CIRPÉE
Centre interuniversitaire sur le risque, les politiques économiques et l'emploi

Cahier de recherche/Working Paper **03-44**

Quantifying the Uncertainty about the Fit of a New Keynesian Pricing Model: Extended Version

André Kurmann

Novembre/November 2003

Kurmann: Department of Economics, Université du Québec à Montréal, P.O. Box 8888, Downtown station, Montréal (QC) H3C 3P8, Canada. *Phone*: 514-987-3000 ext. 3503; *Fax*: 514-987-8494
kurmann.andre@uqam.ca

A shorter version of this paper is forthcoming in the *Journal of Monetary Economics*. Special thanks to Bob King for invaluable guidance and support. Thanks also to Marianne Baxter, Mike Dotsey, Jinill Kim, Chris Otrok, Mark Watson and Alex Wolman for helpful suggestions on this research. Earlier versions of this paper have been presented at Bocconi University, Emory University, Michigan State University, Université de Montréal, UQAM, the University of Virginia, the Federal Reserve Bank of Boston, the Federal Reserve Bank of Richmond, the Bank of England, the Swiss National Bank, the Bank of Canada workshop on the Phillips Curve and the European Economic Association's 2002 annual meeting. I thank the seminar participants for their comments. Financial support by the Bankard Fund of the University of Virginia and the Swiss National Science Foundation is gratefully acknowledge.

Abstract: Recent studies by Gali and Gertler (1999) and Sbordone (2002) conclude that a theoretical inflation series implied by the forward-looking New Keynesian pricing model of Calvo (1983) fits post-1960 U.S. inflation closely. Their theoretical inflation series is conditional on (i) a reduced-form forecasting process for real marginal cost; and (ii) the calibration of the structural pricing equation implied by the Calvo model. The present paper shows that both of these determinants are surrounded by considerable uncertainty. When quantifying the impact of this uncertainty on theoretical inflation, I find that we can no longer say whether the Calvo model explains observed inflation dynamics very well or very poorly.

Keywords: Inflation, New Keynesian pricing, real marginal cost

JEL Classification: E31, E32, E37

1 Introduction

New Keynesian (NK) pricing models have emerged as the dominant theoretical attempt to explain the dynamics of inflation and its interaction with real aggregates. One of the most popular versions of NK pricing is derived from Calvo (1983) and implies that inflation is the present-value of current and expected future real marginal cost. Recently, studies by Gali and Gertler (1999) and Sbordone (2002) – referenced henceforth as GGS – conclude that if real marginal cost is measured by labor income share, the Calvo model provides a good approximation of post-1960 U.S. inflation. Besides reporting estimates that are broadly consistent with the theory, GGS illustrate the goodness-of-fit of the model with a theoretical (present-value) inflation series that is computed conditional on a reduced-form vector autoregressive (VAR) forecasting process for the expected future marginal cost terms. For their dataset, this theoretical inflation series tracks observed inflation closely.

The objective of this paper is to quantify the uncertainty about theoretical inflation in order to assess how much confidence we should have in the suggested good fit of the Calvo model. The reason for undertaking this task is straightforward. Theoretical inflation as calculated by GGS is a series of point estimates that depends on at least three important assumptions: (i) the assumption that the VAR coefficient estimates of the forecasting process for future marginal cost represent the true population values; (ii) the assumption that future real marginal cost terms are well forecasted from the information contained in the variables of the VAR process; and (iii) the assumption that the calibration of the structural pricing equation implied by the Calvo model is correct.[1]

To assess the robustness of the fit of theoretical inflation with respect to each of these assumptions, the paper proceeds in four steps. First, I compute a benchmark example of theoretical inflation conditional on a bivariate VAR forecasting process in labor income share and inflation that is very similar to the one reported by GGS. Using the same dataset as Gali and Gertler, I find that the resulting series of theoretical inflation fits observed inflation remarkably well. The estimated correlation coefficient between the two series is 0.97 and the volatility of theoretical inflation relative to observed inflation is 0.78.

Second, I present a bootstrap approach in order to quantify the uncertainty about the fit of theoretical inflation that is due to imprecisely estimated coefficients in the VAR forecasting process for marginal cost. When applied to the benchmark example, I find that the bootstrapped 90% confidence interval for the correlation coefficient between theoretical and observed inflation is large and extends from 0.40 to 0.99. Hence, taking into account the imprecision in the estimated VAR coefficients uncovers an important source of uncertainty about the fit of the Calvo model.

Third, I assess the robustness of the goodness-of-fit of theoretical inflation with respect to alternative specifications of the forecasting process. In particular, there is no specific reason to believe that a bivariate VAR in labor income share and inflation provides a good approximation of how markets

[1] Another important assumption is that real marginal cost is itself well approximated by labor income share, as GGS propose. For the purpose of this study, I will disregard this issue and uphold that labor income share *is* the correct measure of real marginal cost.

forecast future labor income share (i.e. real marginal cost). I show that computing theoretical inflation conditional on a VAR in labor income share and unit labor cost – as Sbordone proposes – does not change the results of the benchmark example substantially. Then, I invoke statistical selection criteria to motivate two alternative forecasting processes: a univariate AR process in labor income share only; and an expanded VAR that contains an additional set of prominent macro-economic aggregates. When computing theoretical inflation conditional on either one of these specifications, I find that the Calvo model fails to track observed inflation. The point estimate of the correlation coefficient between theoretical and observed inflation drops to 0.51 conditional on the AR forecasting process, while it falls to 0.55 for the expanded VAR (with a 90% confidence interval ranging from -0.54 to 0.84). This dramatic change in goodness-of-fit highlights that the empirical promise of the Calvo model crucially hinges on our assumptions about how markets forecast real marginal cost.

Fourth and finally, I evaluate the robustness of theoretical inflation to changing the calibration of the slope coefficient on real marginal cost in the Calvo pricing equation. This coefficient is, among other things, a function of the average degree of price rigidity in the economy and the elasticity of firms' real marginal cost with respect to their output. Micro surveys offer little guidance about the values of these two parameters and hence, the calibration of the slope coefficient remains an open question. Analytical developments reveal that the correlation coefficient between theoretical and observed inflation is not affected by the value of this slope coefficient but that it plays an important role for the volatility of theoretical inflation relative to the volatility of observed inflation. In particular, when calibrating the slope parameter such that it corresponds to an average price fixity of four quarters and assuming that firm-specific real marginal cost is inelastic – as proposed by Yun (1996) in the traditional version of the Calvo model – theoretical inflation is between 2.5 and 4 times more volatile than observed inflation (depending on the choice of forecasting process). Concurrently, for an alternative version of the pricing equation proposed by Sbordone that implies some degree of elasticity for firm-specific real marginal cost, the estimated relative volatility of theoretical inflation becomes more reasonable.

In sum, the analysis of this paper makes clear that the fit of theoretical inflation with observed inflation is surrounded by a great deal of uncertainty. Hence, we cannot say with any degree of confidence whether the Calvo model explains U.S. inflation very poorly or very well. This result represents a cautionary note about the positive conclusions by GGS. It suggests that we first need to work on a more precise understanding about the cyclical interaction of marginal cost with other macroeconomic aggregates as well as about the calibration of important firm-specific parameters, before we can reasonably assess the empirical relevance of the Calvo model.

2 The fit of the Calvo model with U.S. inflation

This section provides a brief overview of how to compute theoretical inflation from the Calvo pricing model. Following, I present a benchmark example that results in a theoretical inflation series similar to the ones reported by GGS.

2.1 Theoretical inflation

The Calvo pricing model is built on two central assumptions: (i) firms have market power; and (ii) firms face costs of changing prices. Following Blanchard and Kiyotaki (1987), market power is implemented by modelling firms as monopolistic competitors. As for the cost of changing prices, Calvo's framework simply posits that each firm has an exogenous probability $1-\kappa$ in any given period that it may change its price and consequently a probability κ that it must keep its price unchanged. This probability is independent of the number of periods since the firm was allowed its last price revision.

Given these assumptions, the average number of periods a firm keeps its price fixed equals $1/(1-\kappa)$, and the dynamics of inflation can be described by a simple log-linearized equation of the form:[2]

$$\pi_t = \beta\gamma_y E_t \pi_{t+1} + \phi\psi_t, \tag{1}$$

which can be rewritten in present-value form as

$$\pi_t = \phi \sum_{i=0}^{\infty} (\beta\gamma_y)^i E_t \psi_{t+i} \; . \tag{2}$$

The variables π_t and ψ_t represent percentage deviations of inflation and (average) real marginal cost from their respective steady states. The parameter β is the discount factor; γ_y is the steady state value of real output growth; and ϕ is a composite of different structural parameters that Woodford (2003) derives as[3]

$$\phi = \frac{(1-\kappa)(1-\kappa\beta\gamma_y)}{\kappa(1+\eta\mu)}. \tag{3}$$

In this expression, κ is the fraction of firms that cannot adjust their price in a given period; μ is the elasticity of substitution between differentiated goods; and η is the elasticity of the firm's real marginal cost with respect to its output. This last parameter η is a function of maintained assumptions about factor markets. GGS impose specific restrictions in order to calibrate η and I will return to discussing these restrictions in Section 5.

Aside from obtaining estimates of the structural parameters that are consistent with micro evidence, an important question is how well the Calvo model fits observed inflation dynamics. To this end, a theoretical inflation series can be computed using a VAR projection method that was first proposed by Campbell and Shiller (1987) in the context of the expectations theory of interest rates. The principal idea behind this method is that present-value equations such as (2) contain unknown expectational elements, which need to be forecasted in terms of available information in order to obtain an empirically operational expression. The derivation of this expression goes as follows. Consider a subset $\omega_t = [z_t \; z_{t-1} ... z_{t-p+1}]'$ of the market's full information set Ω_t (i.e. $\omega_t \subseteq \Omega_t$), where z_t is an n-variable vector of information available at date t but not at date $t-1$. Let z_t contain current real

[2] The simplicity of this pricing equation is the main reason for the popularity of the Calvo model and explains why it has been used in many DSGE models. See for example King and Wolman (1996), Yun (1996), Rotemberg and Woodford (1997), or McCallum and Nelson (1998).

[3] See the appendix for a derivation.

marginal cost (i.e. $\psi_t \in z_t$). Furthermore, assume that the dynamics of the np elements in ω_t are well described by a VAR process expressed in companion form as

$$\omega_t = M\omega_{t-1} + e_t. \tag{4}$$

Then, the law of iterated expectations implies that multiperiod forecasts of real marginal cost conditional on information ω_t are computed as

$$E[E_t\psi_{t+i}]|\omega_t = E[E\psi_{t+i}|\Omega_t]|\omega_t = E[\psi_{t+i}|\omega_t] = h_\psi M^i \omega_t \ , \tag{5}$$

where h_ψ is a $1 \times np$ selection vector that singles out the forecast for real marginal cost (e.g. if ψ_t takes the first position in z_t, then $h_\psi = [1\ 0\ 0...0]$). Finally, we map these forecasts into the present value representation of the Calvo pricing model (2) to obtain the following closed-form expression

$$\pi'_t = \phi \sum_{i=0}^{\infty} (\beta\gamma_y)^i E\left[\psi_{t+i}|\omega_t\right] = \phi h_\psi [I - \beta\gamma_y M]^{-1}\omega_t, \tag{6}$$

where I is an $np \times np$ identity matrix.

Analogous to Campbell and Shiller, π'_t is defined as *theoretical inflation.* It is the model-based path of inflation conditional on VAR forecasts of real marginal cost from information ω_t and represents a measure of how well the model fits the observed dynamics of inflation.[4] The advantage of computing theoretical inflation conditional on a reduced-form forecasting process is that it does not involve making assumptions about the structure of the rest of the economy (for example, about household preferences). In other words, the hope is that by taking as given the predicted path of future labor income share rather than deriving it from an explicit structural framework, it may be easier to identify failures that are specific to the proposed pricing model.

Under the null that the Calvo model is true and the additional assumption that π_t is part (or a linear combination) of the econometrician's information set ω_t (with both real marginal cost and inflation being perfectly observable), theoretical inflation equals the observed rate of inflation; i.e. $\pi_t = \pi'_t$. A direct implication of this equality is that under the null, observed inflation and theoretical inflation are perfectly correlated and have the same standard deviation. The correlation coefficient $\rho(\pi_t, \pi'_t)$ and the ratio of the standard deviations $\Delta(\pi_t, \pi'_t) \equiv \sigma(\pi_t)/\sigma(\pi'_t)$ therefore provide a set of statistics that summarize how well the Calvo model fits observed inflation dynamics.

An important point to emphasize about the Campbell and Shiller approach to computing theoretical inflation is that the coefficients of the VAR companion matrix M in (4) are left *unrestricted* and can therefore be estimated from ordinary least squares (OLS). As I discuss more completely in Kurmann (2003), this implies that we consider the Calvo model not as the true description of inflation

[4]The fit of theoretical inflation with observed inflation is by no means the only metric to evaluate how well the model can explain the data. For example, one could alternatively consider the forecasting performance of the model, or compare model-based correlation functions with their empirical counterparts (see for example Fuhrer and Moore, 1995; and Sbordone, 2002). The conclusions reached from these goodness-of-fit measures would be the same than the ones reached here.

but merely as an approximation. Concurrently, computing theoretical inflation under the null would necessitate imposing rational expectations cross-equation restrictions that ensure consistency of the VAR with the dynamics of inflation and labor income share as implied by the Calvo model.[5]

2.2 A benchmark example

A key issue in evaluating NK pricing models concerns the measurement of (average) real marginal cost. In principle, real marginal cost is some weighted average of marginal factor costs such as the real wage paid for hiring an additional unit of labor and the rental price of acquiring an incremental unit of capital. However, reliable data on the aggregate rental price of capital does not exist, which explains why researchers have adopted alternative methods for measuring real marginal cost.

Following Bils (1987), GGS invoke a framework where firms act as price takers in the labor market and are subject to a production function that is log-linear in the different inputs. Under these two assumptions, cost minimization implies that average real marginal cost is proportional to labor income share; i.e. $\psi_t = w_t - (y_t - n_t) = s_t$ in log-linearized form, where s_t denotes labor income share, w_t is the real wage, y_t is real output, and n_t stands for employment. Hence, the empirical specification of the Calvo pricing equation becomes

$$\pi_t = \beta\gamma_y E_t \pi_{t+1} + \phi s_t. \tag{7}$$

In their work, GGS estimate $\beta\gamma_y$ and ϕ for quarterly U.S. data between 1960 and 1997. Despite different estimation techniques, their results are very similar and paint a promising picture about the Calvo model's explanation of U.S. inflation.[6] For example, a representative estimate of (7) taken from the battery of results of Gali and Gertler's study is the one where $\beta\gamma_y$ is restricted to unity[7]

$$\pi_t = E_t \pi_{t+1} + \underset{(0.007)}{0.035}\, s_t \ , \tag{8}$$

where standard errors are in parenthesis here and below. The estimate of ϕ is positive and significantly different from zero, thus supporting the theory that real marginal cost is driving inflation. Furthermore, $\hat{\phi} = 0.035$ together with $\beta\gamma_y = 1$ implies between 2.5 and 6 quarters of average price rigidity, depending on the assumptions about ϕ in (3). This range is roughly consistent with the evidence from micro

[5] As Campbell and Shiller argue in their paper, not imposing these constraints is sensible because statistical tests of the cross-equation restrictions may be "...highly sensitive to deviations – so sensitive, in fact, that they may obscure some of the merits of the theory" [page 1080].

[6] Gali and Gertler estimate (7) via General Methods of Moments. Sbordone estimates $\beta\gamma_y$ and ϕ using an alternative formulation of the Calvo pricing equation that relates the (log-linearized) price level to the discounted present value of future expected unit labor cost (labor income share times the aggregate price level). Her estimation approach consists of minimizing the variance of the distance between the theoretical (model-based) price level and the observed price level with respect to the structural parameters. The theoretical price level is computed from a present-value representation of the Calvo model, taking as given the estimates of an unrestricted VAR forecasting process.

[7] Note that Gali and Gertler actually set $\gamma_y = 1$ and thus impose $\beta = 1$, which is the upper bound of theoretically admissible values for the discount factor.

surveys.[8] Finally, GGS test for the importance of the forward-looking characteristic of the Calvo model by estimating a pricing equation that allows for lagged price information to influence current inflation. Both studies find that expectations about future inflation matter greatly and generally dominate the lagged price information.

Taken together, the promising estimates by GGS suggest that the Calvo pricing model provides a good approximation of U.S. inflation dynamics conditional on real marginal cost being measured with labor income share. This finding stands in contrast with earlier studies by Fuhrer and Moore (1995) and others who approximate real marginal cost by the real output gap. These studies find that (i) the real output gap is not a significant determinant of inflation; and (ii) the purely forward-looking characteristic of the Calvo model cannot explain the substantial degree of inflation persistence that we observe in the data.[9] As Fuhrer and Moore emphasize, the Calvo model conditional on the real output gap fails because it implies a strong *positive* correlation between the current real output gap and lagged inflation. Yet, exactly the opposite pattern is observed in the data where the real output gap and lagged inflation are *negatively* correlated with each other.[10] Concurrently, as GGS point out, the Calvo model conditional on labor income is more successful because labor income share and inflation in the data are *positively* correlated at both leads and lags. GGS thus conclude that the main empirical difficulty is not in explaining inflation with a purely forward-looking pricing model, but in reconciling the behavior of output with the behavior of real marginal cost.[11]

To compute the benchmark example of theoretical inflation, I will adopt $\beta\gamma_y = 1$ and $\phi = 0.035$ as the calibration of the Calvo pricing model. Section 5 will return to assessing the robustness of the results reported here with respect to alternative values of ϕ. For the specification of the VAR process to forecast labor income share (i.e. real marginal cost), I follow Gali and Gertler and specify the information set ω_t as a vector of current and lagged labor income share and inflation.[12] Furthermore, I assume as in GGS that both inflation and labor income share are stationary (see the next section for more discussion). Using Gali and Gertler's quarterly U.S. dataset over the same sample 1960:1-1997:4, I choose to include four lags of each series; i.e. $\omega_t = [s_t \;\; \pi_t \;\; s_{t-1} \;\; \pi_{t-1} \; s_{t-2} \; \pi_{t-2} \; s_{t-3} \; \pi_{t-3}]'$.[13]

[8] See Taylor (1998) who concludes that the average degree of price fixity should not exceed 4 quarters. Section 5 will discuss the robustness of the results with respect to different calibrations of ϕ.

[9] Proxying real marginal cost with the real output gap leads to the following form of the Calvo pricing equation (1):

$$\pi_t = \beta\gamma_y E_t\pi_{t+1} + \phi\eta x_t,$$

where x_t stands for the real output gap and η is the output gap elasticity of real marginal cost. This expression is commonly referred to as the "New Keynesian Phillips curve".

[10] A similar point is made more recently by Estrella and Fuhrer (2002).

[11] A companion paper by Gali, Gertler and Lopez-Salido (2001) evaluates the Calvo specification for several European countries. The results are equally promising if not better.

[12] The data have been kindly provided by Jordi Gali. See the appendix for details.

[13] While the Aikake information criterion (AIC) selects an optimal lag number of three, I decided to adding one more lag so that there is no statistically significant evidence of either serial correlation or heteroscedasticity in the VAR residuals (absence of both serial correlation and heteroscedasticity are important regularity conditions for the bootstrap approach introduced in Section 3).

Table 1 reports the OLS estimates for the coefficients of this VAR (tables are reported at the end of the text). The large and highly significant coefficient on s_{t-1} in the labor income share equation and on π_{t-1} in the inflation equation illustrate the sluggish behavior of the two variables in the data. Also note that almost none of the other coefficient estimates differ significantly from zero. In particular, the role of lagged inflation in the labor income share equation is very imprecisely estimated, which is a finding that will figure importantly in the analysis of this paper.

With the coefficient estimates of the VAR forecasting process at hand and the coefficients of the Calvo pricing equation set to $\beta\gamma_y = 1$ and $\phi = 0.035$, theoretical inflation π'_t is easily computed from equation (6). Figure 1 shows the plots for theoretical inflation and observed inflation. The results

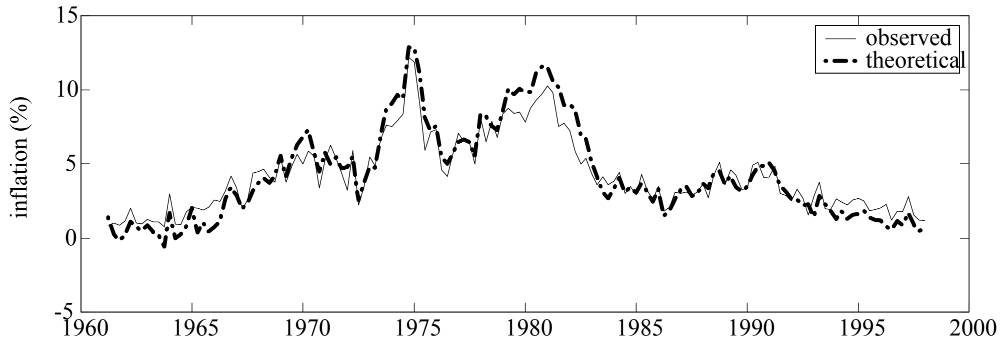

Figure 1: Fit of theoretical inflation computed from benchmark VAR

are remarkable: although the series of theoretical inflation does not capture every wiggle of the data, it overall tracks observed inflation well. This visual impression is tellingly summarized by the high estimated correlation coefficient of $\hat{\rho} = 0.97$ and a standard deviations ratio of $\hat{\Delta} = 0.78$. Furthermore, note that the fit of π'_t with π_t appears to be even better than in Gali and Gertler's study (their Figure 2) and also comes very close to the results reported by Sbordone (Figure 2b of her study).[14]

2.3 Behind the close fit

The close fit between theoretical and observed inflation underlines the conclusion by GGS that the Calvo model provides a good approximation of U.S. inflation dynamics. At the same time, it is important to realize that the calculation of theoretical inflation is conditional on a variety of assumptions. One is that labor income share (i.e. real marginal cost) is well forecasted by a VAR in four lags of labor income share and inflation. A second assumption is that both the coefficient estimates of the VAR and the calibrations chosen for $\beta\gamma_y$ and ϕ represent their true population values. In other words, the discussed series of theoretical inflation and thus its correlation and standard deviation relative

[14]Gali and Gertler actually compute their series of π'_t from a hybrid variant of the Calvo pricing model that has a lagged inflation term tagged on to it. The better results of the benchmark example here suggest that for the sample under consideration, adding a small lagged inflation term rather worsens than improves the (informal) fit of the model with the data.

to observed inflation (ρ and Δ) are mere point estimates. Hence, it is unclear how much confidence we can have in the close fit of theoretical inflation with observed inflation suggested by both the benchmark example above and the evidence reported in GGS.

The rest of the paper proceeds in three distinct steps to quantify this uncertainty about the goodness-of-fit. First, I compute the sample distributions of the correlation coefficient ρ and the variance ratio Δ as a function of the VAR coefficient estimates. Second, I assess the sensitivity of ρ and Δ to alternative specifications of the forecasting process (i.e. alternative information sets ω_t). Third, I evaluate the robustness with respect to the structural parameter ϕ.[15]

3 Uncertainty due to imprecision in the VAR coefficients

Table 1 reveals that most of the coefficients of the benchmark VAR forecasting process used above are imprecisely estimated. To quantify the impact of this imprecision on the theoretical inflation series and thus on the goodness-of-fit of the model, I use a bias-corrected bootstrap approach along the lines proposed by Kilian (1998a).

3.1 Motivating the bias-corrected bootstrap

Boostrapping in the present context consists of (i) generating many artificial series of the variables in the VAR from the estimated coefficients in $\hat{M}$ and the residuals $\hat{e}_t$ as if they were population values; (ii) estimating new VAR coefficients $\hat{M}^*$ from the simulated data, which in turn imply a simulated series of theoretical inflation (with $\beta\gamma_y$ and ϕ taken as given); and (iii) computing the correlation coefficient $\hat{\rho}^* = \hat{\rho}^*(\beta\gamma_y, \phi, \hat{M}^*)$ and the variance ratio $\hat{\Delta}^* = \hat{\Delta}^*(\beta\gamma_y, \phi, \hat{M}^*)$ for each of the simulation runs. The distribution of $\hat{\rho}$ and $\hat{\Delta}$ is then inferred from the series of simulated $\hat{\rho}^*$ and $\hat{\Delta}^*$.

The advantage of bootstrapping the distribution is that it respects by definition the boundedness of the statistics of interest (i.e. $-1 < \rho < 1$ and $\Delta > 0$). Furthermore, the bootstrap allows for skewness because it does not impose symmetry.[16] However, as Kilian (1998a) showed for the case of impulse response estimates, a standard bootstrap may perform poorly when it is used to compute distributions of statistics that are nonlinear functions of VAR coefficients (such as ρ and Δ). In fact, OLS coefficient estimates of autoregressive processes systematically suffer from small-sample bias.[17] As a result, the

[15] Campbell and Shiller's method of computing theoretical inflation is a two-step approach in the sense that the forecasting VAR is estimated separately from the pricing parameters. Therefore, no statistical link exists between the VAR coefficients and $\beta\gamma_y$, ϕ, which precludes us from computing the sample distribution of ρ and Δ jointly as a function of all the parameters that determine theoretical inflation.

[16] By contrast, traditional methods such as the asymptotic normal approximation that relies on a delta expansion impose symmetry and do not take into account the boundedness. Confidence intervals for $\hat{\rho}$ and $\hat{\Delta}$ infered from such methods may therefore be highly inappropriate.

[17] A standard assumption of OLS is that the regressors X_t are independent of regression error u_s for all t and s. Now, consider an autoregressive process $y_t = X_t b + u_t$, where X_t contains lagged values of y. Even if we assume that u_t and X_t are independent of each other, it will not be the case that u_t is independent of X_{t+1}, which means that OLS estimates of b will be biased in small samples. See Hamilton (1994), chapter 8.

small sample distributions of statistics that are nonlinear functions of these autoregressive coefficients (i.e. ρ and Δ) are likely to be biased, and correcting for median bias in these nonlinear statistics ignores the fact that their distributions are not scale invariant.

Given the difficulties with directly correcting for the bias in nonlinear statistics, Kilian proposes an adapted bootstrap algorithm that removes the bias prior to simulation. The idea is to replace the simulated VAR coefficient estimates $\hat{M}^*$ by bias-corrected estimates $\bar{M}^*$ before computing $\hat{\rho}^*$ and $\hat{\Delta}^*$, i.e. to bootstrap $\hat{\rho}^*(\beta\gamma_y, \phi, \bar{M}^*)$ and $\hat{\Delta}^*(\beta\gamma_y, \phi, \bar{M}^*)$ rather than $\hat{\rho}^*(\beta\gamma_y, \phi, \hat{M}^*)$ and $\hat{\Delta}^*(\beta\gamma_y, \phi, \hat{M}^*)$.[18] In addition, a second bias-correction is necessary because the OLS estimates $\hat{M}$ are themselves systematically biased away from their population value. Consequently, the coefficients in $\hat{M}$ cannot be considered good approximations of the population coefficients M and should not be used to generate artificial data series from which to estimate $\hat{M}^*$. To preserve the validity of the bootstrap, we thus need to bias-correct the point estimates $\hat{M}$ *prior* to simulating dataseries such that the bias-corrected simulated coefficients $\bar{M}^*$ are approximately unbiased estimators of the population coefficients M.

Aside from the double bias-correction, the employed bootstrap algorithm takes into account of lag-order uncertainty in the simulated VAR as suggested by Kilian (1998b), and applies Stine's (1987) block method to randomly select starting values for each bootstrap simulation.[19]

3.2 Preliminaries

The previously discussed bias-correction of the VAR estimates $\hat{M}$ has a direct effect on the estimated series of theoretical inflation. As Figure 2 shows, the bias-correction hardly affects the comovement

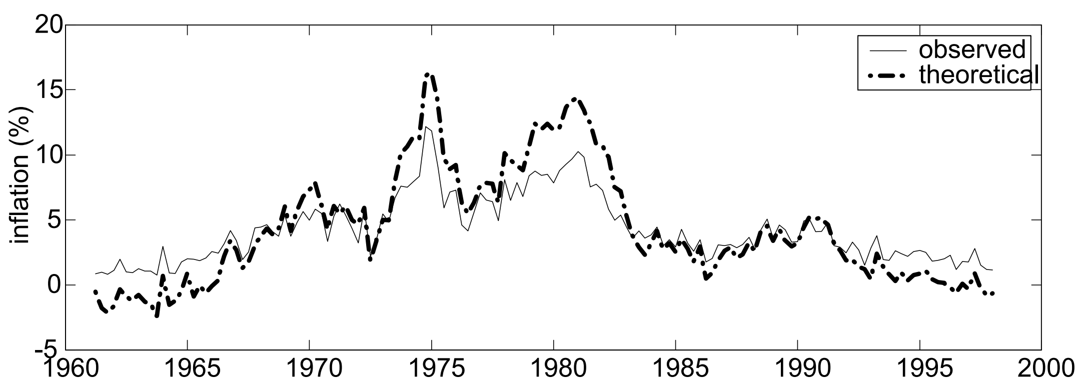

Figure 2: Fit of theoretical inflation computed from bias-corrected benchmark VAR

between theoretical inflation and observed inflation (the point estimate $\hat{\rho}$ increases slightly from 0.97

[18] Since the bias in $\hat{\rho}^*$ and $\hat{\Delta}^*$ may arise not only from bias in $\hat{M}^*$ but also because of the nonlinear nature of the two statistics, this procedure will not in general produce unbiased estimates. However, under the assumption that it yields a bootstrap that is unbiased on average, the sample distribution is likely to be a good approximation. Furthermore, it is important to point out that Kilian's approach only corrects for first-order bias.

[19] Refer to the appendix for a description of the different steps in the bootstrap.

to 0.98) However, the impact on the volatility of theoretical inflation is noticeable, with the estimated standard deviations ratio decreasing by roughly 30% from 0.79 to 0.57. This sizable negative impact on the volatility of theoretical inflation is surprising given that none of the bias-corrections exceeds 0.015. It is a first indicator that the fit of theoretical with observed inflation is sensitive even to small changes in the underlying VAR coefficients.

The accuracy of the bootstrap approach to computing the sample distributions of ρ and Δ relies on two important assumptions: (i) the errors e_t of the VAR are homoscedastic and serially uncorrelated; (ii) the variables in the VAR are stationary. As for the first assumption, I find no statistical evidence of serial correlation and heteroscedasticity for neither the labor income share nor the inflation equation of the (bias-corrected) bivariate VAR(4) of the benchmark example.[20] With regards to the stationarity assumption for labor income share and inflation, the null of a unit root can be rejected for labor income share but I am unable to reject the same null for inflation.[21] Despite the unit root problem for inflation, which is a common finding in macroeconomics (see for example King and Watson, 1994), I choose to remain with the stationarity assumption for inflation because of three reasons. First, the existence of a steady state for inflation (i.e. stationarity) is one of the main assumptions underlying the derivation of the Calvo pricing equation. Second, bootstrapping the distributions of ρ and Δ requires simulating inflation series from the VAR and thus, inflation (either directly or as some linear combination of other variables) needs to be included in the information set. Third, unit root tests are known for their low power in small samples, which is one of the reasons why the (non-)stationarity of inflation remains an open question.[22]

[20] To test for serial correlation, I carry out a Breusch-Godfrey Lagrange Multiplier (BGLM) test, which is based on a regression of the residuals of the equation on the original regressors and lags of the residuals. The BGLM test statistic is computed as the number of observations times the uncentered R^2 of the regression. Under the null of no serial correlation, this statistic is distributed as χ^2_p, where p is the number lags of the residuals in the regression (see Godfrey, 1988 for details). The p-values for the labor income share equation and the inflation equation of the (bias-corrected) benchmark VAR(4) are 0.496 and 0.177, respectively. Thus, the null of no serial correlation cannot be rejected.

To test for heteroscedasticity, I use White's (1980) test that is based on a regression of the squares of the residuals on all the regressors. The test statistic is the same than for the BGLM test. Under the joint null that the errors are both homoscedastic and independent of the regressors (and that the linear specification of the regression is correct), this statistic is distributed as χ^2_n where n is the number of coefficients in the test regression. The p-values for the labor income share equation and the inflation equation of the (bias-corrected) benchmark VAR(4) are 0.101 and 0.142, respectively. Thus, the joint null of homoscedasticity and independence of the errors cannot be rejected at high significance levels.

[21] The results are based on an Augmented Dickey-Fuller (ADF) test, which consists of regressing the first difference of the variable to be tested on its lagged level and several lags of first differences. The null of a unit root corresponds to a zero coefficient on the lagged level of the variable. For three lags of first differences in the regression, the ADF test statistic for labor income share is 3.29, which means that the null can be rejected at the 5% level. For inflation, the ADF test statistic equals 1.74, which means that we cannot reject the null of a unit root for plausible significance levels.

[22] Another potential explanation for the low power of the unit root test is that inflation – while stationary over the business cycle – has undergone important shifts over the sample. Thus, stationarity tests that allow for structural breaks may be more likely to reject the null of a unit root.

3.3 Results

Based on the bias-corrected estimates of the benchmark VAR, I bootstrap the bias-corrected sample distribution of $\hat{\rho}$ and $\hat{\Delta}$. Figure 3a displays the resulting density of the correlation coefficient. Although visual inspection suggests that most of the probability mass is concentrated about the high point estimate of $\hat{\rho} = 0.98$, a closer look at the distribution tells a different story. For example, the 90% confidence interval extends from 0.40 to 0.99, which means that there is a lot of uncertainty about the comovement between theoretical and actual inflation. The uncertainty about the relative volatility between observed and theoretical inflation is even more severe. As Figure 3b shows, the sample distribution of the variance ratio is very disperse, with a 90% confidence interval that spans from 0.01 to 1.57. Furthermore, roughly half of the probability mass is located between 0 and 0.6, implying that there is about a 50% chance that theoretical inflation is at least one and half times as volatile as observed inflation!

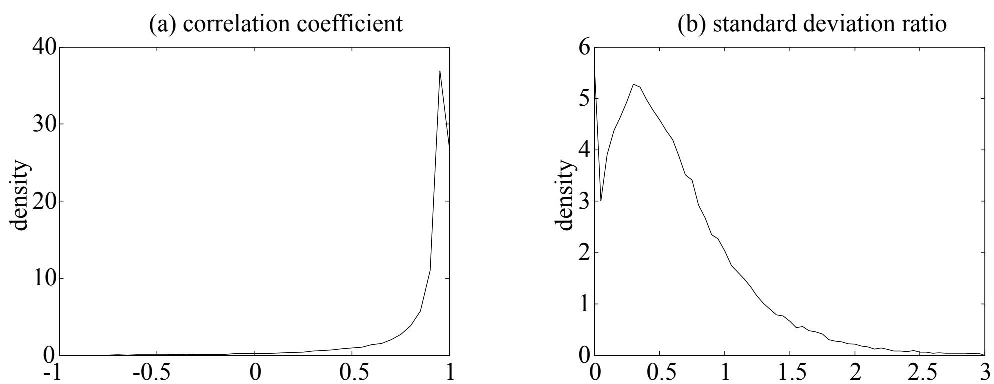

Figure 3: Uncertainty of fit due to imprecisely estimated VAR coefficients (benchmark case)

The large confidence intervals for both $\hat{\rho}$ and $\hat{\Delta}$ illustrate that once the sampling imprecision of the estimated VAR coefficients is taken into account, we do not know whether theoretical inflation tracks actual inflation almost perfectly or very poorly. Hence, even if we disregard the issue of whether the specification of the forecasting process is appropriate or whether the structural pricing equation is correctly calibrated, the present finding highlights that there is a great amount of uncertainty about the fit of theoretical inflation with the data.

4 Robustness to alternative forecasting processes

As discussed above, a maintained assumption behind the computation of theoretical inflation with the Campbell and Shiller method is that the Calvo model represents only an approximation of true inflation dynamics. This implies that the econometrician's information set ω_t is just a subset of the unknown vector of information Ω_t that markets use to forecast labor income share.[23] But then, there is no particular reason to believe that ω_t should consist of current and lagged values of labor income share and inflation only. In this section, I therefore assess the robustness of the fit between theoretical and observed inflation to alternative specifications of the forecasting process. First, I investigate whether using a VAR in labor income share and unit labor cost – as Sbordone proposes – alters the results. Second, I invoke statistical selection criteria to motivate both a reduction and an expansion of the information contained in ω_t. I then compute theoretical inflation both for the reduced and the expanded forecasting specification.

4.1 Replacing inflation with unit labor cost changes

In her paper, Sbordone computes theoretical inflation conditional on a bivariate VAR in labor income share and the first difference of unit labor cost $\triangle ulc_t$, where ulc_t is defined in logs as $s_t + p_t$. Note that $\triangle ulc_t$ can be rewritten as $\triangle ulc_t = \triangle s_t + \pi_t$, which means that her forecasting process also contains information about inflation.

Using the same sample of quarterly data as before, the Aikake information criteria selects an optimal lag length of four for the (bias-corrected) VAR in s_t and $\triangle ulc_t$. Thus, $\omega_t = [s_t \;\; \triangle ulc_t \;\; s_{t-1} \; \triangle ulc_{t-1} ... \; s_{t-4} \; \triangle ulc_{t-4}]'$. Table 2 reports the unrestricted (bias-corrected) coefficient estimates. Analogous to Gali and Gertler's benchmark VAR process in labor income share and inflation, s_{t-1} is the only significant determinant in the labor income share equation, while lagged inflation (now contained in $\triangle ulc$) remains a highly imprecise predictor of labor income share. Concurrently, s_{t-1} and $\triangle ulc_{t-1}$ are the only significant determinants in the unit labor cost equation. With regards to the statistical properties of the error terms, neither the null of homoscedastic errors nor the null of serially uncorrelated errors can be rejected.[24]

Figure 4 displays the path of theoretical inflation computed conditional on this (bias-corrected) VAR(4) in labor income share and unit labor cost changes (keeping $\beta\gamma_y = 1$ and $\phi = 0.035$ as before). Its fit with observed inflation is better than the fit of theoretical inflation computed from the (bias-corrected) benchmark VAR in labor income share and inflation (Figure 2). In particular, theoretical inflation overpredicts observed inflation by less, with the estimated variance ratio increasing to $\hat{\Delta} = 0.74$. In turn, the comovement between theoretical and observed inflation remains strong with an estimated correlation coefficient of $\hat{\rho} = 0.98$.

However, a similar picture than before emerges with respect the uncertainty about the correlation

[23] The explanation for this result is subtle. Under the null, $\pi_t = \phi \sum_{i=0}^{\infty} (\beta\gamma_y)^i E_t \psi_{t+i}$ holds exactly, which implies that inflation embodies all information that markets use to forecast real marginal cost. Hence, as long as $\pi_t \in \omega_t$, it must be that ω_t contains all *relevant* information to forecast future real marginal cost; i.e. $E[\sum_{i=0}^{\infty} (\beta\gamma_y)^i mc_{t+i} | \Omega_t] = E[\sum_{i=0}^{\infty} (\beta\gamma_y)^i mc_{t+i} | \omega_t]$. However, under the alternative that the model is not exactly true, π_t no longer embodies all relevant information about expected future real marginal cost and ω_t becomes a subset of Ω_t only.

[24] The BGLM p-value for the labor income share equation is 0.719, and 0.514 for the equation of unit labor cost changes. The p-values of the White test statistics are 0.104 and 0.315 for the labor income share equation and the unit labor cost changes equation, respectively.

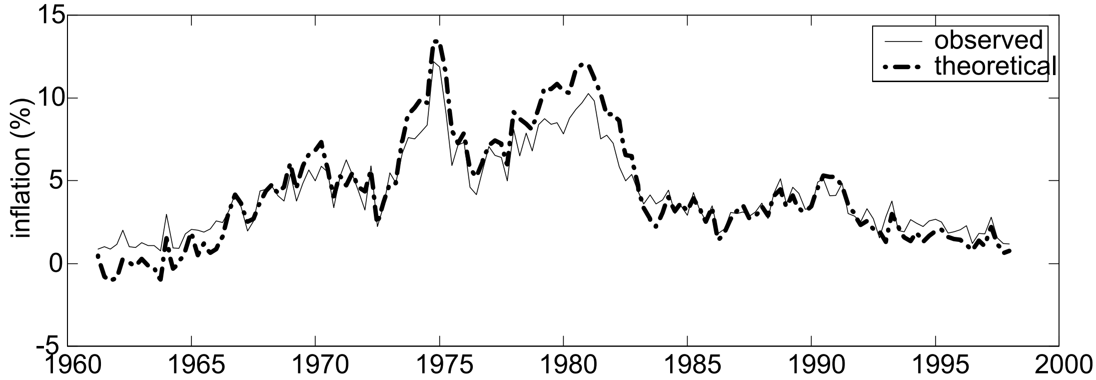

Figure 4: Fit of theoretical inflation computed from (bias-corrected) VAR in s and Δulc.

coefficient and the variance ratio. Figure 5 display the bootstrapped sampling distribution of $\hat{\rho}$ and $\hat{\Delta}$. While visual inspection suggests again that the distribution of $\hat{\rho}$ is concentrated about its large

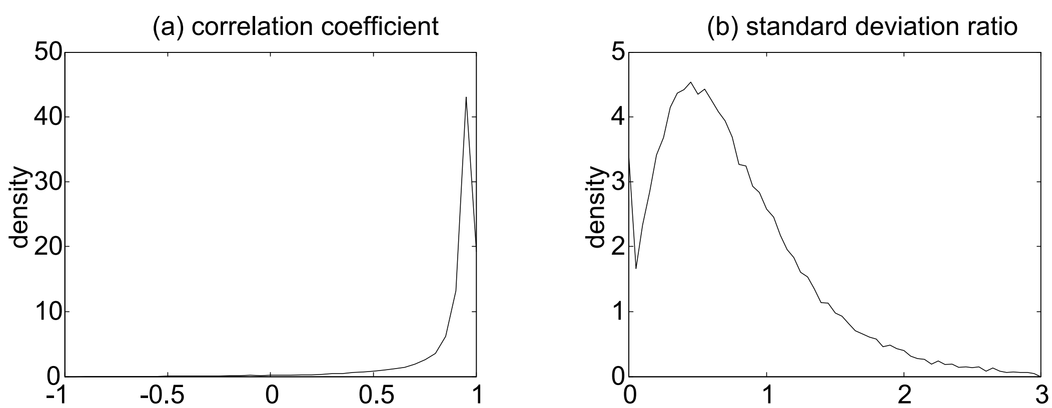

Figure 5: Uncertainty due to imprecisely estimated VAR coefficients (VAR in s and Δulc)

point estimate of $\hat{\rho} = 0.98$, the 90% confidence interval remains large, extending from 0.45 to 0.99. Furthermore, the sample distribution of the variance ratio becomes even more disperse than for the benchmark VAR case, with a 90% confidence interval ranging from 0.07 to 1.83. In sum, Sbordone's alternative VAR forecasting process in labor income share and changes of unit labor cost improves the estimated fit of theoretical inflation with observed inflation. Yet, it is incapable to reduce the substantial uncertainty that surrounds these estimates.

4.2 The crucial but uncertain role of inflation in the forecasting process

The discussion of the coefficient estimates of both the benchmark VAR(4) in Table 1 and the alternative VAR(4) in Table 2 reveal that information about past inflation is insignificant in the labor income share equation. To assess whether inflation is statistically useful in predicting future labor income share, I apply a Granger F-test. For a lag length of four, I find that the null hypothesis of π *does not Granger cause s* can only be rejected at a marginal confidence level of 0.22. Likewise, for the same lag length, the null of $\triangle ulc$ *does not Granger cause s* can only be rejected at a marginal confidence level of 0.18. An econometrician may thus conclude that inflation (respectively changes in unit labor cost) should not be part of the information set used to predict labor income share and instead specify a univariate autoregressive AR(4) forecasting process in labor income share alone.[25]

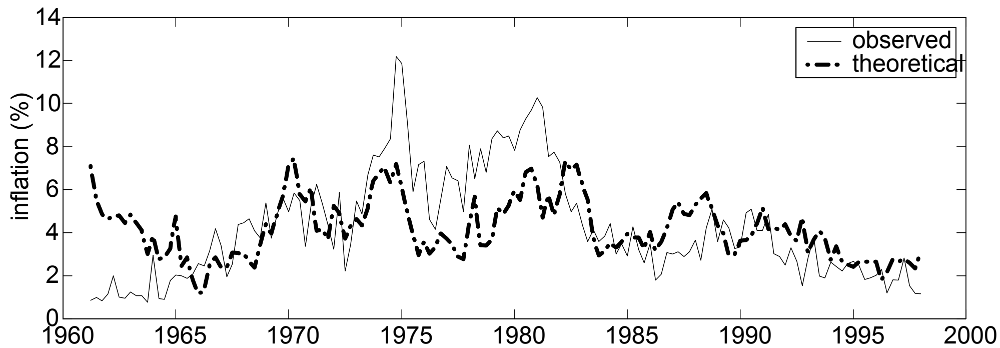

Figure 6: Lack of robustness with respect to univariate AR in s

Figure 6 displays the series of theoretical inflation conditional on this (bias-corrected) AR(4) forecasting process in labor income share (keeping $\beta\gamma_y = 1$ and $\phi = 0.035$ as before). Excluding inflation from the information set has dramatic effects. When labor income share is forecasted by lags of labor income share alone, the tight fit between theoretical inflation and observed inflation all but breaks down, with the estimated correlation coefficient dropping to $\hat{\rho} = 0.51$ and the estimated ratio of the standard deviations increasing to $\hat{\Delta} = 1.79$.[26]

The breakdown in fit highlights that theoretical inflation is highly sensitive to whether the forecasting process for labor income share includes information about inflation or not. The result also provides an intuitive explanation for the great degree of uncertainty about ρ and Δ uncovered above. If information about inflation is useful in predicting labor income share, then a forecasting process in labor income share and inflation is justified. In terms of point estimates, the Calvo model conditional on such a bivariate VAR process provides a good approximation of observed inflation dynamics. By

[25] The estimated dynamics of this AR(4) process in labor income share are:

$$s_t = \underset{(0.083)}{0.906}\, s_{t-1} - \underset{(0.113)}{0.032}\, s_{t-2} - \underset{(0.119)}{0.129}\, s_{t-3} - \underset{(0.082)}{0.045}\, s_{t-4} + \hat{e}_{s,t}$$

Note that all but the coefficient estimate on the first lag of labor income share are insignificant. However, I will use this AR(4) instead of an AR process with fewer lags because I want to isolate the impact on theoretical inflation of dropping inflation (respectively unit labor cost) from the forecasting process.

[26] Note that it is impossible to bootstrap the sample distribution of $\hat{\rho}$ and $\hat{\Delta}$ if theoretical inflation is computed from a forecasting process that does not contain explicit information about inflation. This is because bootstrapping the distribution of $\hat{\rho}$ and $\hat{\Delta}$ necessitates simulated inflation series, which can only be generated if the forecasting model (from which I bootstrap) contains inflation.

contrast, if inflation is an inappropriate proxy and therefore does not help forecasting labor income share, the evolution of labor income share should rather be described by an univariate process. In this case, the model does a poor job explaining observed inflation dynamics. Since the role of inflation in the forecasting of labor income share is so uncertain, theoretical inflation could fit observed inflation either well or badly, which is equivalent to saying that the confidence intervals of $\hat{\rho}$ and $\hat{\Delta}$ are very wide.

4.3 Expanding the forecasting process

Equivalently, one may ask whether including information other than labor income share and inflation improves the prediction of future labor income share. To examine this question, I consider augmenting the bivariate VAR of the benchmark example as well as the alternative VAR of Sbordone with the following set of variables: changes in employment ($\triangle n$), changes in real wages ($\triangle w$), the difference between output and consumption ($y - c$), the difference between output and investment ($y - i$), the first difference in the nominal stock of money ($\triangle M$), and the spread between long and short-term Treasury bill rates ($R^L - R^S$).

The choice of variables is motivated by the belief that they additional contain information about labor market conditions, economic activity in general, and the stance of monetary policy.[27] Moreover, block-exogeneity likelihood ratio tests, reported in Table 3a and 3b, reveal that the combination of these variables is significant at the 5% level in improving the prediction of labor income share for both the benchmark VAR and Sbordone's alternative VAR.[28] The expanded information set to forecast labor income share that I will consider is thus $z_t = [s_t,\, \pi_t,\, \triangle n_t,\, \triangle w_t,\, y_t - c_t,\, y_t - i_t,\, \triangle M_t,\, R_t^L - R_t^S]'$ and lags thereof, respectively $z_t = [s_t,\, \triangle ulc_t,\, \triangle n_t,\, \triangle w_t,\, y_t - c_t,\, y_t - i_t,\, \triangle M_t,\, R_t^L - R_t^S]'$ and lags thereof.[29]

Figure 7a displays the theoretical inflation series conditional on the (bias-corrected) extended VAR that is built from s_t and π_t (again keeping $\beta\gamma_y = 1$ and $\phi = 0.035$). Figure 7b displays the theoretical inflation series conditional on the (bias-corrected) extended VAR that is built from s_t

[27] The mixture of first differences and combinations of different variables is motivated by a host of empirical evidence about stochastic trends and cointegration characteristics (King, Plosser, Stock and Watson, 1991; or King and Watson, 1996). The additional series are identical with the ones used by Stock and Watson (1999). I thank Mark Watson for making them available on his website. See the appendix for a description of the data.

[28] A word of caution about this statistical selection procedure is in order, however. Block-exogeneity tests as well as the Granger F-tests are only concerned with (in-sample) forecasting performance one period ahead. The test results thus provide only an imperfect account of how much the addition of variables helps in improving the *discounted sum of forecasts* of labor income share (which is what determines theoretical inflation). To assess the impact of adding variables on multi-period forecasts, I also computed Theil's inequality coefficients (not reported in the tables). All the added variables improve the forecasting of labor income share for 4 and 8 quarters ahead. However, to the author's knowledge, no statistical test criteria exists for Theil's inequality coefficient.

[29] The AIC selects an optimal lag length of one for the extended VAR based on s_t and π_t. For the extended VAR based on s_t and Δulc_t, the optimal lag length is two. For the sake of improving the properties of the residuals (absence of both serial correlation and heteroscedasticity), I choose a lag length of two for both specifications; i.e. $\omega_t = [z_t,\, z_{t-1}]'$.

and Δulc_t. In both cases, theoretical inflation dramatically fails at replicating the large swings of

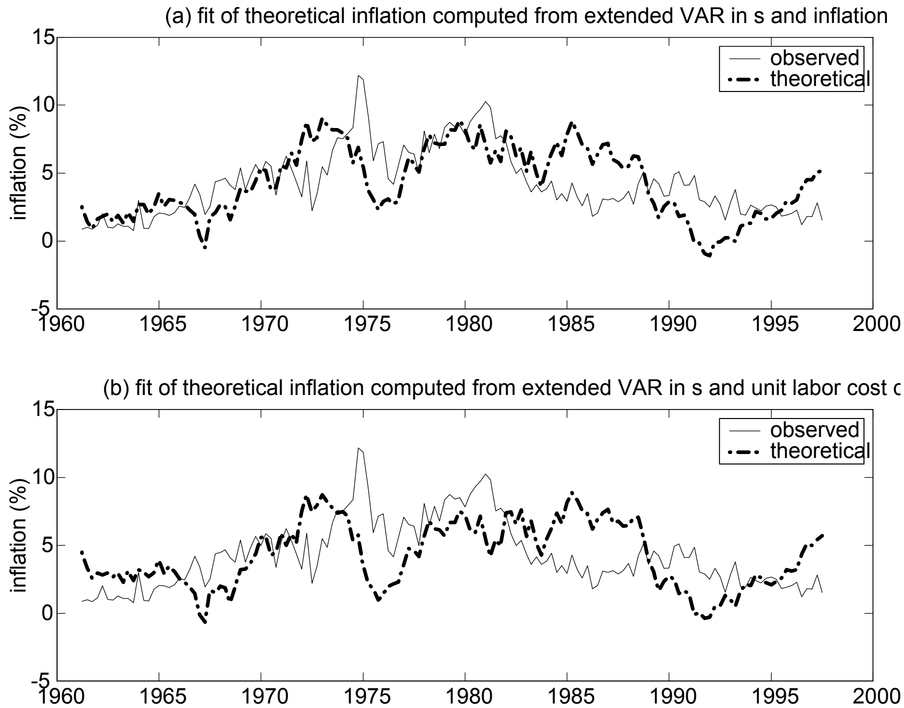

Figure 7: Lack of robustness with respect to extended VAR process

observed inflation in the 1970s and cannot account for the dynamics of observed inflation after 1983. As a result, the estimated correlation coefficient drops to $\hat{\rho} = 0.55$ in Figure 7a, and to $\hat{\rho} = 0.32$ in Figure 7b. Concurrently, the relative volatility of theoretical inflation improves in both cases, with an estimated standard deviation ratio of $\hat{\Delta} = 0.95$ in Figure 7a and $\hat{\Delta} = 1.03$ in Figure 7b. However, this improvement is of no great interest given the weak correlation coefficients.

Adding more variables in the VAR forecasting process also leads to a substantial increase in the dispersion of the correlation coefficient. For theoretical inflation computed from the extended VAR that is built from s_t and π_t, the 90% confidence interval of the correlation coefficient extends from -0.54 to 0.84. For theoretical inflation computed from the extended VAR that is built from s_t and Δulc_t, the same 90% confidence interval spans from -0.42 to 0.69. Concurrently, the dispersion of the variance ratio remains essentially unaffected: the 90% confidence interval ranges from 0.002 to 1.33 for the case built from s_t and π_t, respectively 0.01 to 1.52 for the case built from s_t and Δulc_t.

To some extent, the decrease in goodness-of-fit when inflation is removed from the forecasting process, respectively when more variables are added is intuitive. In the benchmark example with the bivariate VAR forecasting process, current and lagged inflation terms are a very important (yet uncertain) component in the equation for theoretical inflation (6).[30] Since inflation is a highly persistent process in the data, it should therefore not come as a great surprise that theoretical inflation is highly

[30]For example, in the equation for theoretical inflation of the (bias-corrected) benchmark case, the sum of coefficients

correlated with observed inflation. But then, removing current and lagged inflation from the formula for theoretical inflation means that we decrease the comovement between π_t' and π_t. By the same token, adding more variables to the formula for theoretical inflation will "water down" the importance of the current and lagged inflation terms on theoretical inflation. Since all of these additional variables are less than perfectly correlated with inflation (unlike inflation with itself), the fit of π_t' with π_t is likely to suffer.

The findings in this section highlight that the performance of the Calvo model in terms of tracking observed inflation depends crucially on the specification of the forecasting process for labor income share. Hence, while the Campbell and Shiller method of evaluating present-value models conditional on a reduced-form VAR process saves us from taking a stand about the rest of the economy, it still requires us to make empirical assumptions that heavily affect the fit of the model with observed data.

5 Robustness with respect to the marginal cost coefficient

A final set of determinants that affect theoretical inflation are the coefficients $\beta\gamma_y$ and ϕ of the Calvo pricing equation. The value of $\beta\gamma_y = 1$ appears appropriate because β should equal the reciprocal of the long-run level of the real interest rate, which is itself linked inversely to the long-run growth rate of output γ_y. However, considerable uncertainty surrounds the value of ϕ – the coefficient on real marginal cost in the pricing equation. Expression (3) shows that ϕ depends on (i) the degree of price rigidity κ; (ii) the elasticity of substitution μ; and (iii) the elasticity of the firm's real marginal cost with respect to its output η.

As for the first issue, Taylor (1998) notes in his survey that micro studies uncovered a great deal of heterogeneity about price setting across different industries. Hence, it is difficult to pin down the average frequency of price adjustment. Secondly, the calibration of μ determines the average markup that firms charge, which we can try to match to the average markup in the data. Finally, η depends on assumptions about factor markets. The traditional and most commonly used version of the Calvo model assumes that all factors including capital are traded in perfectly competitive markets and can be reallocated across firms instantaneously at no cost (see Yun, 1996; or Gali and Gertler, 1999). Under the additional assumption that firms produce with constant returns to scale technology, real marginal cost is independent of the level of output; i.e. $\eta = 0$. Consequently, the definition of ϕ becomes

$$\phi = \frac{(1-\kappa)(1-\kappa\beta\gamma_y)}{\kappa} > 0. \tag{9}$$

Alternatively, if we assume – as proposed by Sbordone (2002) – that capital stocks are predetermined for every firm, it is possible show that $\eta = \alpha/(1-\alpha)$, where α is the share of capital in the production

on current and lagged inflation terms is 1.76, while the sum of coefficients of current and lagged labor income share terms is only 0.03. Hence, theoretical inflation is mainly driven by inflation data, and not labor income share data.

function.[31] In this case

$$\phi = \left(\frac{(1-\kappa)(1-\kappa\beta\gamma_y)}{\kappa}\right)\left(\frac{1-\alpha}{1-\alpha+\alpha\mu}\right) > 0. \tag{10}$$

Given the uncertainty about both the average degree price rigidity and the definition of ϕ, it is important to evaluate how robust the fit of theoretical inflation with observed inflation is to changes in ϕ. Reconsider the definition of the correlation coefficient and the standard deviations ratio, $\rho(\pi_t, \pi_t') = E[\pi_t \pi_t']/\sqrt{E[\pi_t^2]E[\pi_t'^2]}$ and $\Delta(\pi_t, \pi_t') = \sqrt{E[\pi_t^2]/E[\pi_t'^2]}$. Using the definition for theoretical inflation (6) and writing $\pi_t \equiv h_\pi \omega_t$ (where h_π is a selection vector for inflation similar to h_ψ for real marginal cost), these two expressions become

$$\rho(\pi_t, \pi_t') = \frac{\phi h_{mc} A \Sigma_\omega h_\pi'}{\phi\sqrt{h_{mc} A \Sigma_\omega A' h_{mc}' h_\pi \Sigma_\omega h_\pi'}} = \frac{h_{mc} A \Sigma_\omega h_\pi'}{\sqrt{h_{mc} A \Sigma_\omega A' h_{mc}' h_\pi \Sigma_\omega h_\pi'}} \tag{11}$$

and

$$\Delta(\pi_t, \pi_t') = \frac{1}{\phi}\sqrt{\frac{h_\pi \Sigma_\omega h_\pi'}{h_\psi A \Sigma_\omega A' h_\psi'}}, \tag{12}$$

where $A = [I - \beta\gamma_y M]^{-1}$ and $\Sigma_\omega = E[\omega_t \omega_t']$. Two aspects are apparent from these formulae. First, ϕ cancels out of the definition of ρ; i.e. the correlation of theoretical inflation with observed inflation is independent of the degree of price fixity in the economy and the assumed market and production structure in the Calvo model. Second, the standard deviation ratio Δ depends inversely on ϕ. From the above definitions of ϕ, we therefore know that the smaller the degree of price fixity in the economy (i.e. the smaller κ), the larger ϕ and thus the larger the volatility of theoretical inflation relative to the volatility of observed inflation (i.e. the smaller Δ).

Table 4 reports the quantitative impact of changing ϕ for both the scenario where theoretical inflation is computed conditional on the bivariate VAR of the benchmark example and the scenario where theoretical inflation is computed conditional on the expanded VAR that includes s_t and π_t (in both cases, I leave $\beta\gamma_y = 1$).[32] The first case considered is the reference used so far, $\phi = 0.035$. For the traditional definition of ϕ in (9), this value implies an average degree of price rigidity of 5.87 quarters, which is too large compared to evidence from micro-studies.[33] Concurrently, for the definition proposed by Sbordone in (10), the value of $\phi = 0.035$ together with a calibration of $\alpha = 0.40$ and $\mu = 10$ implies a more plausible average degree of price fixity of 2.5 quarters.[34] The second case considered is $\phi = 0.083$, which implies an average price rigidity of 4 quarters for the traditional definition in (9), respectively 1.8 quarters for Sbordone's definition when $\alpha = 0.40$ and $\mu = 10$.

[31] See the appendix for the mathematics behind the different values of η.

[32] For the sake of conciseness, I do not report the results for the VAR that involve Δulc_t. The results are very similar.

[33] Taylor concludes that *"...it would be inaccurate and misleading to build a model in which the average frequency of price* [or wage] *adjustment is longer than one year"* [page 23].

[34] These values for α and μ are quite standard (see for example Basu, 1996 for the value of $\mu = 10$, which implies a steady state markup of price over marginal cost of 11% in the Calvo model). Lowering α or μ would increase the implied price fixity.

The results in the table illustrate that the standard deviations ratio is highly sensitive even to this relatively small change in ϕ. The point estimate of Δ drops from 0.57 to 0.24 for the bivariate VAR, respectively from 0.95 to 0.40 for the extended VAR. In other words, the volatility of theoretical inflation increases by more than 100% relative to the volatility of observed inflation with the consequence that theoretical inflation for both the bivariate VAR of the benchmark and the extended VAR fails to track actual inflation.[35] In addition, while the bootstrapped 90% confidence intervals for these new values of Δ remain considerable, they do by far not include the theoretical value $\Delta = 1$ of the null that the model is correct.

The sensitivity of theoretical inflation with respect to ϕ is disconcerting given that we have no precise knowledge about neither the degree of price fixity in the economy nor the firm-specific elasticity of marginal cost. Furthermore, the dramatic increase of the relative volatility of theoretical inflation when $\phi = 0.083$ is bad news for the traditional version of the Calvo model with instantaneous capital allocation. This value of ϕ still implies an average price fixity of 4 quarters, which is the upper bound of admissible price rigidity according to Taylor. Hence, once we calibrate κ to a reasonable (but still high) degree of price rigidity, the traditional Calvo model implies a theoretical inflation series that is much too volatile. By contrast, if we adopt the alternative definition of ϕ proposed by Sbordone, this conclusion is not necessary since the benchmark calibration of $\phi = 0.035$ implies a degree of price fixity of 2.5 quarters, which is well within reasonable bounds.

6 Conclusion

The results of this paper illustrate that the "good fit" of theoretical inflation with observed inflation reported in GGS is surrounded by a great deal of uncertainty. Hence, we cannot conclude whether the Calvo model explains U.S. inflation dynamics very poorly or very well. Instead, the results suggest that important research on the determinants of firms' cost and market structure is necessary before we can assess the empirical relevance of NK pricing models that link inflation to expected future real marginal cost terms.

On a more general level, this paper also highlights that while Campbell and Shiller's method of evaluating present-value models has the advantage that we do not need to take a stand about the structure of the rest of the economy, it does not save us from making empirical assumptions that may severely affect the fit of the model with observed data. In particular, theoretical series computed with the Campbell-Shiller method appear to be highly sensitive to the set of information used to approximate expectations. Furthermore, a great variety of dynamic stochastic theories imply present-value relationships (i.e. Euler equations containing forward-looking expectations terms) that resemble the form of the Calvo pricing equation.[36] Many of these theories have been taken to the data using a VAR forecasting approach similar to the one discussed here. To the author's knowledge, however, very

[35] See the extended version of the paper for a graph that further illustrates this lack of fit.

[36] Prominent examples include the permanent income hypothesis of consumption, the present-value model of stock prices or the expectations theories of interest rates and exchange rates.

few (and incomplete) attempts have bee made to systematically quantify the uncertainty about the theoretical data series that the respective models imply. The analysis presented in this paper offers a starting point to do so.

References

[1] Basu, S., 1996, Procyclical Productivity: Increasing Returns or Cyclical Utilization? Quarterly Journal of Economics 111, 719-751.

[2] Berkowitz, J. and L. Kilian, 2000, Recent Developments in Bootstrapping Time Series, Econometric Reviews 19.

[3] Bils, M., 1987, The Cyclical Behavior of Marginal Cost and Prices, American Econonomic Review 77, 838-857.

[4] Blanchard, O.J. and N. Kiyotaki, 1987, Monopolistic Competition and the Effect of Aggregate Demand, American Economic Review, 77.

[5] Calvo, G.A., 1983, Staggered Prices in an Utility-Maximizing Framework, Journal of Monetary Economics 12, 383-398.

[6] Campbell, J.Y. and R.J. Shiller, 1987, Cointegration and Tests of Present-Value Models, Journal of Political Economy 95, 1062-1088.

[7] Estrella, A. and J. C. Fuhrer, 1998, Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational Expectations Models, American Economic Review 92, 1013-1028..

[8] Fuhrer, J.C. and G. Moore, 1995, Inflation Persistence, Quarterly Journal of Economics 110, 127-159.

[9] Gali, J. and M. Gertler, 1999, Inflation Dynamics: A Structural Econometric Analysis, Journal of Monetary Economics 44, 195-222.

[10] Gali, J., M. Gertler and D. J. Lopez-Salido, 2001, European Inflation Dynamics, European Economic Review 45, 1237-1270.

[11] Godfrey, L.G., 1988, Specification Tests in Econometrics (Cambridge University Press, Cambridge).

[12] Hamilton, J. D., 1994, Time Series Analysis, (Princeton University Press, Princeton).

[13] Kilian, L., 1998a, Small-Sample Confidence Intervals for Impulse Response Functions, Review of Economics and Statistics 80, 218-230.

[14] Kilian, L., 1998b, Accounting for Lag Order Uncertainty in Autoregressions: The Endogenous Lag Order Bootstrap Algorithm, Journal of Time Series Analysis 19, 531-548.

[15] King, R.G., C.I. Plosser, J.H. Stock and M.W. Watson, 1991, Stochastic Trends and Economic Fluctuations, American Economic Review 81, 819-840.

[16] King, R.G. and M. W. Watson, 1994, The Post-War U.S. Phillips Curve: A Revisionist Econometric History, Carnegie-Rochester Series of Public Policy, 41.

[17] King, R.G. and M.W. Watson, 1996, Money, Prices, Interest Rates and the Business Cycle, Review of Economics and Statistics 78, 35-53.

[18] King, R. G. and A. L. Wolman, 1996, Inflation Targeting in a St. Louis Model of the 21st Century, Federal Reserve Bank of St. Louis Review, 78.

[19] Kurmann, A., 2003, Maximum Likelihood Estimation of Dynamic Stochastic Theories with an Application to New Keynesian Pricing, Université du Québec à Montréal, unpublished.

[20] McCallum, B.T. and E. Nelson, 1998, Performance of Operational Policy Rules in an Estimated Semiclassical Structural Model, in: J.B. Taylor, ed., Monetary Policy Rules (Chicago University Press, Chicago).

[21] Pope, A. L., 1990, Biases of Estimators in Multivariate Autoregressions, Journal of Time Series Analysis, 11.

[22] Sbordone, A.M., 2002, Prices and Unit Labor Costs: Testing Models of Pricing, Journal of Monetary Economics 49, 265-292.

[23] Stine, R.A., 1987, Estimating Properties of Autoregressive Forecasts, Journal of the American Statistical Association 82, 1072-1078.

[24] Stock, J.H. and M.W. Watson, 1999, Business Cycle Fluctuations in U.S. Macroeconomic Time Series, in: J.B. Taylor and M. Woodford, eds., Handbook of Macroeconomics, Vol. 1A (North-Holland, Amsterdam).

[25] Rotemberg, J. J. and M. Woodford, 1997, An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy, in National Bureau of Economic Research Macroeconomics Annual, (MIT Press: Cambridge).

[26] Taylor, J.B., 1998, Staggered Wage and Prices in Macroeconomics, in: J.B. Taylor and M. Woodford, eds., Handbook of Macroeconomics, Vol. 1B (North-Holland, Amsterdam).

[27] White, Halbert, 1980, A Heteroskedasticity-Consistent Covariance Matrix and a Direct Test for Heteroskedasticity, Econometrica, 48.

[28] Woodford, M., 2003, Interest and Prices (Princeton University Press, Princeton).

[29] Yun, T., 1996, Nominal Price Rigidity, Money Supply Enodogeneity, and Business Cycles, Journal of Monetary Economics 37, 345-370.

A Appendix

A.1 Data

All data used in this paper are quarterly U.S. time series from DRI Basic Economics (formerly Citibase). The exact definition of each variable is provided in the following table (DRI mnemonics in the second column):

Variable	Definition	Description	Source
P	ln(gdp)-ln(gdpq)	implicit GDP deflator (total)	GG
s	ln(lblcpu)-ln(lbgdpu)	labor income share (nonfarm business)	GG
y	ln(gdpuq)	real GDP (nonfarm business)	GG
c	ln(gcq)	real personal consumption expenditures (total)	SW
i	ln(gifq)	real private fixed investment (total)	SW
n	ln(lbmnu)	hours of all persons (nonfarm business)	GG
w	ln(lbcpu)-ln(lbgdpu)	real wage rate (nonfarm business)	GG
M	ln(fm2)	nominal money stock M2	GG
R^S	fygm3	3 months T-bill rate, in annual %	GG
R^L	fygt10	10 year T-bill rate, in annual %	SW

where GG indicates that the datasource is Gali and Gertler (1999), and where SW indicates that the datasource is Stock and Watson (1999). The sample period for the data by GG is 1947:1-1998:2 and the sample period for the data by SW is 1947:1-1997:2. To keep my results comparable to the results by GGS, I only consider the sample 1961:1-1997:4 for the bivariate VAR forecasting process, and 1961:1-1997:2 for the expanded VAR (not including lags needed for the VAR and to compute inflation from the GDP deflator).[37]

A.2 Derivation of Calvo pricing equation

A.2.1 Underlying monopolistic competition setup

The starting point of the Calvo pricing model is a monopolistic competition setup, based on that in Blanchard and Kiyotaki (1987), in which households value differentiated products and firms have the power to set prices. Assume that there is a continuum of goods ($y_t(z), z \in [0,1]$), which can be aggregated into an index y_t according to

$$y_t = \left[\int_0^1 y_t(z)^{(\mu-1)/\mu} dz \right]^{\mu/(\mu-1)}. \tag{13}$$

[37] Note that I chose to use the data from GG and SW instead of considering more recent data from DRI Basic Economics because important revisions have occured for the labor income share series between 1998 (the approximate collection date by GG and SW) and 2001.

Cost minimization on the part of consumers who value the consumption aggregate, but not its separate components, implies that demand for the zth good takes the form

$$y_t(z) = (P_t(z)/P_t)^{-\mu} y_t, \tag{14}$$

where P_t is an index of the cost of buying a unit of c_t

$$P_t \equiv \left[\int_0^1 P_t(z)^{1-\mu} dz \right]^{\frac{1}{1-\mu}}. \tag{15}$$

The parameter $\mu > 1$ is the (constant) elasticity of substitution between the differentiated goods. The closer μ is to unity, the smaller the degree of competition.

A.2.2 Other aspects of household's behavior

Households directly own all factors of production and rent these to firms. Households also own a diversified portfolio of claims to the profits earned by the monopolistically competitive firms. The utility value of these profits depends on the real shadow price of a unit of the consumption aggregate, which is denoted as λ_t.

A.2.3 Firms' revenues and costs

At any level of output, the firm seeks to minimize its cost. These costs include its payments to labor (possibly of a specialized sort within its local market), its cost of capital and other factor payments including those for energy and materials. Define the nominal cost function of the firm as $\Phi(y_t(z), s_t)$, where s_t is a vector of state variables that can affect these costs but do not depend on z. We write the comparable real cost function as $\phi(y_t(z), s_t) = \Phi(y_t(z), s_t)/P(s_t)$. It is convenient to have an expression for nominal marginal cost

$$\Psi_t(z) = \frac{\partial \Phi(y_t(z), s_t)}{\partial y_t(z)}$$

and for real marginal cost

$$\psi_t(z) = \frac{\partial \phi(y_t(z), s_t)}{\partial y_t(z)}$$

Firm z's real profit is thus

$$q_t(z) = \left[\frac{P_t(z)}{P_t} y_t(z) - \phi(y_t(z), s_t) \right].$$

A.2.4 Optimal pricing when prices are flexible

Consider first the good z monopolist's profit maximization problem under the assumption that the nominal price can be set each period after learning the state of demand and cost. Supposing that the monopolist sets the nominal price to maximize its real profits $q(\cdot)$, then its efficiency condition can be written as

$$\begin{aligned}
\frac{\partial q(y_t(z), s_t)}{\partial P_t(z)} &= \left[\frac{1}{P_t} y_t(z) + \frac{P_t(z)}{P_t}\frac{\partial y_t(z)}{\partial P_t(z)}\right] - \left[\frac{\partial \phi(y_t(z), s_t)}{\partial P_t(z)}\right] = 0 \\
&= \left[\frac{1}{P_t} y_t(z) + \frac{P_t(z)}{P_t}\frac{\partial y_t(z)}{\partial P_t(z)}\right] - \left[\frac{\partial \phi(y_t(z), s_t)}{\partial y_t(z)}\frac{\partial y_t(z)}{\partial P_t(z)}\right] = 0 \\
&= \left[\frac{1}{P_t} y_t(z) + \frac{P_t(z)}{P_t}\frac{\partial y_t(z)}{\partial P_t(z)}\right] - \left[\psi_t(z)\frac{\partial y_t(z)}{\partial P_t(z)}\right] = 0
\end{aligned}$$

This expression equates a nontraditional notion of marginal revenue $\left[\frac{1}{P_t} y_t(z) + \frac{P_t(z)}{P_t}\frac{\partial y_t(z)}{\partial P_t(z)}\right]$, that from altering the product price, with the cost consequences of doing so, $\left[\psi_t(z)\frac{\partial y_t(z)}{\partial P_t(z)}\right]$. Plugging in the demand function $y_t(z) = P_t(z)^{-\mu} P_t^{\mu} y_t$ and its derivative $\frac{\partial y_t(z)}{\partial P_t(z)} = -\mu P_t(z)^{-\mu-1} P_t^{\mu} y_t$, we find that:

$$(1-\mu)\frac{1}{P_t}\left(\frac{P_t(z)}{P_t}\right)^{-\mu} y_t + \mu \frac{1}{P_t(z)}\left(\frac{P_t(z)}{P_t}\right)^{-\mu} y_t \psi_t(z) = 0$$

Solving this equation, we obtain a standard result for a monopolist facing a constant elasticity demand curve, which is that the firm sets its relative price as a *fixed* markup over its real marginal cost

$$\frac{P_t(z)}{P_t} = \frac{\mu}{\mu-1}\psi_t(z)$$

where $\frac{\mu}{\mu-1}$ indicates the size of that markup. Equivalently, the firm sets its nominal product price as a fixed markup over its nominal marginal cost, $P_t(z) = \frac{\mu}{\mu-1}[P_t \psi_t(z)] = \frac{\mu}{\mu-1}\Psi_t(z)$.

A.2.5 Some equilibrium analysis

We are interested in economies in which all firms have the same marginal cost functions, which we impose as a symmetry condition. Since we also assumed that all firms faced the same demand functions, then all firms will have the same output level and choose the same optimal price P^*. In turn, this price will also be equal to the general price level,

$$P_t = \left[\int_0^1 P_t(z)^{1-\mu} dz\right]^{\frac{1}{1-\mu}} = \left[\int_0^1 P_t^{*1-\mu} dz\right]^{\frac{1}{1-\mu}} = P_t^*$$

In the flexible price version of the model, the equilibrium relative price therefore always equals unity.

A.2.6 Optimal pricing when prices are sticky

In Calvo's framework, it is simply assumed that in each period t, every firm z faces a probability $1-\kappa$ that it adjusts its price, disregarding of the number of periods the price has been kept unchanged. Hence, the fraction ω_j of firms that charge the same price in period t than in period $t-j$ is

$$\omega_j = \kappa^j (1-\kappa), \text{ for } j = 0, 1, 2, ..., J-1.$$

In addition, it is straightforward to show that the average number of periods for which a firm's price remains fixed equals $\frac{1}{1-\kappa}$.

Adjusting firms set their new optimal price such as to maximize the discounted sum of expected profits, where the discount factor takes into account the probability that the firm may not readjust its price for several periods to come. The optimality condition is

$$0 = \sum_{j=0}^{\infty} (\beta\kappa)^j E_t \left(\frac{\lambda_{t+j}}{\lambda_t} \frac{\partial q_{t+j,t}}{\partial P_t^*} \right)$$

If this condition were not to hold, then a slight increase or a slight decrease in the price could raise the price-adjusting firm's value.

To turn this optimality condition into a more explicit expression, take the profit derivative of the flexible price case above and write it for period $t+j$

$$\frac{\partial q_{t+j,t}}{\partial P_t^*} = (1-\mu)\frac{1}{P_{t+j}} \left(\frac{P_t^*}{P_{t+j}} \right)^{-\mu} y_{t+j} + \mu \frac{1}{P_t^*} \left(\frac{P_t^*}{P_{t+j}} \right)^{-\mu} y_{t+j} \psi_{t+j,t}$$

where $\psi_{t+j,t}$ is the real marginal cost at period $t+j$ of a firm that charges P_t^*. Note that – depending on the maintained assumptions about technology and input factor markets – this real marginal cost may be different (the value, not the function) from the real marginal cost of firms that adjusted last in a period other than t. This is because marginal cost depends on the amount of goods produced, which in turn depends on the relative price of the firm.[38] Plugging these profit derivatives into the efficiency condition and rearranging terms, we obtain the optimal pricing formula:

$$P_t^* = \frac{\mu}{\mu-1} \frac{\sum_{j=0}^{\infty} (\beta\kappa)^j E_t \left[\lambda_{t+j} {P_{t+j}}^{\mu} y_{t+j} \psi_{t+j,t} \right]}{\sum_{j=0}^{\infty} (\beta\kappa)^j E_t \left[\lambda_{t+j} {P_{t+j}}^{\mu-1} y_{t+j} \right]}. \tag{16}$$

This equation is identical to the one reported, for example, in Yun (1996). Furthermore, the aggregate price level can be expressed as:

$$P_t^{1-\mu} = (1-\kappa)(P_t^*)^{1-\mu} + \kappa P_{t-1}^{1-\mu}. \tag{17}$$

A.2.7 Loglinearizations

Consider an inflationary steady state where nominal prices and wages are growing at gross rate g. Then, the steady state relative price of adjusting firms is given by

$$\frac{P}{P^*} = \left[\sum_{j=0}^{\infty} \omega_j g^{(\mu-1)j} \right]^{\frac{1}{1-\mu}}.$$

[38] An important assumption in this line of argument is that firms are willing to satisfy any demand at the price they charge (even if this involved selling below marginal cost). This assumption is admittedly unrealistic. However, taking into account such discontinuities would very much complicate the analysis of the model.

To loglinearize P_t, it is convenient to rewrite the aggregate price level expression as

$$1 = \left[\sum_{j=0}^{\infty} \omega_j (g^{(\mu-1)j})(\frac{g^j P^*_{t-j}}{P_t})^{1-\mu}\right] = \left[\sum_{j=0}^{\infty} \omega_j (g^{(\mu-1)j}) x^{1-\mu}_{t-j,t}\right],$$

where $x_{t-j,t} = (\frac{g^j P^*_{t-j}}{P_t})$. The idea behind rewriting the expression in this way is that we are left only with variables that are constant in a steady state where the price level is growing at gross rate g.

Now, take logs on each side of this expression, which yields

$$0 = \log \left[\sum_{j=0}^{\infty} \omega_j (g^{\mu-1)j}) x^{1-\mu}_{t-j,t}\right]$$

and perform a first-order Taylor approximation around the steady state. We obtain

$$0 = \sum_{j=0}^{\infty} \omega_j (g^{\mu-1)j}) \hat{x}_{t-j,t},$$

where $\hat{x}_{t-j,t}$ represents a percentage deviation from its steady state (i.e. $\hat{x}_{t-j,t} = \log x_{t-j,t} - \log x \cong (x_{t-j,t} - x)/x$). Furthermore, it will turn out to simplify matters greatly if we assume that the steady state of inflation is zero; i.e. $g = 1$ (thus, prices are stationary variables). Then, we can write $\hat{x}_{t-j,t} = \hat{P}^*_{t-j} - \hat{P}_t$ and the log-linearized aggregate price formula becomes

$$\hat{P}_t = (1-\kappa) \sum_{j=0}^{\infty} \kappa^j \hat{P}^*_{t-j}.$$

Equivalently, we can invoke some lag operator techniques to rewrite this expression as

$$\begin{aligned}
\hat{P}_t &= (1-\kappa) \sum_{j=0}^{\infty} \kappa^j L^j \hat{P}^*_t \\
\hat{P}_t &= \frac{(1-\kappa)}{(1-\kappa L)} \hat{P}^*_t \\
(1-\kappa L)\hat{P}_t &= (1-\kappa)\hat{P}^*_t \\
\hat{P}_t &= (1-\kappa)\hat{P}^*_t + \kappa \hat{P}_{t-1} \qquad (18)
\end{aligned}$$

Finally, defining the percentage deviations in the *relative* optimal price from its steady state as $\hat{p}^*_t = \hat{P}^*_t - \hat{P}_t$, we have

$$\hat{p}^*_t = \frac{\kappa}{1-\kappa} \hat{\pi}_t$$

where $\hat{\pi}_t = \hat{P}_t - \hat{P}_{t-1}$ represents percentage deviations in the rate of inflation from its steady state (which we assumed to equal zero above).

Now, consider the optimal price setting equation (16). To loglinearize it, rewrite it as

$$\frac{P^*_t}{P_t} = \frac{\mu}{\mu-1} \frac{\sum_{j=0}^{\infty} (\beta\kappa)^j E_t \left[\lambda_{t+j} \left(\frac{P_{t+j}}{P_t}\right)^{\mu} y_{t+j} \psi_{t+j,t}\right]}{\sum_{j=0}^{\infty} (\beta\kappa)^j E_t \left[\lambda_{t+j} \left(\frac{P_{t+j}}{P_t}\right)^{\mu-1} y_{t+j}\right]}$$

$$
\begin{aligned}
&= \frac{\mu}{\mu-1}\frac{\sum_{j=0}^{\infty}(\beta\kappa g^{\mu})^j E_t\left[\lambda_{t+j}\left(\frac{P_{t+j}}{g^j P_t}\right)^{\mu} y_{t+j}\psi_{t+j,t}\right]}{\sum_{j=0}^{\infty}(\beta\kappa g^{\mu-1})^j E_t\left[\lambda_{t+j}\left(\frac{P_{t+j}}{g^j P_t}\right)^{\mu-1} y_{t+j}\right]} \\
&= \frac{\sum_{j=0}^{\infty}(\beta\kappa g^{\mu})^j E_t\left[\lambda_{t+j}x_{t+j,t}^{\mu}c_{t+j}\psi_{t+j,t}\right]}{\sum_{j=0}^{\infty}(\beta\kappa g^{\mu-1})^j E_t\left[\lambda_{t+j}x_{t+j,t}^{\mu}c_{t+j}\right]} = \frac{\mu}{\mu-1}\frac{N_t}{D_t},
\end{aligned}
$$

Taking logs, we obtain

$$
\log P_t^* - \log P_t = \log\left(\frac{\mu}{\mu-1}\right) + \log N_t - \log D_t,
$$

or in percentage deviations from their respective steady states

$$
\hat{P}_t^{*\prime} - \hat{P}_t = \hat{p}_t^* = \hat{N}_t - \hat{D}_t
$$

where

$$
\begin{aligned}
\hat{N}_t &= \frac{1}{N_{ss}}\sum_{j=0}^{\infty}(\beta\gamma_y\kappa g^{\mu})^j \lambda_{ss}x_{ss}^{\mu}y_{ss}\psi_{ss}E_t\left[\hat{\lambda}_{t+j} + \mu\hat{x}_{t+j,t} + \hat{y}_{t+j} + \hat{\psi}_{t+j,t}\right] \\
\hat{D}_t &= \frac{1}{D_{ss}}\sum_{j=0}^{\infty}(\beta\gamma_y\kappa g^{\mu-1})^j \lambda_{ss}x_{ss}^{\mu-1}y_{ss}E_t\left[\hat{\lambda}_{t+j} + (\mu-1)\hat{x}_{t+j,t} + \hat{y}_{t+j}\right]
\end{aligned}
$$

and

$$
\begin{aligned}
N_{ss} &= \frac{\lambda_{ss}x_{ss}^{\mu}y_{ss}\psi_{ss}}{1-\beta\gamma_y\kappa g^{\mu}} \\
D_{ss} &= \frac{\lambda_{ss}x_{ss}^{\mu-1}y_{ss}}{1-\beta\gamma_y\kappa g^{\mu-1}}
\end{aligned}
$$

Finally, we assume again that $g = 1$. Many of the terms cancel out and we obtain the following loglinearized optimal pricing equation

$$
\begin{aligned}
\hat{p}_t^* &= (1-\beta\gamma_y\kappa)\sum_{j=0}^{\infty}(\beta\gamma_y\kappa)^j E_t\left[\hat{x}_{t+j,t} + \hat{\psi}_{t+j,t}\right] \\
&= (1-\beta\gamma_y\kappa)\sum_{j=0}^{\infty}(\beta\gamma_y\kappa)^j E_t\left[\hat{P}_{t+j} - \hat{P}_t + \hat{\psi}_{t+j,t}\right]. \qquad (19)
\end{aligned}
$$

This expression is exactly the same than the one reported in Goodfriend and King (1997) or Sbordone (2001).

A.2.8 The New Keynesian Pricing equation

The loglinearized aggregate price equation and the loglinearized optimal pricing equation can be combined to arrive at an expression linking inflation and *average* real marginal cost that is commonly referred to as the New Keynesian pricing equation. The first step in deriving this equation is to link firm-specific real marginal cost $\hat{\psi}_{t+j,t}$ to *firm-wide average* real marginal cost $\hat{\psi}_{t+j}$. We can write[39]

$$\hat{\psi}_{t+j} = (1-\kappa)\hat{\psi}_{t+j,t+j} + \kappa\bar{\psi}_{t+j,t+j-1}$$

where $\bar{\psi}_{t+j,t+j-1}$ is the marginal cost of a firm that charges the average price P_{t+j-1}in period $t+j$. Next, we linearize $\hat{\psi}_{t+j,t+j}$ and $\bar{\psi}_{t+j,t+j-1}$ around the period t optimal price level P_t^*

$$\begin{aligned}\hat{\psi}_{t+j,t+j}|_{\hat{P}_t^*} &= \hat{\psi}_{t+j,t} + \left[\frac{\partial\hat{\psi}_{t+j,t}}{\partial\hat{y}_{t+j,t}}\frac{\partial\hat{y}_{t+j,t}}{\partial P_t^*}\right](\hat{P}_{t+j}^* - \hat{P}_t^*) \\ \hat{\psi}_{t+j,t+j}|_{\hat{P}_t^*} &= \hat{\psi}_{t+j,t} + \eta\mu(\hat{P}_{t+j}^* - \hat{P}_t^*)\end{aligned}$$

and

$$\begin{aligned}\bar{\psi}_{t+j,t+j-1}|_{\hat{P}_t^*} &= \hat{\psi}_{t+j,t} + \left[\frac{\partial\hat{\psi}_{t+j,t}}{\partial\hat{y}_{t+j,t}}\frac{\partial\hat{y}_{t+j,t}}{\partial P_t^*}\right](\hat{P}_{t+j-1} - \hat{P}_t^*) \\ \bar{\psi}_{t+j,t+j-1}|_{\hat{P}_t^*} &= \hat{\psi}_{t+j,t} + \eta\mu(\hat{P}_{t+j-1} - \hat{P}_t^*)\end{aligned}$$

In this notation, η is the elasticity of the firm's marginal cost with respect to its output, and μ is the elasticity of the firm's demand with respect to its price. Using these two equations in the expression for $\hat{\psi}_{t+j}$, we obtain after some rearrangement

$$\hat{\psi}_{t+j,t} = \hat{\psi}_{t+j} + \eta\mu(\hat{P}_t^* - \hat{P}_{t+j})$$

Now, we can substitute for $\hat{\psi}_{t+j,t}$ in the optimal pricing equation and solve for the NK pricing equation

$$\begin{aligned}\frac{\kappa}{1-\kappa}\hat{\pi}_t &= (1-\beta\kappa)\sum_{j=0}^{\infty}(\beta\kappa)^j E_t\left[\hat{P}_{t+j} + \hat{\psi}_{t+j} + \eta\mu(\hat{P}_t^* - \hat{P}_{t+j})\right] - \hat{P}_t \\ \frac{\kappa}{1-\kappa}(1+\eta\mu)\hat{\pi}_t &= (1-\beta\kappa)\sum_{j=0}^{\infty}(\beta\kappa)^j E_t\left[(1+\eta\mu)\hat{P}_{t+j} + \hat{\psi}_{t+j}\right] - (1+\eta\mu)\hat{P}_t \\ \frac{\kappa}{1-\kappa}(1+\eta\mu)\hat{\pi}_t &= (1-\beta\kappa)\sum_{j=0}^{\infty}(\beta\kappa)^j E_t L^{-j}\left[(1+\eta\mu)\hat{P}_t + \hat{\psi}_t\right] - (1+\eta\mu)\hat{P}_t \\ (1-\beta\kappa E_t L^{-1})(1+\eta\mu)\left[\frac{\kappa}{1-\kappa}\hat{\pi}_t + \hat{P}_t\right] &= (1-\beta\kappa)\left[(1+\eta\mu)\hat{P}_t + \hat{\psi}_t\right]\end{aligned}$$

Rearranging this equation and canceling out different terms, we obtain

$$\hat{\pi}_t = \beta E_t\hat{\pi}_{t+1} + \frac{(1-\beta\kappa)(1-\kappa)}{\kappa}\frac{1}{1+\eta\mu}\hat{\psi}_t,$$

which is the NK pricing equation as reported by Woodford (2003).

[39] This derivation is similar to the one provided in Woodford (2003).

A.2.9 Special cases

Suppose that (i) input markets are perfectly competitive; (ii) input factors can be reallocated instantaneously at no cost; and (iii) technology is loglinear in its input factors. Under these assumptions, real marginal cost is a function of factor prices, which are exogenous at the firm level. Hence, the elasticity of marginal cost with respect to output is zero; i.e. $\eta = 0$ and the NK pricing equation becomes

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \frac{(1-\kappa)(1-\kappa\beta)}{\kappa}\hat{\psi}_t.$$

This expression is identical to the one derived by Goodfriend and King (1997), Gali and Gertler (1999), and many others.

The second special case is the one proposed by Sbordone (2002), where technology is assumed to be Cobb-Douglas and capital cannot be reallocated among firms within a period. Minimizing total cost under these assumptions leads to the following first order conditions

$$\begin{aligned} w_t &= \psi_t (1-\alpha) A_t \left(k_t/n_t\right)^{\alpha} \\ y_t &= A_t {n_t}^{(1-\alpha)} {k_t}^{\alpha}, \end{aligned}$$

where the second condition is simply the production function and replaces the optimality condition with respect to capital (since k is fixed within a period). Solving for ψ_t yields

$$\psi_t = B {y_t}^{\alpha/(1-\alpha)},$$

where

$$B \equiv \frac{1}{w_t(1-\alpha)} {A_t}^{-1/(1-\alpha)} {y_t}^{\alpha/(1-\alpha)} {k_t}^{-\alpha/(1-\alpha)}.$$

Hence

$$\frac{\partial \psi_t}{\partial y_t}\frac{y_t}{\psi_t} = \frac{\alpha}{(1-\alpha)}.$$

Replacing η by this expression, we get the following NK pricing equation

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \frac{(1-\kappa)(1-\kappa\beta)}{\kappa}\frac{1-\alpha}{1-\alpha+\eta\mu}\hat{\psi}_t \tag{20}$$

which is the same than equation (2.8) in Sbordone (2002).

A.3 Bootstrap algorithm

The bootstrap algorithm used to compute the sample distribution and confidence intervals of $\hat{\rho}$ and $\hat{\Delta}$ incorporates the bias-correction mechanism by Kilian (1998a) and takes into account of the lag order uncertainty in the VAR forecasting process as in Kilian (1998b). It can be described as follows.[40]

[40] With the exception of a few changes, the description of the algorithm is taken from Kilian (1998a) and Kilian (1998b).

1. Determine the optimal lag order $\hat{p}$ by an appropriate selection criteria and estimate the VAR($\hat{p}$) process, written in companion form as $\omega_t = M\omega_{t-1}+e_t$, by means of OLS. Consider the estimates $\hat{M}$ and the vector of residuals $\{\hat{e}_t\}$ as the data-generating process.[41]

2. Generate N artificial series $\{\omega_t^*\}$ based on the recursion:

$$\omega_t^* = \hat{M}\omega_{t-1}^* + \hat{e}_t^*, \tag{21}$$

where $\hat{e}_t^*$ is a random draw with replacement from the vector of residuals $\{\hat{e}_t\}$.[42]

3. For each artificial series, fit a VAR(p) and estimate simulated coefficients $\hat{M}_i^*$ by means of OLS.

4. Approximate the OLS small-sample bias term $\Psi = E(\hat{M} - M)$ by $\hat{\Psi} = 1/N \sum_{i=1}^N \hat{M}_i^* - \hat{M}$.[43]

5. Construct the bias-corrected coefficient estimate $\bar{M} = \hat{M} - \hat{\Psi}$. Compute the absolute value of the largest root of $\bar{M}$. Denote this quantity by $m(\bar{M})$. If $m(\bar{M}) \geq 1$, let $\hat{\Psi}_{i+1} = \delta_i \hat{\Psi}_i$ and $\delta_{i+1} = \delta_i - 0.01$. Set $\bar{M} = \bar{M}_i$ after iterating on $\bar{M}_i = \hat{M} - \hat{\Psi}_i$, $i = 1, 2, ...$until $m(\bar{M}_i) < 1$. The purpose of this adjustment is to avoid that the bias-corrected companion matrix $\bar{M}$ implies a non-stationary VAR process. As Kilian (1998a) explains, the adjustment has no effect asymptotically since it does not restrict the parameter space of the OLS coefficient estimates $\hat{M}$ themselves but only its bias estimates.

6. Replace $\hat{M}$ with $\bar{M}$ in equation (21) and generate N new artificial series $\{\omega_t^*\}$ from this bias-corrected data-generating process.

7. For each artificial series, determine the optimal lag order $\hat{p}^*$ by the same selection criteria as in Step 1. Fit a VAR($\hat{p}^*$) to the artificial series in question and compute the simulated estimate $\hat{M}^*$ by OLS. Then construct the corresponding bias-corrected simulated coefficient estimate $\bar{M}^* = \hat{M}^* - \hat{\Psi}^*$ where the bias estimate $\hat{\Psi}^*$ is computed as in Steps 3 and 4 but for the appropriate lag order $\hat{p}^*$.[44]

   For each artificial series, compute the simulated series of theoretical inflation $\{\pi_t'^*\}$ as in formula (6) taking the structural parameter values β and λ as given, i.e.:

$$\pi_t'^* = \phi h_\psi [I - \beta\gamma_y \bar{M}^*]^{-1} \omega_t^*.$$

[41] In order to ensure zero mean and desired variance, all the residuals $\hat{e}_t$ are rescaled according to $\hat{e}_t' = 1/T \sum_{t=1}^T \hat{e}_t \sqrt{T/(T-2p)}$. See Berkowitz and Kilian (2000) for a discussion.

[42] To initialize the recursion, I randomly select p initial values using the block method by Stine (1987).

[43] This approximation amounts to assuming that the bias is constant in the neighbourhood of $\hat{M}$. See Kilian (1998a) for more discussion. Concurrently, we could have used the closed-form solution by Pope (1990) to compute the bias. However, simulations revealed that the bootstrap method to compute the bias was more accurate.

[44] Note that to estimate this mean bias $\hat{\Psi}^*$ would require nesting a separate bootstrap inside each of the bootstrap loops. Instead, to reduce the computational requirements, I compute the mean bias $\hat{\Psi}^*$ for each possible lag order $\hat{p}^* = 1...\hat{P}^*$ apart in an initial bootstrap (with $\hat{P}^*$ being an arbitrary upper bound). As Kilian (1998a) shows, this short-cut is justified asymptotically.

From this simulated series of theoretical inflation and corresponding simulated series of inflation taken from ω_t^*, compute the simulated correlation coefficient $\rho(\pi_t^*, \pi_t'^*)$ and the simulated variance ratio $\Delta(\pi_t^*, \pi_t'^*)$.

8. The simulated sample distribution of $\hat{\rho}$ and $\hat{\theta}$ consists of the (ordered) N simulations $\rho(\pi_t^*, \pi_t'^*)$ and $\Delta(\pi_t^*, \pi_t'^*)$, respectively. The α and $1-\alpha$ percentile interval endpoints can be read off of this distribution.

Table 1
Unrestricted VAR estimates of benchmark example

	s_{t-1}	π_{t-1}	s_{t-2}	π_{t-2}	s_{t-3}	π_{t-3}	s_{t-4}	π_{t-4}	R^2
s_t	0.876	0.008	0.004	-0.077	-0.017	0.204	-0.074	0.150	0.781
	(0.084)	(0.252)	(0.112)	(0.294)	(0.112)	(0.291)	(0.084)	(0.250)	
π_t	0.073	0.632	-0.059	0.042	-0.009	0.211	-0.007	0.048	0.825
	(0.029)	(0.086)	(0.038)	(0.100)	(0.038)	(0.099)	(0.029)	(0.085)	

Notes: This table reports the coefficient estimates for the OLS regressions of U.S. labor income share and inflation on lags thereof. The sample period is 1960:1-1997:4. Standard errors are shown in brackets.

Table 2
Unrestricted estimates for VAR(4) in labor income share and unit labor cost changes

	s_{t-1}	Δulc_{t-1}	s_{t-2}	Δulc_{t-2}	s_{t-3}	Δulc_{t-3}	s_{t-4}	Δulc_{t-4}	R^2
s_t	0.856	0.025	0.094	-0.065	-0.292	0.208	0.137	0.955	0.782
	(0.275)	(0.245)	(0.496)	(0.293)	(0.494)	(0.253)	(0.258)	(0.081)	
Δulc_t	-0.723	0.676	0.628	-0.011	-0.468	0.429	0.350	0.123	0.362
	(0.307)	(0.274)	(0.555)	(0.327)	(0.552)	(0.283)	(0.288)	(0.090)	

Notes: This table reports the (bias-corrected) coefficient estimates for the OLS regressions of U.S. labor income share and unit labor cost changes on lags thereof. The sample period is 1960:1-1997:4.

Table 3a
Block exogeneity tests for extending the bivariate benchmark VAR in labor income share and inflation

Test for 1 additional variable (4 lags)

	dn	dw	y-c	y-i	dM	spread
Likelihood ratio	26.9226**	12.243	32.2305**	19.9441*	18.4357*	12.147
Degrees of freedom	8	8	8	8	8	8

Test for several additional variables (4 lags of each)

	dn, dw	y-c, y-i	dM, spread	dn,dw,y-c,y-i	dn,dw,y-c,y-i, dM,spread
Likelihood ratio	40.6528**	42.2442**	32.0625**	71.1806**	86.8998**
Degrees of freedom	16	16	16	32	48

Notes: * = significance at 5% level, ** = signficance at 1% level.
The null hypothesis of the tests is H0: the forecasting performance of the bivariate VAR(4) in labor income share and inflation is not improved by the addition of any of these (combination of) variables. The sample period is 1960:1-1997:2.

Table 3b
Block exogeneity tests for extending the alternative VAR in labor income share and unit labor cost changes

Test for 1 additional variable (4 lags)

	dn	dw	y-c	y-i	dM	spread
Likelihood ratio	26.414**	11.258	26.720**	19.087*	15.500*	10.986
Degrees of freedom	8	8	8	8	8	8

Test for several additional variables (4 lags of each)

	dn, dw	y-c, y-i	dM, spread	dn,dw,y-c,y-i	dn,dw,y-c,y-i, dM,spread
Likelihood ratio	38.377**	35.955**	26.565*	64.153**	78.162**
Degrees of freedom	16	16	16	32	48

Notes: * = significance at 5% level, ** = signficance at 1% level.
The null hypothesis of the tests is H0: the forecasting performance of the bivariate VAR(4) in labor income share and unit labor cost changes is not improved by the addition of any of these (combination of) variables. The sample period is 1960:1-1997:2.

Table 4
The uncertain fit of the Calvo pricing model

	Marginal cost coefficient ϕ			
	$\phi = 0.035$		$\phi = 0.083$	
Model	Implied average price fixity (quarters)			
Traditional Calvo	5.9		4.0	
Calvo without capital mobility ($\mu = 10$ / $\alpha = 0.4$)	2.5		1.8	
	benchmark VAR	extended VAR	benchmark VAR	extended VAR
correlation coefficient	0.978 (0.402, 0.990)	0.550 (-0.542, 0.839)	0.978 (0.402, 0.990)	0.550 (-0.528, 0.843)
standard deviations ratio	0.567 (0.009, 1.570)	0.948 (0.002, 1.326	0.239 (0.004,0.662)	0.400 (0.001, 0.564)

Notes: This table reports the correlation coefficient between theoretical inflation and observed U.S. inflation as well as the standard deviation ratio of theoretical inflation relative to observed inflation for different marginal cost coefficients and different VAR forecasting processes. The sample period is 1960:1-1997:2. The numbers in brackets report the 90% confidence interval of the respective point estimates. All numbers are bias-corrected (see text).