_{t-1} + \Omega_\epsilon^* \epsilon_t.$$

Combining with the updating equations

$$a_t = (I_{n_a} - K_m H_m') a_{t-1} + K_m z_t - K_m \Omega_s^* s_{t-1}$$

we can express the model in the following form

$$\begin{bmatrix} I_{n_z} & 0 & 0 & 0 \\ -K_m & I_{n_a} & 0 & 0 \\ 0 & 0 & I_{n_s} & 0 \\ 0 & 0 & 0 & I_{n_a} \end{bmatrix} \begin{bmatrix} z_t \\ a_t \\ s_t \\ a_{t-1} \end{bmatrix} = \begin{bmatrix} 0 & C_0 H_m' & \Omega_s^* & 0 \\ 0 & I_{n_a} - K_m H_m' & -K_m \Omega_s^* & 0 \\ 0 & 0 & \Xi & 0 \\ 0 & I_{n_a} & 0 & 0 \end{bmatrix} \begin{bmatrix} z_{t-1} \\ a_{t-1} \\ s_{t-1} \\ a_{t-2} \end{bmatrix} + \begin{bmatrix} \Omega_\epsilon^* \\ 0 \\ I_{n_s} \\ 0 \end{bmatrix} \epsilon_t.$$

9.3 Optimal Policy

We consider the model with i.i.d. shocks. Write the model in compact notation *as*

$$\bar{A}_0 z_t = \bar{A}_{11} \hat{E}_t^i \sum_{T=t}^{\infty} \beta^{T-t} z_{T+1} + \bar{A}_{12} \hat{E}_t^i \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} z_{T+1} + \bar{A}_2 \epsilon_t$$

which on evaluating expectations can be written as:

$$\bar{A}_0 z_t = \bar{B}_0 a_{t-1} + \bar{A}_2 \epsilon_t$$

where

$$\begin{aligned} z_t &= \begin{bmatrix} \pi_t \\ x_t \\ i_t \end{bmatrix} \\ \bar{A}_0 &= \begin{bmatrix} 1 & -\kappa & 0 \\ 0 & 1 & \sigma^{-1} \end{bmatrix} \\ \bar{A}_{11} &= \begin{bmatrix} 0 & 0 & 0 \\ \sigma^{-1} & 1 - \beta & -\sigma^{-1}\beta \end{bmatrix} \\ \bar{A}_{12} &= \begin{bmatrix} (1 - \alpha)\beta & \kappa\alpha\beta & 0 \\ 0 & 0 & 0 \end{bmatrix} \\ \bar{A}_2 &= \begin{bmatrix} 0 & \kappa & 0 \\ \sigma^{-1} & 0 & 0 \end{bmatrix}. \end{aligned}$$

and

$$\bar{B}_0 = \frac{1}{1 - \beta} \bar{A}_{11} + \frac{1}{1 - \alpha\beta} \bar{A}_{12}$$

The policymaker's problem can then be written as

$$\min_{y_t, z_t, a_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \begin{array}{l} \frac{1}{2} z_t' W z_t + \\ \underbrace{\Lambda'_{1,t}}_{1 \times (n_z - 1)} \left[-\bar{A}_0 z_t + \bar{B}_0 a_{t-1} + \bar{A}_2 \epsilon_t \right] + \\ \underbrace{\Lambda'_{2,t}}_{1 \times n_a} \left[-a_t + (I_{n_a} - K_m) a_{t-1} + K_m z_t \right] + \end{array} \right\}$$

which has the first-order conditions

$$\frac{1}{2}z'_t(W + W') - \Lambda'_{1,t}\bar{A}_0 + \Lambda'_{2,t}K_m = 0$$

$$-\Lambda'_{2,t} + \beta\mathbb{E}_t\Lambda'_{1,t+1}\bar{B}_0 + \beta\mathbb{E}_t\Lambda'_{2,t+1}(I_{n_a} - K_m) = 0.$$

Transposing, and using the fact that W is diagonal, provides

$$-\bar{A}'_0\Lambda_{1,t} + K'\Lambda_{2,t} + Wz_t = 0$$

$$\Lambda_{2,t} = \beta\bar{B}'_0E_t\Lambda_{1,t+1} + \beta(I_{n_a} - K_m)'E_t\Lambda_{2,t+1}.$$

The evolution of the economy is then described by the following equations

$$-\bar{A}'_0\Lambda_{1,t} + K'\Lambda_{2,t} + Wz_t = 0$$

$$\Lambda_{2,t} = \beta\bar{B}'_0\mathbb{E}_t\Lambda_{1,t+1} + \beta(I_{n_a} - K_m)'\mathbb{E}_t\Lambda_{2,t+1}$$

$$\bar{A}_0z_t = \bar{B}_0a_{t-1} + \bar{A}_2\epsilon_t$$

$$a_t = (I_{n_a} - K_m)a_{t-1} + K_mz_t.$$

Table 1: Baseline Model

| Parameters | Distribution | Prior | | Posterior | |
|----------------------|--------------|-------|------|-----------|---------------|
| | | Mean | Std | Mode | 95% interval |
| α | Beta | 0.66 | 0.08 | 0.66 | 0.61;0.70 |
| σ | Gamma | 1.00 | 0.30 | 1.24 | 0.87;1.71 |
| ϕ_π | Gamma | 2.00 | 0.15 | 2.05 | 1.83;2.28 |
| ϕ_x | Gamma | 0.12 | 0.10 | 0.84 | 0.51;1.64 |
| ϕ_μ | Normal | 0 | 1 | 0.65 | 0.28;1.43 |
| ρ_r | Beta | 0.70 | 0.12 | 0.89 | 0.91;0.93 |
| ρ_μ | Beta | 0.70 | 0.12 | 0.82 | 0.89;0.93 |
| $\rho_{\bar{i}}$ | Beta | 0.70 | 0.12 | 0.76 | 0.59;0.70 |
| σ_r | Igamma | 0.5 | 4 | 0.14 | 0.13;0.18 |
| σ_μ | Igamma | 0.5 | 4 | 0.89 | 0.81;1.02 |
| $\sigma_{\bar{i}}$ | Igamma | 0.5 | 4 | 0.79 | 0.65;1.09 |
| $\rho_{r\mu}$ | normal | 0 | 5 | -0.19 | -0.24;-0.15 |
| p_{11} | Igamma | 0.1 | 5 | 0.09 | 0.078;0.11 |
| p_{22} | Igamma | 0.1 | 5 | 0.02 | 0.015;0.054 |
| p_{21} | Normal | 0 | 5 | -0.01 | -0.02; -0.007 |
| λ_{iy} | Normal | 0 | 5 | -0.05 | -0.081;-0.011 |
| $4\sigma_{1Q,\pi}$ | Igamma | 0.1 | 1 | 0.72 | 0.64;0.82 |
| $4\sigma_{1Q,i}$ | Igamma | 0.1 | 1 | 0.68 | 0.61;0.78 |
| $4\sigma_{4Q,\pi}$ | Igamma | 0.1 | 1 | 0.21 | 0.18;0.25 |
| $4\sigma_{4Q,i}$ | Igamma | 0.1 | 1 | 0.43 | 0.38;0.50 |
| $4\sigma_{510Y,\pi}$ | Igamma | 0.1 | 1 | 0.28 | 0.24;0.36 |
| $4\sigma_{510Y,i}$ | Igamma | 0.1 | 1 | 0.29 | 0.24;0.38 |

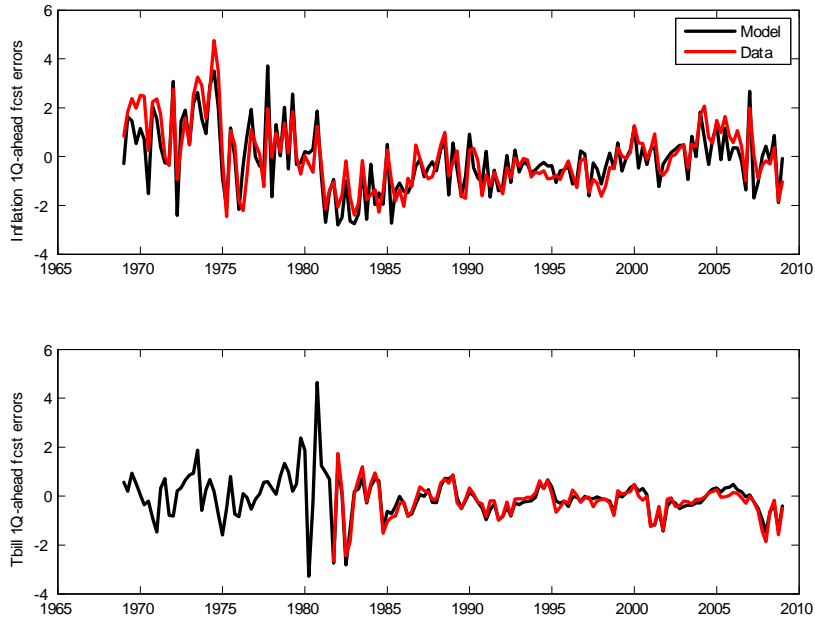


Figure 6: **One-quarter ahead forecast errors.** Top panel shows the data and model predictions for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.

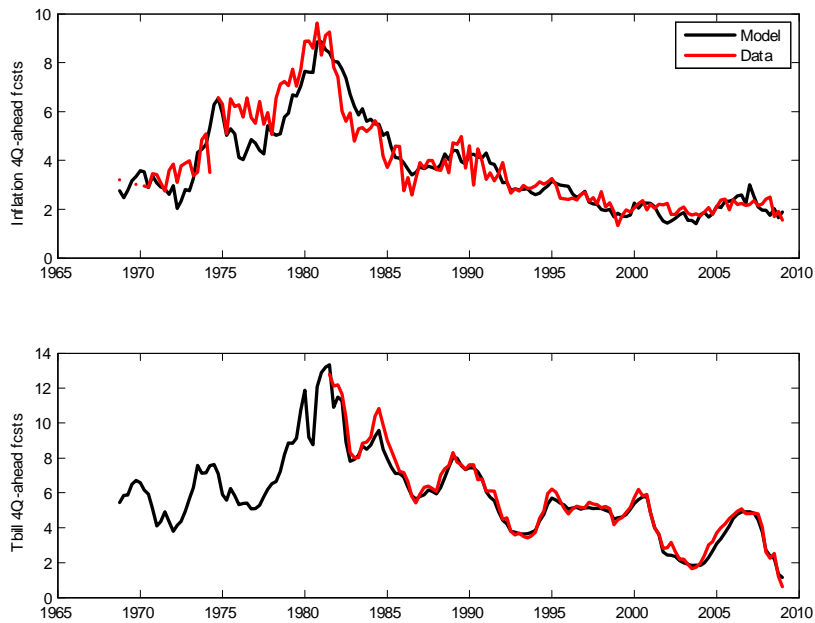


Figure 7: **Four-quarter ahead forecast errors.** Top panel shows the data and model predictions for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.

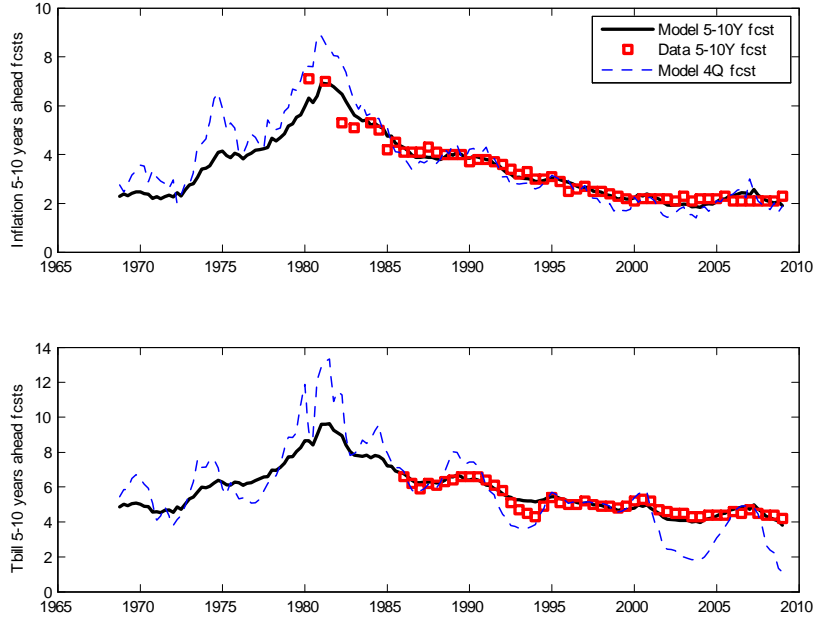


Figure 8: **Five-to-ten year forecast errors.** Top panel shows the data and model predictions at both the 4 quarter and 5-10 year horizon for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.

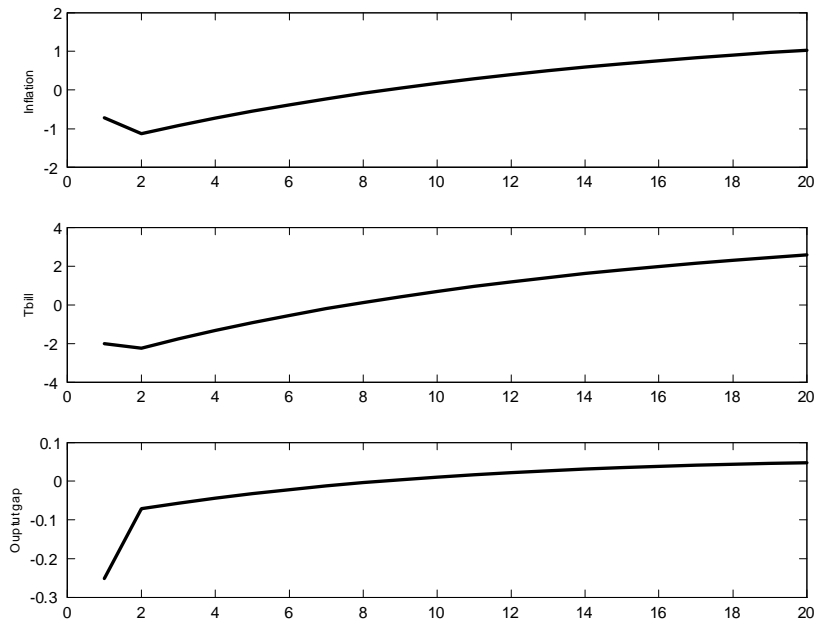


Figure 9: Impulse response functions for a natural-rate disturbance and a Taylor Rule.

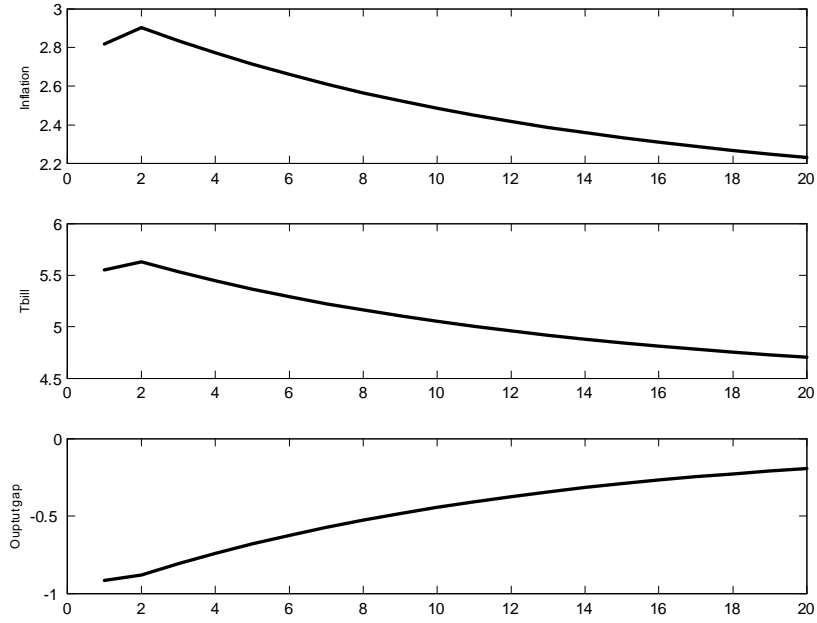


Figure 10: Impulse response functions for a cost-push disturbance and a Taylor Rule.

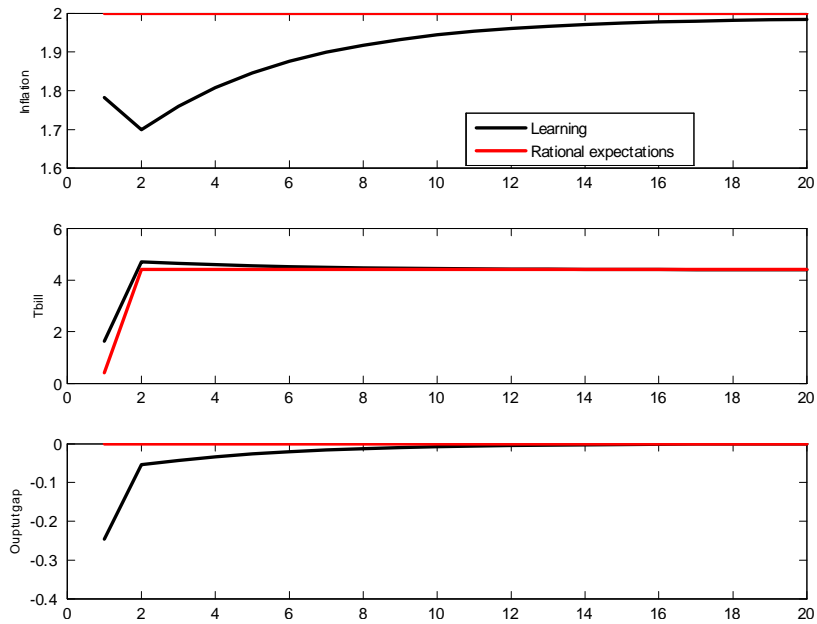


Figure 11: Impulse response for a natural-rate shock under optimal policy and diagonal gain beliefs. Learning is the black line; rational expectations the red line.

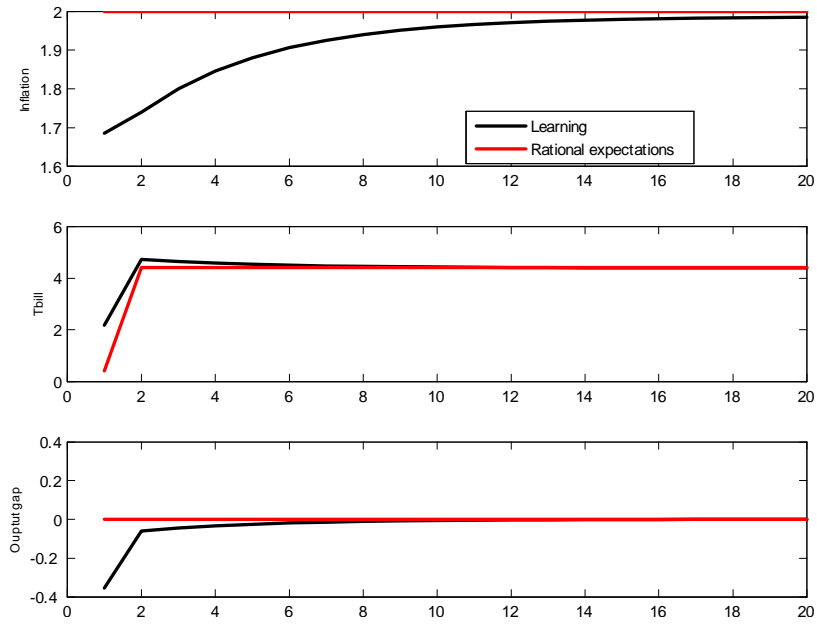


Figure 12: Impulse response for a natural-rate shock under optimal policy and low weight on output gap stabilization. Learning is the black line; rational expectations the red line.

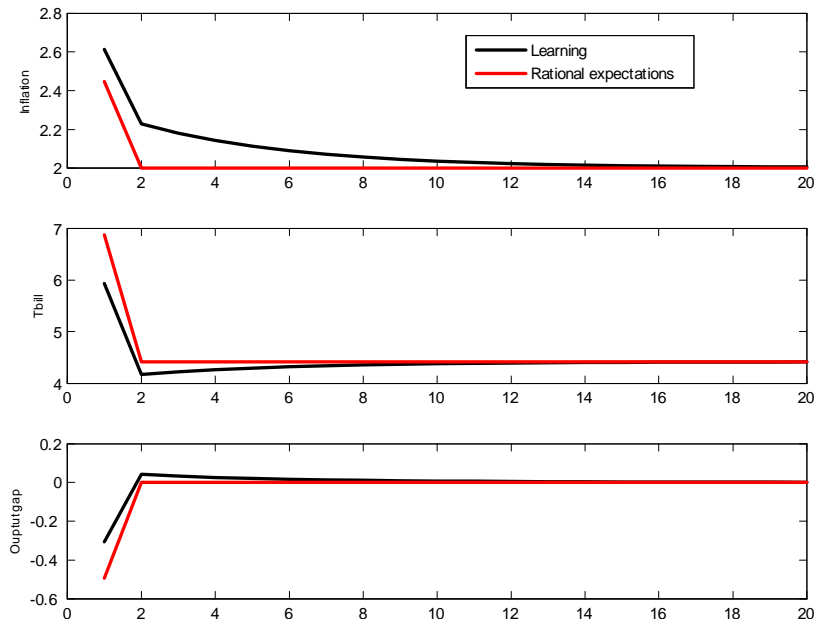


Figure 13: Impulse response for a cost-push shock under optimal policy. Learning is the black line; rational expectations the red line.

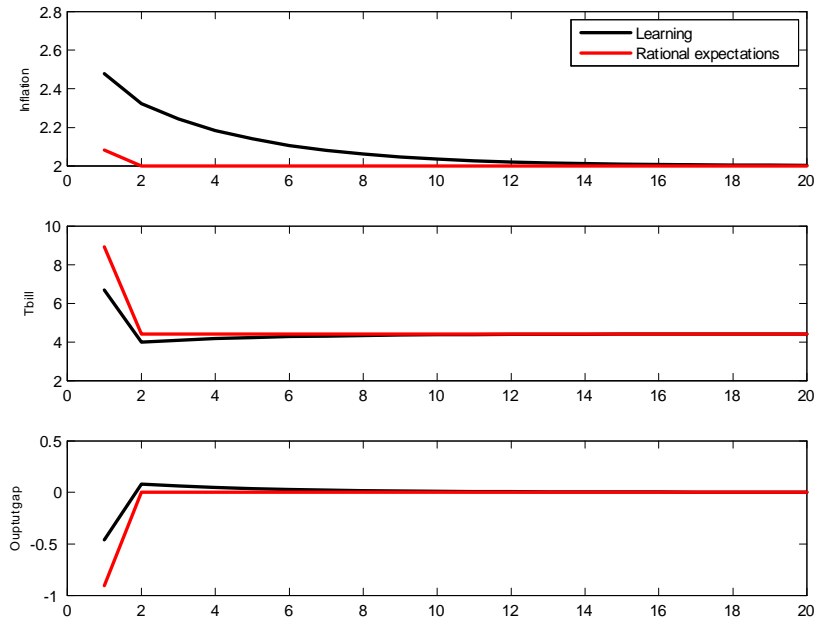


Figure 14: Impulse response for a cost-push shock under optimal policy with low weight on output gap stabilization. Learning is the black line; rational expectations the red line.

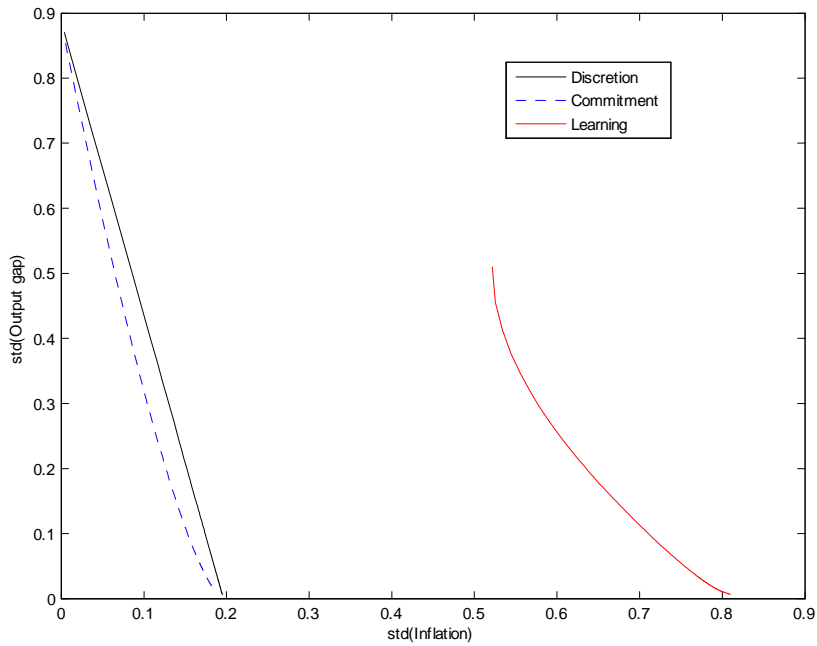


Figure 15: Taylor Frontier under optimal commitment, optimal discretion and optimal policy with long-term drift in expectations.

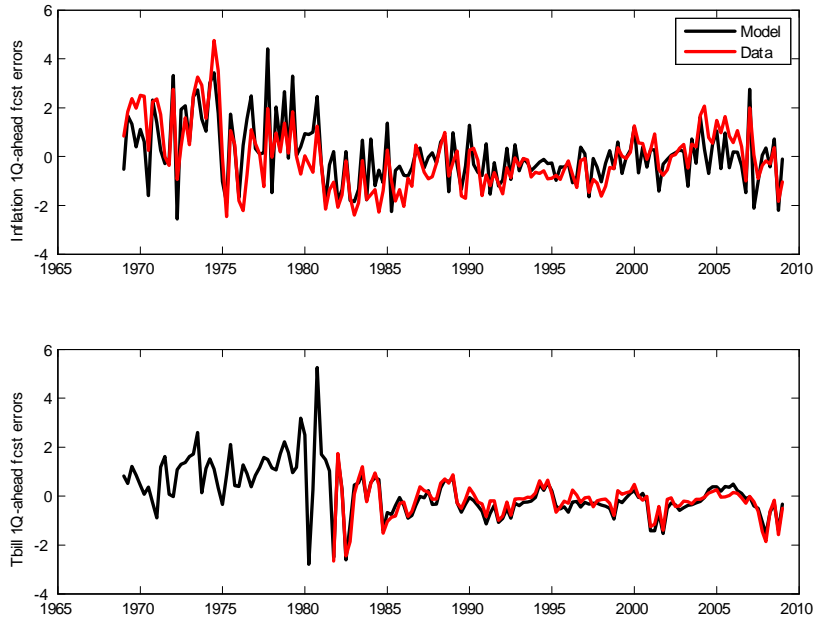


Figure 16: **One-quarter ahead forecast errors with constant gain beliefs.** Top panel shows the data and model predictions for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.

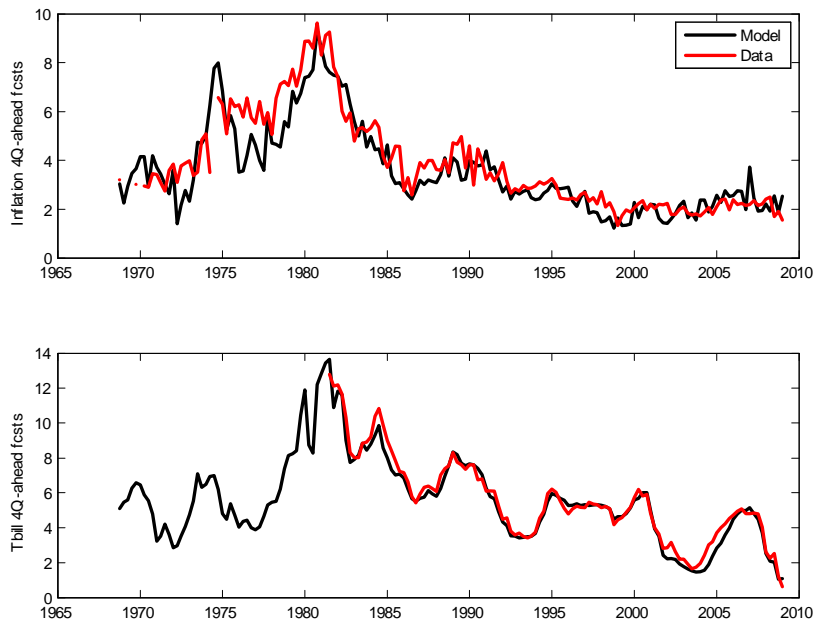


Figure 17: **Four-quarter-ahead forecast errors with constant gain beliefs.** Top panel shows the data and model predictions for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.

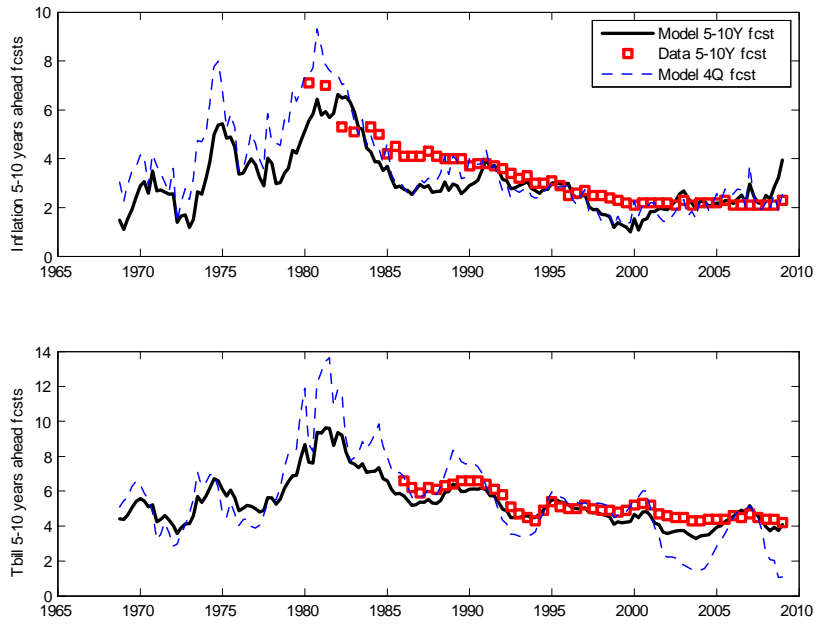


Figure 18: **Five-to-ten year forecast errors with constant gain beliefs.** Top panel shows the data and model predictions at both the 4 quarter and 5-10 year horizon for the GDP deflator. The bottom panels gives corresponding values for 3-month Treasury Bill.