Volume Title: NBER Macroeconomics Annual 2007, Volume 22

Volume Author/Editor: Daron Acemoglu, Kenneth Rogoff and Michael Woodford, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 978-0-226-00202-6

Volume URL: http://www.nber.org/books/acem07-1

Conference Date: March 30-31, 2007

Publication Date: June 2008

Chapter Title: Comment on "How Structural Are Structural Parameters?"

Chapter Author: Timothy Cogley

Chapter URL: http://www.nber.org/chapters/c4088

Chapter pages in book: (p. 139 - 147)

Comment

Timothy Cogley, University of California, Davis

1 Introduction

Lucas (1976) has stated that Cowles Commission models are not structural, but he did not say that DSGE models are. Under what conditions would that be the case?

Imagine that we had a doubly long sample with a single, unanticipated once-and-for-all policy shift in the middle. To make things simple, neglect learning and assume immediate convergence to a rational-expectations equilibrium after the regime change. Also suppose that the date of the change is known, so that we can cleanly estimate parameters of a DSGE approximating model using the two subsamples.

Subject to regularity conditions, a Bayesian consistency theorem states that parameter estimates converge in probability to the argmin of the Kullback-Liebler information criterion (KLIC).[1] This probability limit is also known as the *pseudo-true value.* When the approximating model is correctly specified, the pseudo-true values coincide with the true parameters. Otherwise, there is an asymptotic bias: the estimates converge to something, but not to the true parameters.

If the approximating model were correctly specified, it would be structural in the sense of Lucas. In that case, the estimates would converge to the truth in both of the long subsamples. Hence, after a policy shift, parameters of the estimated policy rule would change, but estimates of preference and technology parameters would remain the same. Since the latter are invariant, the model could be trusted to evaluate the consequences of alternative policies. This is why many economists favor DSGE modeling as a strategy for overcoming the Lucas critique.

For a misspecified approximating model, invariance of preference and technology parameters is not guaranteed. After a policy shift, the estimated parameters would adjust to re-optimize the KLIC, and there is no reason to expect that the adjustment would be confined to policy rule parameters, with private sector parameters remaining at the same values as in the first subsample. All the parameters are free to adjust to re-optimize fit, and typically they will do so.[2] In order for estimates of preference and technology parameters to remain invariant, the asymptotic bias would have to be unaffected by the policy shift. This is a statistical property of the model, not grounded in the economic bedrock of preference and technology conditions. If we take seriously that DSGE models are approximations, the best we can hope for is that preference and technology estimates won't change too much.

Relatively little work has been done to examine whether existing DSGE models are structural. This is where Fernández-Villaverde and Rubio-Ramírez enter. They examine a dynamic new Keynesian model of the kind now becoming popular in central banks and question whether its parameters are invariant to the shifts in monetary policy experienced in the United States over the last 50 years. They do so by allowing some of the model parameters to drift, interpreting drift in private sector parameters as evidence that a model is not structural. Some of their results point to changes in the monetary policy rule, corroborating findings of Clarida, Galí, and Gertler (2000), Lubick and Schorfheide (2004), and others. Having documented shifts in monetary policy, they go on to examine the invariance of Calvo pricing parameters. They document drift in the frequency of price and wage adjustment as well as in indexation parameters, and they relate the patterns of drift to changes in trend inflation. For the most part, the patterns intuitively make sense.[3] They conclude that the Calvo pricing model is not structural for the range of monetary policies experienced in the United States, and they recommend against using it for policy analysis.

The paper is technically very sharp, and it substantially advances econometric methods for estimating drifting-parameter models. In particular, as far as I know, FVRR are the first to estimate a DSGE model with drifting parameters. The main challenge in estimating such a model involves evaluating the likelihood function. The model can be cast as a nonlinear state-space model, and the likelihood function can be expressed using the prediction-error decomposition. Unlike a linear state-space model, however, the Kalman filter cannot be used to evaluate the terms in the prediction-error decomposition. Building on their own earlier research, FVRR substitute a nonlinear filtering algorithm known as a *particle filter.* Although some kinks remain to be worked out, their particle-filtering algorithm seems very promising.

Although the paper is technically very impressive, I am a bit skeptical about their substantive conclusions. In particular, I question their interpretation of drift in Calvo pricing parameters. In part, that is because they adopt a one-at-a-time approach to modeling parameter drift. In addition, other evidence exists that is more supportive of invariant Calvo parameters for the United States. The remainder of my comment focuses on these issues.

2 Are Calvo Pricing Parameters Structural?

For examining the invariance of Calvo pricing parameters, the ideal modeling strategy would allow policy and Calvo parameters to drift jointly. If we found drift in policy parameters but not in Calvo parameters, we could conclude that the latter were invariant with respect to shifts in the former. On the other hand, if Calvo parameters drifted as well, we would conclude that they are not structural.

This is not what FVRR do. Instead, they adopt a one-at-a-time approach to modeling parameter drift. That is, they estimate a sequence of approximating models, each of which involves a single drifting parameter while holding all other parameters constant. In one model, the drifting parameter is the long-run inflation target in the central bank's policy rule. In others, it is one of the policy-rule feedback parameters. In still others, one of the Calvo wage- or price-adjustment or indexation parameters is free to vary, with the policy-rule parameters held constant.

This approach is problematic because it does not reliably identify the source of time variation. For example, suppose that Clarida et al. and Lubick and Schorfheide are correct—that an important shift in monetary policy occurred around the time of the Volcker disinflation. They say that the Federal Reserve failed to satisfy the Taylor principle in the 1970s but did satisfy it after the early 1980s.[4] Among other things, this change in policy would alter inflation-gap persistence, presumably making inflation revert more quickly to the Fed's long-run target.[5]

Suppose, however, that the approximating model held the policy parameters constant and allowed one of the Calvo parameters—say the price indexation parameter—to drift. Since the policy parameters are held constant, the estimates would fail to detect the change in policy and the model would not identify the true source of time-varying inflation persistence. On the other hand, the indexation parameter is free to drift, and changes in the degree of indexation after inflation-gap persistence. Since this feature of the approximating model can fit time-varying inflation persistence, the indexation parameter is likely to drift to compensate for misspecification of the policy rule. Fernández-Villaverde and Rubio-Ramírez would interpret this as evidence that the Calvo model is not structural. But in this example, a Calvo parameter drifts because the true source of time variation has been shut off.

My point is not to argue for this scenario. I just want to provide an example to question FVRR's reading of the evidence. I would not jump to the conclusion that Calvo models fail to be invariant to shifts in U.S. monetary policy. That question remains open.

Indeed, there is something odd about FVRR's scenarios. If I understand correctly, their chief concern is that Calvo parameters might change in response to a shift in the monetary policy rule. But in their models, whenever Calvo parameters are free to drift, policy parameters are held constant. This makes it hard to attribute drift in Calvo parameters to changes in the monetary policy rule. Their exercise is probably better interpreted as a way to look for evidence of state-dependent pricing within a stable monetary regime. That is also interesting, provided that one believes the Fed's policy rule has not changed.

This criticism is a bit churlish, because FVRR are tackling a hard problem and are constrained by what is computationally feasible. The ideal strategy—allowing for joint drift in policy and Calvo parameters—is computationally very intensive and seems to be beyond the current state of the art. With particle-filtering algorithms, convergence is hard to achieve in models with high-dimensional state vectors, and that is why they allow only one parameter at a time to drift. Liu and West (2001) also report that convergence is sometimes difficult to attain when simultaneously filtering for hidden states and estimating unknown, constant hyperparameters. Convergence is also finicky when a model's nonlinear cross-equation restrictions permit more than one solution. Thus, a number of technical challenges remain. Nevertheless, FVRR are the masters of the craft within economics, and I look forward to their future work.

Since the one-parameter-at-a-time approach is a shortcut, one might ask about other shortcuts. For instance, Boivin and Giannoni (2005), Canova (2005), and Schorfheide (2007) look for evidence of parameter shifts by estimating constant-parameter DSGE models for various subsamples. Cogley and Sbordone (2005) estimate a reduced-form VAR with drifting parameters and then estimate Calvo pricing parameters from the model's cross-equation restrictions. Contrary to FVRR, Sbordone and I find that a constant-parameter Calvo model can be reconciled with time-variation in the law of motion for inflation and other variables. I want to conclude by reviewing that evidence.

One of the main findings of reduced-form studies such as Cogley and Sargent (2005) is that trend inflation drifts, starting low in the 1960s, rising during the Great Inflation, and falling after the Volcker disinflation. Sbordone and I interpret this in terms of shifts in the Fed's long-run target for inflation, and we incorporate this into the Calvo model. With drift in the central bank's inflation target, the New Keynesian Phillips curve (NKPC) becomes

$$\hat{\pi}_t \approx \tilde{\rho}_t \hat{\pi}_{t-1} + \zeta_t \hat{s}_t + b_{1t} E_t \hat{\pi}_{t+1} + b_{2t} E_t \sum_{j=2}^{\infty} \varphi_{1t}^{j-1} \hat{\pi}_{t+j} + u_t, \quad (1)$$

where π_t represents inflation, s_t is real marginal cost, and u_t is a white noise error. Hat variables represent log differences from time-varying trends, measured using the reduced-form, time-varying-parameter VAR (see our paper for details). The NKPC parameters $\tilde{\rho}_t$, ζ_t, b_{1t}, and φ_{1t} are functions of trend inflation and the Calvo parameters, $\psi = (\alpha, \theta, \rho)$. The parameter α represents the frequency of price adjustment, θ is the elasticity of substitution in the Dixit-Stiglitz aggregator, and ρ measures the extent to which nonoptimizing firms mechanically index to past inflation. The introduction of drifting trend inflation alters the standard NKPC in two ways: additional forward-looking inflation terms appear on the right-hand side of (1), and the NKPC parameters drift through time. Note that the NKPC parameters will drift even if the Calvo parameters are constant. We are concerned about invariance of the Calvo parameters.

We estimate the model in two steps. In the first, we estimate a time-varying parameter VAR for inflation and real marginal cost. In the second, following Sbordone (2002, 2005), we estimate Calvo parameters to satisfy the model's cross-equation restrictions. To derive those restrictions, write the VAR in companion form

$$z_t = \mu_t + A_t z_{t-1} + \varepsilon_{zt}, \quad (2)$$

where z_t stacks current and lagged values of inflation and marginal cost, μ_t includes the VAR intercepts, and A_t contains the VAR autoregressive parameters. Next, take the reduced-form conditional expectation of the inflation gap,

$$E(\hat{\pi}_t \mid \hat{z}_{t-1}) = e_{\pi}' A_t \hat{z}_{t-1}, \quad (3)$$

where e_π is a selector vector. Then take the conditional expectation for the inflation gap from the NKPC,

$$E(\hat{\pi}_t \mid \hat{z}_{t-1}) \approx [\tilde{\rho} e'_\pi + \zeta_t e'_s A_t + b_{1t} e'_\pi A_t^2 + b_{2t} e'_\pi \varphi_{1t}(I - \varphi_{1t} A_t)^{-1} A_t^3] \hat{z}_{t-1}, \tag{4}$$
$$\equiv g(\mu_t, A_t, \psi) = \hat{z}_{t-1},$$

where e_s is another selector vector. After equating the two and insisting that they hold for all realizations, we obtain a vector of cross-equation restrictions involving the Calvo parameters ψ and the drifting VAR parameters μ_t and A_t:

$$F_t(\mu_t, A_t, \psi) = e'_\pi A_t - g(\mu_t, A_t, \psi) = 0. \tag{5}$$

In principle, these restrictions should hold for all dates. To keep the problem to a manageable dimension, we focus on five representative years—1961, 1978, 1983, 1995, and 2005—which span the variety of monetary experience in the sample. The year 1961 is drawn from an initial period of low and stable inflation. The height of the Great Inflation is represented by 1978, when both trend inflation and the degree of persistence were close to their maxima. The Volcker disinflation is represented by 1983, a key turning point in postwar U.S. monetary history. The final two years, 1995 and 2003, are drawn from the Greenspan era, a mature, low-inflation environment. The first was chosen to represent the preemptive Greenspan, the second reflects his later wait-and-see approach.

We stack the restrictions at each date into a vector,

$$\mathcal{F}(\cdot) = (F'_{1961}\, F'_{1978}\, F'_{1983}\, F'_{1995}\, F'_{2003})', \tag{6}$$

and then estimate the Calvo parameters α, ρ, θ by minimizing $\mathcal{F}(\cdot)'\mathcal{F}(\cdot)$, subject to constraints that they lie within economically sensible regions.[6] For each draw in the posterior sample for the VAR parameters μ, A, we calculate best-fitting values for α, ρ, thus deducing a distribution for the Calvo parameters from the posterior for the VAR parameters.[7] The results are summarized in table 2C1.1.

The estimates of α, ρ, and θ are all economically sensible. The estimate of θ implies a steady state markup of about 11 percent, which is in line with other estimates in the literature. The estimate of the indexation parameter ρ sits on the lower bound of zero. Although this contrasts with much of the empirical literature on the NKPC, we regard it as a virtue

Table 2C1.1
Estimates of Calvo Parameters

	α	ρ	θ
Median	0.602	0	9.97
Median absolute deviation	0.048	0	0.90

for two reasons. One is that the microfoundations for indexation are less compelling than for other elements of the Calvo model. The other is that there is no support in micro data for indexation. When $\rho > 0$, the model implies that every firm would change price every quarter, some optimally rebalancing marginal benefit and marginal cost, others mechanically marking up prices in accordance with the indexation rule. Yet Bils and Klenow (2004) report that approximately 75 percent of prices remain unchanged every month. That would not happen if firms indexed mechanically to past inflation.

Finally, with $\rho = 0$, our estimate of α implies a median duration of prices of 1.36 quarters, or 4.1 months, a value consistent with microeconomic evidence on the frequency of price adjustment. For example, before adjusting for sales, Bils and Klenow report a median duration of 4.4 months.[8] Our estimate from macroeconomic data therefore accords well with the conclusions they draw from microeconomic data.

Most importantly, our constant-parameter Calvo model fists the time-varying dynamics of inflation and marginal cost quite well. Following Sbordone's earlier work, we assess the model's fit by comparing the expected inflation gap implied by the NKPC with the expected inflation gap estimated by the unconstrained VAR. The VAR inflation forecast is given by equation (3), while the NKPC forecast is defined by the right-hand side of equation (4). The distance between the two forecasts measures the extent to which the model's cross-equation restrictions are violated. Figure 2C1.1 portrays the two series, plotting the VAR forecast as a solid line and the NKPC forecast as a dotted line.

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

Finally, Sbordone and I also estimated versions of the NKPC that allow Calvo parameters to vary across dates. We found no compelling evidence of changes in α, ρ, or θ, or that variation in the Calvo parameters improves the fit of the model. Thus, we concluded that a constant-parameter version of the Calvo model can be reconciled with postwar U.S. monetary experience.

Why do we get a different answer? One important difference between our model and that of FVRR is that drift in Calvo parameters is not the sole source of time variation in the law of motion for inflation, as it is in their model. Once drift in trend inflation is accounted for, there seems to

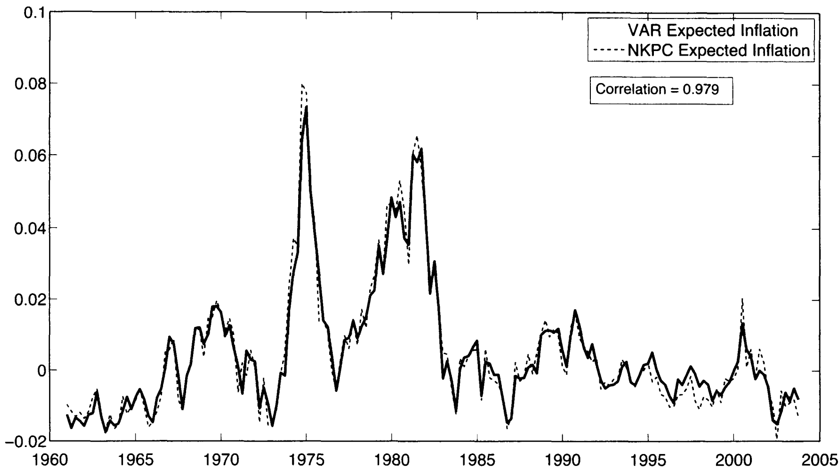

Figure 2C1.1
Assessing the Cross-Equation Restrictions

be less need for drift in Calvo parameters. But there are a number of other important differences as well, and sorting this out will require more work.

Acknowledgment

I am grateful to Argia Sbordone for many discussions of the issues addressed herein.

Endnotes

1. See Gelman, Carlin, Stern, and Rubin 2000, appendix B.

2. This is one way to interpret the subsample estimates of Schorfheide (2007).

3. There are a few puzzling results. For instance, actual inflation exceeds target inflation from the late 1960s until the early 1980s. It is hard to imagine that a purposeful central bank would allow a positive inflation gap to remain open for fifteen years.

4. The evidence on this is mixed, but let's assume this for the sake of argument. Orphanides (2001) and Sims and Zha (2006)provide other perspectives.

5. Cogley and Sargent (2007) report empirical evidence that the inflation gap was less persistent after the Volcker disinflation.

6. We restrict $\alpha \in (0,1)$, $\rho \in (0,1)$, $\theta \in (1, \infty)$.

7. This is not a Bayesian posterior for the Calvo parameters. It is a change of variables with respect to the reduced-form posterior.

8. Nakamura and Steinsson (2006) report comparable numbers. Their disagreement with Bils and Klenow concerns sales-adjusted durations.

References

Bils, M., and P. J. Klenow. 2004. Some evidence on the importance of sticky prices. *Journal of Political Economy* 112:947–85.

Boivin, J., and M. Gianonni. 2006. Has monetary policy become more effective? *Review of Economics and Statistics* 88 (3): 445–62.

Canova, F. 2005. Monetary policy and the evolution of the U.S. economy. Unpublished manuscript, Universitat Pompeu Fabra.

Clarida, R., J. Galí, and M. Gertler. 2000. Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics* 115:147–80.

Cogley, T., and T. J. Sargent. 2005. Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics* 8:262–302.

———. 2007. Inflation-gap persistence in the U.S. Unpublished manuscript. UC Davis and NYU.

Cogley, T. and A. M. Sbordone. 2005. A search for a structural Phillips curve. Unpublished manuscript, UC Davis and Federal Reserve Bank of New York.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 2000. *Bayesian data analysis.* London: Chapman and Hall.

Liu, J., and M. West. 2001. Combined parameter and state estimation in simulation-based filtering. In *Sequential Monte Carlo methods in practice,* 199–224. New York: Springer.

Lubick, T., and F. Schorfheide. 2004. Testing for indeterminacy: An application to U.S. monetary policy. *American Economic Review* 94:190–217.

Lucas, R. E., Jr. 1976. Econometric policy evaluation: A critique. In *The Phillips curve and labor markets,* vol. 1, ed. K. Brunner and A. Meltzer, Carnegie-Rochester Series on Public Policy, 19–46. Amsterdam: North-Holland.

Nakamura, E., and J. Steinsson. 2006. Five facts about prices. Unpublished manuscript, Harvard University.

Orphanides, A. 2001. Monetary policy rules based on real-time data. *American Economic Review* 91:964–85.

Sbordone, A. M. 2002. Price and unit labor costs: A new test of price stickiness. *Journal of Monetary Economics* 49:265–92.

———. 2005. Do expected future marginal costs drive inflation dynamics? *Journal of Monetary Economics* 52:1183–97.

Schorfheide, F. 2007. Comment on "How structural are structural parameters." Unpublished manuscript, University of Pennsylvania.

Sims, C. A., and T. Zha. 2006. Were there regime switches in U.S. monetary policy? *American Economic Review* 96:54–81.