

WHAT EXPLAINS THE GREAT MODERATION IN THE U.S.? A STRUCTURAL ANALYSIS

Fabio Canova

ICREA and Universitat Pompeu Fabra

Abstract

This paper investigates what has caused output and inflation volatility to fall in the U.S. using a small scale structural model using Bayesian techniques and rolling samples. There are instabilities in the posterior of the parameters describing the private sector, the policy rule, and the variance of the shocks. Results are robust to the specification of the policy rule. Changes in the parameters describing the private sector are the largest, but those of the policy rule and the covariance matrix of the shocks explain the changes most. (JEL: E52, E47, C53)

1. Introduction

Many researchers, including Blanchard and Simon (2000), McConnell and Perez Quiroz (2000), and Stock and Watson (2002), have documented a marked decline in the variance of real activity and the variance and the persistence of inflation in the United States since the early 1980s. Although some have questioned the statistical significance of the reported changes (see Canova and Gambetti 2007a; or Pivetta and Reis 2007), there is agreement among macroeconomists that the nature and the causes of these changes should be carefully investigated.

Taylor (1998), Sargent (1999), Clarida, Gali, and Gertler (2000), and Lubik and Schorfheide (2004), among others, have attributed the fall to a permanent alteration in the weight that inflation receives in the objective function of the monetary authority. The popular version of the story runs as follows: The increase in inflation in the 1970s occurred because the authorities believed there was an exploitable trade-off between inflation and output. Because output was low following the two oil shocks, the temptation to inflate, in order to bring output at or

The editor in charge of this paper was Roberto Perotti.

Acknowledgments: I would like to thank two editors of the journal, two anonymous referees, and the participants of seminars at a number of universities and central banks for comments and constructive suggestions. The financial support of the Spanish Ministry of Education through the grant SEJ-2004-21682-E and of CREI is acknowledged. The author is also associated with CEPR, CREI, and AMeN.

E-mail address: Canova: fabio.canova@upf.edu

above its potential level, was strong. Between keeping inflation low (and output low) or inflation high (and output high), the monetary authorities systematically chose the latter option. Hence, inflation in the long run turned out to be higher while output simply settled to its potential level. Since the 1980s, however, the perception of the output–inflation trade-off has changed—the Fed learned that it was not exploitable and concentrated instead on the objective of fighting inflation. A low-inflation regime ensued, and the predictability of monetary policy contributed to make the macroeconomic environment less volatile and the swings in inflation and output more unpredictable.

Although prevalent, this view is not fully shared within the profession. Some researchers claim that monetary policy has not displayed any permanent switch; that the same policy rule characterizes most of the post-WWII experience; that monetary policy has little influence on output; and that good luck is responsible for the changes (see, e.g., Bernanke and Mihov 1998; Leeper and Zha 2003; Hanson 2006; Sims and Zha 2006). Others have suggested “real” reasons to explain the volatility fall (see, e.g., Ireland 1999; McConnell and Perez Quiroz 2000; Gordon 2005; or Campbell and Hercowitz 2006).

Some progress has been made in the investigation of these issues using empirical models where coefficients are allowed to vary over time. Cogley and Sargent (2001, 2005), who used a reduced-form time-varying coefficient VAR, find evidence that supports the causation story running from monetary policy changes to changes in the rest of the economy. Primiceri (2005), Sims and Zha (2006), and Canova and Gambetti (2007a), who estimate structural time-varying coefficients VARs, find little posterior support for this hypothesis. Because structural VARs only use a minimal amount of the restrictions implied by the current generation of DSGE models, one may wonder how truly structural the estimated relationships are. For example, Ireland (2001), Lubik and Schorfheide (2004), and Boivin and Giannoni (2006), who explicitly condition their analyses on a small-scale dynamic stochastic general equilibrium (DSGE) model, do find evidence of policy instability.

This paper provides new evidence on the causes of output and inflation volatility changes recursively estimating a small-scale DSGE model with Bayesian techniques. Recursive estimation provides a shortcut to more complicated analyses that allow for varying taste, technology, and policy parameters into a structural model, but requires estimation of second-order approximations to the solution and much more time-consuming posterior simulators (see Fernandez Villaverde and Rubio Ramirez 2007). Also, compared with analyses where subsamples are arbitrarily chosen, a recursive approach allows us to obtain more solid evidence on the nature of the time variations. Because the volatility of output (inflation) displays a U-shaped (inverted U-shaped) pattern, conclusions may crucially depend on the selected break point. Bayesian methods have inferential and computational advantages over traditional maximum likelihood techniques when dealing

with models that are a “false” description of the data generating process. This is important because, despite recent attempts to make them more realistic, DSGEs are still highly stylized; important relationships are modeled with black-box frictions; and ad-hoc shocks are used to dynamically span the probabilistic space of the data. In these situations, unrestricted maximum likelihood estimates are often unreasonable, and asymptotic standard errors constructed assuming that the model is “true” (under the null) are uninterpretable. Posterior estimates are meaningful even for models displaying such features. A Bayesian framework is also preferable to an approach that obtains estimates of the structural parameters by matching a subset of impulse responses in two respects: All the information of the model is efficiently used, and the trade-off between identifiability and non-linearities is dealt with in a more transparent way (see, e.g., Canova and Sala 2006).

Rather than searching for the best empirical model, we take a standard specification that is popular in the theoretical literature and show what it tells us about the causes of the changes experienced in the United States. We first consider standard subsample analysis; then we estimate the model a number of times, using overlapping samples that span a 20-year window over the period 1955–2002, and analyze the evolution of the posterior distributions of the structural parameters. Our analysis is geared to shed light on two issues. First, we would like to know which parameters, if any, are drifting over time. Second, we would like to know which variation has contributed most to the observed changes in the volatility of output and inflation.

Although it is common to examine this latter question via counterfactuals in which parameters from different subsamples are switched (see, e.g., Boivin and Giannoni 2006), this practice violates a basic principle underlying the Lucas critique—agents are unaware that changes may repeatedly occur—and thus fails to provide a reliable answer. Our approach is to estimate unrestricted and restricted specifications, and then to examine by how much the fit of the model changes and the consequences of restricting some parameters on fraction of output and inflation variabilities that are explained by the model.

We find instabilities in all the parameters of interest. Consistent with the common wisdom, the inflation coefficient in the policy rule increases if the sample includes only the years after 1982. However, changes are relatively small and often insignificant. The parameters describing the private sector also change, and the variations are significantly larger. Finally, the covariance matrix of the shocks changes over time and the adjustments are broadly in line with those reported in the VAR literature. These results are robust to the choice of policy rule: A rule that makes the interest rate respond to output growth rather than to the output gap, or to future rather than current developments in the economy, produces qualitatively similar results.

We show that changes in the parameters of the policy rule and the covariance matrix of the shocks are the most important in accounting for the changes in the volatility of output and inflation: Restricting them to be unchanged over the samples makes the fit of the model drop dramatically and the decline in volatility disappear. Interestingly, restricting the parameters of the policy rule imply a much higher inflation volatility, whereas restricting the standard deviations of the shocks increases the variance of output by a factor of 10. Hence, the changes in the volatility of the two variables may have different causes.

In sum, it appears that both the “good policy” and the “good luck” hypotheses have some support in the data. However, only by combining the two explanations do we account for the decline in the variability of real activity and inflation over time.

The rest of the paper is organized as follows. Section 2 describes the model and the estimation technique. Section 3 presents the basic results and a few robustness exercises. Section 4 compares our results to those in the literature, and Section 5 studies what explains the observed changes in the volatility of output and inflation. Section 6 concludes.

2. The Model and the Estimation Approach

The model we consider is a standard three-equations New-Keynesian model, consisting of a log-linearized Euler equation, a forward-looking Phillips curve, and a monetary policy rule. The system in log-linear form is

$$x_t = E_t(x_{t+1}) - \frac{1}{\varphi}(i_t - E_t\pi_{t+1}) + e_{1t}, \quad (1)$$

$$\pi_t = \beta E_t\pi_{t+1} + \kappa x_t + e_{2t}, \quad (2)$$

$$i_t = \psi_r i_{t-1} + (1 - \psi_r)(\psi_\pi \pi_t + \psi_x x_t) + e_{3t}, \quad (3)$$

where β is the discount factor, φ is the coefficient of relative risk aversion, κ is a parameter regulating the slope of Phillips curve, and $(\psi_r, \psi_\pi, \psi_x)$ are policy parameters. Here x_t is the output gap, π_t the inflation rate, and i_t the nominal interest rate. The shocks attached to each equation may not be structural in the sense that they may represent linear combinations of primitive disturbances to the economy. We assume

$$e_{1t} = \rho_1 e_{1t-1} + v_{1t}, \quad (4)$$

$$e_{2t} = a_{12} e_{1t} + \rho_2 e_{2t-1} + v_{2t}, \quad (5)$$

$$e_{3t} = a_{13} e_{1t} + v_{3t}, \quad (6)$$

where ρ_1, ρ_2 capture the persistence of the shocks and a_{12}, a_{13} the cross-equation effects, and where v_{jt} are mean zero processes with variance σ_j^2 , $j = 1, 2, 3$.

A system of equations like (1)–(3) can be obtained from a standard dynamics stochastic general equilibrium model with sticky prices, monopolistic competition, and preferences that are additive in consumption and leisure provided that labor is the only productive factor (see, e.g., Clarida, Galí, and Gertler 1999). The specification of the policy rule is consistent with the ideas that (i) the monetary authority observes current values of the output gap and of inflation when deciding the current interest rate and (ii) policy changes are smooth, in the sense that interest rate movements may be persistent. The specification for the error terms reflects that the expected level of potential output is omitted from the estimated specification but that the monetary authorities may pay attention to potential output changes when making decisions (see also An and Schorfheide 2007). The AR(1) assumption on e_{1t} and e_{2t} , on the other hand, is quite standard.

Throughout this paper we use a statistically computed measure of the output gap rather than the deviation of output from the level obtained in the flexible price equilibrium. We chose this approach for two reasons. First, this choice ensures comparability with previous work. Second, a flexible price measure that does not take capital accumulation into account is likely to be misspecified and this may potentially distort inference.

Several authors, including Smets and Wouters (2003), Rabanal and Rubio Ramirez (2005), and others have specified more complicated and realistic structures, that allow for additional shocks and frictions. But rather than augmenting the specification with bells and whistles to generate a better-fitting model, we perform our exercises using a simple and internally consistent specification that is close to those used in theoretical discussions.

The model contains 13 parameters: six have some structural interpretation $\alpha_1 = (\beta, \varphi, \kappa, \psi_r, \psi_x, \psi_\pi)$ and seven are auxiliary $\alpha_2 = (\rho_1, \rho_2, a_{12}, a_{13}, \sigma_1^2, \sigma_2^2, \sigma_3^2)$. Our exercise is geared to obtaining posterior distributions for $\alpha_T = (\alpha_{1T}, \alpha_{2T})$ over different samples T and to compare the time-series properties of their posterior distributions. Our system can be solved using standard first-order log-linear methods. The solution has a state-space format

$$y_{1t+1} = A_1(\alpha)y_{1t} + A_2(\alpha)v_{t+1}, \quad (7)$$

$$y_{2t} = A_3(\alpha)y_{1t}, \quad (8)$$

where $y_{2t} = [\pi_t, x_t, i_t]$, $y_{1t} = [\pi_{t-1}, x_{t-1}, i_{t-1}, e_{1t}, e_{2t}, e_{3t}]$, and the matrices $A_i(\alpha)$, $i = 1, 2, 3$ are nonlinear functions of the structural parameters α .

Bayesian estimation of equation (8) is simple. Given some α , we compute the likelihood of the model, denoted by $f(y_T|\alpha)$, by means of the Kalman filter and the prediction error decomposition. Then, for any specification of the prior distribution, denoted by $g(\alpha)$, the posterior distribution for the parameters is $g(\alpha|y_T) = (g(\alpha)f(y_T|\alpha))/f(y)$. The analytical computation of the posterior is impossible in our setup because the denominator of the expression,

$f(y)$, can be obtained only by integrating $g(\alpha)f(y_T|\alpha)$ with respect to α , a 13-dimensional vector. To obtain numerically a sequence from this unknown posterior, we employ a Metropolis algorithm. Roughly speaking, given α_0 and a transition function satisfying regularity conditions, we can produce a sequence from the unknown posterior, iterating on this transition function, after discarding an initial set of draws. We choose a standard random walk transition with jumps taken from a normal distribution centered at zero and with covariance matrix equal to a scaled version of the Hessian at the mode. The scale is sample dependent and is chosen to ensure that an appropriate number of draws (between 20% and 50%) is accepted. For each sample we draw five chains of 50,000 elements each and check convergence using standard CUMSUM methods. Posterior distributions are constructed using the last 5,000 draws from each of the chains.

We assume that the prior distribution can be factored as $g(\alpha) = \prod_{i=1}^{13} g(\alpha_i)$. Prior distributions are selected according to the following rule: Gamma distributions are used for parameters that must be positive, beta distributions for parameters that must lie in an interval, and normal distributions for all other parameters. This implies that φ , κ and σ_j^2 , $j = 1, 2, 3$ have gamma priors, that β , ψ_r , ρ_1 , ρ_2 have beta priors, and that the other parameters have normal priors except for ψ_π , whose prior is truncated below 1.0. The mean and the standard deviation of these distributions are shown in Table 1.

Clarida, Galí, and Gertler (2000) and Lubik and Schorfheide (2004), among others, have emphasized the potential importance of indeterminacies for characterizing the U.S. experience over the last 35 years. Because our prior distribution for the inflation coefficient in the policy rule is truncated at 1.0, no indeterminacy is allowed. Therefore, the changes we emphasize are changes within a determinate regime rather than changes across regimes. Canova and Gambetti (2007) showed that the dynamics induced by this model under indeterminacy (continuity solution) can be reasonably matched in a system where only determinate equilibria are considered. Hence, considering only determinate equilibria is less restrictive than it may first appear. Also, because our samples cut across periods with potentially different regimes, our prior assumption that the inflation policy coefficient is no less than 1.0 on average is not inconsistent with the possibility that, in particular periods of the sample, such a restriction is not satisfied.

The means of the priors are located around standard calibrated values and the mean for κ reflects a priori knowledge about its underlying components. The selected standard deviations imply proper but noninformative densities over a range of economically reasonable parameter values. We select “loose” priors in order to minimize subjective information and to allow the posterior to move away from the prior if the data is informative. Because we maintain the same prior in every subsample, differences in the location and shape of the posterior will tell us how much the likelihood evolves over time.

TABLE 1. Prior and posterior moments, basic rule.

	Prior		Posterior							
			1955–2002		1955–1979		1979–2002		1982–2002	
			Mean	Std.	Mean	HPI	Mean	HPI	Mean	HPI
φ	2.0	0.35	0.47	[0.42,0.52]	3.04	[2.67,3.51]	0.87	[0.25,0.41]	0.74	[0.45,1.01]
β	0.98	0.01	0.98	[0.98,0.99]	0.979	[0.97,0.99]	0.97	[0.95,0.99]	0.97	[0.95,0.99]
κ	2.0	1.0	4.15	[4.10,4.22]	3.42	[3.07,3.93]	3.90	[3.84,4.22]	3.90	[3.42,4.22]
ψ_r	0.8	0.25	0.98	[0.98,0.99]	0.99	[0.99,0.99]	0.98	[0.98,0.99]	0.98	[0.98,0.99]
ψ_x	0.5	0.25	0.02	[0.00,0.08]	0.16	[0.01,0.50]	0.03	[-0.00,0.01]	0.01	[0.00,0.02]
ψ_π	1.3	0.5	1.71	[1.65,1.76]	1.44	[1.42,1.50]	1.54	[1.23,1.53]	1.57	[1.51,1.58]
ρ_1	0.85	0.25	0.97	[0.97,0.98]	0.94	[0.92,0.94]	0.97	[0.97,0.97]	0.95	[0.94,0.96]
ρ_2	0.85	0.25	0.99	[0.99,0.99]	0.99	[0.99,0.99]	0.99	[0.99,0.99]	0.99	[0.99,0.99]
σ_1	0.01	0.50	0.012	[0.011,0.012]	0.021	[0.016,0.029]	0.012	[0.012,0.012]	0.012	[0.012,0.012]
σ_2	0.01	0.50	0.079	[0.071,0.081]	0.047	[0.044,0.061]	0.071	[0.067,0.087]	0.076	[0.065,0.078]
σ_3	0.01	0.50	0.156	[0.156,0.167]	0.031	[0.012,0.048]	0.131	[0.132,0.203]	0.146	[0.117,0.160]
a_{12}	0.0	0.25	-0.33	[-0.66,-0.31]	0.01	[-0.48,0.72]	-0.00	[-0.40,0.32]	-0.01	[-0.46,0.23]
a_{13}	0.0	0.25	0.15	[0.27,0.53]	0.52	[0.10,0.87]	0.02	[-0.22,0.20]	-0.10	[-0.51,0.31]
Acceptance rate			0.28		0.55		0.21		0.38	
Convergence after			25,000	draws	29,000	draws	30,000	draws	29,000	draws

The data is quarterly for the sample 1955:Q1–2002:Q1 and it is the same as in Ireland (2004). The output gap is proxied by gross domestic product (GDP) in deviation from a linear trend, inflation is measured as quarter-on-quarter log changes in the consumer price index (CPI), and the nominal interest rate we use is the Federal funds rate. Because output is linearly detrended once and for the whole sample, trend breaks cannot explain the changes we are interested in.

We estimate the model over a number of samples. We start from the [1955:Q1–1974:Q4] sample and then repeat the estimation by moving the starting and ending dates by 4 years so as to keep the size of the window constant at 20 years. Keeping a fixed window size is important for minimizing the differences produced by the varying precision of the estimates. The last subsample is [1983:Q1–2002:Q1], which means that we produce eight posterior distributions for the parameters.

3. Results

Before we describe the estimation results, we plot in Figure 1 the variance (in percentage terms) of the three variables in eight different samples. This plot may help to better explain the reasons of our study and the estimates we obtain.

Three features of Figure 1 are important. First, there is a fall in the variance of inflation only for samples that start no sooner than 1982. Samples that include any year preceding 1982 display a variance that is much higher and roughly unchanged. Second, the variance of the output gap is U-shaped, with the flex point represented by the 1967–1986 sample. This means that, for appropriately selected samples, one can claim that the variability of the output gap has fallen or risen over time (compare, e.g., the 1959–1978 and 1983–2002 samples with 1963–1982 and 1983–2002 samples). In general, the absence of a once-and-for-all break makes the rolling analysis more informative than subsample exercises when studying reasons for the changes. Third, the variance of the nominal interest rate shows an inverted U-shaped pattern. Interesting, the pre-1979 and post-1982 volatilities are almost identical; whereas any sample that includes part or all of the Volker experiment yields a much higher volatility of the nominal interest rate. Once again, our rolling analysis may shed some light for why this pattern emerges.

3.1. Evidence for Subsample Estimation

We start by presenting results for the 1955:Q1–2002:Q1 sample and for three subsamples commonly employed in the literature (1955:Q1–1979:Q2, 1979:Q3–2002:Q1, and 1982:Q4–2002:Q1). We are interested in two issues: we want to assess in which dimension the structural system changes to cope with the time profile of the volatility of output and inflation documented in Figure 1; and to see

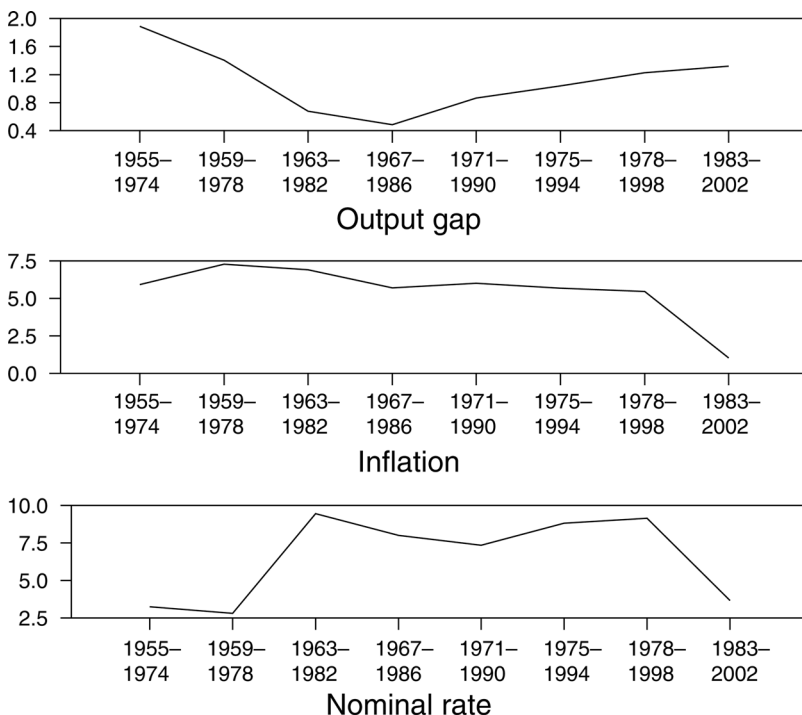


FIGURE 1. Variances in eight different samples.

how distorted inference becomes when potential heterogeneities in the process generating the data are not accounted for.

Table 1 presents the posterior mean and the highest 95% posterior interval (HPI) for each of the parameters in each sample. This measure, which corresponds to classical confidence intervals, tells us where 95% of the mass of the posterior distribution is located. For distributions that are skewed or multimodal, the HPI need not contain the posterior mean (which is precisely the case for certain subsamples) or could have disjoint pieces.

There are several interesting aspects of Table 1 that are worth emphasizing. First, the samples are informative for all parameters of interest. In fact, the location changes and the spreads of the posteriors are smaller than those of the prior. Therefore, the identification problems that Canova and Sala (2006) have highlighted in the context of this model appear to be less dramatic with the selected parameterization. The mean estimate of κ is typically larger than the estimates available in the literature (which are of the order of 0.5). We also can obtain such mean estimates of κ if estimation is performed conditional on $a_{12} = 0$, so one

may conjecture that either misspecification or the impossibility to separate κ and a_{12} is responsible for the differences.

Second, splitting the sample in two changes the point estimates of the policy parameters, with full sample estimates being closer to the 1982–2002 estimates. Cross-sample variations in ψ_x and ψ_r are small or insignificant. However, consistent with the conventional wisdom, the second subsamples are characterized by a higher ψ_π and the differences (at least for the latest subsample) are statistically significant: HPIs for the 1955–1979 and 1982–2002 samples do not overlap.

Third, two of the parameters characterizing the private sector, the risk aversion coefficient φ and the Phillips curve trade-off κ , exhibit considerable changes. The point estimate of φ dramatically drops in the last two subsamples, and the HPIs do not overlap with that of the first sample; the point estimate of κ increases, but the uncertainty around the point estimate is sufficiently large to make changes a-posteriori insignificant.

Fourth, using as the second subsample 1979–2002 or 1982–2002 makes little difference for the point estimates we obtain. However, HPIs do change: Excluding the 1979–1982 period makes the posterior intervals for ψ_π , σ_2 , σ_3 smaller and those of φ and κ larger. Because excluding the Volker experiment from the sample makes information on the location of the private sector parameters weaker, one must conclude that it is the information present in this period that identifies the location of the posterior distribution of these parameters.

Fifth, the covariance matrix of the shocks displays considerable changes. The standard deviation of the shocks to the Euler equation is larger in the first subsample, whereas the standard deviations of the shocks to the other two equations are larger after 1979. Excluding the 1979–1982 period does not change the point estimates of the standard deviations but their HPIs are centered around a lower value if estimation starts at 1982. The covariance terms have HPIs that are entirely on one side of zero for the full sample but generally so within the subsamples. Hence, the statistical correlation one finds in the full sample may be spurious.

Sixth, as Table 2 shows, the model underestimates the variance of the output gap regardless of the sample and it overestimates the variance of inflation by a substantial amount even 10 times in some samples when one uses posterior mean estimates to compute the variabilities implied by the model. Moreover, the actual values of output variability are in the upper tail of the estimated posterior distribution of output variability of any sample, whereas the actual values of the inflation variance are in the very low tail of the estimated posterior distribution of the inflation variance. In other words, the specification is too simple to jointly account for the variability of the two variables. Nevertheless, the model captures the fall in variability across subsamples: When going from the 1955–1979 sample to the 1979–2002 or 1982–2002 sample the estimated variance of output falls by half (from 1.61 to 0.60 or 0.80) and the estimated variance of inflation falls by about two-thirds (from 47.9 to 12.7 or 16.3).

Campbell and Hercowitz (2006) have suggested that changes in the credit constraints faced by consumers in the early 1980s could account for the fall in inflation and output volatilities observed after that date. In their model, volatility drops because labor supply (and therefore real activity) is extremely sensitive to shocks when credit constraints are binding but much less so when constraints are relaxed. In our model labor supply decisions are absent; hence such an effect is unmeasured. Nevertheless, changes in the risk aversion coefficient could play a similar role. In Section 5 we study whether variations in the elasticity of the output gap to real interest rate changes can account for part of the volatility changes.

Arias, Hansen, and Ohanian (2006) have argued that, to account for the fall in output volatility, one need not change the parameters of the model across subsamples, but instead simply allow the variance of the Solow residuals to be reduced over time. Comparisons are difficult because we use a different model, but our results seem to tell a different story. Given that the parameters of the private sector have changed, the variance of the Phillips curve shock increases rather than decreases after 1979 to fit the evidence.

The increase in the standard deviation of the shock to the interest rate equation may appear odd. However, one should be careful when comparing our estimates to those present in the literature, because disturbances in the model do not necessarily have a structural interpretation. We ran a VAR with the same three variables, computed the standard deviation of the reduced-form residuals of the interest rate

TABLE 2. Data and estimated posterior variabilities.

	1955–2002	1955–1979	1979–2002	1982–2002
$\text{Var}^d(y)$	5.76	2.16	1.79	1.38
$\text{Var}^d(\pi)$	5.95	6.92	4.75	1.06
			Basic rule	
Mean posterior $\text{Var}(y)$	3.55	1.61	0.60	0.80
Percentile where $\text{Var}^d(y)$ lies	83	80	89	90
Mean posterior $\text{Var}(\pi)$	74.13	47.90	12.72	16.39
Percentile where $\text{Var}^d(\pi)$ lies	01	05	13	02
			Output growth rule	
Mean posterior $\text{Var}(y)$	0.20	0.18	1.25	1.08
Percentile where $\text{Var}^d(y)$ lies	100	99	80	82
Mean posterior $\text{Var}(\pi)$	7.50	3.50	26.92	31.84
Percentile where $\text{Var}^d(\pi)$ lies	34	97	03	02
			Forward rule	
Mean posterior $\text{Var}(y)$	4.34	0.27	0.11	0.11
Percentile where $\text{Var}^d(y)$ lies	72	97	98	99
Mean posterior $\text{Var}(\pi)$	13.44	5.47	4.12	2.81
Percentile where $\text{Var}^d(\pi)$ lies	17	60	54	28

equation in the three subsamples, and found a similar pattern. Hence, to fit the time path for the endogenous variables, we need a combination of changes in the parameters of the model. Given the pattern for the variance of interest rates presented in Figure 1, this combination must include an increase in the volatility of the residuals of the interest rate equation.

Overall, the analysis of this section has highlighted two important conclusions. First, changes in the parameters of the policy rule do occur but their magnitude is smaller than usually claimed in the literature. Second, coefficients describing the private sector and the standard deviation of the shocks display large and significant changes. Next we examine whether these conclusions remain valid when posterior distributions are obtained over rolling samples with a window of 20 years.

3.2. Evidence from Rolling Estimation

We have argued that arbitrarily breaking the sample in two is less than ideal for what we want to investigate. Two reasons make the results of that approach potentially difficult to interpret. First, using fixed subsamples forces all the relationships of the model to break at the same date—clearly violating what we have displayed in Figure 1—and this may induce important biases. Second, the patterns displayed by the level and variability of output and interest rates do not fit well into the null of stability nor the alternative of a permanent jump. Therefore, the conclusions drawn would be highly sensitive to the choice of break date. Our rolling estimation approach does not entirely solve these problems. Accounting for them would require estimation techniques that allow structural parameters to be fully time varying. Nevertheless, by comparing posterior estimates over different samples, we can provide a more robust characterization of the changes observed over the last 35 years than by simply using (fixed) subsample analysis.

In Figure 2, we plot the posterior mean (straight line) and HPI estimates (dashed lines) obtained in the eight subsamples for the parameters of interest. The figure confirms and qualifies the conclusions we have previously reached. There are considerable variations in both the posterior mean and the posterior HPI for the coefficient of relative risk aversion φ over samples. Although variations are present in samples that include years before 1982, it is only after that date that the fall becomes considerable and significant. The Phillips curve trade-off κ is increasing over time in a manner consistent with the previous analysis: The trend is clear but the posterior significance of the changes is low.

The Phillips curve trade-off in more structural versions of the model that we consider is typically regulated by a (nonlinear) function of four parameters: The coefficient of relative risk aversion, the inverse elasticity of labor supply, the discount factor, and the price stickiness parameter. Because, at least in micro

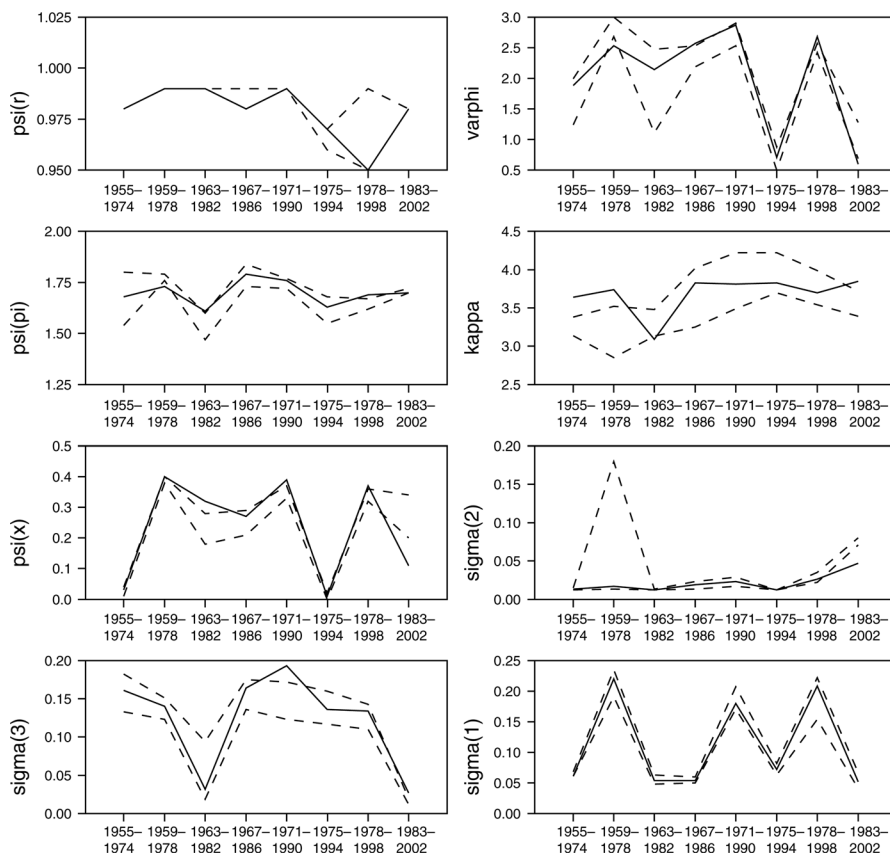


FIGURE 2. Posterior mean estimates and HPI, eight different samples.

studies, there is little evidence that the price stickiness parameter has changed over time, and because estimates of the discount factor do not display important variations over the samples, one must conclude that variations in the intertemporal elasticity of labor supply must counteract variations in the risk aversion coefficient and so produce the mild trend we observe in κ . It is tempting to associate this trend with the changes that the U.S. labor market experienced over the period (higher female participation, larger number of migrant workers, etc.). However, one should realize that more general specifications of the model, for example, with decreasing returns to scale production, produce more complicated Phillips curve trade-off where additional parameters enter. Rather than asking our model to explain features it was not designed to accommodate, we leave for future research the question of what drives the trend in κ .

The parameters of the policy rule display minor variations across samples: HPIS for different samples almost always overlap, except for the coefficients on the output gap. Consistent with the analysis of Bernanke and Mihov (1998), the patterns evident in Figure 2 square well with the notion that none of the three policy coefficients has permanently shifted over time. Also, consistent with Canova and Gambetti (2007a), our recursive posterior analysis shows that the policy rule during the tenures of Burns and Greenspan were not too different. Assuming that the policy rule represents the actual policy well over the entire period, HPIS for the policy coefficients in the earlier and the later samples overlap.

Fourth, the standard deviation of two of the three disturbances (σ_2 , σ_3) display considerable variations over subsamples. Because the covariance parameters (not shown here) also exhibit this feature, it is the entire covariance structure of the disturbances that is significantly altered over time.

Given these results, the temptation is strong to associate variations in the variance of the output gap over time with changes in φ and to associate variations in the variances of inflation and interest rates with changes in the covariance structure of the disturbances. In order to make the link more transparent, in Section 5 we estimate restricted systems in which certain parameters are fixed at their 1955–1974 mean value. By comparing restricted and unrestricted estimates, one can obtain a formal indication of what parameters contributed most to the variations in the variance of output and inflation.

3.3. Robustness

The model we employ is rather standard. However, the specification of its details may be subject to some debate. In particular, although we have chosen to work with a policy rule where the nominal interest rate depends on the current output gap and current inflation, a policy rule specified in terms of current output growth and current inflation is probably equally reasonable. Furthermore, some literature (see, e.g., Clarida, Gali, and Gertler 2000; Boivin and Giannoni 2006) specify a forward-looking rule where the current interest rate responds to future expected changes in the output gap and in inflation. Would our main conclusions change if one of these alternative rules were used? Evidence on this issue is in Tables 3 and 4, which report for two alternative rules posterior means and HPIS for the full sample and the three subsamples presented in Table 1. Results obtained using rolling samples are comparable and not presented.

It is well known that the statistical output gap proxy we use is subject to a large amount of measurement error. Consequently, estimates of the structural parameters may fail to have the correct magnitudes because a large amount of measurement error is present in each sample. We have already argued that model-based measures of the output gap are not viable because they are typically obtained

disregarding the role of the capital stock. Because Orphanides (2004) has argued that measurement errors are significantly reduced if output growth is used in place of the output gap, it is worth investigating what happens to our estimates when we employ this new policy equation.

Table 3 indicates that the coefficient of relative risk aversion φ still falls but now the fall is much more limited in magnitude. Nevertheless, HPIs for the 1955–1979 and 1982–2002 samples do not overlap. Estimates of the Phillips curve trade-off κ increase (as with the output gap measure), but now the magnitude of the increase is much larger and the HPIs for the 1955–1979 and 1982–2002 samples do not overlap. As a consequence of these changes, the standard deviation of the shocks to the first two equations shows a pattern that is opposite to the one in Table 1: The standard deviation of the disturbance to the Euler equation slightly increases, whereas that of the Phillips curve decreases. Surprisingly, the coefficient on inflation in the policy rule falls when we move from the pre-1979 to the post-1979 samples and the fall is significant. Taken at face value, this implies that the policy rule has become less aggressive in responding to inflation since the beginning of the 1980s. This may be due to inflation expectations being much less volatile in the 1980s (thereby necessitating a smaller coefficient for inflation stabilization); however, one should also recognize that differences across samples may reflect model misspecifications. The time profile of the standard deviation of the disturbance of the interest rate equation suggests that this is probably true. Hence, despite their large size, changes in the policy parameters account for little of the variation in interest rates.

Table 3 indicates that even with the output growth rule the model has a hard time mimicking the variability of the output gap and inflation, regardless of whether we use mean estimates or the percentiles where the actual values lie. As in the previous case, the variance of the output gap is underestimated and the variance of inflation is typically overestimated but here the magnitude of the discrepancy is larger with the former than with the latter. However, the estimates we obtain imply no volatility moderation.

With a policy rule that reacts to expected changes in the output gap and inflation, results are roughly similar. As seen in Table 4, φ falls as we move from the earlier to the later part of the sample, but changes are smaller and the HPI of different samples overlap. Here κ increases over time, making the Phillips curve trade-off flatter; and changes are a posteriori significant. φ_π increases in the later part of the sample and the increase is now significant a posteriori. With this specification of the policy rule, the mean estimate is on the boundary of the stability region in the first subsample, suggesting that the likelihood of the data is quite sensitive to the specification of the policy rule. Once again, changes in the parameters of the policy rule account for little of the variations in the nominal interest rates; evidently changes in the standard deviation of the disturbance seem to do most of the work in matching the time path of the variance of interest rates.

TABLE 3. Prior and posterior moments, output growth rule.

	Prior		Posterior											
			1955–2002			1955–1979			1979–2002			1982–2002		
	Mean	Std.	Mean	HPI	draws	Mean	HPI	draws	Mean	HPI	draws	Mean	HPI	draws
φ	2.0	0.35	3.05	[2.58,3.56]	20,000	3.02	[2.61,3.22]	20,000	2.61	[2.02,3.06]	21,000	2.27	[1.80,2.52]	18,000
β	0.98	0.01	0.98	[0.96,0.99]	21,000	0.98	[0.97,0.99]	21,000	0.98	[0.97,0.99]	21,000	0.97	[0.98,0.99]	18,000
κ	2.0	1.0	4.11	[4.02,4.22]	20,000	2.66	[2.28,3.17]	20,000	2.97	[2.58,3.38]	21,000	3.90	[3.29,3.91]	18,000
φ_r	0.8	0.25	0.98	[0.98,0.99]	20,000	0.99	[0.99,0.99]	20,000	0.99	[0.99,0.99]	21,000	0.99	[0.99,0.99]	18,000
φ_x	0.5	0.25	0.52	[0.14,0.99]	21,000	0.47	[0.08,0.81]	21,000	0.42	[-0.10,0.86]	21,000	0.52	[0.12,1.01]	18,000
φ_π	1.3	0.5	1.42	[1.32,1.43]	20,000	2.22	[1.87,2.64]	20,000	1.22	[1.11,1.34]	21,000	1.15	[1.07,1.22]	18,000
ρ_x	0.85	0.25	0.88	[0.88,0.90]	20,000	0.83	[0.80,0.84]	20,000	0.76	[0.75,0.78]	21,000	0.71	[0.68,0.73]	18,000
ρ_p	0.85	0.25	0.98	[0.98,0.99]	20,000	0.99	[0.99,0.99]	20,000	0.99	[0.99,0.99]	21,000	0.99	[0.99,0.99]	18,000
σ_1	0.01	0.50	0.012	[0.012,0.012]	21,000	0.042	[0.040,0.045]	21,000	0.037	[0.029,0.045]	21,000	0.046	[0.037,0.055]	18,000
σ_2	0.01	0.50	0.052	[0.047,0.060]	21,000	0.019	[0.015,0.020]	21,000	0.020	[0.016,0.024]	21,000	0.015	[0.013,0.019]	18,000
σ_3	0.01	0.50	0.122	[0.106,0.130]	21,000	0.062	[0.051,0.086]	21,000	0.101	[0.085,0.118]	21,000	0.095	[0.081,0.109]	18,000
a_{12}	0.0	0.25	1.39	[0.88,1.59]	21,000	1.52	[1.42,1.59]	20,000	1.47	[1.36,1.59]	21,000	1.42	[1.30,1.59]	18,000
Acceptance rate			0.35			0.25			0.43			0.36		
Convergence after			21,000	draws	20,000	21,000	draws	21,000	draws	21,000	draws	18,000	draws	18,000

TABLE 4. Prior and posterior moments, forward rule.

	Prior		Posterior											
			1955–2002			1955–1979			1979–2002			1982–2002		
	Mean	Std.	Mean	HPI	Mean	HPI	Mean	HPI	Mean	HPI	Mean	HPI		
φ	2.0	0.35	2.56	[1.81,2.90]	2.36	[2.61,3.22]	2.71	[2.63,2.93]	2.05	[1.62,2.74]	0.98	[0.97,0.99]		
β	0.98	0.01	0.98	[0.97,0.98]	0.97	[0.97,0.99]	0.98	[0.98,0.98]	0.98	[0.97,0.99]	0.98	[0.97,0.99]		
κ	2.0	1.0	3.27	[2.79,3.67]	2.93	[2.28,3.17]	3.97	[3.88,4.22]	3.48	[3.18,3.66]				
φ_r	0.8	0.25	0.98	[0.98,0.99]	0.99	[0.99,0.99]	0.99	[0.99,0.99]	0.99	[0.99,0.99]	0.99	[0.99,0.99]		
φ_x	0.5	0.25	0.21	[0.06,0.28]	-0.01	[0.08,0.81]	0.32	[0.29,0.39]	-0.01	[-0.02,0.00]				
φ_π	1.3	0.5	1.34	[1.00,1.46]	1.02	[1.00,1.05]	1.69	[1.65,1.72]	1.52	[1.24,1.76]				
ρ_x	0.85	0.25	0.93	[0.93,0.95]	0.95	[0.94,0.95]	0.91	[0.90,0.91]	0.91	[0.89,0.92]				
ρ_p	0.85	0.25	0.99	[0.99,0.99]	0.99	[0.99,0.99]	0.96	[0.96,0.97]	0.98	[0.98,0.98]				
σ_1^2	0.01	0.50	0.027	[0.020,0.032]	0.012	[0.012,0.013]	0.027	[0.022,0.030]	0.012	[0.012,0.014]				
σ_2^2	0.01	0.50	0.028	[0.018,0.052]	0.061	[0.063,0.063]	0.161	[0.128,0.193]	0.066	[0.049,0.068]				
σ_3^2	0.01	0.50	0.015	[0.012,0.015]	0.026	[0.024,0.031]	0.046	[0.041,0.055]	0.035	[0.032,0.043]				
a_{12}	0.0	0.25	-1.05	[-1.59,-0.66]	0.13	[-0.37,0.03]	0.07	[0.03,0.07]	0.10	[0.02,0.07]				
a_{13}	0.0	0.25	1.25	[0.79,1.46]	0.13	[0.25,0.35]	0.13	[0.10,0.13]	0.10	[0.10,0.14]				
Acceptance rate			0.50		0.25		0.30		0.43					
Convergence after			22,000	draws	21,000	draws	19,000	draws	19,000	draws				

A specification with a forward rule appears to be better in matching inflation variability than the original one but worse in matching output variability (see Table 2). Moreover, although the forward specification can reproduce the fall in the variances of the two variables in the last two subsamples, the fall in inflation in the 1982–2002 sample is small relative to the one actually observed in the data.

Overall, these alternative policy rules produce results that are qualitatively similar to those obtained in the baseline case as far as trends in crucial parameter estimates are concerned. However, they seem to face more important problems in matching either the level or the decline in the volatility of output and inflation over the subsamples.

4. A Comparison with the Literature

Our findings may seem puzzling relative to what it is currently available in the literature, except perhaps for Gordon (2005). It is therefore worthwhile to discuss how our results are different and what can account for these differences.

To start with, we would like to point out three facts. First, our structural estimation does find an increase in the inflation coefficient of the policy rule when moving from a sample including the 1970s to a sample that excludes those years. What we show is that the variations are not statistically large relative to variations in other parameters. Second, time variations in parameters other than the policy ones are often detected when the model is estimated using system-wide methods (see, e.g., Ireland 2001; Boivin and Giannoni 2006), but they are left undiscussed. Third, direct structural estimation typically leads to conclusions that are different from those obtained by estimating structural VARs with or without time-varying coefficients; the former finds changes mainly in the parameters of the model, the latter mainly in the covariance matrix of shocks. In some cases this occurs because variations in the standard deviations of the shocks cannot be identified with the chosen objective function (see, e.g., Boivin and Giannoni 2006); in other cases, because specification choices impose a particular structure on the estimated structural shocks. Our findings, which are obtained conditioning on a model, are consistent with the VAR evidence.

In contrast to Clarida, Galí, and Gertler (2000), who use single-equation structural estimation, we take a system-wide estimation approach and use Bayesian rather than classical techniques. Although the second difference may be of minor importance because the priors are sufficiently noninformative over the ranges we choose, the first one is important. Single-equation methods may produce a distorted view of the structural relationships when important endogeneities are present (see also Lubik and Schorfheide 2004). In addition, they neither take system-wide relationships into account, nor do they use the cross-equation restrictions present in the model. Hence, single-equation methods are inefficient.

In comparison with Boivin and Giannoni (2006), who use a minimum distance estimator to obtain parameter estimates, our approach has the advantage of enabling a better identification of the structural relationships. Canova and Sala (2006) have shown that minimum distance estimators, when used to derive the parameters of a New Keynesian model from responses to monetary shocks, face severe identification problems: The objective function is very flat and ridges are present. This means that variations in the coefficients identified by this procedure could be points of equivalent height on this surface or could represent variations linked to variations in the other parameters. On the contrary, the likelihood function of the system is much more peaked and displays relationships among the parameters that are more easily disentangled (see also Linde 2005). Two additional reasons may explain the different findings.

First, Boivin and Giannoni (2006) adjust the estimated specification in order to achieve the best possible fit—endowing the theoretical model with ad hoc exogenous frictions and searching among the (forward) specifications of the policy rule for the one that best fit the interest rate data—whereas we take a textbook specification and do no preliminary data mining exercises. Table 4 shows that it is possible to roughly reproduce the pattern of point estimates they obtain with a one-period forward-looking rule and no ad-hoc frictions. However, Boivin and Giannoni neglect the fact that pretesting downsizes the standard errors of their estimates. Thus, changes that are a posteriori insignificant may look artificially significant. Second, the counterfactual exercises of Boivin and Giannoni are subject to the Lucas critique—agents behave as if there will never be a structural break and then when the break occurs, they learn immediately that there will never be a break in the future—the exercises we conduct in Section 5 are largely free of these problems. As a matter of fact, the majority of the counterfactual exercises performed in the literature suffer from various types of inconsistencies that make results uninterpretable. For example, the practice of switching coefficients and variances across samples does not take into account the correlation structure of estimates and the fact that the parameters/variances estimates obtained in a sample may be in the tails of the estimated distribution of parameters/variances estimates in another sample.

With regard to Lubik and Schorfheide (2004), who also employ system-wide methods and Bayesian estimation on a model similar to ours, two important differences should be mentioned. First, the policy rule they estimate uses output growth. As shown in Table 2, this choice has some consequences for the results but does not change the main features of the conclusions. The second difference is that they allow for indeterminacy (and sunspots) in the estimation, whereas we don't. Consequently, our work is a complement to, rather than a substitute for, theirs.

Finally, several papers have estimated structural VARs with or without time variations in the coefficients (see, e.g., Cogley and Sargent 2001, 2005; Canova

and Gambetti 2007a; Sims and Zha 2006). Although most are concerned with estimates of the policy rules and of the monetary policy shock, some papers have tried to estimate sources of variations in systems that have similarities to the model in Section 2 (see, e.g., Gambetti, Pappa, and Canova 2008). Our results help explain some of their findings. For example, the large impact that supply shocks have in explaining the time profile of the volatility and persistence of output is consistent with the time profile of the estimates of the coefficient of relative risk aversion in the log-linearized Euler equation. Moreover, the fall in the standard deviation of the supply shocks over time that they find is due, in part, to the change in the Phillips curve trade-off κ .

5. What Change Explains the Great Moderation?

The analysis so far has documented the presence of generalized parameter instabilities over the samples under consideration and has shown that variations in only certain parameters are statistically significant. In this section we ask which of these changes contributes most to the changes in the variance of output and inflation documented in Figure 1. In particular, suppose we repeat estimation over the 1983–2002 subsample while fixing some parameters at their 1955–1974 posterior mean estimates. Would the fit of the model change? Would the model reproduce the fall in the volatility of output and inflation observed in the sample?

When examining which feature of the model is responsible for the Great Moderation, one typically performs counterfactual exercises in which parameters for different subsamples are switched and interesting statistics are recomputed under these alternative parameter values. As we have mentioned, although prevalent in the literature (see, e.g., Stock and Watson 2002; Boivin and Giannoni 2006), these exercises cannot credibly answer the question that concerns us. Our approach, which allows unrestricted parameters to be readjusted in the estimation, can provide a more reasonable scenario for evaluating the economic consequences of parameter changes.

Table 5 reports estimates of the variance of the output gap and inflation obtained in the unrestricted specification and in three restricted specifications for which the parameters describing the private sector behavior (φ, κ), the policy rule ($\varphi_x, \varphi_r, \varphi_\pi$), and the standard deviations of the shocks ($\sigma_1, \sigma_2, \sigma_3$) are in turn constrained to have a prior mean equal to the posterior mean of the 1955–1974 subsample and a small variance (0.0001). Table 5 also reports the posterior probability of each model and the risk of matching the variance of output and inflation of the unrestricted model with each restricted specification.

Posterior probabilities are computed using the prior probability of each restricted specification (set to one-third) and their marginal likelihood. The marginal likelihood is a synthetic measure of fit that is comparable to \bar{R}^2 in linear models: A higher marginal likelihood obtains if a model fits the data better,

TABLE 5. Posterior moments, probabilities and risk. Restricted and unrestricted specifications, 1983–2002 sample.

	Unrestricted	Restricting φ, κ	Restricting $\psi_x, \psi_r \varphi_\pi$	Restricting σ_j^2
$\text{Var}^d(y)$	1.38			
$\text{Var}^d(\pi)$	1.06			
$\text{Var}(y)$	0.80	0.05	2.98	4.57
$\text{Var}(\pi)$	16.39	1.27	47.04	14.97
Posterior				
Probability		0.999	3.0e-23	2.7e-82
Risk		21.19	2.0e-21	1.2e-81

given a common prior, or if the prior of one model is closer to the likelihood, given a common likelihood. Because the experiments we conduct involve changing both the likelihood and the prior of the parameters, the marginal likelihood is altered through both channels. We compute marginal likelihoods using a modified harmonic mean estimator and 10 chains of parameter draws (see, e.g., Geweke 1998).

The risk measure is computed by comparing the volatilities of output and inflation produced by each restricted specification to the ones of the unrestricted specification under an absolute loss function, equally weighting the two volatilities by the posterior probability of each restricted specification. This type of measure, popularized in Schorfheide (2000), is useful for comparing models that are likely to be misspecified and thus may have very low posterior probability. We also computed a risk measure using two alternative loss functions: a quadratic loss function and a loss that asymmetrically weights only positive deviations from the volatilities of the basic specification. The results we present are robust with respect to these choices. To interpret the risk measure, observe that if time variations in one set of parameters are relatively unimportant (important), then the posterior probability of the restricted specification will be high (low) and the risk relatively high (low).

Table 5 indicates that variations in the parameters of the private sector induce changes that are opposite to those we are hoping for. In fact, had we kept them fixed at the posterior mean value estimated over the 1955–1974 sample, then the fall in the variance of output and inflation implied by the model would have been much larger than in the unrestricted case. As a result, changes in φ and κ cannot be the drivers of the Great Moderation. Restricting the parameters of the policy rule to their 1955–1974 posterior mean values implies that the variability of output and inflation would have counterfactually increased rather than decreased over the 1982–2002 sample. Hence, the Great Moderation would not have occurred if policy parameters were invariant over the sample. Finally, restrictions on the volatility of the shocks have minor effects on the variance of inflation but considerable effects on the variance of output. Thus, the fall in the variances of output

and inflation may have distinct causes. Output volatility declines owing to a combination of causes, of which the most important is the fall in the standard deviation of the shocks. The fall in the variance of inflation, on the other hand, appears to be largely due to changes in the parameters of the policy rule.

How can one explain the extreme posterior probabilities of Table 5? To start with, one should observe that Euler equation shocks are those with the lowest variability in all the samples. Therefore, given a shock of this type, changing the coefficient of relative risk aversion produces only small changes in the volatility of the output gap and, given the changes in the Phillips curve trade-off, this implies small variations in the volatility of inflation. In contrast, small variations in the coefficients of the policy rule imply considerably different covariance matrices of the shocks and therefore large effects on the volatility of output and inflation. Finally, fixing the standard deviation of the shocks forces the parameters of the policy rule to change dramatically (for example, the output coefficient rises from 0.11 in the unrestricted specification to 0.99 when we fix the standard deviation of the shocks), and this has important consequences for the volatility of output produced by the model.

In sum: Changes in the parameters of the policy rule and the variability of the shocks are crucial to understanding the Great Moderation; relatively speaking, posterior probabilities and risk measures suggest that time variations in the standard deviation of the shocks are the most important cause of the observed variations.

The results of this section should be seen as a warning against taking the results of statistical estimation at face value. Variations that are statistically large may produce only small economic consequences and small statistical variations like those experienced in the parameters of the policy rule may generate significant economic implications because of their effects on the covariance structure of the shocks.

6. Conclusions

This paper recursively estimates a conventional small-scale DSGE model using U.S. post-WWII data and Bayesian techniques. The model belongs to the class of New-Keynesian structures that have been extensively used in the current literature for welfare and policy analyses. Bayesian techniques are preferable to standard likelihood methods or to indirect inference (impulse response matching) exercises, especially for models like the one we consider, which are clearly false and misspecified. We show that the model and the methodology are useful tools for understanding the nature of the changes responsible for the Great Moderation episode.

We estimate the model a number of times, using a different starting date but keeping the window size fixed, and analyze the role of changes in the private sector

parameters, in the coefficients of the policy rule, and in the covariance structure of the shocks. We find that changes over time in the parameters of the private sector are the largest and the most significant: They tend to make the output gap more elastic in response to changes in the real rate and to make inflation more reactive to marginal costs. Changes in the covariance structure of the shocks are also considerable, whereas changes in the coefficients of the policy rule are small and a posteriori insignificant.

Nevertheless, when we analyze which of these changes explain best the Great Moderation episode, we find that the changes in the parameters of the private sector alone cannot generate the observed fall in the variance of output and inflation whereas changes in the parameters of the policy rule and the covariance of the shocks can. We also show that the fall in variances of output and inflation appear to have different causes, suggesting that the quest for one common explanation to both facts is probably misplaced.

These results stand midway relative to those in the literature. As in the structural VAR analyses (Primiceri 2005; Sims and Zha 2006; Canova and Gambetti 2007a), we find evidence that the shocks hitting the economy have considerably changed over time. Also, consistent with the analyses of McConnell and Perez Quiroz (2000), Gordon (2005), and Campbell and Hercowitz (2006), we detect statistical changes in the parameters of the private sector but these changes do little to explain the Great Moderation. Finally, although policy parameters change little, they seem to matter quite a lot.

Our work has a number of limitations which we would like to spell out in detail. As we have mentioned, our analysis imposes the restriction that only a determinate equilibrium is present within each sample. This is relatively common in the literature (see, e.g., Rabanal and Rubio Ramirez 2005; or Fernandez Villaverde and Rubio Ramirez 2007) and, for the rolling analysis we perform, the restriction is probably less important than one would initially think. An obvious extension of what we have done here would be to allow for indeterminacies in every subsample and then to check whether rolling analysis would confirm or disprove our conclusions.

Second, although our estimation approach is convenient, it imposes a form of irrationality on agents' behavior. In fact, the analysis implicitly assumes that agents have rational expectations within each sample where estimation is conducted but not over the entire sample: Agents never take into account that changes in the structural parameters may occur. In order to fully take this into account, one would need to employ the techniques recently developed by Fernandez Villaverde and Rubio Ramirez (2007), which use higher-order approximations to agents' decision rules and more complicated Monte Carlo techniques. This option, however, requires considerable computational time even in a model with only three equations.

Third, as we have argued in the Introduction, the model is taken “off-the-shelf” and is not optimized to fit the data in any sense. Hence, there is always the possibility that results are driven by misspecification, omitted variables, or shocks. To fully understand the sources of the Great Moderation, one should probably employ a larger-scale model that fits the data better than the simple specification we consider. Such an extension is relatively straightforward to undertake but again would require considerable computational time.

Finally, although it is common to look at the U.S. and only at output and inflation, there are obvious reasons to ask whether other variables display similar behavior and whether common explanations for the international patterns documented (e.g., in Stock and Watson 2002 or Canova et al. 2007), could be found. A cross-country perspective would be fundamental to understanding the source of variations because we know quite a bit about policy changes and the dates at which they occurred in countries other than the United States. We leave all these issues for future research.

References

- An, Sungbae, and Frank Schorfheide (2007). “Bayesian Analysis of DSGE Models.” *Econometrics Reviews*, 26, 113–172.
- Arias, Andres, Gary Hansen, and Lee Ohanian (2006). “Why Have Business Cycle Fluctuations Become Less Volatile?” NBER Working Paper No. 12079.
- Bernanke, Ben, and Ian Mihov (1998). “Measuring Monetary Policy.” *The Quarterly Journal of Economics*, 113, 869–902.
- Blanchard, Olivier, and John Simon (2000). “The Long and Large Decline in U.S. Output Volatility.” *Brookings Papers on Economic Activity*, 1, 135–147.
- Boivin, Jean, and Marc Giannoni (2006). “Has Monetary Policy Become Less Powerful?” *Review of Economics and Statistics*, 88(3), 445–462.
- Campbell, Jeffrey, and Zvi Hercowitz (2006). “The Role of Collateralized Household Debt In Macroeconomic Stabilization.” Working paper, Federal Reserve Bank of Chicago.
- Canova, Fabio, and Luca Gambetti (2007a). “Structural Changes in US Economy: Is There a Role for Monetary Policy?” *Journal of Economic Dynamics and Control*, 33, 477–490.
- Canova, Fabio, and Luca Gambetti (2007). “Do Inflation Expectations Matter? The Great Moderation Revised?” Available at <www.econ.upf.edu/crei/people/canova>.
- Canova, Fabio, and Luca Sala (2006). “Back to Square One: Identification Issues in DSGE Models.” Available at <www.econ.upf.edu/crei/people/canova>.
- Clarida, Richard, Jordi Galí, and Mark Gertler (1999). “The Science of Monetary Policy: A New-Keynesian Perspective.” *Journal of Economic Literature*, 37, 1661–1707.
- Clarida, Richard, Jordi Galí, and Mark Gertler (2000). “Monetary Policy Rule and Macroeconomic Stability: Evidence and Some Theory.” *Quarterly Journal of Economics*, CXV, 147–180.
- Cogley, Timothy, and Thomas Sargent (2001). “Evolving Post-World War II U.S. Inflation Dynamics.” *NBER Macroeconomic Annual*, 16, 331–373.
- Cogley, Timothy, and Thomas Sargent (2005). “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.” *Review of Economic Dynamics*, 8, 262–302.
- Fernandez Villaverde, Jesus, and Juan Rubio Ramirez (2007). “How Structural are Structural Parameters?” forthcoming, *NBER Macroeconomic Annual*.

- Gambetti, Luca, Paraskevi Pappa, and Fabio Canova (2008). "The Structural Dynamics of US Output and Inflation: What Explains the Changes?" *Journal of Money, Credit and Banking*, 40, 369–388.
- Geweke, John (1998). "Using Simulation Methods for Bayesian Econometric Models: Inference, Development and Communication." Federal Reserve Bank of Minneapolis Staff Report 249.
- Gordon, Robert (2005). "What Explains the Decline in US Business Cycle Volatility?" NBER Working Paper No. 11777.
- Hanson, Michael (2006). "Varying Monetary Policy Regimes: A Vector Autoregressive Investigation." *Journal of Business and Economics*, 58 (5–6), 407–427.
- Ireland, Peter (1999). "Does the Time Consistent Problem Explain the Behavior of Inflation in the United States?" *Journal of Monetary Economics*, 44, 279–292.
- Ireland, Peter (2001). "Sticky Price Models of the Business Cycle: Specification and Stability." *Journal of Monetary Economics*, 47, 3–18.
- Ireland, Peter (2004). "Technology Shocks in the New Keynesian Model." *Review of Economics and Statistics*, 86, 923–936.
- Leeper, Eric, and Tao Zha (2003). "Modest Policy Interventions." *Journal of Monetary Economics*, 50, 1673–1700.
- Linde, Jesper (2005). "Estimating New Keynesian Phillips Curve: A Full Information Maximum Likelihood." *Journal of Monetary Economics*, 52, 1135–1149.
- Lubik, Thomas, and Frank Schorfheide (2004). "Testing for Indeterminacy: With an Application to US Monetary Policy." *American Economic Review*, 94, 190–217.
- McConnell, Michelle, and Gabriel Perez Quiroz (2000). "Output Fluctuations in the US: What Has Changed Since the Early 1980s?" *American Economic Review*, 90, 1464–1476.
- Orphanides, Athanasios (2004). "Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches." *Journal of Money, Banking and Credit*, 36, 151–175.
- Pivetta, Fredric, and Ricardo Reis (2007). "The Persistence of Inflation in the US." *Journal of Economic Dynamics and Control*, 31, 1326–1358.
- Primiceri, Giorgio (2005). "Time Varying Structural VAR and Monetary Policy." *Review of Economic Studies*, 72, 453–472.
- Rabanal, Pau, and Juan Rubio Ramirez (2005). "Comparing New Keynesian Models of the Business Cycle: A Bayesian Approach." *Journal of Monetary Economics*, 52, 1150–1169.
- Sargent, Thomas (1999). *The Conquest of American Inflation*. Princeton University Press.
- Schorfheide, Frank (2000). "Loss Function Based Evaluation of DSGE Models." *Journal of Applied Econometrics*, 15, 645–670.
- Sims, Christopher, and Tao Zha (2006). "Were There Regime Switches in US Monetary Policy?" *American Economic Review*, 96, 54–81.
- Smets, Frank, and Raf Wouters (2003). "An Estimated DSGE Model for the Euro Area." *Journal of the European Economic Association*, 1, 1123–1175.
- Stock, James, and Mark Watson (2002). "Has the Business Cycle Changed and Why?" *NBER Macroeconomic Annual*, 17, 159–218.
- Taylor, John (1998). "Monetary Policy Guidelines for Unemployment and Inflation Stability." In *Inflation, Unemployment and Monetary Policy*, edited by R. Solow and J. Taylor, NBER.