Comparing New Keynesian models of the business cycle: A Bayesian approach ☆

Pau Rabanal[a], Juan F. Rubio-Ramírez[b],*

[a]*International Monetary Fund, 700 19th Street NW, Washington, DC 20431, USA*
[b]*Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, NE, Atlanta, GA 30309, USA*

Abstract

The baseline New Keynesian model cannot replicate the observed persistence in inflation, output, and real wages for sensible parameter values. As a result, several extensions have been suggested to improve its fit to the data. We use a Bayesian approach to estimate and compare the baseline sticky price model of Calvo's [1983. Staggered prices in a utility maximizing framework. Journal of Monetary Economics 12, 383–398.] and three extensions. Our empirical results are as follows. First, we find that adding price indexation improves the fit of Calvo's [1983. Staggered prices in a utility maximizing framework. Journal of Monetary Economics 12, 383–398.] model. Second, models with both staggered price and wage setting dominate models with only price rigidities. Third, introducing wage indexation does not significantly improve the fit. Fourth, all model estimates suggest a high degree of price stickiness. Fifth, the estimates of labor supply elasticity are higher in models with both staggered price and wage contracts. Finally, the estimated inflation parameters of the Taylor rule are stable across models.
© 2005 Elsevier B.V. All rights reserved.

JEL classification: C11; C15; E31; E32

Keywords: Nominal rigidities; Indexation; Bayesian econometrics; Model comparison

1. Introduction

Dynamic general equilibrium models with nominal rigidities ("New Keynesian" models) have become increasingly popular in the analysis of monetary policy. However, the baseline sticky price model does not generate the persistence in inflation, output, and real wages that we observe in the data unless implausible levels of nominal rigidity are assumed.[1] As a result, several extensions to the baseline sticky price model have been considered to improve its fit to the data. Despite these extensions, the existing literature lacks a formal comparison between competing alternatives.

In this paper, we fill this gap; we use a Bayesian approach to estimate and compare the baseline sticky price model of Calvo (1983) and three extensions. The first extension introduces price indexation to last period's inflation rate. Introducing price indexation results in a lagged inflation term in the price equation and, therefore, a better fit of inflation persistence. In the second extension, we add staggered wage contracts to the baseline sticky price model as in Erceg et al. (2000). As Galí and Gertler (1999) point out, in a pure forward-looking model, inflation persistence is driven by the sluggish adjustment of real marginal costs. Adding sticky nominal wages to the sticky price model delivers sticky *real* wages. Since inflation is a discounted stream of real marginal costs, dampening the real marginal cost movement restrains inflation fluctuation, generating more persistence. The final extension adds wage indexation to the Erceg et al. (2000) model.

On the estimation side, we combine priors and the likelihood function to obtain the posterior distribution of the structural parameters. We use the Kalman filter to evaluate the likelihood function of a log-linear approximation of the model and the Metropolis-Hastings algorithm to draw from the posterior distribution. Then, we use the marginal likelihood to compare the four models. In doing so, we are able to determine how much each additional model extension helps in explaining the data, and we are able to compare the models. An advantage of the marginal likelihood criterion is that it penalizes overparametrization. Therefore, models with more rigidities do not necessarily rank better if the extra rigidity does not sufficiently help in explaining the data.

Although we are not aware of any formal work comparing different New Keynesian models, various approaches have been used to estimate the structural parameters of some extensions of the baseline sticky price model. Galí and Gertler (1999) and Sbordone (2002) used minimum distance methods to estimate price and/or wage-setting equations separately. Kim (2000) and Ireland (2001) pursued maximum likelihood estimation in a general equilibrium framework. Rotemberg and Woodford (1997) and Christiano et al. (2005) minimized the distance between a structural VAR and the models' predicted impulse responses to a monetary shock. Finally, Smets and Wouters (2003) used a Bayesian approach to estimate a New Keynesian model using a "synthetic" data set for the Euro area.

We view our paper as a complement to previous approaches. Like Kim (2000) and Ireland (2001), we use a likelihood approach to estimate the structural parameters of the model. Taken further, we can easily perform model comparison of models using a Bayesian approach. Christiano et al. (2005) approach allows a better understanding of the

---

[1]For instance, Fuhrer and Moore (1995) show that a sticky wage model generates persistence in the price level but not in the inflation rate. Chari et al. (2000) point out that models with nominal rigidities do not generate enough persistence in output following a monetary shock.

models' implications for the transmission mechanism of monetary policy. However, looking at the overall fit and comparing different alternatives is an imperative exercise to evaluate the models' performances.

We take a Bayesian approach for several reasons. First, it takes advantage of the general equilibrium approach. As discussed in Leeper and Zha (2000), estimation of reduced-form equations or partial equilibrium models suffers from identification problems. Second, the Bayesian approach outperforms GMM and maximum likelihood in small samples. Third, it does not rely on the identification scheme of the VAR but does follow the likelihood principle (see Berger and Wolpert, 1998). In addition, Fernández-Villaverde and Rubio-Ramírez (2004) show that, even in the case of misspecified models, Bayesian estimation and model comparison are consistent.

The main results of this paper are as follows. First, adding price indexation to the baseline sticky price model clearly improves the fit. This result holds because introducing price indexation results in a lagged inflation term in the price equation and therefore, it better fits inflation persistence. Second, Erceg et al. (2000) model dominates the baseline sticky price model, even if we consider price indexation. This occurs because Erceg et al. (2000) model is better able to match the autocorrelation in the real wage. Finally, adding wage indexation to the Erceg et al. (2000) model does not substantially improve the fit to the data.

Other empirical results worth mentioning include a high degree of price stickiness across models and smaller estimates of the elasticity of labor supply in models with staggered wage contracts. Also, none of the models match the degree of autocorrelation in the nominal interest rate and the estimated inflation parameters of the Taylor rule are stable across models and in accord with prior studies.

The remainder of the paper is organized as follows. In Section 2 we discuss the dynamics of each model, with particular attention placed on the price- and wage-setting equations. In Section 3 we explain the data, the likelihood function, and the priors. In Section 4 we present and discuss the results, leaving Section 5 for concluding remarks.

2. The models

In this section we describe the four models. Our baseline model is a sticky price model as in Calvo (1983) (we refer to this model as BSP). We extend this baseline model in three different ways. First, we allow for indexation in prices to last period's inflation rate (henceforth INDP). Second, in accordance with Erceg et al. (2000), we introduce staggered wage contracts (henceforth EHL). Finally, we allow for both staggered wage contracts and indexation in wages to last period's inflation rate (henceforth INDW).

Since these four models are well known in the literature, we detail only the equations that describe the linear dynamics of each model. These equations are obtained by taking a log-linear approximation of the first-order conditions around the steady state. An accurate description of the various models can be found in the appendix of Rabanal and Rubio-Ramírez (2005). Throughout this paper, the lower-case variables denote log-deviations from steady-state values.

The rest of section describes the set of equations that is common to the four models and discusses the price- and wage-setting equations, which are different for each model.

2.1. Common equations

First, we have the Euler equation, which relates output growth with the real rate of interest in the following way:

$$y_t = E_t y_{t+1} - \sigma(r_t - E_t \Delta p_{t+1} + E_t g_{t+1} - g_t), \tag{1}$$

where y_t denotes output, r_t is the nominal interest rate, g_t is the preference shifter shock, p_t is the price level, σ is the elasticity of intertemporal substitution, Δ is the first difference operator, and E_t is the conditional expectation operator with information up to time t.

The production and the real marginal cost of production functions are defined by

$$y_t = a_t + (1 - \delta)n_t \quad \text{and} \quad mc_t = w_t - p_t + n_t - y_t, \tag{2}$$

where a_t is a technology shock, n_t is the amount of hours worked, mc_t is the real marginal cost, w_t is the nominal wage, and δ is the capital share of output.

The desired marginal rate of substitution (mrs_t) between consumption and hours takes the form

$$mrs_t = \frac{1}{\sigma} y_t + \gamma n_t - g_t, \tag{3}$$

where γ is the inverse elasticity of labor supply with respect to real wages.

We use the following specification for the Taylor rule:

$$r_t = \rho_r r_{t-1} + (1 - \rho_r)[\gamma_\pi \Delta p_t + \gamma_y y_t] + z_t, \tag{4}$$

where γ_π and γ_y are the long-run responses of the monetary authority to deviations of inflation and output from their steady-state values, and z_t is the monetary shock, to be defined below. We include an interest rate smoothing parameter, ρ_r, following recent empirical work (as in Clarida et al., 2000).

In order to close the model, we link real wage growth, nominal wage growth, and price inflation in the following way:

$$w_t - p_t = w_{t-1} - p_{t-1} + \Delta w_t - \Delta p_t. \tag{5}$$

We specify the shocks as

$$a_t = \rho_a a_{t-1} + \varepsilon_t^a,$$
$$g_t = \rho_g g_{t-1} + \varepsilon_t^g,$$
$$z_t = \varepsilon_t^z \quad \text{and}$$
$$\lambda_t = \varepsilon_t^\lambda,$$

where each innovation ε_t^i follows a *normal* $(0, \sigma_i^2)$ distribution, for $i = a, g, z, \lambda$, and innovations are uncorrelated with each other.

Below, we discuss the price- and wage-setting equations, which are different across models.

2.2. Baseline sticky price model (BSP)

The optimal choice of price setters under the Calvo-type restriction delivers the following forward-looking equation for price inflation (Δp_t):

$$\Delta p_t = \beta E_t \Delta p_{t+1} + \kappa_p(mc_t + \lambda_t), \tag{6}$$

where $\kappa_p = (1 - \delta)(1 - \theta_p\beta)(1 - \theta_p)/\{\theta_p[1 + \delta(\bar{\varepsilon} - 1)]\}$ and $\bar{\varepsilon} = \bar{\lambda}/(\bar{\lambda} - 1)$ is the steady-state value of ε.

Eq. (6) is the so-called New Keynesian Phillips curve, which relates current inflation to expectations of future inflation, the real marginal cost, and the price-markup shock. It denotes the forward-looking behavior of the firms in response to the Calvo-type restriction.

Since wages are flexible, the usual condition that real wages equal the desired marginal rate of substitution is satisfied. Therefore

$$w_t - p_t = mrs_t \tag{7}$$

holds.

2.3. Model with sticky prices and price indexation (INDP)

In this case, Eq. (6) is replaced by

$$\Delta p_t = \gamma_b \Delta p_{t-1} + \gamma_f E_t \Delta p_{t+1} + \kappa'_p(mc_t + \lambda_t), \tag{6'}$$

where $\kappa'_p = \kappa_p/(1 + \omega\beta)$, $\gamma_b = \omega/(1 + \omega\beta)$, and $\gamma_f = \beta/(1 + \omega\beta)$, and ω is the degree of price indexation to last period's inflation. The wage-setting equation remains the same (7).

2.4. Model with sticky prices and wages (EHL)

In this case, both price and wage inflation behave in a forward-looking manner. The price inflation equation is given by (6). Introducing the Calvo-type wage restriction delivers the following process for the nominal wage growth equation (Δw_t), replacing (7):

$$\Delta w_t = \beta E_t \Delta w_{t+1} + \kappa_w[mrs_t - (w_t - p_t)], \tag{7'}$$

where $\kappa_w = (1 - \theta_w)(1 - \beta\theta_w)/[\theta_w(1 + \phi\gamma)]$. With staggered wage-setting, it no longer holds that workers remain on their desired labor supply schedule all the time. Hence, the driving force of current nominal wage growth is the expected nominal wage growth as well as the distance between the desired marginal rate of substitution and the real wage.

2.5. Model with sticky prices, wages, and wage indexation (INDW)

Expanding on EHL, the nominal wage growth equation (7') incorporates indexation:

$$\Delta w_t - \alpha\Delta p_{t-1} = \beta E_t \Delta w_{t+1} - \alpha\beta\Delta p_t + \kappa_w[mrs_t - (w_t - p_t)], \tag{7''}$$

where α is the degree of wage indexation to last period's inflation.

3. Empirical analysis

Section 3 outlines how to draw from the posterior distribution of the models' structural parameters and evaluate the marginal likelihood of the data implied by each model. First, we describe the data we want to explain. Second, we write the likelihood function of the data implied by each of the models. Third, we describe the prior distribution of the parameters. Finally, we briefly explain how to draw from the posterior distribution.

3.1. The data

We explain the joint behavior of price inflation, real wages, interest rates, and real output for the United States at a quarterly frequency. The sample period is 1960:01 to 2001:04. The series for output, prices, and wages come from the Bureau of Labor Statistics. Let $\{d_t\}_{t=1}^T$ represent the series of observables. We use "output for the nonfarm business sector" as a measure of output and its associated price deflator as a measure of prices. We use "hourly compensation for the nonfarm business sector" as nominal wages. Finally, we use the federal funds rate as the relevant instrument for monetary policy. This last series comes from the FRED data base that uses the Board of Governors of the Federal Reserve System as a source. We demean all variables and detrend the real wage and output series using a quadratic trend.

3.2. The likelihood function

Let $\psi = (\sigma, \theta_p, \theta_w, \beta, \phi, \alpha, \gamma_y, \gamma_\pi, \rho_r, \rho_a, \rho_g, \delta, \bar{\lambda}, \gamma, \sigma_a, \sigma_m, \sigma_g, \sigma_\lambda)'$ be the vector of parameters that describe preferences and technology of each model. We use standard solution methods to solve the system of equations (1)–(7). Then, we write the solution in state-space form and use the Kalman filter to evaluate the likelihood of each of the models. Let $L(\{d_t\}_{t=1}^T | \psi, m)$ be the likelihood function of model m.

3.3. The priors

Table 1 presents the prior distribution of the parameters. The inverse of the elasticity of intertemporal substitution, σ^{-1}, follows a gamma distribution. This assumption implies a positive support for σ^{-1}. Given our hyperparameters, we assume a prior mean of 2.5 and a standard deviation of 1.76. We also assume a gamma distribution for the average duration of prices.[2] Thus, the average duration of prices has a prior mean of 3 and a prior standard deviation of 1.42. This assumption reflects the informal evidence presented in Taylor (1999).

Regarding the Taylor rule coefficients, because we do not impose nonnegativity restrictions, we assume normal distributions. We set the mean of γ_π to 1.5 and that of γ_y to 0.125, which are Taylor's original estimates.[3] We assume a normal distribution for the inverse of the elasticity of the labor supply, γ, centered at 1 and with a standard deviation of 0.5. The interest rate smoothing coefficient, ρ_r, has a uniform prior between $[0, 1)$. We choose prior uniform distributions between $[0, 1)$ for the autoregresive parameter of the technology and preference shocks and for all standard deviations. We make this choice for two reasons: we do not have strong prior information about the standard deviations of the innovations; and lower values of estimated σ_λ necessitate higher values of estimated κ_p to explain the observed inflation volatility. Since there is a negative relationship between κ_p and θ_p, higher values of κ_p result in lower values of the estimated θ_p. Therefore, truncation of σ_λ can result in underestimation of θ_p. We want to preclude the

---

[2]Since we need to keep the probability of the Calvo lottery between 0 and 1, we formulate the prior in terms of the parameter $1/(1 - \theta_p) - 1$.

[3]Taylor (1993) uses annual data, while we use quarterly data. Therefore, we multiply the prior mean of γ_y by four to compare it to Taylor's results.

Table 1
Prior and posterior distributions for the parameters

	Prior distribution		Posterior distribution			
			BSP	INDP	EHL	INDW
		Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)
$\frac{1}{1-\theta_p}$	Gamma(2, 1) + 1	3.00 (1.42)	4.49 (0.43)	6.07 (0.14)	4.37 (0.35)	4.18 (0.22)
$\frac{1}{1-\theta_w}$	Gamma(3, 1) + 1	4.00 (1.71)	1 (−)	1 (−)	2.72 (0.27)	2.31 (0.17)
ω	Uniform[0, 1]	0.5 (0.28)	− (−)	0.76 (0.06)	− (−)	− (−)
α	Uniform[0, 1]	0.5 (0.28)	− (−)	− (−)	− (−)	0.25 (0.07)
γ_π	Normal(1.5, 0.25)	1.5 (0.25)	1.02 (0.02)	1.08 (0.04)	1.08 (0.09)	1.13 (0.12)
γ_y	Normal(0.125, 0.125)	0.125 (0.125)	0.10 (0.03)	0.09 (0.03)	0.26 (0.06)	0.27 (0.06)
ρ_r	Uniform[0, 1]	0.5 (0.28)	0.46 (0.04)	0.39 (0.05)	0.74 (0.02)	0.78 (0.02)
σ^{-1}	Gamma(2, 1.25)	2.5 (1.76)	6.67 (3.30)	8.33 (2.20)	8.33 (2.50)	7.69 (2.20)
γ	Normal(1, 0.5)	1.0 (0.5)	0.46 (0.09)	0.49 (0.10)	1.74 (0.29)	2.22 (0.33)
ρ_a	Uniform[0, 1]	0.5 (0.28)	0.85 (0.03)	0.83 (0.03)	0.74 (0.05)	0.73 (0.05)
ρ_g	Uniform[0, 1]	0.5 (0.28)	0.80 (0.03)	0.85 (0.03)	0.82 (0.03)	0.82 (0.03)
$\sigma_a(\%)$	Uniform[0, 1]	50.0 (28.0)	1.15 (0.23)	1.15 (0.23)	3.88 (1.09)	3.79 (0.89)
$\sigma_m(\%)$	Uniform[0, 1]	50.0 (28.0)	0.43 (0.03)	0.47 (0.04)	0.33 (0.02)	0.34 (0.02)
$\sigma_\lambda(\%)$	Uniform[0, 1]	50.0 (28.0)	40.82 (8.42)	52.94 (5.62)	31.67 (5.32)	25.56 (2.41)
$\sigma_g(\%)$	Uniform[0, 1]	50.0 (28.0)	8.46 (2.08)	10.40 (2.47)	11.88 (3.28)	10.73 (2.10)
$\log(\hat{L})$			−	64.20	146.91	149.57

underestimation of θ_p and be symmetric on the prior assumptions for all four standard deviations; therefore, we opt for high prior upper bounds on all four standard deviations.

In the BSP model, wages are flexible, and there is no price indexation. Therefore, we set θ_w, α, and ω to zero. In the INDP model, while we maintain θ_w and α equal to zero, we assume a prior uniform distribution between 0 and 1 for the price indexation parameter, ω. In the EHL model, we set the two indexation parameters, α and ω, to zero, and we establish a gamma distribution for the prior duration of wages with a mean of four quarters and standard deviation of 1.71. We assume a uniform distribution between 0 and 1 for the prior distribution of the wage indexation parameter, α. Finally, we limit the support of all parameters to the region where the model has a unique, stable solution.[4]

We impose dogmatic priors over the parameters β, δ, ϕ, and $\bar{\varepsilon}$. Because we do not consider capital, we have difficulty estimating β and δ. Similarly, there is an identification problem between the probability of the Calvo lottery, θ_p, and the mean of the price markup, $\bar{\varepsilon}$.[5] Therefore, it is impossible to identify θ_p and $\bar{\varepsilon}$ at the same time. This problem

---

[4]We use an appropriate normalizing constant to ensure that the prior is a proper density.

[5]The slope of the Phillips curve, κ_p, is the only equation containing $\bar{\varepsilon}$ and θ_p.

also emerges between θ_w and ϕ. The values we use ($\beta = 0.99$, $\delta = 0.36$, $\phi = 6$ and $\bar{\varepsilon} = 6$) are quite conventional in the literature.

3.4. Drawing from the posterior and model comparison

Let $M = \{BSP, INDP, EHL, INDW\}$ be the set of models that we wish to compare and let $m \in M$. In the Bayesian approach, the main parameter estimation tool is the parameters' posterior distribution of model m given the data, $p(\psi|\{d_t\}_{t=1}^T, m)$, while the main model comparison apparatus is the marginal likelihood of model m, $L(\{d_t\}_{t=1}^T|m)$.

The posterior distribution is proportional to the product of the likelihood function and the prior. Given our priors and the likelihood functions implied by the models, we are unable to obtain a closed-form solution for the posterior distributions. However, we are able to evaluate both expressions numerically. Therefore, we use the random walk Metropolis-Hastings algorithm to obtain 2,000,000 draws from each model's posterior distribution. We use the draws to estimate the moments of the posterior distributions. The marginal likelihood is equal to the integral of the likelihood function across the parameter space using the prior as the weighting function. Since, we are unable to obtain the marginal likelihood's closed-form, we follow Geweke (1998) to estimate it. As shown in Fernández-Villaverde and Rubio-Ramírez (2004), if $m^* \in M$ is the best model under the Kullback-Leibler distance, then for any other model $n \in M$, the Bayes factor of model n over model m, $L(\{d_t\}_{t=1}^T|n)/L(\{d_t\}_{t=1}^T|m^*)$, converges to zero as T increases. Hence, we focus on the Bayes factor as a tool to determine which model best explains the joint behavior of our four variables. We provide a more careful description of the Metropolis-Hastings algorithm and Geweke's (1998) procedure in the appendix of Rabanal and Rubio-Ramírez (2005).

4. Findings

In this section we present our findings. First, we present our posterior moments estimates for each of the four models. Second, we display the estimates of the marginal likelihood for each of the models. Third, in order to check the robustness of our results, we recompute the posterior moments estimates and the marginal likelihoods using 1982:04 to 2001:04 data. Finally, we analyze the persistence that each of the models generates and compare the results with the persistence observed in the data.

4.1. Posterior distributions and moments

The last four columns of Table 1 present the mean and the standard deviation of the posterior distributions of the parameters for the four models. The fourth column of Table 1 presents the estimates for the BSP model. The posterior mean of the average duration of price contracts is 4.49 quarters, which, by any standard, does not imply a too-long price duration.[6] In the BSP model, wages are flexible, so we fix the parameter on wage duration to be one. For the coefficients of the Taylor rule, the coefficient on inflation is close to one,

---

[6]Our results depend on the values of β and $\bar{\varepsilon}$. However, for a reasonable range of values, the average duration of prices does not change significantly.

with a small posterior standard deviation. We find a coefficient on the output gap and the interest rate smoothing parameter similar to those reported by Clarida et al. (2000).

The fifth column of Table 1 reports the results of the INDP model. In this case, we see that the coefficient on price indexation is high, with a posterior mean value close to 0.76. The average duration of price contracts increases to 6.07 quarters. Note that this result does not imply a rejection of the forward-looking nature of price inflation versus a pure backward-looking specification of inflation. In order to obtain pure backward-looking behavior in this model, we would need to estimate high indexation and price flexibility. Clearly, the two conditions are not met. When looking at the reduced-form values that result from our estimates, we obtain coefficients close to one-half for both the forward- and the backward-looking component of inflation. Hence, the parameter estimates favor a hybrid specification for price inflation. The estimates of the Taylor rule for the INDP model are similar to those for the BSP model.

We present the estimates of the EHL model in the sixth column of Table 1. The estimated average duration of price contracts is 4.37 quarters, similar to that estimated under BSP. A surprising result is the low estimated average duration of wage contracts. We obtain an average duration of less than 3 quarters (the point estimate is 2.72). Given our priors, we expected wages to be fixed for longer periods of time than prices.[7] However, we should stress that wage flexibility is rejected, as indicated from the posterior standard deviation of the wage duration parameter. The estimated output gap and lagged interest rate coefficients on the Taylor rule differ from those estimated for the models with flexible wages. Both coefficients are higher. On the other hand, the estimated coefficient on inflation remains close to one. The last column of Table 1 presents the estimates of the INDW model. The parameter of wage indexation is 0.25. The estimated price and wage durations are lower (4.18 and 2.31 quarters, respectively) than in the EHL model. The parameters estimated for the Taylor rule are similar to those estimated for EHL.

The remainder of the estimated parameters is as follows. For the elasticity of intertemporal substitution, σ, we obtain estimates that range between 0.12 and 0.15. These values are similar to those usually reported in the literature. See, for instance, in DeJong et al. (2000) and Basu and Kimball (2000).[8] The estimates of the elasticity of labor supply, γ^{-1}, are smaller in models with staggered wage contracts (and closer to values suggested in studies using microdata, see Altonji, 1986). We estimate values of γ close to 0.5 for the models with flexible wages (BSP and INDP) and close to 2 for the models with staggered wage contracts (EHL and INDW).[9] When wages are flexible the marginal rate of substitution between consumption and labor equals the real wage. Therefore, large values of this elasticity are needed to match the observed real wage fluctuations. In a setting with staggered price and wage, the marginal rate of substitution does not need to be equal to the real wage. Hence, it is not necessary to rely on high estimates of the elasticity of labor supply in order to match the data. In all four models we obtain high correlation coefficients for the technology and preference-shifter shocks.

---

[7]There are interactions between ϕ and the duration of wage contracts. It is difficult to obtain a higher duration of wage contracts for reasonable values of ϕ. Similarly, Christiano et al. (2005) find an average duration of wage contracts of 3 quarters.

[8]Methods that minimize the distance between model-based and VAR impulse responses rely on higher values for σ. See, for instance, Rotemberg and Woodford (1997).

[9]We do not obtain the downward-sloping labor supply schedule of Sbordone (2002).

The estimated posterior mean for σ_λ is always larger than 25% (being 52.94% in the case of the INDP model). This result is important for two reasons. First, it shows that these models imply a large estimated volatility of price markups. Second, it justifies the choice of a higher upper bound on the prior distribution of σ_λ relative to the other standard deviations. As a comparison, all other standard deviation estimates are lower than 12%. In addition, the estimated posterior means for σ_λ for EHL and INDW are lower than the estimated value for BSP and INDP. We believe that this difference reflects higher endogenous inflation persistence in models with any type of wage rigidity.

Therefore, we reach the following conclusions. First, data clearly provide support for an average duration of price contracts between four and seven quarters and a average duration of wage contracts of less than three quarters. Second, price indexation is more important than wage indexation. Third, the estimated Taylor rule coefficients for inflation remain stable across models and very close to one. Finally, the estimates of the elasticity of labor supply, γ^{-1}, are smaller in models with sticky wages.

4.2. Model comparison

Which model best explains the behavior of our data set? The last row of Table 1 reports the difference between the log marginal likelihood of each model with respect to the log marginal likelihood of BSP.[10] The Bayes factor clearly favors INDW and EHL over BSP and INDP; i.e., the data favor models with both price and wage stickiness over models with only price rigidities. As we explain in Section 4.4, models with price and wage stickiness are able to match the autocorrelation in the real wage more closely. Additionally, the log marginal likelihood difference between INDP and BSP is large. Hence, the data favor price indexation. Introducing price indexation results in a lagged inflation term in the price equation and, therefore, a better fit of inflation persistence. Finally, the log marginal likelihood difference between INDW and EHL is less than three. As suggested by Jeffreys (1961), this difference cannot be accepted as decisive evidence in favor of one model over the other. Therefore, adding wage indexation to price and wage stickiness does not improve the ability of the model to explain the data.

The log marginal likelihood difference between INDP and BSP is 64.20. This result suggests that in order to choose BSP over INDP, we need a prior probability over BSP 7.6×10^{27} (= exp(64.20)) times larger than our prior probability over INDP. We believe that this factor is too large; therefore, we conclude that price indexation improves the baseline sticky price model considerably.

How does the inclusion of sticky wages to the baseline sticky price model compare to the inclusion of price indexation? The log marginal likelihood difference between EHL and INDP is 82.7. This result implies that in order to choose INDP over EHL, we need a prior probability over INDP 8.3×10^{35} times larger than our prior over EHL. Since this is a similar factor as our BSP to INDP factor, we conclude that EHL outperforms INDP.

How much does wage indexation add to EHL? In this case we would need only to have a prior probability over EHL 14.3 (= exp(2.66)) times larger than our prior over INDW in order to choose EHL. Since this factor is not as large as the one reported before, we conclude that wage indexation does not improve the ability of the EHL model to explain the data.

---

[10]Since we are only interested in ranking the models, the relative log marginal likelihood is sufficient.

Table 2
Prior and posterior distributions for the parameters (sample period 1982:04 to 2001:04)

	Prior distribution		Posterior distribution			
			BSP	INDP	EHL	INDW
		Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)
$\frac{1}{1-\theta_p}$	Gamma(2, 1) + 1	3.00 (1.42)	7.71 (0.47)	7.58 (0.53)	6.20 (0.80)	6.39 (1.03)
$\frac{1}{1-\theta_w}$	Gamma(3, 1) + 1	4.00 (1.71)	1 (–)	1 (–)	2.37 (0.39)	2.50 (0.54)
ω	Uniform[0, 1]	0.5 (0.28)	– (–)	0.54 (0.09)	– (–)	– (–)
α	Uniform[0, 1]	0.5 (0.28)	– (–)	– (–)	– (–)	0.28 (0.08)
γ_π	Normal(1.5, 0.25)	1.5 (0.25)	1.33 (0.15)	1.40 (0.14)	1.39 (0.17)	1.42 (0.17)
γ_y	Normal(0.125, 0.125)	0.125 (0.125)	0.27 (0.05)	0.25 (0.05)	0.53 (0.12)	0.53 (0.13)
ρ_r	Uniform[0, 1]	0.5 (0.28)	0.64 (0.04)	0.61 (0.04)	0.75 (0.03)	0.76 (0.03)
σ^{-1}	Gamma(2, 1.25)	2.5 (1.76)	5.26 (2.30)	6.67 (2.10)	4.55 (2.80)	5.26 (2.20)
γ	Normal(1, 0.5)	1.0 (0.5)	0.80 (0.19)	0.89 (0.24)	1.68 (0.40)	1.64 (0.38)
ρ_a	Uniform[0, 1]	0.5 (0.28)	0.92 (0.02)	0.88 (0.02)	0.83 (0.04)	0.82 (0.04)
ρ_g	Uniform[0, 1]	0.5 (0.28)	0.67 (0.16)	0.65 (0.16)	4.10 (2.02)	5.19 (3.10)
$\sigma_a(\%)$	Uniform[0, 1]	50.0 (28.0)	1.15 (0.23)	1.15 (0.23)	3.88 (1.09)	3.79 (0.89)
$\sigma_m(\%)$	Uniform[0, 1]	50.0 (28.0)	0.20 (0.02)	0.20 (0.02)	0.18 (0.02)	0.18 (0.02)
$\sigma_\lambda(\%)$	Uniform[0, 1]	50.0 (28.0)	55.56 (5.51)	44.71 (5.15)	40.49 (10.50)	42.13 (13.87)
$\sigma_g(\%)$	Uniform[0, 1]	50.0 (28.0)	5.38 (0.16)	6.47 (0.19)	5.94 (2.31)	6.47 (2.56)
$\log(\hat{L})$			–	13.41	32.94	34.56

It is natural to ask why a richer model (INDW) does *not* rank better than a simpler model (EHL). The reason is simple: richer models have many more hyperparameters, and the Bayes factor discriminates against them. This "built-in" Ockham's razor is a final and attractive feature of the Bayes factor that embodies a strong preference for parsimonious modeling.

4.3. Using a more appropriate sample for the Taylor rule

A strong feature of the estimates presented in Table 1 is that the estimated coefficient on the reaction of the Taylor rule to price inflation is extremely close to one. Since we require a Taylor rule that induces a unique and stationary solution, values for this parameter that are less than one are ruled out by our priors. In this subsection we discuss the results of reducing our sample period to starting in 1982:04. The choice to reduce the sample period reflects the fact that the Fed shifted its operating procedure at that time toward using a target for the federal funds rate (see Clarida et al., 2000; Bernanke and Mihov, 1998).

Results of the new estimates are presented in Table 2. The main differences with respect to the full sample estimates are: (i) the parameter on the reaction of the Taylor rule to price
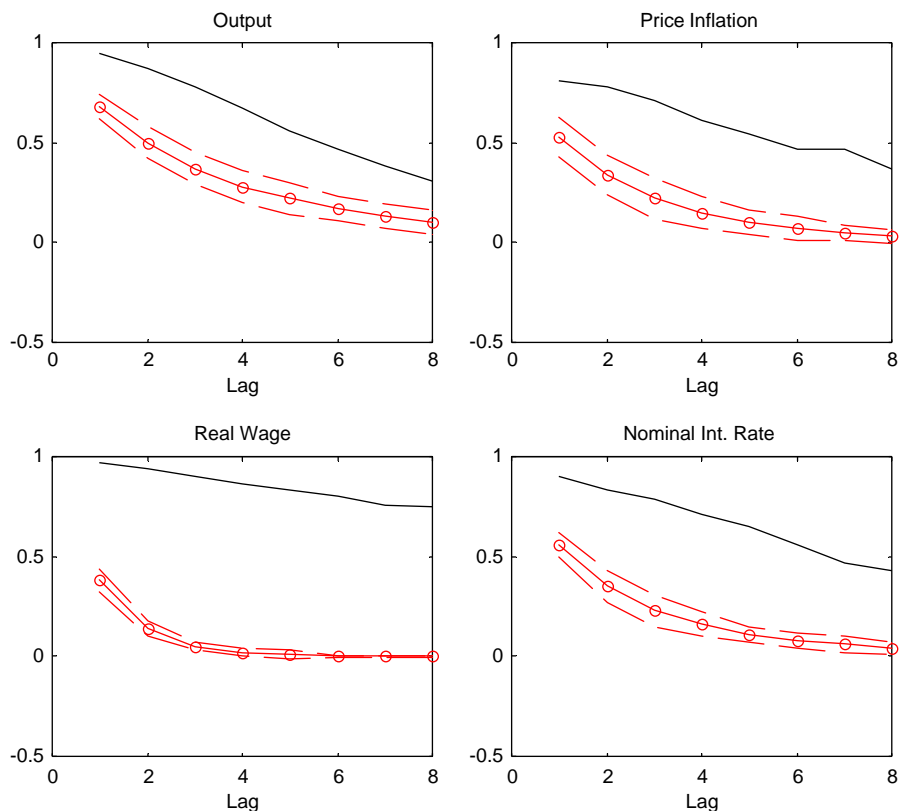
Fig. 1. Autocorrelations, BSP model. Circle = mean posterior, dashed lines = ±2SD posterior, solid line = US Data.

inflation is estimated to be higher—in the range of 1.3–1.5; (ii) the parameter on the reaction of the Taylor rule to the output gap is also estimated to be higher—in the range of 0.25–0.53; (iii) the coefficient on price indexation decreases to 0.54; (iv) the degree of wage indexation increases to 0.28; and (v) the average duration of both price and wage contracts increases slightly.

The remaining parameters of the models do not change significantly from what we estimated using the full sample. We estimate smaller standard deviations for productivity, money and preference shocks, reflecting the significantly lower volatility of the macrovariables during this period. Finally, we obtain the same qualitative results regarding the use of the marginal likelihood to determine which model best explains the data.

4.4. Persistence

An important shortcoming of the baseline sticky price model is its inability to generate enough persistence in the endogenous variables when facing exogenous shocks. In Figs. 1–3, we compare the implied autocorrelation functions of output, price inflation, real
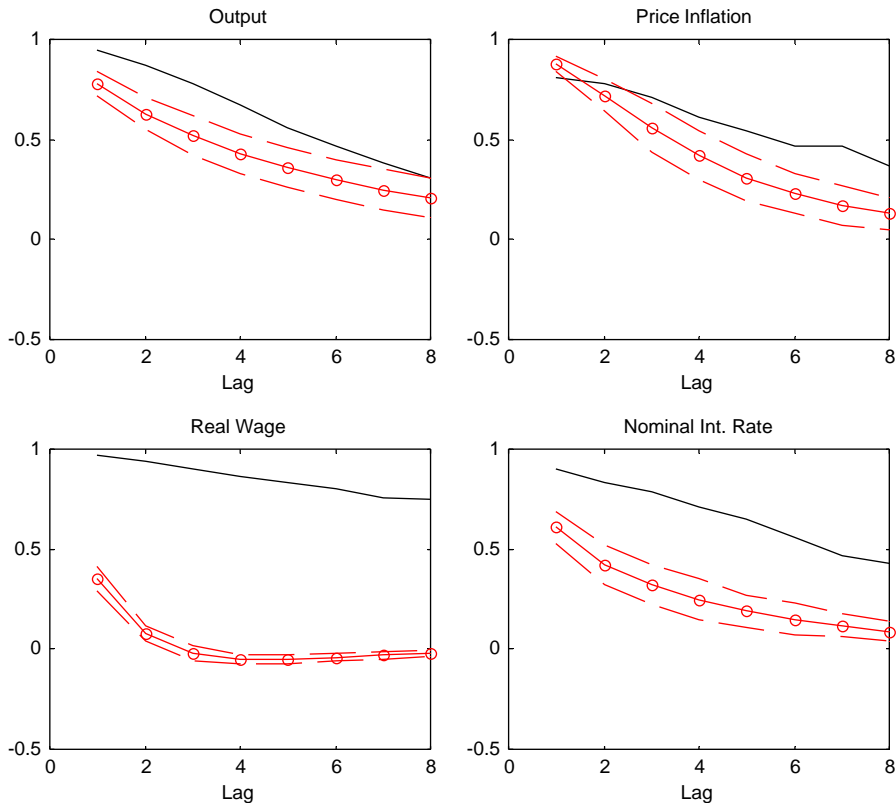
Fig. 2. Autocorrelations, INDP model. Circle = mean posterior, dashed lines = ±2SD posterior, solid line = US Data.

wages, and nominal interest rates of the BSP, INDP, and EHL models to U.S. data for the period 1960:01 to 2001:04.[11] We present the posterior mean and bands of two posterior standard deviations for each autocorrelation function.

In Fig. 1, we present the autocorrelation functions of the BSP model. Overall, the following picture emerges. It is not possible to match the autocorrelation function of real wages and the nominal interest rate; however, the autocorrelation functions of inflation and output match closely to the data. In Fig. 2, the autocorrelation functions are shown for the INDP model. The fit improves greatly for the behavior of inflation persistence. Hence, price indexation is important in order to explain inflation persistence. However, the model still does not explain the observed autocorrelation of real wages and the nominal interest rate. In Fig. 3, we present the implied persistence of the EHL model. In this case we observe less robust results with respect to the INDP model in matching the persistence of price inflation. We have to keep in mind that the EHL model preserves the pure forward-looking behavior of the price- and wage-setting equations, while INDP relies on lagged inflation. Meanwhile, the implied persistence of the real wage greatly increases closer to the

---

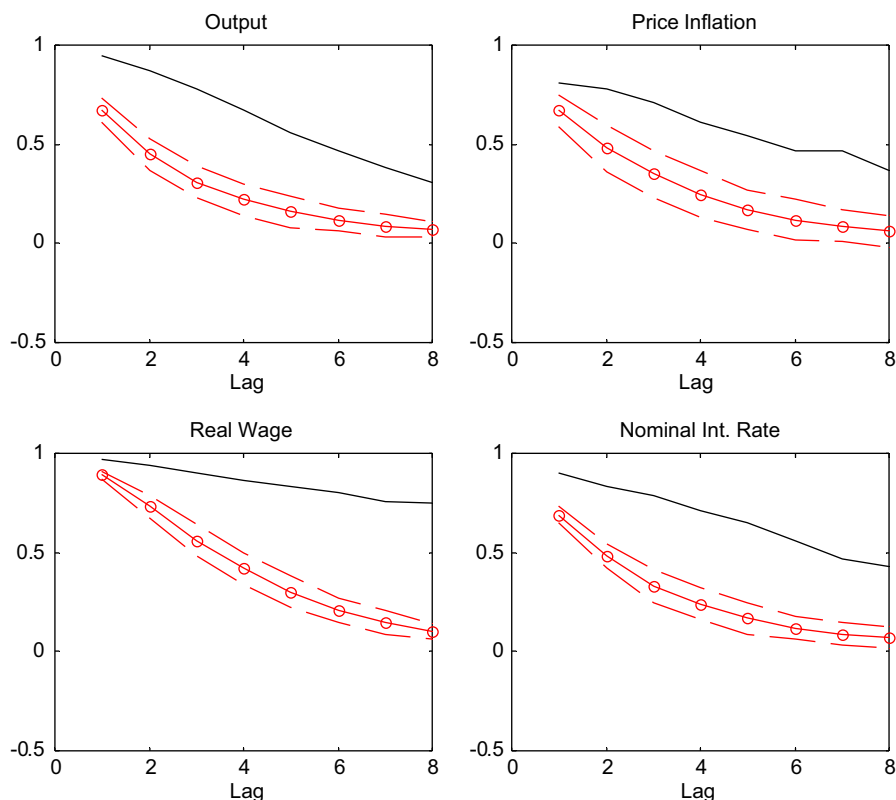[11]The autocorrelation functions of INDW look very much like those of EHL, so they are not reported.

Fig. 3. Autocorrelations, EHL model. Circle = mean posterior, dashed lines = ±2SD posterior, solid line = US Data.

observed one. Hence, in order to explain inflation and real wage persistence we need models with both staggered price and staggered wage contracts. Finally, the EHL model is still not able to match the observed persistence in the nominal interest rate.

Hence, we conclude that (i) while adding price indexation to the baseline sticky price model helps to explain the persistence of price inflation, it does not help to explain the persistence of real wages; (ii) in order to match the persistence of real wages, we need to consider both staggered price and staggered wage contracts; and (iii) none of the considered models can match the observed nominal interest rate autocorrelation function.

It is important to note the convenience of using the marginal likelihood to compare between alternative models. In order to discriminate among models using model-based and observed autocorrelations, we would need to specify (i) a *distance* to measure the difference between estimated and observed autocorrelations and (ii) a *loss function* that would determine which autocorrelations, and which lags, are the most important to match. The marginal likelihood criterion solves these two needs. A gratifying result is that the model that obtains the highest marginal likelihood, EHL, seems to match the data best.

5. Concluding remarks

In this paper, we use a Bayesian approach to estimate and compare the baseline sticky price model of Calvo (1983) and three extensions. We find the following (i) adding price indexation to the baseline sticky price model clearly improves the fit; (ii) models with both price and wage staggered contracts dominate models with only price rigidities; (iii) all model estimates suggest a high degree of price stickiness; (iv) the estimates of the elasticity of labor supply are smaller in models with wages rigidities; (v) none of the models match the degree of autocorrelation in the nominal interest rate; and (vi) the estimated inflation parameters of the Taylor rule are stable across models.

We restrict ourselves to estimate simple models that are used in the analysis of monetary policy. In future research, it would be interesting to incorporate factors of interest for policymakers including (i) capital accumulation and investment rigidities; (ii) an explicit credit channel; (iii) other types of labor market rigidities; and (iv) exchange rates, international trade, and other open economy aspects.

References

Altonji, J.G., 1986. Intertemporal substitution in labor supply: evidence from micro data. The Journal of Political Economy 94, S176–S215.

Basu, S., Kimball, M.S., 2000. Long-run labor supply and the elasticity of intertemporal substitution. Mimeo, University of Michigan.

Berger, J., Wolpert, R., 1998. The Likelihood Principle: A Review and Generalizations, second ed. Institute of Mathematical Statistics, Hayward.

Bernanke, B.S., Mihov, I., 1998. Measuring monetary policy. Quarterly Journal of Economics 113, 869–902.

Calvo, G., 1983. Staggered prices in a utility maximizing framework. Journal of Monetary Economics 12, 383–398.

Chari, V.V., Kehoe, P., McGrattan, E., 2000. Sticky price models of the business cycle: can the contract multiplier solve the persistence problem? Econometrica 68, 1151–1181.

Christiano, L., Eichenbaum, M., Evans, C., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. Journal of Political Economy, forthcoming.

Clarida, R., Galí, J., Gertler, M., 2000. Monetary policy rules and macroeconomic stability: evidence and some theory. Quarterly Journal of Economics 115, 147–180.

DeJong, D.N., Ingram, B.F., Whiteman, C.H., 2000. A Bayesian approach to dynamic macroeconomics. Journal of Econometrics 98, 203–223.

Erceg, C.J., Henderson, D.W., Levin, A.T., 2000. Optimal monetary policy with staggered wage and price contracts. Journal of Monetary Economics 46, 281–313.

Fernández-Villaverde, J., Rubio-Ramírez, J.F., 2004. Comparing dynamic equilibrium economies to data: a Bayesian approach. Journal of Econometrics 123, 153–187.

Fuhrer, J.C., Moore, G., 1995. Inflation persistence. Quarterly Journal of Economics 110, 127–160.

Galí, J., Gertler, M., 1999. Inflation dynamics: a structural econometric analysis. Journal of Monetary Economics 44, 195–222.

Geweke, J., 1998. Using Simulation Methods for Bayesian Econometric Models: Inference, Development and Communication, Federal Reserve Bank of Minneapolis Staff Report 249.

Ireland, P., 2001. Sticky price models of the business cycle: specification and stability. Journal of Monetary Economics 47, 3–18.

Jeffreys, H., 1961. Theory of Probability, third ed. Oxford University Press, Oxford.

Kim, J., 2000. Constructing and estimating a realistic optimizing model of monetary policy. Journal of Monetary Economics 45, 329–359.

Leeper, E., Zha, T., 2000. Assessing Simple Policy Rules: A View from a Complete Macro Model, Federal Reserve Bank of Atlanta Working Paper 2000-19.

Rabanal, P., Rubio-Ramírez, J.F., 2005. Comparing New Keynesian Models of the Business Cycle: A Bayesian approach. Federal Reserve Bank of Atlanta Working Paper 2001-22 (Revised version).

Rotemberg, J.J., Woodford, M., 1997. An optimization-based econometric framework for the evaluation of monetary policy. In: Bernanke, B.S., Rotemberg, J.J. (Eds.), NBER Macroeconomics Annual. The MIT Press, Cambridge, pp. 297–346.

Sbordone, A., 2002. An Optimizing Model of U.S. Wage and Price Dynamics, Rutgers University Working Paper 2001-10.

Smets, F., Wouters, R., 2003. An estimated stochastic dynamic general equilibrium model for the euro area. Journal of the European Economic Association 1, 1123–1175.

Taylor, J., 1993. Discretion versus policy rules in practice. Carnegie-Rochester Series on Public Policy 39, 195–214.

Taylor, J., 1999. Staggered price and wage-setting in macroeconomics. In: Taylor, J., Woodford, M. (Eds.), Handbook of Macroeconomics. North-Holland, Amsterdam, pp. 1009–1050.