

Optimal Disinflation Under Learning*

Timothy Cogley Christian Matthes
New York University New York University

Argia M. Sbordone
Federal Reserve Bank of New York

April 2010
(Preliminary and Incomplete)

Abstract

We examine how to design a disinflation when agents must learn the new policy rule. Erceg and Levin (2003) demonstrate that learning increases inflation persistence and the sacrifice ratio relative to what would occur under rational expectations. Accordingly, one might think that learning promotes gradualism. On the other hand, a gradual disinflation might retard the rate at which private agents learn, thus prolonging the transition and making it more turbulent. We sort out the competing forces in the context of a simple new Keynesian model. Despite the presence of sticky prices and learning, the optimal policy brings inflation down sharply and converges to the new target from below. This promotes rapid learning and minimizes the duration and turbulence of the transition.

1 Introduction

We examine the problem of a newly-appointed central bank governor who inherits a high average inflation rate from the past. The bank has no official inflation target and lacks the political authority unilaterally to set one, but it has some flexibility in choosing how to implement a vague mandate. We assume that the new governor's preferences differ from his predecessor and, in particular, that he wants to disinflate. What is the optimal policy? Does it imply a sharp or a gradual reduction of inflation?

*For comments and suggestions, we thank Boyan Jovanovic, Thomas Sargent, and seminar participants at Duke, the Federal Reserve Board, and UQAM. The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Sargent (1982) studies an analogous problem in which the central bank not only has a new governor but also is undergoing a fundamental institutional reform. He argues that by suitably changing the rules of the game, the government can persuade the private sector in advance that a low-inflation policy is its best response. In that case, the central bank can engineer a sudden, sharp disinflation at low cost in terms of lost output. Sargent discusses several historical examples that support his theory, emphasizing the institutional changes that establish credibility. Our scenario differs from Sargent's in two important ways. We take institutional reform off the table, assuming instead just a change of personnel. We also take away knowledge of the new policy and assume that the private sector must learn about it from observed outcomes.¹ Our scenario is more like the Volcker disinflation than the end of interwar hyperinflations.

In contrast to Sargent's case studies, the Volcker disinflation was quite costly. Erceg and Levin (2003) and Goodfriend and King (2005) explain the high cost by pointing to a lack of transparency and credibility. Erceg and Levine contend that Volcker's policy lacked transparency, and they develop a model in which the private sector must learn about the central bank's long-run inflation target. In their model, learning makes inflation more persistent relative to what it would be under full information, increases the sacrifice ratio, and produces output losses like those seen in the early 1980s.² Goodfriend and King claim that Volcker's disinflation lacked credibility because no important changes were made in the rules of the game. Because the private sector was initially unconvinced that Volcker would disinflate, the new policy collided with expectations inherited from the old regime and brought about a deep recession.

The analysis of Erceg, Levin, Goodfriend, and King is positive. Their objective is to understand why the Volcker disinflation was costly. In contrast, we address a normative question, viz. what policy is optimal when the private sector is not persuaded in advance and must learn the new policy? Are Sargent's recommendations still wise? Or should the disinflation be more gradual, e.g. as recommended by Gordon (1982)?

We study this question in the context of a dynamic new Keynesian model. We modify the model so that target inflation is positive, as in Ascari (2004) and Sbordone (2007), and we replace rational expectations with learning. We assume that the central bank commits to a simple Taylor-type rule and that it chooses its coefficients by minimizing a discounted quadratic loss function. The private sector learns about the new policy coefficients via Bayesian updating, and the central bank takes learning into account when solving its decision problem.

We find that some policies that work well under rational expectations generate

¹This is tantamount to assuming that the private sector does not know the new governor's preferences.

²See also Orphanides and Williams (2005) and Milani (2007).

catastrophic losses under learning. In particular, implementing the policy that is optimal under full information would be very costly. Taking the learning transition into account is essential when formulating a disinflationary policy.

We also find that an optimal Taylor-type rule reacts only to inflation and not to output. The policy that is optimal under learning pulls inflation sharply downward and overshoots the new long-run target. After the initial decline, inflation converges to the new target from below. This policy promotes rapid learning and results in a small sacrifice ratio. In contrast, reacting to output as well is counterproductive because it retards learning and makes the transition more turbulent. Thus, despite the presence of sticky prices and learning, the optimal policy has more in common with the recommendations of Sargent (1982) than with those of Gordon (1982).

The remainder of the discussion is organized as follows. Section 2 outlines a dynamic new Keynesian model with non-zero trend inflation. Building on Sbordone (2007), we explicitly account for a positive inflation target and work with a more accurate log-linear approximation that tracks the evolution of agents' beliefs.

Section 3 describes how the private sector learns. Our approach differs from much of the macro-learning literature, in particular the branch emanating from Marcet and Sargent (1989a, 1989b) and Evans and Honkapohja (2001). Models in that tradition typically assume that agents use reduced-form statistical representations such as vector autoregressions (VARs) for forecasting. They also commonly assume that agents update parameter estimates by some version of recursive least squares. Our approach to learning retains more of the structure of a DSGE model. The agents who inhabit our model also utilize VARs for forecasting, but their VARs satisfy cross-equation restrictions analogous to those in rational-expectations models. As a consequence, we are able to articulate a tight link between the actual and perceived laws of motion (the ALM and PLM, respectively). Indeed, our equilibrium concept is close to rational expectations. Our agents know the ALM up to the unknown monetary policy parameters, and their PLM is the perceived ALM (i.e., the ALM evaluated at their current estimate of the policy coefficients). Because they know the functional form of the ALM, our agents can use Bayes theorem to update beliefs, efficiently exploiting all available information about the new monetary-policy rule.³

Section 4 summarizes the calibration, and section 5 discusses how we solve the optimal policy problem. Sections 6 and 7 describe optimal policies under full information and learning, respectively. Section 8 concludes by listing directions for future research.

³Conceptually, our approach detaches learning from bounded rationality. Although the latter might turn out to be important for understanding disinflation, we think that developing a framework that disentangles the two is helpful. This paper takes a step in that direction.

2 A dynamic new-Keynesian model with positive target inflation

The timing protocol is critical in learning models, so we begin by describing it. Then we go on to describe the model's structure and our approximation methods, taking beliefs as given. We defer a discussion of how beliefs are updated to section 3.

2.1 The timing protocol

Private agents enter period t with decision rules and beliefs about policy coefficients inherited from $t - 1$. They form expectations and make current-period plans accordingly. At the same time, the central bank sets the systematic part of its instrument rule based on information inherited from $t - 1$. After that, shocks are realized and current-period outcomes are determined. After observing those outcomes, private agents update their estimates. They treat estimated parameters as if they were known with certainty, and they formulate new decision rules, which they carry forward to period $t + 1$.⁴

2.2 The baseline model

Our baseline model is a simple dynamic new Keynesian model modified so that target inflation is positive and that agents take expectations with respect to subjective beliefs. This section presents a log-linearized version of the model. For details on how we arrived at this representation, please see appendix A.

Monetary policy is determined in accordance with a Taylor-type rule. Private-sector behavior is characterized by two blocks of equations, a conventional intertemporal IS curve and an Ascari-Sbordone version of the aggregate supply curve.

2.2.1 Monetary policy

We assume that the central bank commits to a Taylor rule in difference form,

$$i_t - i_{t-1} = \psi_\pi(\pi_{t-1} - \bar{\pi}) + \psi_x(y_{t-1} - y_{t-2}) + \varepsilon_{it}, \quad (1)$$

where i_t is the nominal interest rate, π_t is inflation, y_t is log output, and ε_{it} is an *i.i.d.* policy shock. The timing assumption follows McCallum (1999) and fits

⁴This timing protocol differs slightly from the convention in DSGE models, but it is convenient because it circumvents a simultaneity between the determination of outcomes, the formation of expectations, and the updating of beliefs.

conveniently within the timing protocol described above.⁵ The policy parameters are $\psi = [\bar{\pi}, \psi_\pi, \psi_x, \sigma_i^2]'$, where $\bar{\pi}$ represents the central bank's long-run target for inflation, ψ_π and ψ_x are feedback parameters on the inflation gap and output growth, respectively, and σ_i is the standard deviation of the policy shock.

We adopt this form because others have shown that it performs well in environments like ours. For instance, Coibion and Gorodnichenko (2008) establish that a rule of this form ameliorates indeterminacy problems in Calvo models with positive target inflation. Orphanides and Williams (2007) demonstrate that it performs well under least-squares learning.⁶ More generally, a number of economists (e.g. ???) have argued that the central bank should engage in a high degree of interest smoothing and react to output growth rather than the output gap. Erceg and Levin (2003) also contend that the output growth rate is the appropriate measure to include in an estimated policy reaction function for the U.S.

We assume that private agents know the form of the policy rule but not the policy coefficients. At any given date, the perceived policy rule is

$$i_t - i_{t-1} = \psi_{\pi t-1}(\pi_{t-1} - \bar{\pi}_{t-1}) + \psi_{x t-1}(y_{t-1} - y_{t-2}) + \varepsilon_{it}, \quad (2)$$

where $\psi_{t-1} = [\bar{\pi}_{t-1}, \psi_{\pi t-1}, \psi_{x t-1}, \sigma_{it-1}^2]$ represents the current estimate of ψ . In accordance with our timing protocol, expectations formed at the beginning of period t depend on beliefs based on information available through the end of period $t - 1$. Hence private-sector decisions at t depend on ψ_{t-1} .

The perceived law of motion that we derive below depends on the perceived policy rule (2). The actual law of motion depends on actions taken by the central bank and decisions made by the private sector. Hence the actual law of motion involves both the actual policy (1) and the perceived policy (2).

Finally, we assume that the central bank chooses ψ by minimizing a discounted quadratic loss function,

$$L = E_0 \sum_t \beta^t [(\pi_t - \bar{\pi})^2 + \lambda(y_t - \bar{y})^2], \quad (3)$$

taking private-sector learning into account.⁷ To keep things simple, we first analyze a

⁵McCallum (1999) contends that monetary policy rules should be specified in terms of lagged variables, on the grounds that the Fed lacks good current-quarter information about inflation, output, and other right-hand variables. This is especially relevant for decisions taken early in the quarter, in accordance with our timing protocol.

⁶Orphanides and Williams (2007) postulate that neither the agents nor the central bank know the true structure of the economy, and they replace rational expectations with least-squares learning. They also assume that the central bank cannot observe the natural rates and that the bank estimates them via a simple updating rule. They show that an optimized Taylor rule in difference form dominates optimized standard Taylor-style rules when learning and time-varying natural rates interact. Whether this result applies also to the case of learning with constant natural rates is an open question.

⁷Gaspar et al (2006) distinguish between an unsophisticated central bank - one that accounts for

problem where the central bank arbitrarily sets $\bar{\pi}$ and σ_i and optimizes with respect to ψ_π and ψ_x . At a later stage, we hope to analyze the choice of $\bar{\pi}$ as well.

2.2.2 Approximation methods

We make use of two approximations when solving private-sector decision problems. As usual, the first-order conditions take the form of non-linear expectational difference equations. We follow the standard practice of log-linearizing around a steady state and solving the resulting system of linear expectational difference equations. However, we expand around the agents' perceived steady state in period t rather than around the true steady state.

The true steady state \bar{x} is the deterministic steady state associated with the true policy coefficients ψ . We define the perceived steady state \bar{x}_t as the long-horizon forecast associated with the current estimate ψ_t . The private sector's long-run forecast \bar{x}_t varies through time because changes in the central bank's inflation target have level effects on nominal variables and also on some real variables (Ascari 2004). Since perceptions of $\bar{\pi}$ change as agents update their beliefs, so do their estimates of steady states.

We chose to expand around \bar{x}_t instead of \bar{x} because the first-order conditions for consumers and firms exist in the minds of private agents, and \bar{x}_t better reflects their state of mind at date t . Since private agents eventually learn the true policy coefficients in our model, the perceived steady state \bar{x}_t converges to the full-information steady state \bar{x} as t grows large, but the two differ along the transition path.

Our second approximation involves the assumption that agents treat the current estimate ψ_t as if it were known with certainty when reformulating their decision rules. Kreps (1998) calls this as an 'anticipated-utility' model. In the context of a single-agent decision problem, Cogley and Sargent (2008) compare the resulting decision rules with exact Bayesian decision rules, and they demonstrate that the approximation is quite good as long as the agent is not too risk averse. Like a log-linear approximation, this imposes a form of certainty equivalence, for it implies that decision rules are the same regardless of the degree of parameter uncertainty. This approximation is standard in the macro-learning literature.

2.2.3 A new-Keynesian IS curve

As usual, we assume that a representative household maximizes expected utility subject to a flow budget constraint. Its first-order condition is a conventional consumption Euler equation. After log-linearizing, we derive a version of the new-

the beliefs of the public but not the dynamic process of learning – and a sophisticated central bank that also takes the learning process into account. Our setting corresponds to the latter assumption.

Keynesian *IS* curve,⁸

$$y_t - \bar{y}_{t-1} = E_t^*[(y_{t+1} - \bar{y}_{t-1}) - (i_t - \pi_{t+1} - r) + (g_{t+1} - g)] + \varepsilon_{yt}, \quad (4)$$

where r is the steady-state real interest rate, \bar{y}_{t-1} is the private sector's long-run forecast for output at date t , and g_t and ε_{yt} are shocks.

This equation differs in a number of ways from a textbook *IS* equation. One difference concerns the choice of the expansion point. Our timing protocol makes long-run forecasts at the beginning of period t a function of information through $t - 1$. Hence decisions at date t depend on estimates of steady states formed at the end of $t - 1$. In addition, our anticipated-utility assumption implies that $E_t^* \bar{y}_{t+1} = \bar{y}_{t-1}$, explaining the appearance of \bar{y}_{t-1} on the right-hand side of equations (4).

A second difference concerns the expectation operator E_t^* . Because the *IS* curve exists in the minds of the representative household, E_t^* represents forecasts formed with respect to its perceived law of motion, which is derived below. In contrast, the central bank takes expectations with respect to the actual law of motion, which we denote by E_t . Since the central bank's information set subsumes that of the private sector,⁹ the law of iterated expectations implies $E_t^*(E_t x_{t+j}) = E_t^*(x_{t+j})$ for any random variable x_{t+j} and $j \geq 0$.¹⁰ Because the central bank can reconstruct private forecasts, it also follows that $E_t(E_t^* x_{t+j}) = E_t^*(x_{t+j})$. But $E_t x_{t+j} \neq E_t^* x_{t+j}$.

Finally, two shocks appear in (4), a persistent shock g_t to the growth rate of technology,

$$g_t = (1 - \rho_g) g + \rho_g g_{t-1} + \varepsilon_{gt}, \quad (5)$$

and a white-noise shock ε_{yt} . We introduce the latter to prevent the private sector's learning problem from becoming singular.

2.2.4 A new-Keynesian Phillips curve

We adopt a purely-forward looking version Calvo's (1983) pricing model. A continuum of monopolistically competitive firms produce a variety of differentiated intermediate goods that are sold to a final-goods producer. Firms that produce the intermediate goods reset prices at random intervals. In particular, with probability $1 - \alpha$, an intermediate-goods producer has an opportunity to reset its price, and with probability α its price remains the same. Thus we abstract from indexation, in accordance with the estimates of Cogley and Sbordone (2008). Since pricing and supply decisions depend on the beliefs of private agents, we again log-linearize around

⁸See appendix A for details.

⁹We assume that the central bank knows the private sector's prior over ψ .

¹⁰We assume that both conditional expectations exist.

perceived steady states, obtaining the following block of equations:¹¹

$$\begin{aligned} \pi_t - \bar{\pi}_{t-1} &= \kappa_{t-1}(y_t - \bar{y}_{t-1}) + \beta_{t-1}E_t^*(\pi_{t+1} - \bar{\pi}_{t-1}) + \varsigma_{t-1}(\delta_t - \bar{\delta}_{t-1}) \\ &\quad + \gamma_{1t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}] + u_t + \varepsilon_{\pi t}, \end{aligned} \quad (6)$$

$$\phi_t = \gamma_{2t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}], \quad (7)$$

$$\delta_t - \bar{\delta}_{t-1} = \lambda_{1t-1}(\pi_t - \bar{\pi}_{t-1}) + \lambda_{2t-1}(\delta_{t-1} - \bar{\delta}_{t-1}). \quad (8)$$

This representation differs in four ways from textbook versions of the *NKPC*. First, a variable

$$\delta_t \equiv \log \left(\int_0^1 (p_t(i)/P_t)^{-\theta} di \right), \quad (9)$$

that measures the resource cost induced by cross-sectional price dispersion has first-order effects on inflation and other variables. If target inflation were zero, this would drop out of a first-order expansion.

Second, higher-order leads of inflation appear on the right-hand side of (6). To retain a first-order form, we introduce an intermediate variable ϕ_t that has no interesting economic interpretation and add equation (7). This is simply a device for obtaining a convenient representation.

Third, the *NKPC* parameters depend on both deep parameters and estimates of target inflation $\bar{\pi}_t$:¹²

$$\begin{aligned} \beta_t &= \beta(1 + \bar{\pi}_t), \\ \kappa_t &= (1 + \nu)[1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}][1 - \alpha\beta(1 + \bar{\pi}_t)^\theta]/\alpha(1 + \bar{\pi}_t)^{\theta-1}, \\ \gamma_{1t} &= \beta\bar{\pi}_t[1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}], \\ \gamma_{2t} &= \alpha\beta(1 + \bar{\pi}_t)^{\theta-1}, \\ \varsigma_t &= \nu[1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}][1 - \alpha\beta(1 + \bar{\pi}_t)^\theta]/\alpha(1 + \bar{\pi}_t)^{\theta-1}, \\ \lambda_{1t} &= \alpha\theta\bar{\pi}_t(1 + \bar{\pi}_t)^{\theta-1}/(1 - \alpha(1 + \bar{\pi}_t)^{\theta-1}), \\ \lambda_{2t} &= \alpha(1 + \bar{\pi}_t)^\theta. \end{aligned} \quad (10)$$

The deep parameters are the subjective discount factor β , the probability $1 - \alpha$ that an intermediate-goods producer can reset its price, the elasticity of substitution across varieties θ , and the Frish elasticity of labor supply $1/\nu$. As Cogley and Sbordone (2008) emphasize, the deep parameters are invariant to changes in policy, but the *NKPC* parameters are not. The latter change as beliefs about $\bar{\pi}_t$ are updated.

Finally, we assume two cost-push shocks, a persistent shock u_t that follows an *AR*(1) process,

$$u_t = \rho_u u_{t-1} + \varepsilon_{ut}, \quad (11)$$

and a white-noise shock $\varepsilon_{\pi t}$. As before, we include the latter to prevent a singularity in the learning problem.

¹¹Again, please see appendix A for details.

¹²The *NKPC* parameters collapse to the usual expressions when $\bar{\pi}_t = 0$.

3 Learning about monetary policy

We assume that private agents and central bank know the model of the economy and the form of the policy rule, but that private agents do not know the policy parameters. Instead, they use Bayes theorem to learn about them. This section describes how that is done. We first conjecture a perceived law of motion (PLM) and then derive the actual law of motion (ALM) under the PLM. Then we verify that the PLM is the perceived ALM. Having verified that private agents know the ALM up to the unknown policy coefficients, we use the ALM to derive the likelihood function. We assume that agents combine the likelihood function with their prior over policy parameters and search for the posterior mode.

3.1 The perceived law of motion

By stacking the IS curve (equation 4) and aggregate supply block (equations 6-8) with the exogenous shocks and perceived monetary policy rule (equation (2)), we can represent the private sector's beliefs about the economy as a system of linear expectational difference equations,

$$A_{t-1}S_t = B_{t-1}E_t^*S_{t+1} + C_{t-1}S_{t-1} + D_{t-1}\varepsilon_t. \quad (12)$$

The state vector is $S_t = [\pi_t, \phi_t, \delta_t, u_t, y_t, g_t, y_{t-1}, i_t, 1]'$ and the innovation vector is $\varepsilon_t = [\varepsilon_{\pi t}, \varepsilon_{ut}, \varepsilon_{xt}, \varepsilon_{gt}, \varepsilon_{it}]'$. The matrices A_t, B_t, C_t , and D_t are specified in appendix B. They vary with time because they depend on estimates of the policy coefficients ψ_t and the *NKPC* parameters.

The reduced-form of (12) is a *VAR*(1),

$$S_t = F_{t-1}S_{t-1} + G_{t-1}\varepsilon_t. \quad (13)$$

Under the timing protocol, expectations formed at the beginning of period t depend on $t - 1$ information. Hence

$$E_t^*S_{t+1} = F_{t-1}^2S_{t-1}. \quad (14)$$

After substituting expectations into the structural form (12) and re-arranging terms, we find

$$S_t = A_{t-1}^{-1}[B_{t-1}F_{t-1}^2 + C_{t-1}]S_{t-1} + A_{t-1}^{-1}D_{t-1}\varepsilon_t. \quad (15)$$

By matching coefficients with those in (13), we find $G_{t-1} = A_{t-1}^{-1}D$ and that F_{t-1} solves the matrix quadratic equation

$$B_{t-1}F_{t-1}^2 - A_{t-1}F_{t-1} + C_{t-1} = 0. \quad (16)$$

A solution for F_{t-1} can be computed by solving a generalized eigenvalue problem (Uhlig 1999). Equation (13) is the PLM.

3.2 The actual law of motion

To find the actual law of motion, we stack the IS curve, aggregate supply block, and shocks along with the actual policy rule. This results in another system of expectational difference equations,

$$A_{t-1}S_t = B_{t-1}E_t^*S_{t+1} + C_{at-1}S_{t-1} + D_{t-1}\varepsilon_t. \quad (17)$$

The state and innovation vectors are the same as in (12), as are the matrices A_{t-1} , B_{t-1} , and D_{t-1} . In addition, all rows of C_{at-1} agree with those of C_{t-1} except for the one corresponding to the monetary-policy rule. In that row, the true policy coefficients ψ replace the estimated coefficients ψ_{t-1} . See appendix B for details.

After substituting the PLM forecasts (14) for $E_t^*S_{t+1}$, we obtain a system of backward-looking difference equations,

$$A_{t-1}S_t = (B_{t-1}F_{t-1}^2 + C_{at-1})S_{t-1} + D\varepsilon_t. \quad (18)$$

Private-sector beliefs about monetary policy matter for price-setting and consumption decisions. Those beliefs are encoded in A_{t-1} , B_{t-1} , D_{t-1} , F_{t-1} , and some elements of C_{at} . Outcomes also depend on actions taken by the central bank, which involve the actual policy rule. That is encoded in other elements of C_{at} . We find the ALM under the PLM by premultiplying both sides of (18) by A_{t-1}^{-1} ,

$$S_t = H_{t-1}S_{t-1} + J_{t-1}\varepsilon_t, \quad (19)$$

where

$$\begin{aligned} H_{t-1} &= A_{t-1}^{-1}(B_{t-1}F_{t-1}^2 + C_{at-1}), \\ J_{t-1} &= A_{t-1}^{-1}D. \end{aligned} \quad (20)$$

The equilibrium law of motion is a *VAR* with time-varying parameters and conditional heteroskedasticity, as in Cogley and Sargent (2005) and Primiceri (2006).

An intriguing feature of the equilibrium is that the drifting parameters ψ_t have a lower dimension than the conditional mean parameters $vec(H_{t-1})$. This is qualitatively consistent with a finding of Cogley and Sargent (2005), who reported that drift in an analog to $vec(H_{t-1})$ is confined to a lower dimensional subspace. The conditional variance in (19) differs slightly from their representations, however, so the model does not agree with their identifying restrictions. Another difference is that the model involves temporary drift during a learning transition while their VARs involve perpetual drift.

3.3 The PLM is the perceived ALM

The reduced-forms ALM and PLM are both $VAR(1)$ processes. Furthermore, the reduced-form ALM matrices solve

$$\begin{aligned} H_{t-1} &= A_{t-1}^{-1}(B_{t-1}F_{t-1}^2 + C_{at-1}), \\ J_{t-1} &= A_{t-1}^{-1}D_{t-1}, \end{aligned} \quad (21)$$

while the reduced-form PLM matrices solve

$$\begin{aligned} F_{t-1} &= A_{t-1}^{-1}(B_{t-1}F_{t-1}^2 + C_{t-1}), \\ G_{t-1} &= A_{t-1}^{-1}D_{t-1}. \end{aligned} \quad (22)$$

It follows that $J_{t-1} = G_{t-1}$.

By inspection, one can verify that C_{at-1} and C_{t-1} are identical except for the row corresponding to the monetary-policy rule. In that row, C_{at-1} depends on actual policy coefficients ψ while C_{t-1} depends on perceived policy coefficients ψ_{t-1} . If we were to interview the agents in the model and ask what is their perceived ALM, they would answer by replacing ψ in C_{at-1} with ψ_{t-1} , thus obtaining C_{t-1} . After that substitution, the equation defining H_{t-1} coincides with the condition for F_{t-1} . Since their perceived $H_{t-1} = F_{t-1}$ and $J_{t-1} = G_{t-1}$, it follows that the PLM is the perceived ALM.

3.4 The likelihood function

We collect the observables in a vector $X_t = [\pi_t, u_t, y_t, g_t, i_t]'$. The other elements of the state allow us to express the model in first-order form, but convey no additional information beyond that contained in the history of X_t . Using the prediction-error decomposition, the likelihood function for data through period t can be expressed as¹³

$$p(X^t|\psi) = \prod_{j=1}^t p(X_j|X^{j-1}, \psi). \quad (23)$$

Since the private sector knows the ALM up to the unknown policy parameters, they can use it to evaluate the terms on the right-hand side of (23).

According to the ALM, X_t is conditionally normal with mean and variance

$$\begin{aligned} m_{t|t-1}(\psi) &= e_X H_{t-1}(\psi) S_{t-1}, \\ V_{t|t-1}(\psi) &= e_X J_{t-1} V_\varepsilon(\psi) J_{t-1}' e_X', \end{aligned} \quad (24)$$

where e_X is a selection matrix defined in appendix B, $H_{t-1}(\psi)$ is the ALM conditional-mean array evaluated at a particular value of ψ , and $V_\varepsilon(\psi)$ is the variance of the

¹³The date-zero term drops out because outcomes at date zero were determined under the old regime. Hence X_0 is uninformative for the new policy coefficients.

innovation vector ε_t also evaluated at ψ . It follows that the log-likelihood function is

$$\ln p(X^t|\psi) = -\frac{1}{2} \sum_{j=1}^t \left\{ \ln |V_{j|j-1}(\psi)| - \frac{1}{2} [X_j - m_{j|j-1}(\psi)]' V_{j|j-1}^{-1}(\psi) [X_j - m_{j|j-1}(\psi)] \right\}. \quad (25)$$

3.5 The prior and posterior

We posit that private agents have a prior $p(\psi)$ over the policy coefficients. At each date t , they find the log posterior kernel by summing the log likelihood and log prior. Because of our anticipated-utility assumption, their decisions depend only on a point estimate, not on the entire posterior distribution. Among the various point estimators from which they can choose, we assume they adopt the posterior mode,

$$\psi_t = \arg \max (\ln p(X^t|\psi) + \ln p(\psi)). \quad (26)$$

Roughly speaking, this boils down to recursive nonlinear least squares, tilted in the direction of the prior.

Notice that agents take into account that past outcomes were influenced by past beliefs. They are not recursively estimating a conventional rational-expectations model. By inspecting the ALM and PLM, one can verify that past values of the conditional mean $m_{j|j-1}(\psi)$ and the conditional variance $V_{j|j-1}(\psi)$ depend on past estimate ψ_{j-1} as well as the current candidate ψ . Past estimates are bygones at t and are held constant when agents update the posterior mode.

Notice also that the estimates are based not just on the policy rule, but also on the equations for inflation and output. The agents in our model efficiently exploit all available information about ψ , taking advantage of cross-equation restrictions implied by the ALM to sharpen their estimates. Our approach would reduce to single-equation estimation of the policy rule if the conditional mean and variance of variables other than i_t were independent of ψ and prediction errors for i_t were orthogonal to the others. In our model, changes in ψ alter the conditional mean and variance of inflation and output, and prediction errors are correlated across equations. Hence full-information estimation is more informative. How much the cross-equation restrictions matter is a question we hope to explore later.

4 Calibration

We begin with the following baseline calibration. We take the parameters of the pricing model from estimates in Cogley and Sbordone (2008),

$$\alpha = 0.6, \quad \beta = 0.99, \quad \theta = 10. \quad (27)$$

Notice in particular that we abstract from indexation or other backward-looking pricing influences. We think this is realistic, as it is supported by the estimates in our earlier paper, but eventually we will explore the robustness of our policy recommendations to this assumption.

We calibrate the labor-disutility parameters as

$$\nu = 0.5, \quad \chi = 1. \tag{28}$$

The parameter ν is the inverse of the Frisch elasticity of labor supply – i.e., the elasticity of hours worked with respect to the real wage, keeping constant the marginal utility of consumption. The literature provides a large range of values for this elasticity, typically high in the macro literature (the extreme case in the Rogerson-Hansen lottery model is ∞), and low in the labor literature. Our calibration implies a Frisch elasticity of 2 and represents a compromise between the two literatures. We think this is reasonable given that the model abstracts from wage rigidities.

With respect to parameters governing the shocks, we abstract from average growth, setting $g = 0$. For the persistent shocks u_t and g_t , we take estimates from Cogley, Sargent and Primiceri (2009),

$$\begin{aligned} \rho_u &= 0.4, & 100\sigma_u &= 0.12, \\ \rho_g &= 0.27, & 100\sigma_g &= 0.5. \end{aligned} \tag{29}$$

For the white noise shocks ε_{yt} and $\varepsilon_{\pi t}$ we set

$$\sigma_\pi = \sigma_y = 0.01/4 \tag{30}$$

Finally, we adopt a standard calibration for loss-function parameters. We assume that the new governor arbitrarily sets $\bar{\pi} = 0.005$, which corresponds to a 2 percent annual rate for target inflation. We also assume that the central bank assigns equal weights to annualized inflation and the output gap. Since the model expresses inflation as a quarterly rate, this corresponds to $\lambda = 1/16$.

5 The central bank’s decision problem

The new governor appears at date 0 and formulates a new policy rule. We initially assume that the governor chooses the reaction coefficients ψ_y, ψ_π to minimize expected loss (equation 3), with target inflation $\bar{\pi}$ and the standard deviation of policy shocks σ_i being set exogenously. This is mainly for computational convenience. Later we hope to expand the project to endogenize the latter parameters as well. Also at date 0, private agents formulate a prior over the new policy coefficients. We assume that the central bank can observe or infer their prior, e.g. by surveying private sector expectations. The disinflation commences at date 1.

The central bank’s decision problem depends on the private sector’s prior and also on the initial state of the economy. We describe those elements and before turning to the optimal-policy problem.

5.1 The private sector’s prior

We are interested in a scenario like the end of the Great Inflation. To add an element of realism, we calibrate the prior to match estimates of the policy rule for the period 1965-1979. We assume that the policy rule had the same functional form as (1) during that period, and we estimate $\bar{\pi}$, ψ_π , ψ_y , and σ_i^2 by OLS. In addition, we assume that the policy parameters are independent a priori,

$$p(\psi) = p(\bar{\pi})p(\psi_\pi)p(\psi_y)p(\sigma_i^2), \quad (31)$$

and we adopt gamma priors for each of the components. A gamma specification is convenient because it enforces non-negativity. Hyperparameters for the marginal priors are calibrated so that the prior mode and variance match the OLS point estimate and its variance. In this way, we ensure that prior $p(\psi)$ encodes information from the period leading up to the Volcker disinflation. Table 1 reports the resulting prior mode and standard deviation, and figure 1 portrays the marginal distributions.

Table 1: The Private Sector’s Prior

	$\bar{\pi}$	ψ_π	ψ_y	σ_i
Mode	0.0116	0.043	0.12	0.0033
Standard Deviation	0.013	0.08	0.04	0.01




Figure 1: The Private Sector’s Prior

The prior mode for $\bar{\pi}$ is 0.0116, implying an annualized target-inflation rate of 4.6 percent. The reaction coefficients are both close to zero, with the output coefficient being slightly larger than the inflation coefficient. Policy shocks are large in magnitude and account for a substantial fraction of the total variation in nominal interest. Finally, the prior is disperse, signifying that the old regime lacked transparency.

5.2 Initial conditions

We initialize the economy at the steady state under the old regime. Taking the prior mode as a point estimate, this implies $\pi_0 = 0.0116$, $y_0 = -0.0732$, and $i_0 = 0.0217$, where inflation and nominal interest are expressed as quarterly rates. The initial condition for inflation is lower than that which Volcker confronted. Later we hope to examine how a higher initial value would alter the optimal policy.

5.3 Evaluating expected loss and finding the optimal policy

If the model fell into the linear-quadratic class, the loss function could be evaluated and optimal policy computed using methods developed by Mertens (2009a, 2009b). In our model, the central bank has quadratic preferences, and many elements of the state transition equation are linear, but learning introduces a nonlinear element. Since that element is essential, we prefer to retain it and use other methods for evaluating expected loss.

We proceed numerically. We start by specifying a grid of values for the policy parameters ψ . Then, for each node on the grid, we simulate 100 sample paths for inflation, output, nominal interest, and other variables, updating private-sector estimates ψ_t by numerical maximization at each date. The sample paths are each 25 years long, and we set the terminal loss to zero, representing a decision maker with a long but finite horizon. Along each sample path, we calculate realized loss, and then we average realized loss across sample paths to find expected loss. The optimal rule among this family is the node with smallest expected loss.¹⁴

6 Optimal policy under full information

To highlight the role of learning, we begin by describing the optimal policy under full information. Here we assume that the central bank announces the new policy at date 0, that the private sector regards the announcement as fully credible, and that they form expectations rationally from the outset in accordance with the new rule.

¹⁴An alternative procedure would be to write down a dynamic program and solve it numerically, as in Gaspar, Smets, and Vestin (2006, 2009). This is feasible in models with a low-dimensional state vector, but in our model it runs afoul of the curse of dimensionality.

First we check whether the policy rule delivers a unique non-explosive rational-expectations equilibrium. We can verify numerically that the REE is determinate for all values $\bar{\pi} \in (0, 0.095/4)$, $\psi_\pi \in (0.05, 4)$, and $\psi_y \in (0, 1.5)$. For higher values of $\bar{\pi}$, the equilibrium becomes indeterminate for low values of ψ_π and high values of ψ_y . For instance, when $\bar{\pi} = 0.10/4$, outcomes become indeterminate when $\psi_\pi = 0.05$ and $\psi_y > 1.3$, and the indeterminacy region expands as $\bar{\pi}$ increases.¹⁵ Happily, the troublesome values of $\bar{\pi}$ are outside the region we consider in our disinflation experiment.

Having verified determinacy for the relevant range of $\bar{\pi}$, we focus on a target inflation rate of 2 percent per annum and look for policies that perform well under full information. Figure 2 depicts a contour map for expected loss as a function of ψ_π and ψ_y . Expected loss is normalized by dividing by loss under the optimal rule, so that the level of the contour lines represents the gross deviation from the optimal policy. A relative loss of unity represents the optimal policy, a value slightly above unity represents a policy whose performance is almost as good, and large numbers represent policies that perform poorly.




Figure 2: Relative loss under full information

With target inflation equal to 2 percent per annum, the optimal policy under full information sets $\psi_\pi = 1$ and $\psi_y = 1.1$. Expected loss is quite flat, however, in the neighborhood of the optimum. The lowest contour (the innermost ellipse) is at 1.01, and all parameter combinations that lie inside involve an increase in loss of 1

¹⁵These findings are consistent with the results of Coibion and Gorodnichenko (2008). They find that responding to output growth rather than the output gap helps ensure determinacy, and they report a generalized Taylor principle that “determinacy appears to be guaranteed for any positive and sustainable inflation rate when the Fed responds to both inflation and current output growth by more than one-for-one.” (p. 12)

percent or less. Furthermore, the isoclines are nearly parallel to the ψ_y -axis in the neighborhood of the minimum, implying that expected loss is relatively insensitive to the response coefficient on output. Expected loss is more sensitive to ψ_π , especially as it declines toward zero. The lesson we draw from the figure is that a good policy under full information requires a sufficiently strong response to inflation and that the precise value for the response to output is secondary.

Figure 2 portrays a disinflation under the FI-optimal policy. Recall that the economy is initialized in the steady state of the old regime and that the disinflation commences at date 1. The figure depicts the response of inflation, output, and nominal interest gaps, which are defined as deviations from the steady state of the new regime.¹⁶

The nominal interest rate rises sharply at date 1, causing inflation to decline sharply and overshoot the new target. After that, inflation converges to its new target from below. This rolls back the price level, counteracting the effects of high inflation at date zero. As Woodford (2003) explains, a partial rollback of the price level is a feature of optimal monetary policy under commitment. Intuitively, a credible commitment on the part of the central bank to roll back price increases in the future restrains a firm's incentive to increase its price today. It is interesting that the optimal simple rule under full information also has this feature.




Figure 3: Transition under the FI-optimal policy

The sharp initial increase in the nominal interest rate also causes the output gap to fall below zero. Since inflation and output are both below target at date 1, the central bank sharply cuts the interest rate at date 2. This damps the output

¹⁶Thus, values at date 0 represent the difference between the steady states of the old and new regimes.

loss associated with the disinflation and initiates a saw-tooth pattern in the nominal interest rate. Convergence to the new steady state is rapid, with output and inflation gaps disappearing within 2 years. Quite a lot of nominal interest volatility occurs along the way.

The cumulative output loss is small, especially in comparison with Keynesian models of the 1970s. After 8 periods, inflation is close to its new target, which is roughly 2.6 percentage points below the old target. The cumulative loss in output is also around 2.6 percent. We define the sacrifice ratio as the cumulative loss in output over 8 quarters divided by change in target inflation. Accordingly, the sacrifice ratio is approximately 1 percent of lost output per percentage point of inflation. In contrast, Okun (1978) cites estimates of 10 percentage points of GNP per point of inflation.

The main reason why the sacrifice ratio is small under full information is that the model has no indexation. For the reasons given in Cogley and Sbordone (2008), we believe this is empirically plausible. Although prices are sticky, the absence of indexation means that inflation is weakly persistent. Thus, under full information, Sargent's (1982) arguments go through with slight modification. Learning will complicate matters because it increases inflation persistence and the sacrifice ratio. How this alters optimal policy remains to be seen.

7 Optimal policy under learning

Now we withdraw knowledge of the new policy and assume instead that agents update estimates of the policy coefficients via Bayes' theorem. We first examine what would happen if the central bank were to adopt the FI-optimal policy described above (i.e., $\bar{\pi} = 0.02/4$, $\psi_{\pi} = 1$, $\psi_y = 1.1$, and $\sigma_i = 0.01/4$). Figure 4 portrays average responses of inflation, output, and nominal interest gaps under this policy. Because the learning model is nonlinear, these responses are calculated by taking the ensemble average across 100 sample paths.

As the figure demonstrates, the policy that is optimal under full information is disastrous under learning. Although this policy causes a mild recession when agents are fully informed, it generates catastrophic turbulence when they learn. Furthermore, the picture changes only slightly when the ensemble average is replaced with the sample path associated with the median loss, implying that extreme turbulence occurs along many sample paths, not just a few. This trajectory should not be taken too seriously because no central bank would persist in following this policy for more than a few quarters. Still, it seems safe to conclude that the FI-optimal rule performs badly when agents learn. It follows that accounting for learning is essential when disinflating.




Figure 4: Average responses under the FI-optimal policy.

Figure 5 portrays isoloss contours under learning as function of ψ_π and ψ_y . As before, we normalize by dividing by expected loss under the optimal rule. Unlike the full-information model, in which the contour map was flat over much of the parameter space, here expected loss varies by orders of magnitude, so we depict the natural log of expected loss.

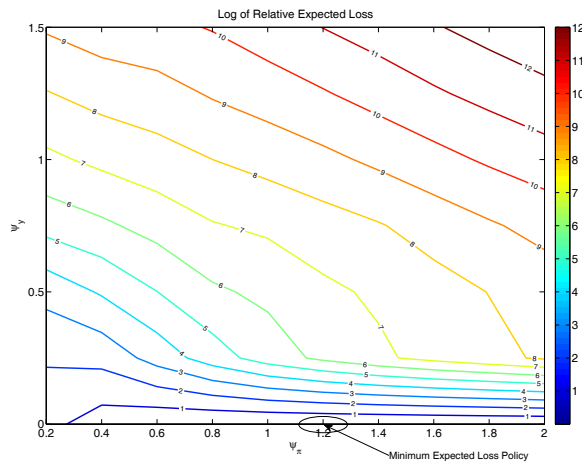


Figure 5: Log relative expected loss under learning

Recall that the FI-optimal policy sets $\psi_\pi = 1$ and $\psi_y = 1.1$. The gradient at that location points steeply downward in the direction of the origin. The combination that minimizes expected loss sets $\psi_\pi = 1.2$ and $\psi_y = 0$. Thus, the optimal Taylor rule reacts only to inflation and disregards movements in output. Indeed, holding ψ_π constant at 1.2, expected loss increases rapidly as ψ_y increases.

As shown in figure 6, turbulence along the transition path is much reduced under this policy.¹⁷ The nominal interest gap increases to 8 percent within 2 quarters, causing output and inflation gaps to turn sharply negative. The inflation gap overshoots the long-run target and reaches a trough of -5.5 percent in quarter 2, at which time the central bank reverses course and begins to reduce the nominal rate. The output gap also reaches a trough in quarter 2, bottoming out at about 3 percent below steady state. Inflation again converges to the new target from below, and the economy reaches the vicinity of its new steady state within 1 year. After 2 years, inflation has declined permanently by 2.6 percentage points, at the cost of a cumulative loss in output of 4.9 percent of a year's GDP. Thus, the sacrifice ratio is 1.9 percentage points of lost output per point of inflation. Although the sacrifice ratio is almost twice as large as under full information, it is only about one-fifth as large as in the models surveyed by Okun (1978).




Figure 6: Average responses under the minimum-expected-loss policy.

Because parameter estimates eventually converge to the true coefficients in our model, the learning economy approaches a full-information economy in the limit. In the full-information economy, expected loss under the learning-optimal policy is roughly 25 percent higher than under the FI-optimal rule (see figure 2). Hence the learning-optimal rule is almost as good in the limit and a lot better along the transition.

Next we explore why the transition is so much less turbulent. In figure 7, we plot mean estimates of the policy coefficients across 100 sample paths. The true coefficients are shown in red while estimates under the learning-optimal rule are portrayed in blue. Because initial beliefs are governed by the prior, date-0 estimates are

¹⁷Responses along the path with median loss resembles the ensemble average, implying that the shape of this path is not sensitive to outliers.

equal to the prior mode for all policies. Learning is rapid under the learning-optimal policy, with estimates converging to the neighborhood of the true policy coefficients after one quarter. From that time forward, the economy behaves essentially as a full-information economy, except with a different initial shock. The mildness of the transition is due to the rapid speed of learning.

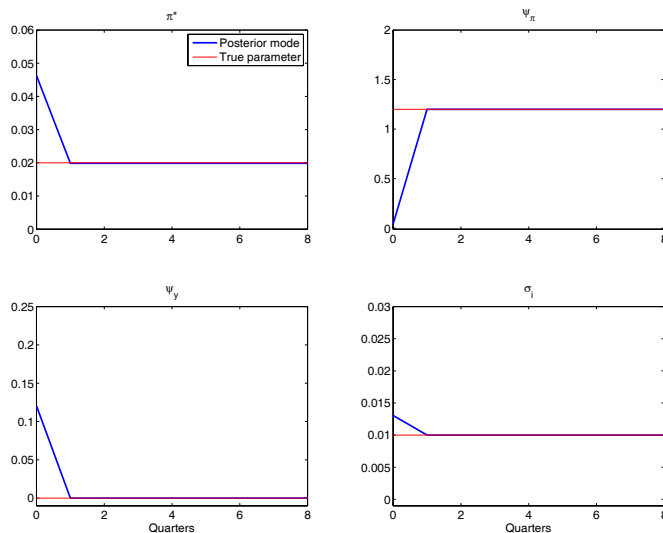


Figure 7: Average of posterior mode estimates when $\psi_\pi = 1.2$ and $\psi_y = 0$

In contrast, learning is slower when the central bank also reacts to output. Figure 8 portrays average parameter estimates under the FI-optimal rule. Estimates of $\bar{\pi}$ decline more slowly and are still near 4 percent per annum after 3 years. Estimates of ψ_π and ψ_y change very little in the first year. Because beliefs about $\bar{\pi}$, ψ_π , and ψ_y are essentially unchanged in the first year, the new policy collides with expectations left over from the old regime, and the transition takes on the characteristics of an adjustment to a contractionary monetary-policy shock. Indeed, under the FI-optimal policy, private agents initially believe the economy is being hit by a sequence of massive policy shocks, and their estimates of σ_i increase accordingly (see the lower right panel). These unexpected movements in the nominal rate make inflation and output gaps highly volatile. That variation eventually identifies ψ_π and ψ_y , but by then the economy is oscillating wildly and the central bank is seriously behind the curve.

Finally, figure 9 explores the consequences of a local departure from the learning-optimal rule, showing average parameter estimates when $\psi_\pi = 1.2$ and $\psi_y = 0.25$. The pattern shown here is closer to that for the FI-optimal rule than for the learning-optimal rule. The lesson we draw is that even a modest response to output retards learning.

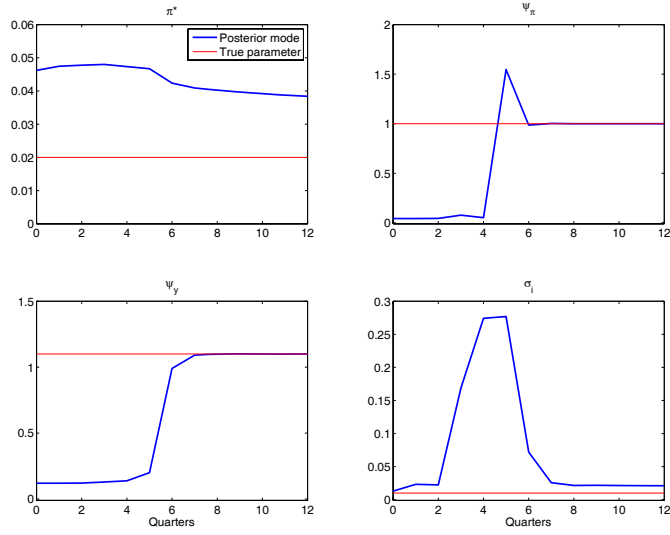


Figure 8: Average estimates when $\psi_\pi = 1$ and $\psi_y = 1.1$




Figure 9: Average estimates when $\psi_\pi = 1.2$ and $\psi_y = 0.25$

8 Conclusion

Learning enhances inflation persistence and increases the sacrifice ratio, making disinflation more costly. Despite that, an efficient disinflation involves a short sharp

shock and a fast recovery. The reason is that the choice of policy affects the speed of learning. In our model, it is optimal for the central bank to move quickly, sharply increasing nominal interest to bring inflation down fast, while at the same time disregarding movements in output. A single-minded focus on inflation speeds learning and promotes a smooth transition. In contrast, reacting to output muddies the signal being sent to the private sector and retards learning. This prolongs the transition and makes it more volatile.

We hope to add a number of extensions as this project moves forward. Work in progress includes endogenizing choice of the inflation target, considering other new Keynesian models that incorporate habit formation and indexation, and comparing Bayesian updating with other forms of learning.

A The baseline model

The demand side of the model consists of equilibrium conditions of a representative household for the optimal choice of consumption and hours of work. The household maximizes expected discounted utility

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\log(C_t) - \chi \frac{H_t^{1+\nu}}{1+\nu} \right), \quad (32)$$

subject to a flow budget constraint

$$E_t(Q_{t,t+1}Z_{t+1}) + P_t C_t = Z_t + W_t H_t + \int_0^1 \Psi_t(i) di, \quad (33)$$

where β is the subjective discount factor, C_t is consumption of the final good, with price P_t . The specification of the period utility – separable in consumption and hours and logarithmic in consumption – guarantees the existence of a balanced growth path. The variable

$$H_t = \int h_t(i) di \quad (34)$$

is an aggregate of the number of hours supplied by the household to firms in the intermediate-goods sector, and W_t is the economy-wide nominal wage. Intermediate goods producers earn profits amounting to $\int_0^1 \Psi_t(i) di$, which they rebate to the household. The variable Z_{t+1} is the state-contingent value of the portfolio of assets held by the household at the beginning of period $t + 1$, and $Q_{t,t+1}$ is a stochastic discount factor.

The first order-conditions for consumption and labor supply are respectively

$$C_t^{-1} = \beta E_t \left[C_{t+1}^{-1} \frac{R_t}{P_{t+1}} P_t \right], \quad (35)$$

and

$$w_t = \chi H_t^\nu C_t, \quad (36)$$

where $R_t = [E_t(Q_{t,t+1})]^{-1}$ is the gross nominal interest rate, and $w_t \equiv W_t/P_t$ is the real wage. Because there is no capital or government, the aggregate resource constraint is simply $C_t = Y_t$.

As we discuss below, the existence of a (possibly non stationary) technological progress Γ_t requires to define stationary variables as ratios to Γ_t , hence we appropriately transform (35) as

$$(C_t/\Gamma_t)^{-1} = \beta E_t \gamma_{t+1}^{-1} (C_{t+1}/\Gamma_{t+1})^{-1} \frac{R_t}{\Pi_{t+1}}, \quad (37)$$

where Π_t is the gross inflation rate: $\Pi_t = P_t/P_{t-1}$. Defining adjusted variables $C_t^a \equiv C_t/\Gamma_t$, $Y_t^a \equiv Y_t/\Gamma_t$ and $\gamma_t = \Gamma_t/\Gamma_{t-1}$, we have

$$(C_t^a)^{-1} = \beta E_t \gamma_{t+1}^{-1} (C_{t+1}^a)^{-1} \frac{R_t}{\Pi_{t+1}} \quad (38)$$

and

$$w_t^a = \chi H_t^\nu C_t^a \quad (39)$$

where $w_t^a \equiv w_t/\Gamma_t$ is the productivity adjusted real wage. The aggregate resource constraint is similarly transformed to give $C_t^a = Y_t^a$.

Imposing market clearing, the two equations are¹⁸

$$(Y_t^a)^{-1} = \beta E_t \gamma_{t+1}^{-1} (Y_{t+1}^a)^{-1} \frac{R_t}{\Pi_{t+1}} \quad (40)$$

and

$$w_t^a = \chi H_t^\nu Y_t^a. \quad (41)$$

The log-linearization of the equilibrium condition (40) gives the dynamic *IS* equation

$$\widehat{Y}_t^a = E_t \left(\widehat{Y}_{t+1}^a + \widehat{\gamma}_{t+1} + \widehat{\gamma}_{yt+1} - (i_t - E_t \pi_{t+1} - r) \right) \quad (42)$$

where $\widehat{Y}_t^a = \log Y_t^a - \log \overline{Y}_t^a$, $i_t = \log R_t$, $\widehat{\gamma}_t = \log \gamma_t - \log \gamma$, $i_t = \log R_t$, and r is the steady state real interest rate. $\gamma_{yt} \equiv \overline{Y}_t^a / \overline{Y}_{t-1}^a$.

The *IS* reported in the main text as eq. (4) is a transformation of (42) where rational expectations are replaced by learning and trend inflation is replaced by agents' perception with date $t - 1$ information, $\overline{\pi}_{t-1}$. Furthermore, since $E_t^* \widehat{\gamma}_{yt+1} = 0$, the term is suppressed. All steady state variables which are functions of trend inflation are similarly denoted with an overbar and subscript $t - 1$. The notation in eq. (4)

¹⁸The nominal interest rate is affected by the non-stationarity of inflation, but its ratio to inflation (and trend inflation) is stationary.

reflects a further simplified notation that avoids the use of superscript a : we set $y_t \equiv \log Y_t^a$, $\bar{y}_t \equiv \log \bar{Y}_t^a$ and $\bar{y}_{t-1} \equiv E_t^* \bar{y}_t$, $g_t \equiv \ln \gamma_t$ and $g \equiv \ln \gamma$.¹⁹

The equilibrium condition (41) is used in the next section to substitute out the real wage in the marginal cost expression of the supply side of the model.

A.1 The supply side

The supply side of the model consists of equilibrium conditions for a continuum of monopolistically competitive firms that produce intermediate goods and a final-good aggregating firm. These equilibrium conditions determine the dynamics of inflation in the model.

The final-good producer combines $y_t(i)$ units of each intermediate good i to produce Y_t units of the final good with technology

$$Y_t = \left[\int_0^1 y_t(i)^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}. \quad (43)$$

where θ is the elasticity of substitution across intermediate goods. The final good producer chooses the intermediate inputs to maximize its profits, taking the price of the final good P_t as given, determining demand schedules

$$y_t(i) = Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\theta}. \quad (44)$$

The zero-profit condition then determines the aggregate price level

$$P_t \equiv \left[\int_0^1 p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}}. \quad (45)$$

Intermediate firm i hires $h_t(i)$ units of labor of type i on an economy-wide competitive market to produce $y_t(i)$ units of intermediate good i with technology

$$y_t(i) = \Gamma_t h_t(i), \quad (46)$$

where Γ_t is an aggregate technological process.

Firms can reset prices at random intervals (we assume a Calvo price-setting mechanism), and we denote by $1 - \alpha$ the probability that an intermediate-goods producer has an opportunity to reset its price. The first order conditions of the optimal price-setting problem²⁰ and the evolution of aggregate prices jointly determine the dynamics of inflation in the model.

¹⁹The term R_t/Π_{t+1} is stationary, and we denote its (log) steady state (which is equal to the steady state value of the ratio of nominal interest rate to trend inflation) by r .

²⁰For simplicity we assume away wedges between the individual firm marginal cost and aggregate marginal costs.

In log-linear form, the supply side can be described by a pair of equations, known as a new Keynesian Phillips curve²¹

$$\begin{aligned}\pi_t - \bar{\pi}_{t-1} &= \tilde{\kappa}_{t-1}(mc_t - \bar{mc}_{t-1}) + \beta_{t-1}E_t^*(\pi_{t+1} - \bar{\pi}_{t-1}) \\ &\quad + \gamma_{1t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}] + u_t + \varepsilon_{\pi t}, \\ \phi_t &= \gamma_{2t-1}E_t^*[(\theta - 1)(\pi_{t+1} - \bar{\pi}_{t-1}) + \phi_{t+1}],\end{aligned}\tag{47}$$

where $\tilde{\kappa}_{t-1} = \kappa_{t-1}/(1 + \nu)$ and the other parameters are defined in expression (10) in the main text.

A.1.1 Marginal costs, output and price dispersion

In order to write the NKPC as a relation between inflation and output, we solve for marginal cost mc_t as function of aggregate output. Given the production function (46), the real marginal cost for the optimizing firm i (which is setting a new price at time t) is

$$mc_t(i) = \frac{W_t/P_t}{\Gamma_t} = w_t^a = mc_t,\tag{48}$$

where the last equality follows from the fact that the marginal cost doesn't depend on firm's i variables.

From equilibrium condition (41) w_t^a is function of aggregate hours, which are obtained by aggregating hours worked in each firm:

$$H_t \equiv \int_0^1 h_t(i) di = \int_0^1 \frac{y_t(i)}{\Gamma_t} di = Y_t^a \int_0^1 \left(\frac{p_t(i)}{P_t}\right)^{-\theta} di = Y_t^a \Delta_t,\tag{49}$$

where we denoted by Δ_t the following measure of price dispersion: $\Delta_t \equiv \int_0^1 \left(\frac{p_t(i)}{P_t}\right)^{-\theta} di$. Alternatively, we can write aggregate output as function of aggregate hours and a measure of price dispersion:

$$Y_t^a = \frac{H_t}{\Delta_t}.\tag{50}$$

Δ_t measures the resource cost induced by price dispersion in the Calvo model, in equilibrium. One can show that $\Delta_t \geq 1$, which implies that in equilibrium more hours are needed to produce the same amount of output (indeed, labor productivity is the inverse of the price dispersion index.) Price dispersion is therefore always a costly distortion in this model. By substituting expressions (41) and (49) in (48) we get

$$mc_t = \chi H_t^\nu Y_t^a = \chi \Delta_t^\nu (Y_t^a)^{1+\nu}.\tag{51}$$

²¹For further detail on the derivation of this Phillips curve, see Cogley and Sbordone (2008). The curve here in that the transformation of inflation into a stationary variable is obtained by dividing current (gross) inflation by perceived (rather than actual) trend inflation. The log-linearization is therefore defined around a point where perceived trend inflation and actual inflation are the same.

This expression shows that price dispersion creates a wedge between marginal costs and output. Substituting out marginal cost, we derive below the log-linear NKPC where both aggregate output and price dispersion are driving variables. Before doing that, we discuss the values of the variables in steady state.

A.1.2 Steady-state relations

From the definition of Δ_t we can derive that²²

$$\Delta_t = (1 - \alpha) (\tilde{p}_t)^{-\theta} + \alpha \Pi_t^\theta \Delta_{t-1},$$

where for ease of notation we indicate by \tilde{p}_t the relative price of the firms that optimizes at t : $\tilde{p}_t \equiv p_t^*(i)/P_t$. Then substituting the value of \tilde{p}_t from the evolution of aggregate prices

$$\tilde{p}_t = \left[\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} \right]^{\frac{1}{1-\theta}}, \quad (52)$$

we get²³

$$\Delta_t = (1 - \alpha) \left(\frac{1 - \alpha \Pi_t^{\theta-1}}{1 - \alpha} \right)^{-\frac{\theta}{1-\theta}} + \alpha \Pi_t^\theta \Delta_{t-1}. \quad (53)$$

From this expression we then obtain a relationship between price dispersion and trend inflation in steady state:

$$\bar{\Delta}_t = \frac{1 - \alpha}{1 - \alpha \bar{\Pi}_t^\theta} \left(\frac{1 - \alpha \bar{\Pi}_t^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta}{\theta-1}}. \quad (54)$$

We can now use the relation between steady state marginal cost and steady state inflation, namely:

$$\bar{m}c_t = \frac{\theta - 1}{\theta} \frac{\left(1 - \alpha \bar{\Pi}_t^{\theta-1}\right)^{\frac{1}{1-\theta}}}{(1 - \alpha)^{\frac{1}{1-\theta}}} \left[\frac{1 - \alpha \beta (\bar{\Pi}_t)^\theta}{1 - \alpha \beta (\bar{\Pi}_t)^{\theta-1}} \right], \quad (55)$$

together with (51) evaluated in steady state,

$$\bar{m}c_t = \chi (\bar{Y}^a_t)^{1+\nu} \bar{\Delta}_t^\nu, \quad (56)$$

to obtain a relationship between inflation and output that should be satisfied in steady state. Equating (55) and (56), substituting $\bar{\Delta}_t$ from (54) and rearranging, we get

$$\bar{Y}^a_t = \left[\frac{\frac{\theta-1}{\theta} \frac{\left(1 - \alpha \bar{\Pi}_t^{\theta-1}\right)^{\frac{1}{1-\theta}}}{(1 - \alpha)^{\frac{1}{1-\theta}}} \left[\frac{1 - \alpha \beta (\bar{\Pi}_t)^\theta}{1 - \alpha \beta (\bar{\Pi}_t)^{\theta-1}} \right]}{\chi \left(\frac{1 - \alpha}{1 - \alpha \bar{\Pi}_t^\theta} \left(\frac{1 - \alpha \bar{\Pi}_t^{\theta-1}}{1 - \alpha} \right)^{\frac{\theta}{\theta-1}} \right)^\nu} \right]^{\frac{1}{1+\nu}} \quad (57)$$

²²This derivation follows Schmitt-Grohe and Uribe (2006, 2007).

²³This expression corresponds to the derivation of Yun (AER, 2005).

which can be interpreted as a “long run” Phillips curve relationship between inflation and output.

A.1.3 Log-linearizations

We start from a log-linear NKPC with marginal cost as forcing variable, and want to transform it into a log-linear Phillips curve in output deviations from steady state: $\widetilde{Y}_t^a \equiv Y_t^a / \overline{Y}_t^a$.

To do that we only need to obtain the log-linearization of (51) around \overline{mc}_t (as defined in (56)), and substitute it into the marginal cost *NKPC*. From (51) we get

$$\widehat{mc}_t = (1 + \nu) \widehat{Y}_t^a + \nu \widehat{\Delta}_t. \quad (58)$$

For the log-linearization of (53), we first let $\widehat{\Delta}_t = \log \Delta_t / \overline{\Delta}_t$ (the non-stationarity of Π_t implies that Δ_t is also non-stationary, but its ratio to trend is by definition stationary).²⁴ Then we log-linearize the resulting expression around a steady state where $\widetilde{\Delta} = \widetilde{\Pi} = 1$, obtaining

$$\widehat{\Delta}_t \simeq \lambda_{1t} \widehat{\Pi}_t + \lambda_{2t} \left(\widehat{\Delta}_{t-1} - \widehat{\gamma}_{\Delta t} \right), \quad (59)$$

where the parameters λ_{1t} and λ_{2t} are defined in the last two rows of (10) in the main text. They are time-varying because they depend on trend inflation. In the main text, for analogy with the other equations, for $\widehat{\Delta}_t$ we use the notation $\delta_t - \overline{\delta}_{t-1}$.

B Arrays for structural representations

The matrices entering the PLM are defined as:

$$A_t = \begin{bmatrix} 1 & 0 & -\varsigma_t & -1 & -\kappa_t & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda_{1t} & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (60)$$

²⁴In first order approximations around a steady state with zero inflation, the variable Δ_t can be ignored (the log deviation $\widehat{\Delta}_t$ would be a first order process with no real consequences for the stationary distribution of the other endogenous variables). But price dispersion must be taken into account if one analyzes economies with trend inflation and imperfect price indexation, as the one in this paper. (see Schmitt-Grohe and Uribe (2007)).

$$B_t = \begin{bmatrix} \beta_t + \gamma_{1t}(\theta - 1) & \gamma_{1t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \gamma_{2t}(\theta - 1) & \gamma_{2t} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (61)$$

$$C_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y \\ 0 & 0 & 0 & 0 & 0 & \rho_g & 0 & 0 & \mu_g \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \psi_{\pi t} & 0 & 0 & 0 & \psi_{yt} & 0 & -\psi_{yt} & 1 & -\psi_{\pi t} \bar{\pi}_t \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (62)$$

$$D_t = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (63)$$

The expressions for the intercepts in C_t are

$$\mu_{\pi t} = [1 - \beta_t - \gamma_{1t}(\theta - 1)]\bar{\pi}_t - \kappa_t \bar{y}_t - \varsigma_t \bar{\delta}_t, \quad (64)$$

$$\mu_{\phi t} = -\gamma_{2t}(\theta - 1)\bar{\pi}_t$$

$$\mu_{\delta t} = (1 - \lambda_{2t})\bar{\delta}_{t-1} - \lambda_{1t}\bar{\pi}_t$$

$$\mu_y = r - g$$

$$\mu_g = (1 - \rho_g)g, \quad (65)$$

where \bar{y}_t and $\bar{\pi}_t$ are private-sector estimates respectively of steady-state output and trend inflation, and r and g are steady state real interest rate and real growth respectively.

The matrices A_t , B_t , and D_t also appear in the ALM. However, C_t is replaced by

$$C_{at} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\pi t} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_{\phi t} \\ 0 & 0 & \lambda_{2t} & 0 & 0 & 0 & 0 & 0 & \mu_{\delta t} \\ 0 & 0 & 0 & \rho_u & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mu_y \\ 0 & 0 & 0 & 0 & 0 & \rho_g & 0 & 0 & \mu_g \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \psi_\pi & 0 & 0 & 0 & \psi_y & 0 & -\psi_y & 1 & -\psi_\pi \bar{\pi} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (66)$$

The selection matrix e_X used to evaluate the likelihood function is defined as

$$e_X = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}. \quad (67)$$

References

- [1] Ascari, G., 2004. Staggered prices and trend inflation: some nuisances. *Review of Economic Dynamics* 7, 642–667.
- [2] Calvo, G. 1983. Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics* 12, 383-398.
- [3] Cogley, T., G. Primiceri, and T.J. Sargent. 2009. Inflation-Gap Persistence in the U.S. Forthcoming *American Economics Journal - Macroeconomics*.
- [4] Cogley, T. and T.J. Sargent. 2005. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262–302.
- [5] Cogley, T. and T.J Sargent, 2008. Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making. *International Economic Review* 49, 185-221.
- [6] Cogley, T. and A.M. Sbordone. 2008. Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve. *American Economic Review* 98, 2101-2126.
- [7] Coibion, O. and Y. Gorodnichenko. 2008. Monetary Policy, Trend Inflation and the Great Moderation: An Alternative Interpretation. *NBER wp 14621*.

- [8] Erceg, C. and Levin, A. 2003. Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50 (4): 915-944.
- [9] Evans, G.W. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press: Princeton, N.J.
- [10] Gaspar, Vito, Frank Smets and David Vestin. 2006. Adaptive Learning, Persistence, and Optimal Monetary Policy. *Journal of the European Economic Association* 4, 376-385.
- [11] Gaspar, Vito, Frank Smets and David Vestin. 2009. Inflation Expectations, Adaptive Learning, and Optimal Monetary Policy. Unpublished manuscript, European Central Bank.
- [12] Goodfriend, M. and R.G. King, 2005. The Incredible Volcker Disinflation. *Journal of Monetary Economics*, 52 (5): 981-1016.
- [13] Gordon, R.J. 1982. Why Stopping Inflation May Be Costly: Evidence from Fourteen Historical Episodes, in *Inflation: Causes and Effects*, edited by Robert Hall, University of Chicago press, pp. 11-40.
- [14] Kreps, D. 1998. Anticipated Utility and Dynamic Choice, in D.P. Jacobs, E. Kalai, and M. Kamien, eds., *Frontiers of Research in Economic Theory*, (Cambridge: Cambridge University Press, 1998), 242-74.
- [15] Marcet, A. and T.J. Sargent. 1989a. Convergence of least-squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48, 337-368.
- [16] Marcet, A. and T.J. Sargent. 1989b. Convergence of least-squares learning in environments with hidden state variables and private information. *Journal of Political Economy* 97, 1306-1322.
- [17] McCallum, B.T., 1999. Issues in the design of monetary policy rules. In: Taylor, J.B., Woodford, M. (Eds.), In: *Handbook of Macroeconomics*, vol. 1C. Elsevier, Amsterdam.
- [18] Mertens, Elmar. 2009a. Managing Beliefs about Monetary Policy under Discretion. Unpublished manuscript, Federal Reserve Board.
- [19] Mertens, Elmar. 2009b. Discreet Commitments and Discretion of Policymakers with Private Information. Unpublished manuscript, Federal Reserve Board.
- [20] Milani, Fabio. 2007. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54 (7): 2065-82

- [21] Okun, A.M., 1978. Efficient disinflationary policies. *American Economic Review Papers and Proceedings* 68, 348–352.
- [22] Orphanides, A. and J.C. Williams, 2002. Robust monetary policy rules with unknown natural rate, *Brookings Papers on Economic Activity*, 2002:2, pp. 63-145.
- [23] Orphanides, A. and J.C. Williams, 2005. Imperfect Knowledge, Inflation Expectations and Monetary Policy, in B. Bernanke and M. Woodford eds, *The Inflation Targeting Debate*
- [24] Orphanides, A. and J.C. Williams, 2007. Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54: 1406-1435.
- [25] Primiceri, G. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72, 821-852.
- [26] Sargent, T.J. 1982. The Ends of Four Big Inflations, in *Inflation: Causes and Effects*, edited by Robert Hall, University of Chicago press, pp. 41-97.
- [27] Sbordone, A.M. 2007. Inflation persistence: Alternative interpretations and policy implications. *Journal of Monetary Economics* 54, 1311–1339
- [28] Schmitt-Grohe, S. and M. Uribe. 2006. Optimal Simple and Implementable Monetary and Fiscal Rules: Expanded version, NBER wp 12402.
- [29] Schmitt-Grohe, S. and M. Uribe. 2007. Optimal Simple and Implementable Monetary and Fiscal Rules. *Journal of Monetary Economics* 54 (6): 1702-1725.
- [30] Yun, Tack, 2005. Optimal Monetary Policy with Relative Price Distorsion. *American Economic Review*, vol 95 (1): 89-109.
- [31] Woodford, M. 2003. Interest and Prices. Princeton University Press: Princeton NJ.