

A Search for a Structural Phillips Curve*

Timothy Cogley
University of California, Davis

Argia M. Sbordone
Federal Reserve Bank of New York

Revised: February 2005

Abstract

The foundation of the New Keynesian Phillips curve is a model of price setting with nominal rigidities which implies that the dynamics of inflation are well explained by the evolution of real marginal costs. The objective of this paper is to analyze whether this is a structurally-invariant relationship. To assess this, we first estimate an unrestricted time-series model for inflation, unit labor costs, and other variables, and present evidence that their joint dynamics are well represented by a vector autoregression with drifting coefficients and volatilities, as in Cogley and Sargent (2004). Then, following Sbordone (2002, 2003), we apply a two-step minimum distance estimator to estimate deep parameters. Taking as given estimates of the unrestricted *VAR*, we estimate parameters of the NKPC by minimizing a quadratic function of the restrictions that the theoretical model imposes on the reduced form. Our results suggest that it is possible to reconcile a constant-parameter NKPC with the drifting-parameter *VAR*, and therefore we argue that the price-setting model is structurally invariant.

JEL Classification: E31

Keywords: Inflation; Phillips curve; time-varying VAR.

*We would like to thank seminar participants at the 2004 Society for Computational Economics Meeting in Amsterdam, the Federal Reserve Bank of New York, Duke University, and the Fall 2004 Macro System Committee Meeting in Baltimore for their comments. 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.

1 Introduction

Much of the modern analysis of inflation is based on the New Keynesian Phillips curve, a model of price setting with nominal rigidities which implies that the dynamics of inflation are well explained by the expected evolution of real marginal costs. A large empirical literature has been devoted to estimating the parameters of this curve, both as a single equation and in the context of general equilibrium models¹. One point of debate concerns whether the model can account for the persistence in inflation which is detected in the data. A common view is that this is possible insofar as a large enough backward-looking component is allowed. However, from a theoretical point of view this is not too satisfactory, since dependence on past inflation is introduced as an ad hoc feature.

Here we reconsider estimates of the New Keynesian Phillips curve in light of recent evidence from reduced form analyses that show significant instability in the parameters of the inflation process. In particular, Cogley and Sargent (2001, 2004) use a vector autoregression model with random-walk coefficients to describe inflation-unemployment dynamics in the U.S. and find strong evidence of coefficient drift. They interpret this as a reflection of the process by which policymakers learn the true model of the economy. A related debate has ensued on whether the more muted response of inflation and output to monetary policy in the 90's is due to a change in the conduct of monetary policy or to a change in the size of the shocks; see Bernanke and Mihov (1998), Stock and Watson (2002), Boivin and Giannoni (2002), and Sims and Zha (2004), among others.²

The question we ask in this paper is whether the NKPC can be regarded as a structural model of inflation dynamics in the sense of Lucas (1976), viz. whether the deep parameters that govern the evolution of inflation are invariant to changes in monetary policy rules, at least over the range experienced after World War II in the U.S.³ Among other things, we investigate whether variation in trend inflation alters estimates of key pricing parameters, how well a constant-parameter version of the NKPC approximates the evolving law of motion for inflation, how the new estimates alter the relative importance of forward- and backward-looking elements in the NKPC, and how the new estimates accord with microeconomic evidence on price changes.

¹Among others: Gali and Gertler (1999), Sbordone (2002, 2003), Kurmann (2002), and Linde (2002) for the U.S., Batini et al. (2002) for the U.K., Gagnon and Khan (2003) for Canada, Gali, Gertler and Lopez-Salido (2000) for the Euro area. For estimates in the context of general equilibrium models see Smets and Wouters (2002), Christiano et al. (2003), and Edge et al. (2003).

²For example, Sims and Zha argue that there is very little evidence for regime switching in the conditional mean parameters, but strong evidence for regime switching in structural disturbances.

³Our analysis does not address whether they are invariant to more extreme interventions. We doubt, for example, that the pricing parameters we estimate from recent U.S. data would well approximate hyperinflationary regimes.

To address these questions, we consider an extension of the discrete-time Calvo (1983) model of staggered price setting, with partial price indexation and strategic complementarities, and consider the form of its approximate solution in the case of non-zero steady-state inflation. This formulation allows us to consider the effects that different policy regimes, which we associate with different levels of trend inflation, have on the relationship between inflation and marginal costs.

Our approach to estimation follows Sbordone (2002, 2003) by exploiting the cross-equation restrictions of the extended Calvo pricing model for a reduced form *VAR*. The wrinkle is that in this paper the reduced form *VAR* has drifting parameters, as in Cogley and Sargent (2004). The estimation is in two steps. In the first, we estimate an unrestricted time series representation for the variables that drive inflation. This is a time-varying *VAR* for inflation, the labor share, GDP growth, and the federal funds rate (expressed on a discount basis), which is estimated as in Cogley and Sargent (2004) with U.S. data from 1960:1 to 2003:4. Then we estimate deep parameters by trying to satisfy the cross-equation restrictions implied by the theoretical model. If we can reconcile a constant-parameter NKPC with the drifting-parameter *VAR*, we say the price-setting model is structurally invariant.

Our estimates point to four conclusions. First, a constant-parameter version of a generalized Calvo model can indeed be reconciled with a drifting-parameter *VAR*. More than that, the model provides an excellent fit to the inflation gap. Second, although there is some weak evidence of changes in the frequency of price adjustment over time, the evidence falls short of statistical significance. Third, our estimates of the backward-looking indexation parameter concentrate on zero, suggesting that a purely forward-looking version of the model fits best. Finally, our estimates of the frequency of price adjustment are not too far from those of Bils and Klenow (2004), so the macro and micro evidence is in accord.

The paper is organized as follows. The next section derives the inflation dynamics for an extended Calvo model and characterizes the cross-equation restrictions that form the basis for the estimation. Section 3 describes the empirical methodology in more detail, section 4 discusses evidence on parameter drift in the *VAR*, and section 5 estimates and assesses the structural parameters. Section 6 concludes.

2 A Calvo model with positive trend inflation

The typical inflation equation derived from the Calvo model is obtained as a log-linear approximation to the equilibrium conditions around a steady state with zero inflation. The model therefore has implications for small fluctuations around the steady state (it links second moments of inflation and real marginal costs). Because we want to investigate the behavior of the model across possibly different policy regimes, and therefore want to allow for shifts in trend inflation, we consider a log-

linear approximation to the equilibrium conditions around a non-zero level of inflation. We show below that, unless there is perfect indexation of prices to the past level of inflation, the Calvo dynamics are more complicated, *and* the pricing model also has predictions for the long-run relationship between trend inflation and marginal cost.⁴

We start with the standard Calvo set-up of monopolistic competition and staggered price setting. We denote by $(1 - \alpha)$ the probability of setting price optimally, with $0 < \alpha < 1$, and we allow the fraction α of firms that do not reoptimize to partially index their price to the inflation level of the previous period. We denote by ϱ the indexation parameter, with $\varrho \in [0, 1]$. Finally, we do not allow capital to be re-allocated instantaneously across firms, and therefore take into account a discrepancy between individual and aggregate marginal costs.

With these assumptions, the equilibrium condition of the price-setting firms is

$$0 = E_t \sum_{j=0}^{\infty} \alpha^j R_{t,t+j} \quad (1)$$

$$\times \left\{ \prod_{k=1}^j \gamma_{y,t+k} \prod_{k=1}^j \pi_{t+k}^{\theta} \prod_{k=0}^{j-1} \pi_{t+k}^{-\varrho(\theta-1)} \left(x_t^{1+\theta\omega} - \frac{\theta}{\theta-1} s_{t+j} \prod_{k=1}^j \pi_{t+k}^{1+\theta\omega} \prod_{k=0}^{j-1} \pi_{t+k}^{-\varrho(1+\theta\omega)} \right) \right\}.$$

while the evolution of aggregate prices is described by⁵

$$1 = \left[(1 - \alpha) x_t^{1-\theta} + \alpha \pi_{t-1}^{\varrho(1-\theta)} \pi_t^{-(1-\theta)} \right]^{\frac{1}{1-\theta}}. \quad (2)$$

The notation is as follows: X_t is the relative price set by the representative optimizing firm, and $x_t = X_t/P_t$ denotes its relative price; S_t is the aggregate nominal marginal cost, and $s_t = S_t/P_t$ denotes real marginal cost; P_t is the aggregate price level, and $\pi_t = P_t/P_{t-1}$ is the gross rate of inflation; $\gamma_{yt} = Y_t/Y_{t-1}$ is the gross rate of output growth, and $R_{t,t+j}$ is a nominal discount factor between time t and $t + j$. In addition to the parameters α and ϱ already introduced, θ is the Dixit-Stiglitz elasticity of substitution among differentiated goods, and ω is the elasticity of marginal cost to firms' own output. The parameter ω enters the equilibrium condition (1) because we assume firm-specific capital: this assumption implies that the marginal cost of the optimizing firm differs from aggregate marginal cost by a function of its relative price, weighted by $\theta\omega$.⁶

Evaluating these two conditions at a steady state with gross inflation rate $\bar{\pi}$, we

⁴Few studies in the literature analyze the policy issues that arise in the context of the Calvo model when one allows for trend inflation. See for ex. Bakhshi et al. (2003), Sahuc (2004) and Ascari (2004).

⁵We provide the main results in the text, and some derivations in the appendix.

⁶The coefficient ω is particularly important because it affects whether there are strategic complementarities in pricing (see Woodford 2003).

get the following relationship between steady-state $\bar{\pi}$ and \bar{s} :

$$(1 - \alpha \bar{\pi}^{(\theta-1)(1-\varrho)})^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\varrho)(1+\omega)}}{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{\theta-\varrho(\theta-1)}} \right) = (1 - \alpha)^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{\theta}{\theta - 1} \right) \bar{s}. \quad (3)$$

Here we have defined by \bar{R} the one-period steady-state discount factor and by $\bar{\gamma}_y$ the steady-state growth rate of output.⁷

The extended Calvo equation is an approximate equilibrium condition obtained by log-linearizing conditions (1) and (2) around a steady-state with inflation $\bar{\pi}$ and then combining the results:

$$\begin{aligned} \hat{\pi}_t &= \tilde{\varrho} \hat{\pi}_{t-1} + \zeta \hat{s}_t + b_1 E_t \hat{\pi}_{t+1} + b_2 E_t \sum_{j=2}^{\infty} \gamma_1^{j-1} \hat{\pi}_{t+j} \\ &\quad + \chi (\gamma_2 - \gamma_1) (P_{\hat{R}_t} + P_{\hat{\gamma}_{yt}}) + u_t. \end{aligned} \quad (4)$$

$$P_{\hat{R}_t} \equiv E_t \sum_{j=0}^{\infty} \gamma_1^j \hat{R}_{t+j,t+j+1}, \quad (5)$$

$$P_{\hat{\gamma}_{yt}} \equiv E_t \sum_{j=0}^{\infty} \gamma_1^j \hat{\gamma}_{y,t+j+1}.$$

With standard notation, hat variables denote log-deviations from steady state values; i.e., for any variable x_t , $\hat{x}_t = \log(x_t/\bar{x})$. We include an error term u_t to account for the fact that this equation is an approximation and to allow for other possible misspecifications.

The coefficients of (4) are functions of the vector of structural parameters $\underline{\psi} = [\alpha \ \tilde{\beta} \ \theta \ \varrho \ \omega \ \bar{\pi}]'$, where $\tilde{\beta}$ is the steady state value of a modified real discount factor,⁸ and

$$\begin{aligned} \tilde{\varrho}(\underline{\psi}) &= \frac{\varrho}{\Delta}, \\ \zeta(\underline{\psi}) &= \frac{\frac{1-\alpha\xi_1}{\alpha\xi_1} \frac{1-\gamma_2}{1+\theta\omega}}{\Delta}, \\ b_1(\underline{\psi}) &= \frac{\frac{1-\alpha\xi_1}{\alpha\xi_1} \phi_1 + \gamma_2}{\Delta}, \\ b_2(\underline{\psi}) &= \frac{\frac{1-\alpha\xi_1}{\alpha\xi_1} \left[\frac{\theta(1-\varrho\gamma_1) + \varrho\gamma_1}{1+\theta\omega} \right] (\gamma_2 - \gamma_1)}{\Delta}, \\ \chi(\underline{\psi}) &= \frac{\frac{1-\alpha\xi_1}{\alpha\xi_1} \frac{1}{1+\theta\omega}}{\Delta}. \end{aligned} \quad (6)$$

⁷As we explain in appendix A, equation (3) involves some additional conditions on α , ϱ , θ , and the steady-state values $\bar{\pi}$, \bar{R} , and $\bar{\gamma}_y$ that are necessary in order that certain present values converge. Our estimates always satisfy those conditions.

⁸The parameter $\tilde{\beta} = \bar{q}\bar{\gamma}_y$, where \bar{q} is the steady state value of $q_{t,t+j}$, a real discount factor between period t and $t+j$. Since $\bar{q} = \bar{R}\bar{\pi}$, one can also write $\tilde{\beta} = \bar{R}\bar{\pi}\bar{\gamma}_y$.

Intermediate symbols used here are

$$\begin{aligned}
\xi_1 &= \bar{\pi}^{(\theta-1)(1-\varrho)}, \\
\xi_2 &= \bar{\pi}^{\theta(1-\varrho)(1+\omega)}, \\
\gamma_1 &= \alpha\tilde{\beta}\xi_1, \\
\gamma_2 &= \alpha\tilde{\beta}\xi_2, \\
\Delta &= 1 + \varrho\gamma_2 - \frac{1 - \alpha\xi_1}{\alpha\xi_1}\phi_0, \\
\phi_0 &= \frac{\varrho\theta(\gamma_1 - (1 + \omega)\gamma_2) - \varrho\gamma_1}{1 + \theta\omega}, \\
\phi_1 &= \frac{1}{1 + \theta\omega} [\gamma_2(1 + \theta\omega) + (\gamma_2 - \gamma_1)(\theta(1 - \varrho\gamma_1) + \varrho\gamma_1)].
\end{aligned} \tag{7}$$

Compared with the standard Calvo equation, obtained as an approximation around a point with zero inflation ($\bar{\pi} = 1$), relationship (4) includes, on the right-hand side, further leads of expected inflation as well as expectations of output growth and the discount rate far into the future. The standard Calvo equation emerges as a special case of (4) when $\bar{\pi} = 1$ (zero steady state inflation), or $\varrho = 1$ (perfect indexation). In that case, $\xi_1 = \xi_2 = 1$, implying $\gamma_1 = \gamma_2 = \alpha\tilde{\beta}$ and causing the terms in $\hat{R}_{t+j,t+j+1}$ and $\hat{\gamma}_{y,t+j}$ to cancel out. The other coefficients collapse to those of the standard Calvo equation,

$$\begin{aligned}
\tilde{\varrho}(\psi) &= \frac{\varrho}{1 + \varrho\tilde{\beta}}, \\
\zeta(\psi) &= \left(\frac{1 - \alpha}{\alpha} \frac{1 - \alpha\tilde{\beta}}{1 + \theta\omega} \right) \frac{1}{1 + \varrho\tilde{\beta}}, \\
b_1(\psi) &= \frac{\tilde{\beta}}{1 + \varrho\tilde{\beta}}, \\
b_2(\psi) &= 0.
\end{aligned} \tag{8}$$

We may draw various implications from a comparison of the coefficients defined in (6) with those defined in (8). For example, the presence of additional terms in equation (4) may create an omitted-variable bias in the estimate of the coefficient of marginal cost in the traditional Calvo equation, should the omitted terms be correlated with the marginal cost term. We comment more on this comparison later.

Here we want to emphasize the fact that the response of inflation to current marginal cost does vary with trend inflation. Indeed, none of the coefficients of the generalized Calvo equation, as defined in (6), are time invariant when trend inflation varies over time (provided $\varrho \neq 1$). But it could still be the case that the underlying parameters of the Calvo model, α , ϱ , and θ , are stable. These parameters govern key behavioral attributes involving the frequency of price adjustment, the extent of

indexation to past inflation, and the elasticity of demand. In the estimation discussed below, we allow the parameters (6) to vary with trend inflation, and we explore the time invariance of α , ϱ , and θ . In particular, we evaluate whether it is still possible, in an environment characterized by a changing level of trend inflation, to fit to the data a Calvo model in which frequency of price adjustment, degree of indexation, and elasticity of demand remain constant.

3 Empirical methodology

The previous section shows that, when derived as an approximate equilibrium condition around a non-zero value for trend inflation, the generalized Calvo model imposes restrictions on both the steady-state values and cyclical components of inflation and real marginal cost. These restrictions are encoded in equation (3) and (4), respectively. In addition, the NKPC parameters are themselves functions of the underlying parameters

$$\psi = [\alpha \quad \tilde{\beta} \quad \theta \quad \varrho \quad \omega \quad \bar{\pi}]', \quad (9)$$

as shown in equation (6). In this section, we explain how to estimate elements of ψ by exploiting conditions (3), (4), and (6). We are particularly interested in α , ϱ , and θ .⁹

Following Sbordone (2002, 2003), we adopt a two-step procedure for estimating these parameters. First we fit a reduced-form *VAR* to summarize the dynamic properties of inflation, real marginal cost, and the other variables that enter the generalized Calvo equation. Then we estimate α , ϱ , and θ by exploiting the cross-equation restrictions that the extended Calvo model implies for the *VAR*. The chief difference from Sbordone (2002, 2003) is that we model the reduced form as a time-varying *VAR*, in order to allow for the possibility of structural breaks. The breaks are manifested as changes in trend inflation, among other things, and our working hypothesis is that they reflect changes in monetary policy.

To illustrate our methodology, we consider first the case where the reduced form model is a *VAR* with constant parameters, and then show its extension to the case of a random coefficients *VAR* model. Suppose the joint representation of the vector time series $x_t = (\pi_t, s_t, R_t, \gamma_{yt})'$ is a *VAR*(p). Then, defining a vector $z_t = (x_t, x_{t-1}, \dots, x_{t-p+1})'$, we can write the law of motion of z_t in companion form as

$$z_t = \mu + Az_{t-1} + \varepsilon_{zt}. \quad (10)$$

From this process, we can express the conditional expectation of the inflation gap as

$$E(\hat{\pi}_t | \hat{z}_{t-1}) = e'_\pi A \hat{z}_{t-1}, \quad (11)$$

⁹As we explain below, ω is calibrated, and $\tilde{\beta}_t$ and $\bar{\pi}_t$ are calculated from the reduced-form estimates.

where we use the notation e_k for a selection vector that picks up variable k in vector z_t (e_k is a column vector with 1 in the position corresponding to variable k , and zero otherwise), and $\hat{z}_t = z_t - \mu_z$, where $\mu_z = (I - A)^{-1}\mu$.¹⁰

The vector z_t also contains all the other variables that drive inflation, so we can use (10) to compute all the conditional expectations that appear on the right-hand side of (4). Furthermore, we can obtain the conditional expectation of the inflation gap according to the model, by projecting the whole right hand side of (4) on \hat{z}_{t-1} ; the resulting expression for the expected inflation gap contains by construction all the restrictions of the theoretical model. Specifically, from expression (4), one obtains¹¹

$$\begin{aligned} E(\hat{\pi}_t | \hat{z}_{t-1}) &= \tilde{\rho} e'_\pi \hat{z}_{t-1} + \zeta e'_s A \hat{z}_{t-1} + b_1 e'_\pi A^2_{t-1} \hat{z}_{t-1} + b_2 e'_\pi \gamma_1 (I - \gamma_1 A)^{-1} A^3_{t-1} \hat{z}_{t-1} \\ &\quad + \chi (\gamma_2 - \gamma_1) e'_R (I - \gamma_1 A)^{-1} A \hat{z}_{t-1} \\ &\quad + \chi (\gamma_2 - \gamma_1) e'_y (I - \gamma_1 A)^{-1} A^2 \hat{z}_{t-1}. \end{aligned} \quad (12)$$

Equating the right-hand sides of (11) and (12), and observing that the equality must hold for any value of \hat{z}_{t-1} , we obtain a set of nonlinear cross-equation restrictions on the companion matrix A ,

$$\begin{aligned} e'_\pi A &= \tilde{\rho} e'_\pi I + \zeta e'_s A + b_1 e'_\pi A^2 + b_2 e'_\pi \gamma_1 (I - \gamma_1 A)^{-1} A^3 \\ &\quad + \chi (\gamma_2 - \gamma_1) e'_R (I - \gamma_1 A)^{-1} A \\ &\quad + \chi (\gamma_2 - \gamma_1) e'_y (I - \gamma_1 A)^{-1} A^2 \end{aligned} \quad (13)$$

The left-hand side of this equation follows from the conditional expectation of inflation implied by the unrestricted reduced-form model, and it reflects relationships that we estimate freely from the data. The right-hand side follows from the conditional expectation implied by the model; it defines a function $g(A, \psi)$ of the deep parameters ψ and the parameters of the VAR. We then define the difference between the ‘data’ and the ‘model’ as

$$F_1(\mu, A, \psi) = e'_\pi A - g(A, \psi). \quad (14)$$

Furthermore, we use equation (3), which relates the steady-state values of inflation and marginal cost, to define a second set of moment conditions,

$$F_2(\mu, A, \psi) = (1 - \alpha \bar{\pi}^{(\theta-1)(1-\varrho)})^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\varrho)(1+\omega)}}{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{\theta-\varrho(\theta-1)}} \right) - (1 - \alpha)^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{\theta}{\theta - 1} \right) \bar{s}, \quad (15)$$

where $\bar{\pi}$ and \bar{s} are computed from the mean values of inflation and real marginal cost, respectively, implied by the VAR. We consolidate the first- and second-moment conditions by defining $F(\mu, A, \psi) = (F'_1 F'_2)'$.

¹⁰This implies that $\hat{z}_t = A \hat{z}_{t-1} + \varepsilon_{zt}$.

¹¹See Appendix A. This expression is obtained by calculating all the expectations in (4) as conditional expectations, given \hat{z}_{t-1} , under the assumption that $E(u_t | \hat{z}_{t-1}) = 0$.

If the model is true, there exists a ψ that satisfies $F(\cdot, \psi) = 0$. Accordingly, we estimate the free elements of ψ by searching for a value that makes $F(\psi)$ as small as possible, where ‘small’ is defined in terms of an unweighted sum-of-squares $F(\cdot, \psi)'F(\cdot, \psi)$. Thus, we estimate ψ by solving an equally-weighted GMM problem,

$$\min_{\psi} F(\cdot, \psi)'F(\cdot, \psi). \quad (16)$$

In what follows, we implement this estimator with a time-varying *VAR*. With drifting parameters, we modify the previous formulas by adding time subscripts to the companion form,

$$z_t = \mu_t + A_t z_{t-1} + \varepsilon_{zt}, \quad (17)$$

and by appropriately redefining the function F as $F_t(\cdot, \psi)$ to represent the restrictions $F(\mu_t, A_t, \psi)$ at a particular date. After selecting a number of representative dates, $t = t_1, \dots, t_n$,¹² we stack the residuals from each date into a long vector

$$\mathcal{F}(\cdot) = [F'_{t_1}, F'_{t_2}, \dots, F'_{t_n}]'. \quad (18)$$

Then we estimate the parameters by minimizing the unweighted sum of squares $\mathcal{F}'\mathcal{F}$.

We estimate two versions of the model, one in which we allow ψ to differ across dates (i.e., in which case the restrictions are $F_t(\cdot) = F(\mu_t, A_t, \psi_t)$), and another in which we hold ψ constant ($F_t(\cdot) = F(\mu_t, A_t, \psi)$). Our objective is to see whether the data support the hypothesis that the parameters of the Calvo model are structurally invariant, i.e., $\psi_t = \psi$.

Further details on the second-stage estimator are provided in section 5. Before illustrating the results, however, we discuss the methodology by which we estimate a time-varying *VAR* and report some evidence on parameter drift.

4 A VAR with drifting parameters

This section documents the time drifting nature of the joint process of inflation and marginal costs. Under the hypothesis of a non-zero level of trend inflation, the dynamics of inflation depend not only on the evolution of marginal costs, but also on the evolution of output growth and the discount rate. We therefore estimate a Bayesian vector autoregression with drifting coefficients and stochastic volatilities for the log of gross inflation, log marginal cost, output growth, and a discount rate.

The methodology for estimating the reduced form follows Cogley and Sargent (2004). We begin by writing the *VAR* as

$$x_t = X_t' \vartheta_t + \varepsilon_{xt}, \quad (19)$$

¹²The problem would be too high-dimensional if we used all the dates.

where ϑ_t denotes a vector of time-varying conditional mean parameters.¹³ In the companion-form notation used above, the matrix A_t refers to the autoregressive parameters in ϑ_t , and the vector μ_t includes the intercepts. As in Cogley and Sargent, ϑ_t is assumed to evolve as a driftless random walk subject to reflecting barriers. Apart from the reflecting barrier, ϑ_t evolves as

$$\vartheta_t = \vartheta_{t-1} + v_t. \quad (20)$$

The innovation v_t is normally distributed, with mean 0 and variance Q . Denoting by ϑ^T the history of VAR parameters from date 1 to T ,

$$\vartheta^T = [\vartheta'_1, \dots, \vartheta'_T]', \quad (21)$$

the driftless random walk component is represented by a joint prior

$$f(\vartheta^T, Q) = f(\vartheta^T|Q) f(Q) = f(Q) \prod_{s=0}^{T-1} f(\vartheta_{s+1}|\vartheta_s, Q). \quad (22)$$

Associated with this is a marginal prior $f(Q)$ that makes Q an inverse-Wishart variate.

The reflecting barrier is encoded in an indicator function, $I(\vartheta^T) = \prod_{s=1}^T I(\vartheta_s)$. The function $I(\vartheta_s)$ takes a value of 0 when the roots of the associated VAR polynomial are inside the unit circle, and it is equal to 1 otherwise. This restriction truncates and renormalizes the random walk prior,

$$p(\vartheta^T, Q) \propto I(\vartheta^T) f(\vartheta^T, Q). \quad (23)$$

This represents a stability condition for the VAR , which rules out explosive representations for the variables in question. Explosive representations might be useful for modeling hyperinflationary economies, but we regard them as implausible for the post World War II U.S.

To allow for stochastic volatility, we assume that the VAR innovations ε_{xt} can be expressed as

$$\varepsilon_{xt} = V_t^{1/2} \xi_t$$

where ξ_t is a standard normal vector, which we assume to be independent of parameters innovation v_t , $E(\xi_t v_s) = 0$, for all t, s . We model V_t as¹⁴

$$V_t = B^{-1} H_t B^{-1'}, \quad (24)$$

¹³ x_t is a $N \times 1$ vector of endogenous variables ($N = 4$ in our case), and $X'_t = I_N \otimes [1 \ x'_{t-l}]$, with x'_{t-l} denoting lagged values of x_t .

¹⁴This is a multivariate version of the stochastic volatility model of Jacquier, Polson, and Rossi (1994).

where H_t is diagonal and B is lower triangular. The diagonal elements of H_t are assumed to be independent, univariate stochastic volatilities that evolve as driftless geometric random walks

$$\ln h_{it} = \ln h_{it-1} + \sigma_i \eta_{it}. \quad (25)$$

The innovations η_{it} have a standard normal distribution, are independently distributed, and are assumed independent of innovations v_t and ξ_t . The random walk specification for h_{it} is chosen to represent permanent shifts in innovation variance, as those emphasized in the literature about the reduction in volatility in US economic time series (see, for example, McConnell and Perez Quiros, 2000).¹⁵

We work with a $VAR(2)$ representation, estimated using data from 1960.Q1 through 2003.Q4. Data from 1954.Q1-1959.Q4 were used to initialize the prior. The posterior distribution for VAR parameters was simulated using Markov Chain Monte Carlo methods. Appendix B sketches the simulation algorithm; for a more extensive discussion, see Cogley and Sargent (2004).

4.1 The data

As noted above, the model depends on the joint behavior of four variables: inflation, real marginal cost, output growth, and a nominal discount factor. Inflation is measured from the implicit GDP deflator, recorded in NIPA table 1.3.4. Output growth is calculated using chain-weighted real GDP, expressed in 2000\$, and seasonally adjusted at an annual rate. This series is recorded in NIPA table 1.3.6. The nominal discount factor is constructed by expressing the federal funds rate on a discount basis. Federal funds data were downloaded from the Federal Reserve Economic Database; they are monthly averages of daily figures and were converted to quarterly values by point-sampling the middle month of each quarter.

That leaves real marginal cost. Under the hypothesis of Cobb-Douglas technology, real marginal cost, s , is proportional to unit labor cost,

$$s = wH/(1 - a)PY = (1 - a)^{-1}ulc, \quad (26)$$

where $1 - a$ is the output elasticity to hours of work in the production function.¹⁶ In previous work, Sbordone (2002, 2003) used an index number constructed by the BLS to measure unit labor cost in the non-farm business sector. That is fine for studying gap relationships, because a change of units does not alter percent differences from a steady state, but here we also exploit a restriction on \bar{s} , and that requires expressing s in its natural units.

¹⁵The factorization in (24) and the log specification in (25) guarantee that V_t is positive definite, while the free parameters in B allow for correlation among the VAR innovations ε_{xt} . The matrix B orthogonalizes ε_{xt} , but it is not an identification scheme.

¹⁶This follows from the fact that the marginal product of labor is proportional to the average product.

To construct such a measure, we compute an index of total compensation in the non-farm business sector from BLS indices of nominal compensation and total hours of work, then translate the result into dollars. Because we lack the right data for the non-farm business sector, we perform the translation using data for private sector labor compensation, which we obtained from table B28 of the *Economic Report of the President* (2004). From that table, we calculated total labor compensation in dollars for 2002;¹⁷ the number for that year comes to \$4978.61 billion. The BLS compensation index is then rescaled so that the new compensation series has that value in 2002. A (log) measure of real unit labor costs ulc is then obtained by subtracting (log of) nominal GDP from (log of) labor compensation. The correlation of this measure of ulc with Sbordone’s original measure is 0.9979, so this is almost entirely just a change in units. The new measure therefore accords very well with the one used in previous work.

Finally, to transform the real unit labor cost (or labor share) into real marginal cost, we subtract the log exponent on labor, $(1 - a)$, which we set equal to 0.7. This also pins down the strategic complementarity parameter ω , for in a model such as this $\omega = a/(1 - a)$. Since a is calibrated when constructing a measure of real marginal cost, ω is no longer free for estimation.

4.2 Calibrating the Priors

Next we describe how the *VAR* priors are calibrated. As in Cogley and Sargent, our guiding principle is to make the priors proper but weakly informative, so that the posterior mainly reflects information in the data. Our settings follow theirs quite closely. We begin by assuming that hyperparameters and initial states are independent across blocks, so that the joint prior can be expressed as the product of marginal priors. Then we separately calibrate each of the marginal priors.

Our prior for ϑ_0 is

$$p(\vartheta_0) \propto I(\vartheta_0)f(\vartheta_0) = I(\vartheta_0)N(\bar{\vartheta}, \bar{P}), \quad (27)$$

where the mean and variance of the Gaussian piece are set by estimating a time-invariant vector autoregression using data from the training sample 1954.Q3-1959.Q4. We set $\bar{\vartheta}$ equal to the point estimate from those regressions and the variance \bar{P} to the asymptotic variance of that estimate.

For the innovation variance Q , we adopt an inverse-Wishart prior,

$$f(Q) = IW(\bar{Q}^{-1}, T_0). \quad (28)$$

¹⁷Column D of table B28 reports private wages and salaries in dollars; to that we add a fraction from column G, supplements to wages and salaries. That fraction was calculated as the ratio of private to total wages. This is an attempt to remove government from column G; the assumption is that the ratio of private to total is the same for wages and salaries as for supplements.

In order to minimize the weight of the prior, the degree-of-freedom parameter T_0 is set to the minimum for which the prior is proper,

$$T_0 = \dim(\theta_t) + 1. \quad (29)$$

To calibrate the scale matrix \bar{Q} , we assume

$$\bar{Q} = \gamma^2 \bar{P} \quad (30)$$

and set $\gamma^2 = 1.25\text{e-}04$. This makes \bar{Q} comparable to the value used in Cogley and Sargent (2004), adjusting for the increased dimension of this model.

The parameters governing stochastic-volatility priors are set as follows. The prior for h_{i0} is log-normal,

$$f(\ln h_{i0}) = N(\ln \bar{h}_i, 10), \quad (31)$$

where \bar{h}_i is the initial estimate of the residual variance of variable i . A variance of 10 on a natural-log scale makes this weakly informative for h_{i0} . The prior for b is also normal with a large variance,

$$f(b) = N(0, 10000 \cdot I_3). \quad (32)$$

Finally, the prior for σ_i^2 is inverse gamma with a single degree of freedom,

$$f(\sigma_i^2) = IG\left(\frac{.01^2}{2}, \frac{1}{2}\right). \quad (33)$$

This also puts a heavy weight on sample information, for (33) does not possess finite moments.

4.3 Evidence on parameter drift

With these priors, the posterior was simulated using the Markov Chain Monte Carlo algorithm outlined in Appendix B. The variables were ordered as $\log \gamma_{yt}$, $\log s_t$, $\log \pi_t$, R_t ; exploring the sensitivity of our results to the ordering is left to future research.

4.3.1 Structure of Drift in ϑ_t

For a first piece of evidence on drift in ϑ_t , we inspect the structure of the innovation variance Q . Recall that this matrix governs the pattern and rate of drift in the conditional mean parameters. Table 1 records the principle components of its posterior mean.

Cogley and Sargent found that patterns of drift in ϑ were highly structured, with Q having only a few non-zero principal components, and the same is true here. The matrix Q is 36×36 , but the posterior mean has only 4 or 5 significant principal components. That means many linear combinations of ϑ are approximately time invariant. In other words, there are stable and unstable subspaces of ϑ .

Table 1
Principal components of Q

	Variance	Cumulative Proportion of $tr(Q)$
PC 1	0.0554	0.637
PC 2	0.0132	0.789
PC 3	0.0065	0.864
PC 4	0.0057	0.930
PC 5	0.0016	0.947
PC 6	0.0011	0.961
PC 7	0.0010	0.972
PC 8	0.0007	0.980
PC 9	0.0005	0.985

This is also illustrated in figure 1, which portrays partial sums of the principal component for $\Delta\vartheta_{t|T}$. It shows rotations of the mean VAR parameters, sorted by degree of time variation. A few move around a lot, the rest are approximately constant throughout the sample.

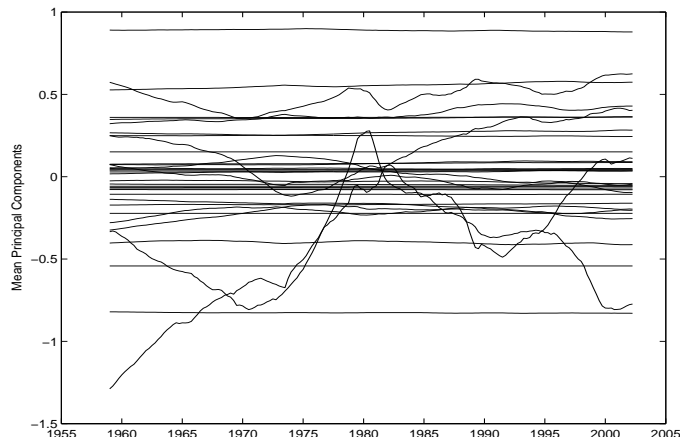


Figure 1: Principal Components of ϑ

From the eigenvectors associated with the first 5 components no obvious pattern or simple interpretation of the factors responsible for the variation in ϑ emerges. Nevertheless, that the drift is structured is an intriguing clue about the source of time variation, for it suggests that many components of a general equilibrium model are likely to be invariant. If changes in monetary policy are indeed behind the drifting components in ϑ , then many other features are likely to be structural. We are curious whether Calvo-pricing parameters are among the invariant features.

4.3.2 Trend Inflation and the Persistence of the Inflation Gap

Next we turn to evidence on trend inflation, $\ln \bar{\pi}_t$, and the inflation gap, $\ln(\pi_t/\bar{\pi}_t)$. Trend inflation is estimated as in Cogley and Sargent by calculating a local-to-date t estimate of mean inflation from the *VAR*,

$$\ln \bar{\pi}_t = e'_\pi (I - A_{t|T})^{-1} \mu_{t|T}. \quad (34)$$

The arrays $\mu_{t|T}$ and $A_{t|T}$ denote posterior mean estimates of the intercepts and autoregressive parameters, respectively. Figure 2 portrays estimates of trend inflation, shown as a red line, and compares it with actual inflation and mean inflation. The latter are recorded in blue and green, respectively, and all are expressed at annual rates

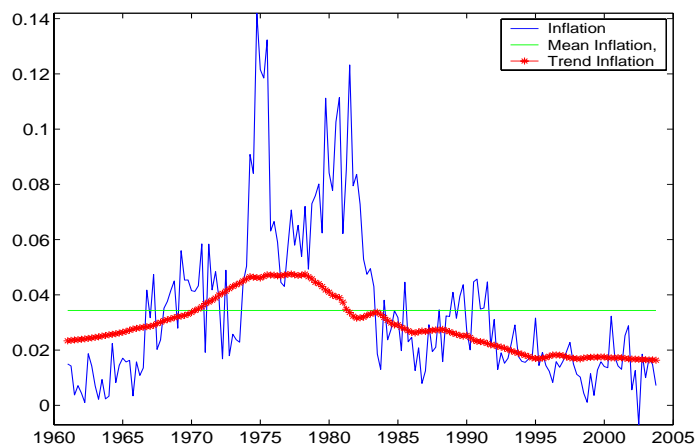


Figure 2: Inflation, Mean Inflation, and Trend Inflation

Two features of the graph are relevant for what comes later. The first, of course, is that trend inflation varies in our sample. We estimate that $\ln \bar{\pi}_t$ rose from 2.3 percent in the early 1960s to roughly 4.75 percent in the 1970s, then fell to around 1.65 percent at the end of the sample. A conventional Calvo model explains inflation gaps, which are usually represented in terms of deviations from a constant mean,¹⁸ but if trend inflation varies, as the data suggest, the appropriate measure of inflation gap is the deviation from its time-varying trend. Accordingly, we aim at modeling a trend-based inflation gap.

The second feature concerns the degree of inflation gap persistence. How the inflation gap is measured – whether as deviations from the mean or from a time-varying trend – matters because that affects the degree of persistence. As the figure

¹⁸In general equilibrium, mean inflation is usually pinned down by the target in the central bank’s policy rule.

illustrates, the mean-based gap is more persistent than the trend-based measure. Notice, for example, the long runs at the beginning, middle, and end of the sample when inflation does not cross the mean. In contrast, inflation crosses the trend line more often, especially after 1985. One of the puzzles in the literature concerns whether conventional Calvo models can generate enough persistence to match mean-based measures of the gap. A backward-looking element is often added to accomplish this. Figure 2 makes us wonder whether this ‘excess persistence’ reflects an exaggeration of the persistence of mean-based gaps rather than a deficiency of persistence in forward-looking models. We comment more on this below.

The figure also suggests that the degree of persistence in the trend-based inflation gap is not constant over the sample. For example, there are also long runs at the beginning and the middle of the sample in which inflation does not cross the trend, while there are many more crossings after 1985. This suggests a decrease in inflation persistence after the Volcker disinflation. Indeed, the first-order autocorrelation for the trend-based inflation gap is 0.75 prior to 1985 and 0.34 thereafter. Changes in inflation persistence may also be part of the resolution of the persistence puzzle. For instance, the ‘excess persistence’ found in time-invariant models may have disappeared from the data after the Volcker disinflation.

Figures 3a and 3b provide another measure of inflation persistence, showing the normalized spectrum of inflation. This is calculated as in Cogley and Sargent (2004) from a local-to-date t approximation to the spectrum for inflation. The normalized spectrum is defined as

$$g_{\pi\pi}(\omega, t) = \frac{2\pi f_{\pi\pi}(\omega, t)}{\int_{-\pi}^{\pi} f_{\pi\pi}(\omega, t) d\omega}, \quad (35)$$

where $f_{\pi\pi}(\omega, t)$ is the instantaneous power spectrum

$$f_{\pi\pi}(\omega, t) = e'_{\pi}(I - A_{t|T}e^{-i\omega})^{-1} \frac{V_{t|T}}{2\pi} (I - A_{t|T}e^{i\omega})^{-1'} e_{\pi}. \quad (36)$$

Once again, the arrays $A_{t|T}$ and $V_{t|T}$ represent posterior means, which are calculated by averaging across the Monte Carlo distribution. In figure 3a, time is plotted on the x -axis, frequency on the y -axis, and power on the z -axis. Figure 3b reports slices along the x -axis for three selected years.

With this normalization,¹⁹ a white noise process has a constant spectrum equal to 1 at all frequencies. Relative to this benchmark, excess power at low frequencies signifies positive autocorrelation or persistence, and deficient power at low frequencies represents negative autocorrelation or anti-persistence. The spectra shown here all have more power at low frequencies than a white noise variate, so there is always positive persistence in the trend-based gap.

¹⁹Notice that we adopt a different normalization than in Cogley and Sargent (2004). Their normalization makes a white noise spectrum equal to $1/2\pi$ at all frequencies.

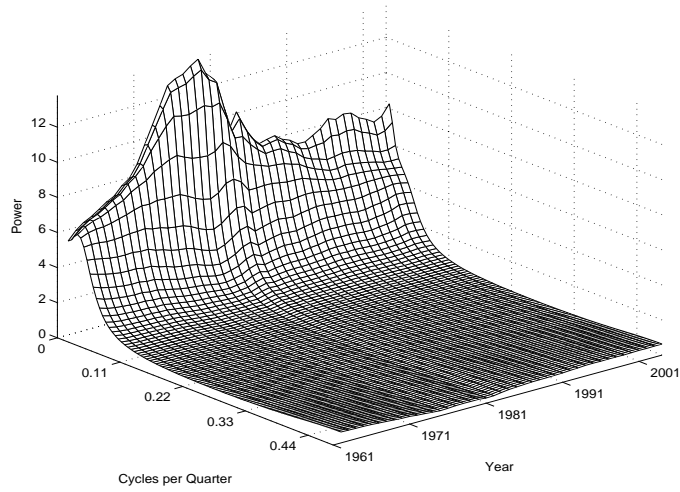


Figure 3a: Normalized Spectrum for Inflation

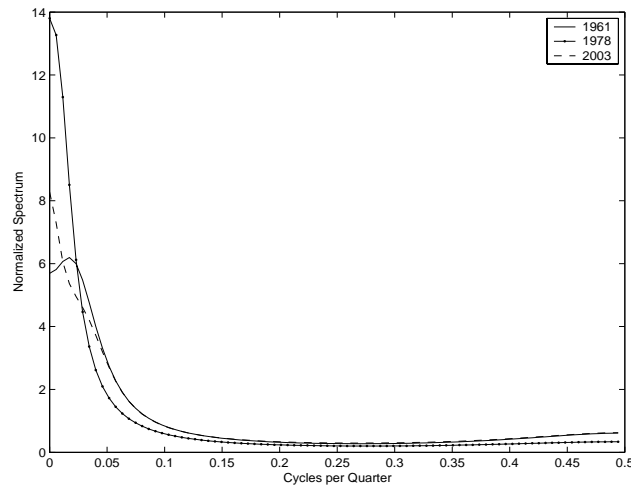


Figure 3b: Normalized Spectrum (selected years)

What varies is the degree of persistence. The rise and fall in low-frequency power signifies a changing degree of autocorrelation. To help interpret the figures, it is convenient to compare them with an $AR(1)$ benchmark, for which the normalized spectrum at zero can be expressed in terms of the autoregressive parameter ρ ,

$$g(0) = (1 + \rho)/(1 - \rho). \quad (37)$$

The normalized spectrum at zero was approximately 6 in the early 1960s, 14 in the late 1970s, and 8 in the 1990s and early 2000s. Those values correspond to autoregressive

roots of 0.71, 0.87, and 0.78, respectively, or half-lives of 2.43, 5.20, and 3.12 quarters. Thus, while there is some variation in inflation persistence, it is not too dramatic.²⁰

5 Estimates of deep parameters

Next we turn to the deep parameters $\psi = [\alpha, \tilde{\beta}, \theta, \varrho, \omega, \bar{\pi}]$ that determine the coefficients of the generalized Calvo equation (4). We estimate deep parameters by searching for values which reconcile that equation with the reduced-form *VAR*.

We begin by noting that three elements of ψ are already determined by other conditions. Trend inflation $\bar{\pi}_t$ is estimated from the reduced-form *VAR* parameters. The value corresponding to the i -th draw from the *VAR* posterior is ²¹

$$\bar{\pi}_t(i) = \exp\left(e'_\pi [I - A_{t|T}(i)]^{-1} \mu_{t|T}(i)\right)$$

where $\mu_{t|T}(i)$ and $A_{t|T}(i)$ represent the i -th draw in the *VAR* posterior sample, and $i = 1, \dots, N_{MC}$, where N_{MC} is the total number of draws in the Monte Carlo sample.

The discount parameter $\tilde{\beta}$ is also a byproduct of *VAR* estimation. Recall that $\tilde{\beta}$ is defined as $\tilde{\beta} = \bar{\gamma}_y \bar{q} = \bar{\gamma}_y \bar{R} \bar{\pi}$, where $\bar{\gamma}_y$ is the steady-state gross rate of output growth and \bar{R} is the steady-state nominal discount factor. Since the latter are also estimated from the *VAR*,

$$\begin{aligned} \bar{\gamma}_{yt}(i) &= \exp\left(e'_y [I - A_{t|T}(i)]^{-1} \mu_{t|T}(i)\right), \\ \bar{R}_t(i) &= e'_R [I - A_{t|T}(i)]^{-1} \mu_{t|T}(i), \end{aligned} \quad (38)$$

that fixes $\tilde{\beta}_t(i) = \bar{\gamma}_{yt}(i) \bar{R}_t(i) \bar{\pi}_t(i)$.

The third parameter that is set in advance is ω , which governs the extent of strategic complementarity. This is pinned down by the condition $\omega = a/(1-a)$, where $1-a$ is the Cobb-Douglas labor elasticity. We calibrated $a = 0.3$ when transforming labor share data into a measure of real marginal cost (see the data description above), and that fixes $\omega = 0.429$.

That leaves three free parameters, α , ϱ , and θ , which we estimate, for every draw ϑ_i , by trying to satisfy the cross-equation restrictions described above. Letting $\psi_i = [\alpha_i, \varrho_i, \theta_i]$, these restrictions are:

$$\begin{aligned} F_{1t}(\psi_i, \mu_{t|T}(i), A_{t|T}(i)) &= e'_\pi [I - b_{1t} A_{t|T}(i) - b_{2t} \gamma_{1t} (I - \gamma_{1t} A_{t|T}(i))^{-1} A_{t|T}^2(i)] A_{t|T}(i) \\ &\quad - \tilde{\varrho}_t e'_\pi I - \zeta_t e'_s A_{t|T}(i) - \chi_t (\gamma_{2t} - \gamma_{1t}) e'_R (I - \gamma_{1t} A_{t|T}(i))^{-1} A_{t|T}(i) \\ &\quad - \chi_t (\gamma_{2t} - \gamma_{1t}) e'_y (I - \gamma_{1t} A_{t|T}(i))^{-1} A_{t|T}^2(i), \end{aligned} \quad (39)$$

²⁰The variation shown here is less pronounced than that reported by Cogley and Sargent, who studied a *VAR* involving different variables.

²¹The *VAR* is estimated for the log of gross inflation, so the local-to-date- t approximation of the mean refers to net inflation. We exponentiate to restore the original units.

$$\begin{aligned}
F_{2t}(\psi_i, \mu_{t|T}(i), A_{t|T}(i)) &= (1 - \alpha \bar{\pi}_t(i)^{(\theta-1)(1-\varrho)})^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{1 - \alpha \bar{R}_t(i) \bar{\gamma}_{yt}(i) \bar{\pi}_t(i)^{1+\theta(1-\varrho)(1+\omega)}}{1 - \alpha \bar{R}_t(i) \bar{\gamma}_{yt}(i) \bar{\pi}_t(i)^{\theta-\varrho(\theta-1)}} \right) \\
&\quad - (1 - \alpha)^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{\theta}{\theta-1} \right) \bar{s}(i)_t,
\end{aligned} \tag{40}$$

$$F_t(\psi_i, \mu_{t|T}(i), A_{t|T}(i)) = [F_{1t}(\cdot)', F_{2t}(\cdot)']'. \tag{41}$$

The parameters $b_{1t}, b_{2t}, \gamma_{1t}, \gamma_{2t}, \tilde{\varrho}_t, \zeta_t,$ and χ_t in (39) are defined as in (6) with $\bar{\pi}_t(i), \tilde{\beta}_t(i),$ and ω set in advance as described above. The moment conditions are indexed by t because they depend on $\mu_{t|T}(i)$ and $A_{t|T}(i),$ which vary through time. Finally, the steady-state value for real marginal cost is also calculated from VAR estimates, as $\bar{s}_t(i) = \exp(e'_s(I - A_{t|T}(i))^{-1} \mu_{t|T}(i)).$

The moment condition $F_t(\cdot)$ has dimension $1 + Np,$ where $N = 4$ is the number of equations in the VAR and $p = 2$ is the number of lags. A complete set of moment conditions for all dates in the sample would therefore have dimension $T(1 + Np).$ Because the sample spans 174 quarters, the complete set of moment conditions would have more than 1500 elements for estimating 3 parameters. That is both intractable and unnecessary. Accordingly, we simplify by selecting 5 representative quarters, 1961.Q3, 1978.Q3, 1983.Q3, 1995.Q3, and 2003.Q3, so that the moment conditions reduce to

$$\mathcal{F}(\cdot) = \begin{bmatrix} F_{1961}(\cdot) \\ F_{1978}(\cdot) \\ F_{1983}(\cdot) \\ F_{1995}(\cdot) \\ F_{2003}(\cdot) \end{bmatrix}. \tag{42}$$

Our selection of quarters is motivated as follows. First, we wanted a relatively small number of dates in order to manage the dimension of the GMM problem. We also wanted to space the dates apart because VAR estimates of μ_t and A_t in adjacent quarters are highly correlated, which would result in high correlation across time in the moment conditions $F_t(\cdot).$ Highly correlated moment conditions would contribute relatively little independent information for estimation and therefore would be close to redundant.

Second, we wanted to span the variety of monetary experience in the sample. Thus, we chose 1961 to represent the initial period of low and stable inflation prior to the Great Inflation. The year 1978 represents the height of the Great Inflation, when both trend inflation and the degree of persistence were close to their maxima. The year 1983 represents the end of the Volcker disinflation, which we regard as a key turning point in postwar US monetary history. This is a point of transition between the high inflation of the 1970s and the period of stability that followed, and expectations may have been unsettled at that time. The final two years, 1995 and 2003 are two points from the Greenspan era, an established low-inflation environment.

The first was chosen to represent a more pre-emptive FOMC, the second reflects a more recent wait-and-see approach.

We emphasize that the dates were chosen based on *a priori* reflection and reasoning, before estimating deep parameters. Exploring the sensitivity of our results to alternative selections would be interesting, provided one does not mine the data too interactively along the way.

With the function \mathcal{F} defined in (42), we estimate the vector of parameters ψ by minimizing the unweighted sum of squares $\mathcal{F}'\mathcal{F}$. As the notation of (39) and (40) indicate, we estimate best-fitting values of ψ for every draw in the posterior sample for the VAR , $\mu_{t|T}(i)$ and $A_{t|T}(i)$. In this way, we obtain a distribution of estimates

$$\psi_i = \arg \min [\mathcal{F}(\cdot)'\mathcal{F}(\cdot)], \quad (43)$$

for $i = 1, \dots, N_{MC}$. where N_{MC} is the number of draws in the Monte Carlo simulation for the first-stage VAR . This allows us to assess how parameter uncertainty in the first-stage VAR matters for estimates of deep parameters. We also estimate best-fitting values of ψ from the posterior mean of VAR estimates, $\mu_{t|T}$ and $A_{t|T}$.²²

In what follows, the median estimate of deep parameters from the distribution (43) is always close to the best-fitting value derived from the posterior VAR mean, but a distribution of estimates is helpful for appraising uncertainty. In effect, we induce a probability distribution over ψ_i by applying a change of variables to the distribution of VAR parameters. The numerical optimizer that we adopt starts from the same initial conditions for each draw and contains no random search elements, so (43) implicitly expresses a deterministic function that uniquely determines the deep parameters as a function of the VAR parameters. Thus, a change-of-variables interpretation is valid. It should be noted that the resulting distribution for ψ_i is not a Bayesian posterior because it follows from the likelihood function for the reduced-form model instead of the structural model. It is in fact a transformation of the posterior for the reduced form parameters ϑ .²³

We estimate two versions of the model, one in which the parameters in ψ are held constant, and another where they are free to differ across dates. In both cases, their values are constrained to lie in the economically meaningful ranges listed in table 2.²⁴ Furthermore, we verify that the parameters satisfy the conditions for existence of a steady state (the inequalities (57) in appendix A).

²²These are defined as follows: $\mu_{t|T} = \frac{1}{N_{MC}} \sum_i^{N_{MC}} \mu_{t|T}(i)$, and $A_{t|T} = \frac{1}{N_{MC}} \sum_i^{N_{MC}} A_{t|T}(i)$.

²³For another approach to this problem, see Hong Li (2004).

²⁴We also considered estimates obtained by minimizing a weighted sum of squares $\mathcal{F}(\cdot)'W\mathcal{F}(\cdot)$. Using the estimates in (43), we calculate the moment condition errors and their covariance, $V_{\mathcal{F}}$. The weighting matrix W is the inverse of that matrix, $W = V_{\mathcal{F}}^{-1}$. Because these weighted estimates do not lead to a gain in precision, based on the median absolute deviation, we report only the unweighted estimates.

Table 2
Admissible Range for Estimates

α	ϱ	θ
$(0, 1)$	$[0, 1]$	$(1, \infty)$

5.1 NKPC with Constant Parameters

Estimates for the constant-parameter case are reported in table 3. Because the distributions are non-normal, we focus on the median and median absolute deviation, respectively, instead of the mean and the standard deviation. All three parameters are economically sensible, the estimates accord well with microeconomic evidence, and they are reasonably precise.

Table 3
Estimates when Calvo Parameters are Constant

	α	ϱ	θ
@VAR mean	0.602	0	10.55
Median	0.602	0	9.97
Median Absolute Deviation	0.048	0	0.90

One especially interesting outcome concerns the indexation parameter, which we estimate at $\varrho = 0$.²⁵ This contrasts with much of the empirical literature based on time-invariant models in which the indexation parameter is estimated as low as 0.2 and as high as 1, and is statistically significant.²⁶ In those models, an important backward-looking component is needed to fit inflation persistence, but that is not the case here. From a purely statistical point of view, a positive coefficient on past inflation may arise from an omitted-variable problem, since the omitted forward-looking terms that belong to the model according to (4), but which are omitted from estimators of standard Calvo models, may be positively correlated with past inflation. Indeed, that is the case when inflation Granger-causes output growth and nominal interest rate.

²⁵To be more precise, 84.2 percent of the estimates lie exactly on the lower bound of 0. The mean estimate is 0.022, and the standard deviation is 0.084. Only 3.3 percent of the estimates lie above 0.2.

²⁶Sbordone (2003) estimates a ϱ ranging from 0.22 to 0.32, depending on the proxy chosen for the marginal cost, in single equation estimates; Smets and Wouters (2002) in a general equilibrium model, estimate a value of approximately 0.6. Giannoni and Woodford (2003) estimate a value close to 1. Other authors, following Gali and Gertler (1999), introduce a role for past inflation assuming the presence of rule-of-thumb firms, instead of through indexation, and also find a significant coefficient on lagged inflation.

More substantially, we believe that allowing for a time-varying trend inflation in the *VAR* reduces the persistence of the gap $\ln(\pi_t/\bar{\pi}_t)$, making it easier to match the data with a purely forward-looking model. In other words, our estimates point to a story in which the need for a backward-looking term arises because of neglect of time-variation in $\ln\bar{\pi}_t$. That neglect creates artificially high inflation persistence in time-invariant *VARs*, and hence a ‘persistence puzzle’ for forward-looking models. In a drifting-parameter environment, however, the inflation gap is less persistent, and a purely forward-looking model is preferred.

Another interesting result concerns the fraction of sticky-price firms, which we estimate at $\alpha = 0.602$ per quarter. In conjunction with the estimate of $\varrho = 0$, this implies a median duration of prices of 1.36 quarters, or 4.1 months,²⁷ a value consistent with microeconomic evidence on the frequency of price adjustment. Bils and Klenow (2004), for example, report a median duration of prices of 4.4 months, which increases to 5.5 months after removing sales price changes, which are only temporary reversals. Our estimate from macroeconomic data therefore accords well with the conclusions they draw from microeconomic data.

In contrast, Calvo specifications estimated from time-invariant *VARs* that require a backward-looking indexation component are grossly inconsistent with their evidence. When $\varrho > 0$, every firm changes price every quarter, some optimally rebalancing marginal benefit and marginal cost, others mechanically marking up prices in accordance with the indexation rule. Unless the optimal rebalancing happened to result in a zero price change or lagged inflation were exactly zero, conditions that are very unlikely, no firm would fail to adjust its nominal price. In a world such as that, Bils and Klenow would not have found that 75 percent of prices remain unchanged each month. We interpret this as additional evidence in support of a purely forward-looking model.

Finally, the estimate of θ implies a steady state markup of about 11 percent, which is in line with other estimates in the literature. For example, this is the same order of magnitude as the markups that Basu (1996) and Basu and Kimball (1997) estimate using sectoral data. With economy-wide data, in the context of general equilibrium models, estimates range from around 6 to 23 percent, depending on the type of frictions in the model. Rotemberg and Woodford (1997) estimate a steady state markup of 15 percent ($\theta \approx 7.8$). Amato and Laubach (2003), in an extended model which include also wage rigidity, estimate a steady-state markup of 19 percent. Edge et al. (2003) find a slightly higher value, 22.7 percent ($\theta = 5.41$). The estimates in Christiano et al. (2003) span a larger range, varying from around 6.35 to 20 percent, depending on details of the model specification. All

²⁷For a purely forward-looking Calvo model, the waiting time to the next price change can be approximated as an exponential random variable (using a continuous approximation), and from that one can calculate that the median waiting time is $-\ln(2)/\ln(\alpha)$. Note that the median waiting time is less than the mean, because an exponential distribution has a long upper tail.

the cited estimates on economy-wide data are obtained by matching theoretical and empirical impulse response functions to monetary shocks. Although obtained through a different estimation strategy, our markup estimate falls within the range found by others.

The model is overidentified, with 3 free parameters to fit 45 elements in $\mathcal{F}(\cdot)$. To test the overidentifying restrictions, we compute a J -statistic,

$$J = \mathcal{F}(\hat{\psi}, \mu_{t|T}, A_{t|T})' \text{Var}(\mathcal{F})^{-1} \mathcal{F}(\hat{\psi}, \mu_{t|T}, A_{t|T}), \quad (44)$$

where $\hat{\psi} = [\hat{\alpha}, \hat{\rho}, \hat{\theta}]$ represent the best-fitting values corresponding to the posterior mean estimates of the VAR parameters, $\mu_{t|T}$ and $A_{t|T}$, and $\text{Var}(\mathcal{F})$ is the variance of $\mathcal{F}(\cdot)$, which we estimate from the sample variance of the moment conditions in the cross section,

$$\text{Var}(\mathcal{F}) = N_{MC}^{-1} \sum_{i=1}^{N_{MC}} \mathcal{F}(\psi_i, \mu_{t|T}(i), A_{t|T}(i)) \mathcal{F}(\psi_i, \mu_{t|T}(i), A_{t|T}(i))'. \quad (45)$$

If $\mathcal{F}(\cdot)$ were approximately normal, J would be approximately chi-square with 42 degrees of freedom.²⁸ We calculate $J = 22.2$, which falls far short of the chi-square critical value. Thus, taken at face value, the model's overidentifying restrictions are not rejected. One should take this with a grain of salt, however, because of the non-normality of the distributions for ψ_i and $\mathcal{F}(\psi_i, \cdot)$. In any case, the J -statistic provides no evidence against the over-identifying restrictions.

A complementary way of evaluating the model involves comparing the expected inflation gap implied by the $NKPC$ with the expected inflation gap estimated by the unconstrained VAR , in the spirit of Campbell and Shiller's (1987) exercise.²⁹ The VAR inflation forecast is given by equation (11), while the $NKPC$ forecast is implicitly defined by the right-hand-side of equation (12), which defines the model's cross-equation restrictions. Thus, the distance between the two forecasts measures the extent to which the cross-equation restrictions are violated. Figure 4 plots the two series, showing VAR forecasts in blue and $NKPC$ forecasts in red.

As the figure shows, $NKPC$ forecasts closely track those of the unrestricted VAR . The correlation between the two series is 0.979, and the deviations are small in magnitude and represent high-frequency twists and turns. Thus the unrestricted VAR satisfies the cross-equation restrictions implied by the $NKPC$.

²⁸There are 45 moment conditions and 3 free parameters.

²⁹We choose to compare inflation forecasts, since eq. (4) doesn't have a unique solution for inflation as a function of real marginal costs.

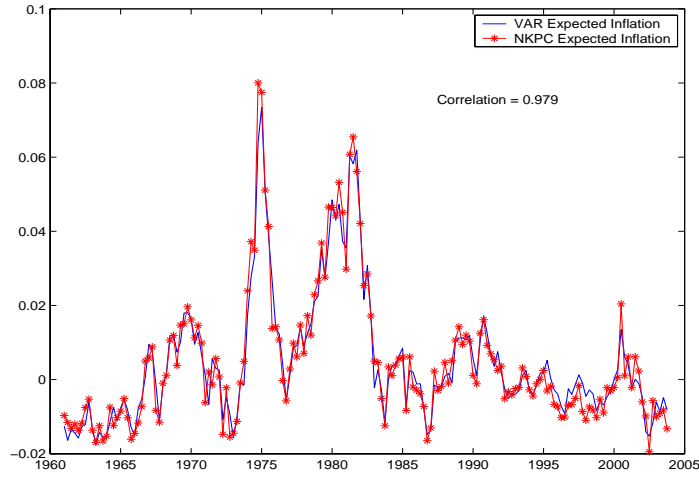


Figure 4: VAR and NKPC Forecasts of Inflation

5.2 NKPC with Variable Parameters

Next we relax the constraint that α , ϱ and θ are constant across dates. When we allow them to vary, we get the estimates recorded in table 4.

Table 4
Estimates when Calvo Parameters are Free to Vary

		α	ϱ	θ
1961	@VAR mean	0.611	0	11.31
	Median	0.590	0	11.32
	Median Absolute Deviation	0.093	0	1.47
1978	@VAR mean	0.566	0	9.60
	Median	0.567	0	10.30
	Median Absolute Deviation	0.053	0	1.18
1983	@VAR mean	0.625	0	12.96
	Median	0.591	0	11.56
	Median Absolute Deviation	0.065	0	1.97
1995	@VAR mean	0.724	0	10.47
	Median	0.672	0	10.71
	Median Absolute Deviation	0.117	0	1.87
2003	@VAR mean	0.734	0	11.05
	Median	0.682	0	10.90
	Median Absolute Deviation	0.124	0	1.90

Once again, we estimate $\rho = 0$ at all the chosen dates. There is, however, some variation in the fraction piling up at zero in various years. This amounted to 74 percent in 1961, 57.5 percent in 1978, 91.4 percent in 1983, 95.3 percent in 1995, and 91 percent in 2003. Thus, support for a purely forward-looking specification is strongest after the Volcker disinflation.

Similarly, the estimates of θ vary a little bit across years, but not a lot. The median point estimates range from a low of 10.30 in 1978 to a high of 11.56 in 1983, values that correspond to mark-ups of 10.8 and 9.5 percent, respectively. These estimates of θ are slightly higher than the median estimate of 9.99 in the constant-parameter version, but they are not dramatically higher.

The estimate of α , the fraction of sticky price firms, also varies slightly across years. Interestingly, this parameter moves in the direction predicted by the New Keynesian theory. For example, Ball, Mankiw, and Romer (1998) say that prices should be more flexible when inflation is high and variable, and less flexible when it is low and stable. Although movements in α are not large (or statistically significant), that is what we find here. Judging by the median point estimates, prices were most flexible (α was smallest) in 1978, when inflation was highest and most variable. For that year, we estimate $\alpha = 0.567$, which implies a median weighting time of 3.66 months to the next price adjustment, a value somewhat lower than what Bils and Klenow estimate.³⁰ Prices were least flexible (α was highest) during the Greenspan era, when inflation was lowest and most stable. For 1995 and 2003, we estimate α equal to 0.672 and 0.682, respectively, which implies a median price duration of roughly 5.3 months. This is somewhat higher than Bils and Klenow's unconditional estimate, but it accords well with what they find after removing sales price changes from their sample.

The next figure provides more detail about the time variation in the estimates. This figure depicts histograms for each of the parameters in various years. The first five rows portray the time-varying estimates, one row for each of the chosen years, and the last row shows the constant-parameter estimates discussed above. Each histogram portrays estimates of α , ρ , and θ for every draw of the *VAR* parameters in the Monte Carlo simulation, that is, 5000 estimates at each date.

There is little evidence here of important time variation in ρ or θ . For ρ , we observe a pile up at zero in all years, as well as in the constant-parameter histogram. The amount of mass at zero varies across years, as noted above, but still there is little evidence of an important indexing or backward-looking component. Similarly, the histograms for θ appear stable across dates, except perhaps for some hard-to-see variation in the long upper tail.

³⁰Their data extend back only to 1995, however, so no contradiction is necessarily implied.

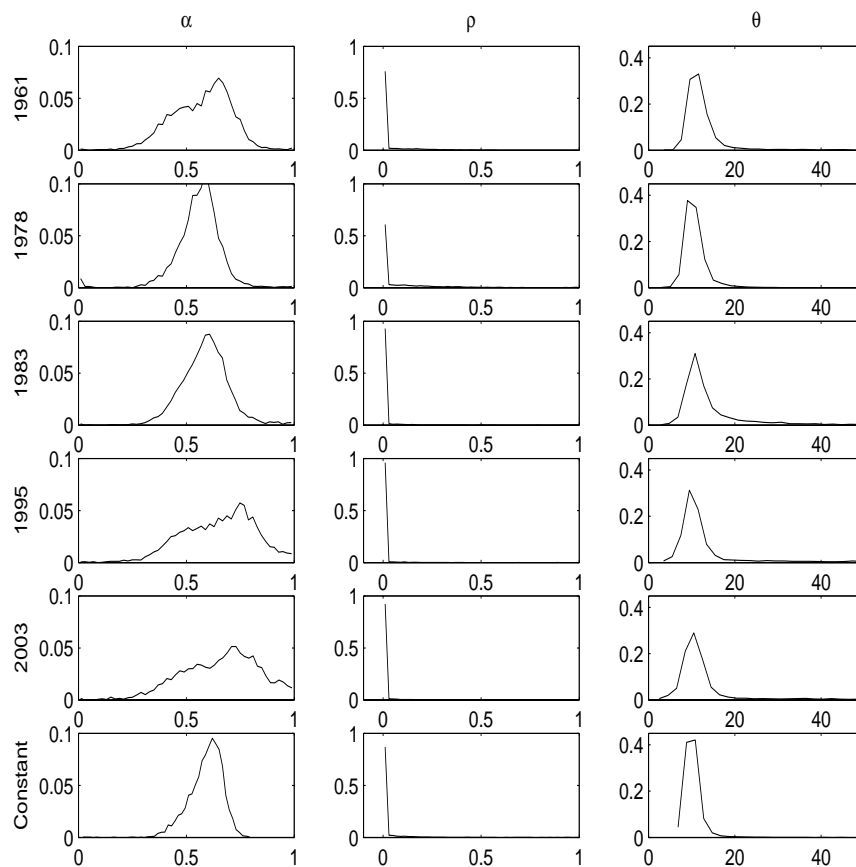


Figure 5 - Histograms for Calvo Parameters

There is slightly more evidence here of changes in α . The histograms for 1995 and 2003 clearly have a different shape than those for 1978 or 1983. Notice, for example, how they are shifted to the right and more disperse than those in earlier years. On the other hand, the histograms for various years also overlap a lot, so it is not clear how strong is the evidence for changes in α .

To dig a bit deeper, we calculated the probability of an increase in α across pairs of years. Recall that we have a panel of estimates α_{it} , $i = 1, \dots, N_{MC}$, and $t = 1961, 1978, 1983, 1995, 2003$. That is, for each of the 5000 sample paths of *VAR* estimates in the Monte Carlo sample, we estimate five α 's, one for each of the chosen years. On each sample path i , we can check whether α increased between various dates. The fraction of sample paths on which α increased is the probability we seek.

Those calculations are reported in the next table. Each entry refers to the probability that α increased from the column date to the row date. For example, the first row shows the probability of an increase between 1961 and 1978, 1961 and 1983, and so on. Numbers smaller than 0.05 or larger than 0.95 may be taken as strong

evidence of shifts in α_t , with numbers close to zero indicating a significant fall in α_t and numbers close to 1 a significant increase.

Table 5
Probability of an Increase in α_t

	1978	1983	1995	2003
1961	0.438	0.532	0.681	0.686
1978		0.625	0.713	0.713
1983			0.663	0.660
1995				0.517

None of the values shown here are strongly significant. Many are not far from 0.5, which says that α was just as likely to fall as to rise. The most significant movements are between 1978 and 1995 or 2003, when we find that α increased on approximately 72 percent of the sample paths. This goes in the right direction, but it falls short of attaining statistical significance at conventional levels. At best, this represents weak evidence of a change in α . If a change did occur, our estimates detect only a vague trace of it.

Table 6 reports analogous calculations for θ . Once again, most of the probabilities are not too far from 0.5, suggesting little evidence of a systematic change.

Table 6
Probability of an Increase in θ_t

	1978	1983	1995	2003
1961	0.335	0.526	0.440	0.458
1978		0.686	0.573	0.587
1983			0.441	0.452
1995				0.517

This result is not surprising. The parameter θ captures the degree of competitiveness and is related to the desired level of mark-up, $\mu = \theta/(\theta - 1)$. Procyclical variations in θ imply countercyclical variations in the desired mark-up, and vice versa, and at a theoretical level, both a countercyclical and a procyclical mark-up can be supported.³¹ At an empirical level, evidence for the U.S. favors countercyclical mark-ups (Bils 1987), while evidence for the U.K. favors procyclical mark-ups (Small 1997).

³¹For example, the model of implicit collusion of Rotemberg and Woodford (1992) implies that the mark-up is a positive function of the ratio of expected future profits to current output, while the customer market model of Phelps and Winter (1970) implies the opposite sign.

It is therefore plausible that variation in trend inflation does not affect the degree of competitiveness one way or the other.³²

Finally, in table 7, we provide an assessment of the probability that $\varrho_t = 0$ in the various periods. Although the median estimate is always zero, the evidence for a purely forward-looking specification is strongest after the Volcker disinflation.

Table 7
Probability $\varrho_t = 0$

1961	1978	1983	1995	2003
0.739	0.576	0.914	0.952	0.909

Overall, the estimates do not point strongly toward variation in α , ϱ , and θ . Over the range of monetary regimes experienced in our sample, the Calvo-pricing parameters appear to be at least approximately invariant to shifts in policy rules. Accordingly, we say the *NKPC* is structural for this class of policy interventions.

6 The effect of positive trend inflation

The traditional *NKPC* is obtained from an approximation around a steady state with zero inflation. In contrast, we estimate a positive and time-varying level of trend inflation in our *VAR* and approximate the local dynamics around that value. In this section, we address how that alters the properties of the *NKPC*.

In figure 5, we show the implied coefficients of the Calvo model, computed as in (6) using the median estimates of α , ϱ , and θ . Dashed lines represent the conventional approximation, which assumes zero trend inflation at all dates, and solid lines represent our approximation, which estimates $\bar{\pi}_t$ from the *VAR*.

The shape of the time-varying *NKPC* parameters is clearly dictated by the dynamics of trend inflation. The parameter ζ , which represents the weight on current marginal cost, varies inversely with $\bar{\pi}$, while the three forward-looking coefficients in (4) vary directly. Thus, as trend inflation rises, the link between current marginal cost and inflation is weakened, and the influence of forward-looking terms is enhanced. This shift in price-setting behavior follows from the fact that positive trend inflation accelerates the rate at which a firm's relative price is eroded when it lacks an opportunity to reoptimize. This makes firms more sensitive to contingencies that may prevail far in the future if their price remains stuck for some time. Thus, relative to the conventional approximation, current costs matter less and anticipations matter more.

³²Khan and Moessner (2003) discuss the relation between competitiveness and trend inflation in the New Keynesian Phillips Curve.

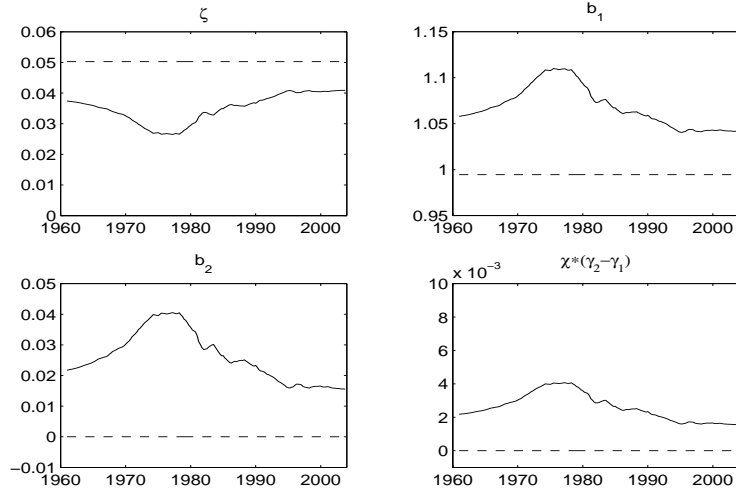


Figure 5: *NKPC* Coefficients

Focusing more closely on the forward-looking coefficients, notice that two of the new terms appearing in (4) – those involving forecasts of output growth and a nominal discount factor – are multiplied by the coefficient $\chi(\gamma_2 - \gamma_1)$. Figure 5 shows that this coefficient is always close to zero,³³ so those terms make a negligible contribution to inflation. In fact, when we omit them from equation (4), *NKPC* expected inflation is virtually the same as that for the complete model shown in figure 4. Thus, the terms $P_{\hat{R}t}$ and $P_{\hat{\gamma}_y t}$ are largely a nuisance and can be neglected without doing too much violence to the theory.

What matters more is how trend inflation alters the coefficients on expected inflation, b_1 and b_2 . Figure 5 shows that b_1 flips from slightly below 1 when trend inflation is zero to around 1.05 or 1.1 for the values of $\bar{\pi}_t$ that we estimate. Similarly, when trend inflation is zero, b_2 is also zero, and multi-step expectations of inflation drop out of equation (4). Those higher-order expectations enter with coefficients of 0.02-0.04 when trend inflation is positive.

As Ascari and Ropele (2004) demonstrate, this shift is so strong that it threatens the determinacy of equilibrium. When trend inflation is zero, we have $b_1 < 1$ and $b_2 = 0$, so we can solve forward to express current inflation in terms of an expected geometric distributed lead of real marginal cost, as in Sbordone (2002, 2003). With positive trend inflation, we can express (4) as

$$E_t [P(L^{-1})\hat{\pi}_t] = E_t(1 - \gamma_1 L^{-1}) \left[\zeta \hat{s}_t + \chi(\gamma_2 - \gamma_1)(P_{\hat{R}t} + P_{\hat{\gamma}_y t}) + u_t \right], \quad (46)$$

where

$$P(L^{-1}) = 1 - (\gamma_1 + b_1)L^{-1} + \gamma_1(b_1 - b_2)L^{-2}, \quad (47)$$

³³This is because $\gamma_2 \simeq \gamma_1$.

when $\varrho = 0$. This polynomial can be factored as

$$P(L^{-1}) = (1 - \lambda_1 L^{-1})(1 - \lambda_2 L^{-1}). \quad (48)$$

Figure 6 portrays λ_1 and λ_2 and shows how they vary with trend inflation. The dashed line also reproduces the value of b_1 that occurs when trend inflation is zero.

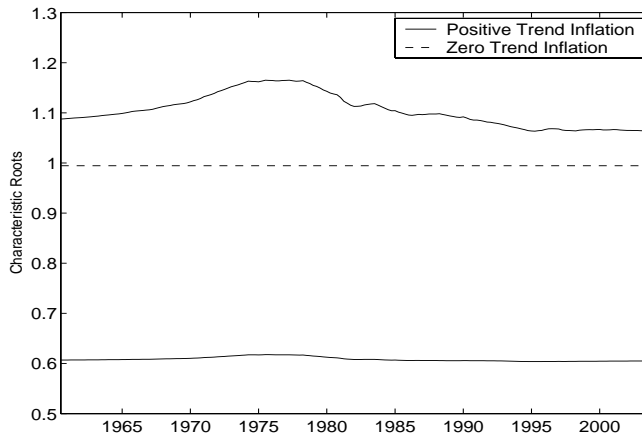


Figure 6: Factorization of $P(L^{-1})$

For our estimates of b_1 , b_2 , and γ_1 , we find $\lambda_1 < 1$ but $\lambda_2 > 1$, which means that a non-explosive forward solution is not guaranteed for arbitrary driving processes. That does not necessarily imply that inflation is indeterminate, for a nonexplosive forward solution could still exist if \hat{s}_{t+j} converged to zero at a faster rate than λ_2^j diverged. The rate of mean reversion in \hat{s}_{t+j} is a property of a general equilibrium, however, and we cannot say much about it in the context of the limited information strategy that we adopt in this paper. Suffice it to say that positive trend inflation diminishes the weight on current marginal cost and increases the weight on future marginal cost, so much so that determinacy of a forward solution is no longer guaranteed. Furthermore, the threat arises even at the low levels of trend inflation experienced in the postwar U.S.

7 Conclusion

In this paper, we address whether the Calvo model of inflation dynamics is structural in the sense of Lucas (1976). In particular, we examine whether its parameters are invariant to shifts in trend inflation, which we associate with different policy regimes.

We first derive the Calvo model as an approximate equilibrium condition around a non-zero steady-state inflation rate and show that its coefficients are nonlinear combinations of deep parameters describing market structure, the pricing mechanism, *and* trend inflation. We estimate deep parameters by exploiting the cross-equation restrictions imposed by the model on a reduced form representation of the data. We model the reduced form as a vector autoregression with time-varying parameters and stochastic volatility, and then ask whether a Calvo-pricing model with constant parameters can be fit to that time-varying reduced form.

We find that a constant-parameter version of the *NKPC* fits very well indeed, closely tracking the *VAR* inflation gap. The estimates are precise, economically sensible, and accord well with microeconomic evidence. In addition, when we allow Calvo-pricing parameters to vary over time, we find little evidence of systematic movements. Thus, the model appears to be structural for policy interventions that may generate shifts in trend inflation of the magnitude of those in our sample.

One important insight that follows from our analysis concerns the importance of backward-looking elements in the model. Our drifting-coefficient *VAR* suggests that trend inflation has been historically quite variable. We believe that measures of the inflation gap that ignore this drift show an artificially high level of persistence, forcing a role for past inflation in the standard Calvo model. In contrast, we show that no indexation or backward-looking component is needed to explain inflation once shifts in trend inflation are properly taken into account. In other words, a purely forward-looking version of the *NKPC* fits post WWII U.S. data very well.

A Appendix A: Derivation of the Calvo equation with trend inflation

The fraction $(1 - \alpha)$ of firms that can set prices optimally choose nominal price X_t (which is not indexed by firms, since each firm that change prices solves the same problem) to maximize expected discounted future profits $\Pi_{t+j} = \Pi(X_t \Psi_{tj}, P_{t+j}, Y_{t+j}(i), Y_{t+j})$

$$\max_{X_t} E_t \sum_j \alpha^j \{R_{t,t+j} \Pi_{t+j}\} \quad (49)$$

subject to their demand constraint

$$Y_{t+j}(i) = Y_{t+j} \left(\frac{X_t \Psi_{tj}}{P_{t+j}} \right)^{-\theta}. \quad (50)$$

$X_t \Psi_{tj} / P_{t+j}$ is the relative price of the firm at $t + j$; $R_{t,t+j}$ is a nominal discount factor between time t and $t + j$; and $Y_t(i)$ is firms' i output. The function Ψ_{tj} captures the fact that individual firms prices that are *not* set optimally evolve according to

$$P_t(i) = \pi_{t-1}^e P_{t-1}(i), \quad (51)$$

and it is therefore defined as

$$\Psi_{tj} = \begin{cases} 1 & j = 0, \\ \prod_{k=0}^{j-1} \pi_{t+k}^\rho & j \geq 1. \end{cases} \quad (52)$$

The FOCs are

$$E_t \sum_{j=0}^{\infty} \alpha^j R_{t,t+j} Y_{t+j} P_{t+j}^\theta X_t^{-\theta-1} \Psi_{tj}^{1-\theta} ((1-\theta)X_t + \theta MC_{t+j,t}(i) \Psi_{tj}^{-1}) = 0. \quad (53)$$

where MC_{t+j} is the *nominal* marginal cost at $t+j$ of the firm that changes its price at t . Dividing through by $Y_t P_t^{\theta+1}$ we can express the equilibrium condition in terms of the (stationary) growth rate of Y , ($\gamma_{y,t} = Y_t/Y_{t-1}$), stationary gross inflation π_t , and stationary relative prices ($x_t = \frac{X_t}{P_t}$). Furthermore, setting $s_{t+j,t}(i) = \frac{MC_{t+j,t}(i)}{P_{t+j}}$, and using the relation between firm's marginal cost and average marginal cost

$$s_{t+j,t}(i) = s_{t+j} x_t^{-\theta\omega} \prod_{k=1}^j \pi_{t+k}^{\theta\omega} \prod_{k=0}^{j-1} \pi_{t+k}^{-\rho\theta\omega}, \quad (54)$$

we obtain expression (1) in the text.

In steady state, (1) is

$$\bar{x}^{(1+\theta\omega)} = \frac{\theta}{\theta-1} \bar{s} \frac{\sum_{j=0}^{\infty} (\alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\rho)(1+\omega)})^j}{\sum_{j=0}^{\infty} (\alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{(\theta-\rho(\theta-1))})^j}. \quad (55)$$

If both $\alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\rho)(1+\omega)}$ and $\alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{(\theta-\rho(\theta-1))}$ are less than 1, the two infinite sums converge, and we obtain

$$\bar{x}^{(1+\theta\omega)} = \frac{\theta}{\theta-1} \left(\frac{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{\theta-\rho(\theta-1)}}{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\rho)(1+\omega)}} \right) \bar{s}. \quad (56)$$

The requirement that the two sums in (55) converge requires that trend inflation must satisfy³⁴

$$\bar{\pi} < \left(\frac{1}{\alpha \bar{R} \bar{\gamma}_y} \right)^{\frac{1}{1+\theta(1-\rho)(1+\omega)}} \quad \text{and} \quad \bar{\pi} < \left(\frac{1}{\alpha \bar{R} \bar{\gamma}_y} \right)^{\frac{1}{\theta-\rho(\theta-1)}}. \quad (57)$$

Combining (56) with the aggregate price condition (2) evaluated at the steady state,

$$\bar{x} = \left(\frac{1 - \alpha \bar{\pi}^{(\theta-1)(1-\rho)}}{1 - \alpha} \right)^{\frac{1}{1-\theta}}, \quad (58)$$

³⁴For any value of $\bar{\pi}$, \bar{R} , and $\bar{\gamma}_y$, there exists values of the pricing parameters for which these inequalities hold. For example, if trend inflation were very high, then $\alpha \doteq 0$ might be needed to satisfy these inequalities. But that makes good economic sense, for the higher is trend inflation the more flexible prices are likely to be. Our estimates always satisfy these bounds.

we get the relationship between steady state $\bar{\pi}$ and \bar{s} :

$$(1 - \alpha \bar{\pi}^{(\theta-1)(1-\varrho)})^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{1+\theta(1-\varrho)(1+\omega)}}{1 - \alpha \bar{R} \bar{\gamma}_y \bar{\pi}^{\theta-\varrho(\theta-1)}} \right) = (1 - \alpha)^{\frac{1+\theta\omega}{1-\theta}} \left(\frac{\theta}{\theta-1} \right) \bar{s}. \quad (59)$$

In the particular case of zero steady-state inflation ($\bar{\pi} = 1$), or perfect indexation ($\varrho = 1$), the expression for the aggregate price level reduces to $\bar{x} = 1$, hence, by (56), $\bar{s} = \frac{\theta-1}{\theta}$.

The log-linearization of the optimal price equation (1) and of the aggregate price evolution (2) around a steady state with inflation $\bar{\pi}$ are respectively

$$\begin{aligned} \hat{x}_t = & \frac{1 - \alpha \tilde{\beta} \xi_2}{1 + \theta\omega} E_t \sum_{j=0}^{\infty} (\alpha \tilde{\beta} \xi_2)^j \\ & \times \left(\hat{R}_{t,t+j} + \hat{s}_{t+j} + \sum_{k=1}^j \hat{\gamma}_{y,t+k} + [1 + \theta(1 + \omega)] \sum_{k=1}^j \hat{\pi}_{t+k} - \varrho\theta(1 + \omega) \sum_{k=0}^{j-1} \hat{\pi}_{t+k} \right) \\ & - \frac{1 - \alpha \tilde{\beta} \xi_1}{1 + \theta\omega} E_t \sum_{j=0}^{\infty} (\alpha \tilde{\beta} \xi_1)^j \left(\hat{R}_{t,t+j} + \sum_{k=1}^j \hat{\gamma}_{y,t+k} + \theta \sum_{k=1}^j \hat{\pi}_{t+k} - \varrho(\theta-1) \sum_{k=0}^{j-1} \hat{\pi}_{t+k} \right), \end{aligned} \quad (60)$$

and

$$\hat{x}_t = \frac{\alpha \xi_1}{1 - \alpha \xi_1} (\hat{\pi}_t - \varrho \hat{\pi}_{t-1}), \quad (61)$$

where the symbols are defined in (7) in the text.

Combining these two equations, simplifying the double sums, and collecting terms, we obtain

$$\begin{aligned} \hat{\pi}_t - \varrho \hat{\pi}_{t-1} = & \frac{1 - \alpha \xi_1}{\alpha \xi_1} \left\{ \frac{1 - \gamma_2}{1 + \theta\omega} \sum_{j=1}^{\infty} \gamma_2^j \hat{s}_{t+j} - \left(\frac{\varrho\theta(1 + \omega)\gamma_2}{1 + \theta\omega} - \frac{\varrho(\theta-1)\gamma_1}{1 + \theta\omega} \right) \hat{\pi}_t \right. \\ & + \frac{1 + \theta(1 + \omega)(1 - \varrho\gamma_2)}{1 + \theta\omega} \sum_{j=1}^{\infty} \gamma_2^j E_t \hat{\pi}_{t+j} - \frac{[\theta(1 - \varrho\gamma_1) + \varrho\gamma_1]}{1 + \theta\omega} \sum_{j=1}^{\infty} \gamma_1^j E_t \hat{\pi}_{t+j} \\ & + \frac{1}{1 + \theta\omega} \left((1 - \gamma_2) \sum_{j=0}^{\infty} \gamma_2^j E_t \hat{R}_{t,t+j} - (1 - \gamma_1) \sum_{j=0}^{\infty} \gamma_1^j E_t \hat{R}_{t,t+j} \right) \\ & \left. + \frac{1}{1 + \theta\omega} \sum_{j=1}^{\infty} (\gamma_2^j - \gamma_1^j) E_t \hat{\gamma}_{y,t+j} \right\} \end{aligned} \quad (62)$$

where γ_1 and γ_2 are also defined in (7) in the text. Finally, we evaluate this expression at $t+1$, multiply it by γ_2 , and subtract its expected value from (62). Collecting terms, we obtain expression (4) in the text.

B Appendix B: Simulating the Posterior Density

Collect the drifting parameters into an array

$$\Theta^T = [\vartheta^T, H^T], \quad (63)$$

and let Ψ denote the static parameters Q , $b = \text{vec}(B)$, and $\sigma = (\sigma_1, \dots, \sigma_N)$. The posterior density,

$$p(\Theta^T, \Psi | X^T), \quad (64)$$

summarizes beliefs about the evolution of the drifting parameters and static hyperparameters. This posterior is simulated via the Markov Chain Monte Carlo algorithm of Cogley and Sargent (2004). They demonstrate that

$$p(\Theta^T, \Psi | X^T) \propto I(\Theta^T) f(\Theta^T, \Psi | X^T), \quad (65)$$

where $f(\Theta^T, \Psi | X^T)$ is the posterior corresponding to the model that does not impose the stability constraint which rules out explosive *VAR* roots. Therefore a sample from $p(\Theta^T, \Psi | X^T)$ can be drawn by simulating $f(\Theta^T, \Psi | X^T)$ and discarding realizations that violate the stability constraint. They also develop a ‘Metropolis within Gibbs’ algorithm for simulating $f(\Theta^T, \Psi | X^T)$ that involves cycling through 5 steps.

1. Sample ϑ^T from $f(\vartheta^T | X^T, H^T, Q, \sigma, b)$ using the forward-filtering, backward-sampling algorithm of Carter and Kohn (1994). This step relies on the Kalman filter and a recursion analogous to the Kalman smoother to update conditional means and variances.
2. Sample Q from $f(Q | X^T, \vartheta^T, H^T, \sigma, b)$. This is a standard inverse-Wishart problem.
3. Sample H^T by cycling through a number of univariate Metropolis chains for $f(h_{it} | h_{-it}, X^T, \vartheta^T, \sigma_i)$, where h_{-it} denotes the rest of the h_{it} vector at dates other than t . This step exploits the stochastic volatility algorithm of Jacquier, Polson, and Rossi (1994).
4. Sample σ from $f(\sigma | X^T, \vartheta^T, H^T, Q, b)$. This is a standard draw from an inverse-gamma density.
5. Sample b from $f(b | X^T, \vartheta^T, H^T, Q, \sigma)$. This is a Bayesian regression, and it is also standard.

The sequence of draws from the conditional submodels forms a Markov Chain that converges to a draw from the joint density, $f(\Theta^T, \Psi | X^T)$. The sample from the unrestricted model can then be transformed into a sample from the restricted model, $p(\Theta^T, \Psi | X^T)$, via rejection sampling. Details of each step and a justification for rejection sampling can be found in the appendices to Cogley and Sargent (2004).

References

- [1] Amato, J. and T. Laubach, 2003. Estimation and control of an optimization-based model with sticky prices and wages, *Journal of Economic Dynamics and Control* 27, 1181-1215.
- [2] Ascari, G., 2004. Staggered prices and trend inflation: some nuisances, *Review of Economic Dynamics* 7, 642-667.
- [3] Ascari, G. and T. Ropele, 2004. Monetary policy under low trend inflation, unpublished.
- [4] Ball, L., N.G. Mankiw and D. Romer, 1988. The new Keynesian macroeconomics and the output-inflation trade-off, *Brookings Papers on Economic Activity* 1, 1-65.
- [5] Bakhshi, H., P. Burriel-Llombart, H. Khan and B. Rudolf, 2003. Endogenous price stickiness, trend inflation and the new Keynesian Phillips curve, Bank of England working paper no. 191.
- [6] Basu, S., 1996. Procyclical productivity: increasing returns or cyclical utilization?, *Quarterly Journal of Economics* CXI, 719-751.
- [7] Basu, S. and M. Kimball, 1997. Cyclical productivity with unobserved input variation, NBER working paper no. 5915.
- [8] Batini, N., B. Jackson and S. Nickell, 2002. An open-economy new Keynesian Phillips curve for the UK, forthcoming, *Journal of Monetary Economics*.
- [9] Bernanke, B. and I. Mihov, 1998. Measuring monetary policy, *Quarterly Journal of Economics* CXIII, 869-902.
- [10] Bils, M., 1987. The cyclical behavior of price and marginal cost, *American Economic Review* 77, 838-855.
- [11] Bils, M. and P.J. Klenow, 2004. Some evidence on the importance of sticky prices, *Journal of Political Economy* 112, 947-985.
- [12] Boivin, J. and M. Giannoni, 2003. Has monetary policy become more effective? NBER working paper no. 9459.
- [13] Calvo, G., 1983. Staggered prices in a utility-maximizing framework, *Journal of Monetary Economics* 12, 383-398.
- [14] Campbell, J.Y. and R.J. Shiller, 1987. Cointegration and tests of present value models, *Journal of Political Economy* 95, 1062-1088.

- [15] Carter, C.K. and R. Kohn, 1994, On Gibbs sampling for state-space models, *Biometrika* 81, 541-553.
- [16] Christiano, L., M. Eichenbaum, and C. Evans, 2003. Nominal rigidities and the dynamic effects of a shock to monetary policy, forthcoming *Journal of Political Economy*.
- [17] Cogley, T. and T. J. Sargent, 2001. Evolving post-World War II U.S. inflation dynamics *NBER Macroeconomics Annual* , vol. 16.
- [18] ———, 2004. Drifts and volatilities: monetary policies and outcomes in the post WWII U.S., forthcoming *Review of Economic Dynamics*.
- [19] *Economic Report of the President* 2004. United States Government Printing Office, Washington DC.
- [20] Edge, R.M., T. Laubach, and J.C. Williams, 2003. The response of wages and prices to technology shocks, Federal Reserve Board Finance and Economics Discussion Series 2004-65.
- [21] Gagnon, E. and H. Khan, 2003. New Phillips curve under alternative production technologies for Canada, the United States, and the Euro area, forthcoming, *European Economic Review*.
- [22] Gali, J. and M. Gertler, 1999. Inflation dynamics: A structural econometric analysis, *Journal of Monetary Economics* 44, 195-222.
- [23] Gali, J., M. Gertler, and J.D. Lopez-Salido, 2000. European inflation dynamics, *European Economic Review* 47, 1237-1270.
- [24] Giannoni, M. and M. Woodford, 2003. Optimal inflation targeting rules, in B. Bernanke and M. Woodford, eds. *Inflation Targeting*, Chicago: University of Chicago Press.
- [25] Jacquier, E., N.G. Polson, and P. Rossi, 1994. Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* 12, 371-418.
- [26] Khan, H. and R. Moessner, 2003. Competitiveness, inflation and monetary policy, manuscript, Bank of England.
- [27] Kurmann, A., 2003. Quantifying the uncertainty about a forward-looking new Keynesian pricing model, forthcoming, *Journal of Monetary Economics*.
- [28] Li, H., 2004. Estimation and testing of Euler equation models with time-varying reduced-form coefficients, unpublished.

- [29] Linde', J. 2003. Estimating new-Keynesian Phillips curves: a full information maximum likelihood approach, unpublished.
- [30] Lucas, R.E., Jr. 1976. Econometric policy evaluation: a critique," in *The Phillips Curve and Labor Markets*, edited by K. Brunner and A. Meltzer, Carnegie-Rochester Series on Public Policy, vol. 1.
- [31] McConnell, M. and G. Perez Quiros, 2000. Output fluctuations in the United States: what has changed since the early '80s? *American Economic Review* 90, 1464-1476.
- [32] Phelps, E.S. and S.G. Winter, 1970. Optimal price policy under atomistic competition, in E.S. Phelps et al., *Microeconomic Foundations of Employment and Inflation Theory*, Norton, New York.
- [33] Rotemberg, J.J. and M. Woodford, 1992. Oligopolistic pricing and the effect of aggregate demand on economic Activity, *Journal of Political Economy* 100, 1153-1207.
- [34] Rotemberg, J.J. and M. Woodford, 1997. An Optimization-based econometric framework for the evaluation of monetary policy, in B. Bernanke and J.J. Rotemberg, *NBER Macroeconomics Annual*, MIT press, Cambridge, pp. 297-346.
- [35] Sahuc, J., 2004. Partial indexation, trend inflation, and the hybrid Phillips curve, unpublished.
- [36] Sbordone, A.M., 2002. Price and unit labor costs: a new test of price stickiness, *Journal of Monetary Economics* 49, 265-292.
- [37] Sbordone, A.M., 2003. A limited information approach to the simultaneous estimation of wage and price dynamics, Rutgers University, working paper no. 2003-23.
- [38] Sims, C. and T. Zha, 2004. Were there regime switches in U.S. monetary policy?, unpublished.
- [39] Small, I., 1997. The cyclicality of mark-ups and profit margins: some evidence for manufacturing and services, Bank of England working paper no. 72.
- [40] Smets, F. and R. Wouters, 2002. Sources of business cycle fluctuations in the U.S.: a Bayesian DSGE approach, unpublished.
- [41] Stock, James H. and Mark W. Watson, 2002, Has the business cycle changed and why? *NBER Macroeconomics Annual* 17, 159-218.

- [42] Woodford, M., 2003. *Interest & Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press: Princeton NJ.