

LEARNING AND TIME-VARYING MACROECONOMIC VOLATILITY

FABIO MILANI

UNIVERSITY OF CALIFORNIA, IRVINE

ABSTRACT. This paper presents a DSGE model in which agents' learning about the economy can endogenously generate time-varying macroeconomic volatility. Economic agents use simple models to form expectations and need to learn the relevant parameters. Their gain coefficient is endogenous and is adjusted according to past forecast errors.

The model is estimated using likelihood-based Bayesian methods. The endogenous gain is jointly estimated with the structural parameters of the system.

The estimation results show that private agents appear to have often switched to constant-gain learning, with a high constant gain, during most of the 1970s and until the early 1980s, while reverting to a decreasing gain later on. As a result, the model can generate a pattern of volatility, which is increasing in the 1970s and falling in the second half of the sample, with a decline that can roughly match the magnitude of the "Great Moderation". The paper also documents how a failure to incorporate learning into the estimation may lead econometricians to spuriously find time-varying volatility in the exogenous shocks, even when these have constant variance by construction.

Keywords: adaptive learning, constant gain, monetary policy, macroeconomic volatility, inflation dynamics.

JEL classification: C11, D84, E30, E50, E52, E58, E66.

Date: This version: May, 2007. First draft: July, 2006.

I would like to thank Bill Branch, Jim Bullard, George Evans, Oscar Jordà, participants at the 'Learning Week' Conference in St. Louis, MO, the Missouri Economics Conference at the University of Missouri, Columbia, and seminar participants at the University of California, Davis and the Federal Reserve of St. Louis for comments and suggestions.

Address for correspondence: Department of Economics, 3151 Social Science Plaza, University of California, Irvine, CA 92697-5100. *Phone:* 949-824-4519. *Fax:* 949-824-2182. *E-mail:* fmilani@uci.edu. *Homepage:* <http://www.socsci.uci.edu/~fmilani>.

1. INTRODUCTION

Several recent studies have documented large changes in the volatility of macroeconomic fluctuations in the US over the post-war period. Kim and Nelson (1999), McConnell and Pérez-Quirós (2000), and Stock and Watson (2002), among several others, have identified a large decline of output growth volatility in the post-1984 sample compared to the previous two decades (the large shift in volatility is commonly referred to as “The Great Moderation”). The reduction in volatility is apparent if one looks at simple measures as the variances of output growth and inflation in the 1950-1980 versus the 1980-2007 samples. Slightly more sophisticated approaches give a similar message: Figure 1, for example, shows the conditional standard deviations from GARCH models for inflation and output gap over time.¹ The conditional standard deviations for both series increase in the 1970s and substantially decline after the early 1980s.

Correctly modeling changes in volatility has been shown to be important for understanding macroeconomic fluctuations. Sims and Zha (2006) find that incorporating regime changes in the volatilities of disturbances in a Bayesian VAR overturns the evidence of large regime switches in US monetary policy. Primiceri (2005), instead, estimates a VAR in which he allows for a continuously changing variance-covariance matrix: he similarly concludes that the role played by the falling volatility of exogenous shocks seems more important than monetary policy changes in explaining the recent behavior of US inflation and unemployment.

The typical estimated DSGE model, however, still commonly assumes that the shocks have maintained constant variance throughout the whole sample (e.g., Smets and Wouters 2003, 2007, Lubik and Schorfheide 2004, and An and Schorfheide 2007). The papers by Justiniano and Primiceri

¹To compute the conditional standard deviation series, I have estimated AR(1) models for inflation and output gap (using the deviation of real GDP from the CBO’s potential GDP series), allowing for a GARCH(1,1) specification for the residuals.

(2006) and Fernandez-Villaverde and Rubio-Ramirez (2007) are the first to relax this assumption. Both papers introduce stochastic volatility in optimizing DSGE models. They find that the volatilities of the shocks have significantly changed over time and that accounting for those variations is important to improve the models' fit.

The existence of time-varying volatility in the economy, therefore, can be now considered an empirical regularity. But what drives the changes in the volatility of macroeconomic fluctuations?

In Justiniano and Primiceri (2006) and Fernandez-Villaverde and Rubio-Ramirez (2007), the changes in volatility are modeled as *exogenous*. But if these are an important feature of the economy as they appear to be, it becomes crucial to strive to understand their potential causes.

This paper takes a step in this direction by presenting a model in which stochastic volatility arises endogenously in the economy. I present a stylized New-Keynesian DSGE model in which agents' learning about the economy has implications for macroeconomic volatility. Economic agents use simple models to form expectations and need to learn the relevant model parameters over time.² Their learning speed is endogenous and depends on previous forecast errors. When the forecast errors are large, the agents become concerned that the economy may be experiencing a structural break and, therefore, they start assigning a larger weight to new information. When the forecast errors are, instead, relatively modest, economic agents remain confident about their model and they are less responsive to new information. The endogenous time-varying learning speed has implications for the volatility of the macroeconomic variables that agents are trying to learn. In this way, agents' learning with an endogenous gain can generate stochastic volatility in the economy.

²See Evans and Honkapohja (2001) for a treatment of several models with adaptive learning.

The learning rule with an endogenously switching gain is in the same spirit as the rule assumed by Marcet and Nicolini (2003), who used a similar mechanism to study hyperinflations. Here, however, the gain is not fixed at a particular value, but estimated from time series data.

This paper is related to the recent work by Branch and Evans (2007), in which they present a framework in which regime changes in volatility arise endogenously. The time-variation in volatility is induced by two channels: agents' parametric learning and the switching between different possible predictors according to their previous forecasting performance. Model uncertainty plays an important role in generating time-varying volatility in their Lucas-type monetary model.³

The paper is also related to a recent paper by Lansing (2006). Lansing presents a New Keynesian Phillips curve with boundedly-rational expectations, which can give rise to time-varying persistence and volatility. In a single-equation setting for inflation, he can derive the optimal variable gain as the fixed point of a nonlinear map that relates the gain to the autocorrelation of inflation changes.⁴ Finally, the paper is related to the extensive literature on adaptive learning in monetary policy models (e.g., Evans and Honkapohja 2001) and, in particular to the papers that study the importance of learning to explain persistence or volatility in macroeconomic data (Orphanides and Williams 2003, 2005a,b, 2006, 2007, Adam 2005, Milani 2007, and Murray 2007).

The simulation results show that time variation in the gain can potentially generate substantial time-varying volatility in the inflation and output

³This paper and the Branch and Evans' approaches should be seen as complementary. A more realistic model, in fact, would possibly include agents that endogenously adjust their gain in response to the previous forecast errors, but that, at the same time, consider different models and switch among them as the performance of one of them becomes superior. This is, however, left for future research. Moreover, learning as in this paper might be seen as a crude way to model economic agents who are concerned about potential changes in the model of the economy, but without having to specify the different possible models or the number of regimes.

⁴Evans and Ramey (2006) also consider the optimal choice of the gain parameter.

gap series. The model is then taken to the data to judge whether changes in the learning process may have been a contributor to the evolution of macroeconomic volatility in the US. The Bayesian approach used in the paper facilitates the joint estimation of the learning gain coefficients together with the structural parameters in the economy. The estimation reveals that the endogenous gain appears to have switched to large constant gain values for most of the 1970s and early 1980s as a consequence of larger forecast errors by private agents in those periods. In the latest two decades, instead, the agents have switched to a decreasing gain. The estimated gain values in the 1970s are large and can justify a sizeable increase in volatility in the period. Simulation of the model, in fact, with the parameters fixed at the posterior mean estimates, implies that under the estimated evolving gain, the economy would observe higher volatility in the 1960s-1970s than later on. The magnitude of the model-implied decline in volatility roughly matches the size of the Great Moderation.

Moreover, the paper shows that even if the economy was subject to structural shocks with constant-variance over the whole sample, a failure to incorporate agents' learning in the estimation would lead econometricians to spuriously find the existence of ARCH/GARCH effects in the model innovations. The paper finally discusses how the evidence of time-varying volatility in the innovations, as they are measured by the econometrician, may itself be the result of monetary policy and, mainly, of the interaction between policy and agents' learning, and not just a matter of luck.

2. THE MODEL

The economy is described by the following New-Keynesian model⁵

$$\pi_t = \beta \widehat{E}_t \pi_{t+1} + \kappa x_t + u_t \quad (2.1)$$

$$x_t = \widehat{E}_t x_{t+1} - \sigma (i_t - \widehat{E}_t \pi_{t+1}) + g_t \quad (2.2)$$

$$i_t = \rho_t i_{t-1} + (1 - \rho_t) (\chi_{\pi,t} \pi_{t-1} + \chi_{x,t} x_{t-1}) + \varepsilon_t \quad (2.3)$$

where π_t denotes inflation, x_t the output gap, and i_t the nominal interest rate; u_t , g_t , and ε_t denote supply, demand, and monetary policy shocks. Equation (2.1) represents the forward-looking New Keynesian Phillips curve that can be derived from the optimizing behavior of monopolistically competitive firms under Calvo price setting or quadratic adjustment costs in nominal prices. Inflation depends on expected inflation in $t + 1$ and on current output gap. The parameter $0 < \beta < 1$ represents the households' discount factor, while κ denotes the slope of the Phillips curve and is an inverse function of the Calvo price stickiness parameter. Equation (2.2) represents the log-linearized intertemporal Euler equation that derives from the households' optimal choice of consumption. The output gap depends on the expected one-period ahead output gap and on the ex-ante real interest rate. The coefficient $\sigma > 0$ represents the intertemporal elasticity of substitution in consumption. Equation (2.3) describes monetary policy. The central bank follows a Taylor-type rule by adjusting its policy instrument, a short-term nominal interest rate, in response to deviations in inflation and the output gap. In light of McCallum's argument that only information up to $t - 1$ might be available in real-time, I assume that the central bank cannot respond to contemporaneous variables, but it responds only to lagged variables. The policy coefficients are allowed to vary over time and,

⁵See Woodford (2003) for a standard derivation.

in particular, they differ between the pre- and post-1979 samples:⁶

$$\rho_t = \begin{cases} \rho_{pre-79} & t < 1979 : 03 \\ \rho_{post-79} & t \geq 1979 : 03 \end{cases}$$

$$\chi_{\pi,t} = \begin{cases} \chi_{\pi,pre-79} & t < 1979 : 03 \\ \chi_{\pi,post-79} & t \geq 1979 : 03 \end{cases}, \quad \chi_{x,t} = \begin{cases} \chi_{x,pre-79} & t < 1979 : 03 \\ \chi_{x,post-79} & t \geq 1979 : 03 \end{cases}.$$

The supply shock u_t may arise endogenously in the model by assuming a time-varying elasticity of substitution among differentiated goods, whereas g_t derives from shocks to preferences or technology, for example (both shocks are assumed to evolve as AR(1) processes). The majority of the papers that focus on the estimation of DSGE models assumes that such shocks maintain constant variance over the whole sample (e.g. Smets and Wouters 2007, An and Schorfheide 2007).

But recent papers have suggested that the changing volatilities of these shocks can be important to understand macroeconomic fluctuations.

State-of-the-art DSGE models cannot endogenously generate time-varying stochastic volatility, but they need to assume that disturbances follow a Markov-Switching process (Laforte 2005) or exogenously assume stochastic volatility (Justiniano and Primiceri 2006 and Fernandez-Villaverde and Rubio-Ramirez 2007). This paper contributes to the literature by showing how stochastic volatility can arise *endogenously* from agents' learning about the economy.

I relax here the assumption of rational expectations. I assume that agents use a linear economic model to form their expectations. The agents do not know the model coefficients and need to learn them over time. Therefore, \hat{E}_t refers to subjective expectations and may differ from the mathematical expectations operator conditional on the true model of the economy E_t .⁷

⁶I assume that a regime switch in policy occurs in 1979, when Paul Volcker begins his term as Chairman of the Federal Reserve (August 1979). Duffy and Engle-Warnick (2005), using nonparametric methods, similarly identify a switch in policy exactly in the third quarter of 1979. Allowing for unknown changes in policy is beyond the scope of this paper.

⁷I have assumed a simple small-scale New-Keynesian model without adding "mechanical" sources of persistence as habit formation in consumption or inflation indexation;

2.1. Learning. Economic agents need to form expectations about future aggregate inflation rates and future output to solve their optimal consumption and price-setting decisions. I assume that agents use a perceived linear model of the economy and they need to learn the relevant coefficients. Therefore, they behave as econometricians by estimating the model and updating their estimates as new data become available. The agents use the following ‘Perceived Law of Motion’ (PLM):

$$Z_t = a_t + b_t Z_{t-1} + \eta_t \quad (2.4)$$

where $Z_t \equiv [\pi_t, x_t, i_t]'$, and a_t, b_t are coefficient vectors and matrices of appropriate dimensions. The agents’ PLM, therefore, is a simple VAR(1) in the endogenous variables π_t, x_t , and i_t . Notice that although the true constants in the model equal zero, agents are not endowed with this information. In this way, they also need to learn the steady-state of the variables. The PLM is, therefore, similar to the Minimum State Variable solution of the system under rational expectations, but with two differences: agents do not know the reduced-form model parameters and they cannot observe the exogenous shocks.⁸ The agents learn the model coefficients according to the following updating equations

$$\hat{\phi}_t = \hat{\phi}_{t-1} + g_{t,y} R_{t-1}^{-1} X_t (Z_t - X_t' \hat{\phi}_{t-1}) \quad (2.5)$$

$$R_t = R_{t-1} + g_{t,y} (X_{t-1} X_{t-1}' - R_{t-1}) \quad (2.6)$$

where $\hat{\phi}_t = (a_t', \text{vec}(b_t)')$ collects in a vector the coefficients, and $X_t \equiv \{1, Z_{t-1}\}_0^{t-1}$ is a matrix of the stacked regressors. The first line describes the updating of the learning rule coefficients, whereas the second describes the updating of the matrix of second moments R_t . The coefficient $g_{t,y}$ denotes

Milani (2006, 2007), in fact, shows that, under learning, those may become redundant as learning is successful in inducing persistence in the model. This simplification is not relevant for the main scope of the paper.

⁸That is, I assume that agents estimate VARs in the endogenous variables, rather than VARMA, as this is a more common practice in econometrics. It seems more realistic to assume that agents do not observe the shocks; the results in the paper, however, do not hinge on this assumption.

the gain, which in the paper will be endogenously determined and time-varying. I allow agents to learn about inflation, output, and interest rates at different rates, letting the gain $g_{t,y}$ differ for $y = \pi_t, x_t, i_t$ (as Branch and Evans 2006 discuss, in fact, if the degree of structural change can be expected to differ across series, the optimal gains should also differ).

The gain endogenously adjusts according to past forecast errors as follows

$$g_{t,y} = \begin{cases} t^{-1} & \text{if } \frac{\sum_{j=0}^J (|y_{t-j} - E_{t-j-1} y_{t-j}|)}{J} < v_t^y \\ \bar{g}_y & \text{if } \frac{\sum_{j=0}^J (|y_{t-j} - E_{t-j-1} y_{t-j}|)}{J} \geq v_t^y, \end{cases} \quad (2.7)$$

where $y = \pi_t, x_t, i_t$. When the average of the past forecast errors (in absolute value) is below a certain threshold v_t^y , the agents use a decreasing gain. When they know the correct model of the economy, with a decreasing gain they can be expected to asymptotically converge to the Rational Expectation Equilibrium (in this model, assuming that the shocks are observed and the gain is always decreasing, the required conditions are derived by Bullard and Mitra 2007). When the average of previous forecast errors is above the threshold v_t , instead, the agents become concerned that the economy may be experiencing a structural break. In the proximity of a structural break, a decreasing gain would be inefficient: the agents therefore switch to a constant gain, which allows them to better track the break by assigning a larger weight to new information. When the forecast errors fall again below the threshold, agents switch back to a decreasing gain, which is initially reset to $\frac{1}{\bar{g}_y + t}$ (rather than t^{-1}).

The endogenous switching gain is in the spirit of the gain assumed by Marcet and Nicolini (2003).⁹ I assume that the threshold v_t^y is given by the mean absolute deviation of historical forecast errors, which is recursively

⁹Although agents' learning with the described endogenous gain is by no means optimal, it can be expected to provide a fairly good approximation to the optimal forecasting behavior of agents who are concerned about possible unknown breaks in the economy, but who do not want to take a stand on the nature or timing of the breaks, or on the existence or number of regimes, and assuming that the agents, in their loss function, are much more concerned about very large forecast errors than relatively small ones.

updated. Notice that the degrees of freedom from this mechanism are the gain coefficients \bar{g}_y , the window length J for past forecast errors, as well as v_t^y . The gain will be estimated from the data, whereas J will be initially fixed (later in the paper I will also treat J as a parameter and estimate its value).

I assume that economic agents dispose of information only up to $t - 1$ when forming expectations for next period. Therefore, economic agents use (2.4) and the updated parameter estimates in (2.5) and (2.6) to form their expectations for $t + 1$ as

$$\hat{E}_{t-1}Z_{t+1} = a_{t-1}(1 + b_{t-1}) + b_{t-1}^2 Z_{t-1}, \quad (2.8)$$

which can be substituted in (2.1) to (2.3) to obtain the Actual Law of Motion of the Economy (ALM).¹⁰

3. ENDOGENOUS GAIN AND ENDOGENOUS TIME-VARYING VOLATILITY

The value of the gain coefficient affects the volatility in the economy. In a simple empirical model of inflation and unemployment dynamics, Orphanides and Williams (2005, 2007), for example, have shown that the volatility of those variables is a positive function of the gain (they consider only gain values between 0.01 and 0.04). I simulate the model (for now with constant, exogenously set, gain values) to show that this is also the case here.¹¹ Figure 2 makes clear that the standard deviations of inflation and the output gap would increase as a function of the constant gain value.

Changes in the gain over time, therefore, may potentially be an important determinant of the observed movements in macroeconomic volatility. To show the potential role of the gain, I turn now to the simulation of the

¹⁰Checking E-Stability in an economy with an endogenous gain as the one presented here is an interesting issue, which is, however, not examined in this paper.

¹¹I fix the following values for the parameters: $\beta = 0.99$, $\kappa = 0.05$, $\sigma = 0.1$, $\rho = 0.95$, $\chi_\pi = 1.5$, $\chi_x = 0.5$, $\rho_u = 0.9$, $\rho_g = 0.9$. Agents use the MSV solution of the system to form expectations. The economy is simulated for 1,000 periods using a grid of constant gain values from 0 to 0.15.

model under an endogenous gain, which is allowed to switch as described in (2.7). Therefore, agents adopt a decreasing gain as long as their forecast errors are ‘small’. They switch to a constant gain when those become larger and above ν_t , the mean absolute deviation of past forecast errors. In the simulation, I assume $\bar{g}_\pi = \bar{g}_x = 0.15$. The choice of such high gain values is, for the moment, purely for descriptive purposes and it is meant to make the effects more striking in the graph (the value will be estimated later in the paper).¹² Figure 3 shows the time-varying endogenous gain together with the rolling standard deviations of inflation and output gap, obtained from a typical simulation. As the gain changes over time, the degree of volatility in the economy also experiences large shifts. The figure displays sizeable time variation in volatility and various episodes characterized by volatility clustering, although the exogenous shocks had constant variance by construction. The persistence of the volatility series and the duration of the clusters obviously depend on the assumed window that agents use to compute past forecast errors; by varying the window size one could in principle mimic a wide range of changing volatility series.¹³ For example, decreasing J to 500 would imply more frequent changes in volatility and shorter clusters, as shown in the upper panels of Figure 4. The lower panels, instead, plot the case when agents only adopt a constant gain (fixed at the lower value of 0.05).

The next section will take the model to the data. The estimation aims to infer the evolution of the endogenous gain from time series observations.

¹²I simulate the economy for 13,000 periods, allowing agents to use a window of 3,000 observations when computing the mean of past forecast errors, and discarding the first 3,000 periods. The large number of observations is again meant to make the time-varying volatility more apparent in the graph. The parameters are: $\beta = 0.99$, $\kappa = 0.05$, $\sigma = 0.1$, $\rho = 0.95$, $\chi_\pi = 1.5$, $\chi_x = 0.5$, $\rho_u = 0.5$, $\rho_g = 0.5$.

¹³Allowing the gain to change in a more ‘continuous’ fashion, rather than abruptly jumping from t^{-1} to \bar{g} would imply more gradual movements in the volatility series. The case of a gradually changing gain, possibly along the lines proposed by Colucci and Valori (2004, 2005), is left for future research.

The simulation can then be repeated in an artificial economy in which the learning process is calibrated to resemble the one estimated from US data.

4. BAYESIAN ESTIMATION

I estimate the model using likelihood-based Bayesian methods. The estimation follows Milani (2007), who extends the techniques reviewed in An and Schorfheide (2007) to allow for near-rational expectations and learning. The vector Θ collects the structural parameters of the model:

$$\Theta = \{\beta, \kappa, \sigma, \rho_t, \chi_{\pi,t}, \chi_{x,t}, \rho_u, \rho_g, \sigma_u, \sigma_g, \sigma_\varepsilon, \bar{g}_\pi, \bar{g}_x, \bar{g}_i\} \quad (4.1)$$

Differently from Milani (2007), the gain coefficient is now endogenous, being allowed to vary over time depending on the magnitude of the past forecast errors that agents make (as made clear by expression 2.7). The gain switches from decreasing (equal to t^{-1}) in ‘stable’ times to constant (\bar{g}_y), when past forecast errors become large and hence suggestive that a break may be occurring. The constant gain coefficient to which agents switch is not fixed to an ad-hoc value, rather its value is jointly estimated with the rest of the model parameters. I use the Metropolis-Hastings algorithm to generate 200,000 draws from the posterior distribution.¹⁴ The likelihood of the system is evaluated at each iteration using the Kalman Filter.¹⁵ I use quarterly US data for the 1960:I-2006:I sample in the estimation to fit the series for inflation, output gap, and nominal interest rates.¹⁶ Data from

¹⁴I discard a burn-in of 40,000 draws. See appendix in Milani (2007) for more details on the estimation.

¹⁵Since stochastic volatility arises endogenously from the adjustment of expectations in the model and it is not assumed, instead, in the exogenous shocks, the estimation can be performed using the Kalman Filter rather than the more computationally-intensive particle filter employed in Fernandez-Villaverde and Rubio-Ramirez (2007).

¹⁶Inflation is defined as the annualized quarterly rate of change of the GDP Implicit Price Deflator, output gap as the log difference between GDP and Potential GDP (Congressional Budget Office estimate), and the federal funds rate is the measure for the nominal interest rate. The series are obtained from FRED, the Federal Reserve Bank of Saint Louis economic database.

the pre-sample period 1954:III-1959:I were, instead, used to initialize the learning algorithm.

4.1. Priors. The priors for the model coefficients are reported in Table 1. Most prior choices follow Milani (2007). To minimize the influence of the priors on the main parameters of interest, I assume a Uniform prior distribution in the $[0, 0.3]$ interval for the constant gain coefficients. I assume a dogmatic prior for β , which is fixed at 0.99, a Gamma prior for σ and κ , a Beta prior for the autoregressive coefficients, and Normal prior distributions for the feedback coefficients to inflation and output gap in the policy rule. I assume for now $J = 4$, i.e. agents care about forecast errors over the previous year (this restriction will be later relaxed) when deciding how much weight to assign to more recent information. I will point out in describing the results the situations in which the priors have important effects on the estimates.

4.2. Empirical Results. Figure 5 shows the evolution of the forecast errors (in absolute value) about inflation, output gap, and the federal funds rate over the sample under the estimated learning rules. Inflation and output were typically harder to predict during the 1970s and until the early 1980s. The forecast errors for both inflation and output gap were on average lower in the 1990s. Monetary policy, instead, was harder to forecast in the late 1960s, in most of the 1970s, and during Volcker's disinflation.¹⁷ Figure 6 shows the episodes in which the rolling means of the absolute forecast errors exceed the updated values of ν_t^y , which imply switches to learning with a constant gain.

Table 2 presents the parameter estimates. The value of the constant gain to which private agents switch when their forecast errors are above threshold is estimated equal to 0.082 for inflation and to 0.073 for output (a very low

¹⁷Best and Milani (2007) study more in detail private agents' expectations and learning about future monetary policies using post-war US data.

gain coefficient is, instead, found for the interest rate equation). Those values are substantially larger than the estimates in Milani (2007), but of course here they refer only to particular periods in the sample.¹⁸ It appears, therefore, that agents, on average, adopt low gain coefficients, but they switch to considerably higher gains in periods of instability. Figure 7 plots the evolution of the time-varying gain coefficients estimated for inflation and the output gap.¹⁹ The learning process for inflation often switches to a constant gain in the 1970s until the early 1980s and it reverts to a decreasing gain shortly after 1985 and for most of the latest part of the sample. Learning about the output gap is also characterized by frequent switches to a constant gain from the 1960s until 1985, and by a decreasing gain for most of the recent two decades (only two switches are identified from 1985 to 2006).

Turning to the other parameters, I estimate $\sigma^{-1} = 5.92$ and $\kappa = 0.022$. The posterior means for the monetary policy rule coefficients indicate a more aggressive response to inflation and a less active response towards the output gap in the second part of the sample than in the first (χ_π goes from 1.37 to 1.53, and χ_x declines from 0.58 to 0.48). The estimated monetary policy rules satisfy the Taylor principle in both sub-samples (a similar result in a model with learning is found in Milani 2006). The posterior distributions for the policy rule parameters, however, are not far from the assumed prior distributions, indicating that the data are, therefore, not very informative about these parameters. The data appear informative, instead, on the values of the gain coefficients. Although Uniform priors were assumed, the

¹⁸The larger gain for inflation than output is consistent with results in Branch and Evans (2007) and Milani (2006, 2007).

¹⁹I focus in this paper on inflation and output gap. I do not try, instead, to explain the time-varying volatility in the Taylor rule equation with learning. The estimated higher volatility of monetary policy shocks in some sub-periods can be more realistically attributed to misspecification of the Taylor rule in the 1979-1982 years and in few other episodes in the 1970s than to a time-varying gain story.

estimation identifies the gains that imply the best fit of the data (Figure 10 shows their posterior distributions).

4.3. Robustness. The results might depend on the particular choice of the number of observations J (in expression 2.7) that agents are assumed to use when computing the average of recent forecast errors. To check for robustness, I repeat the estimation using a longer window, i.e. $J = 20$ (now agents compute the mean absolute forecast errors over the past five years). The estimates are reported in table 3. Figure 8 shows the evolution of the endogenous gain coefficients in this case. The estimated gains equal 0.065 for inflation and 0.064 for output gap. The other estimates are not substantially different.

The window for the mean forecast errors can also be interpreted as a parameter that can be estimated from data. Table 4 reports the results when the estimation is repeated treating J as a free parameter. A gamma distribution with mean 12 and standard deviation 4.9 is assumed as prior for J . The estimated posterior mean is 4 (in the estimation, J is rounded to the closest integer, since agents need to use the previous J periods as described in 2.7), implying that agents care about forecast errors over the previous year (the time-varying gains are therefore similar to those shown in Figure 7). Overall, the results are not too sensitive to the choice of J . The finding of frequent switches to a constant gain coefficient in the 1970s is especially robust to the different J 's. Switches to constant gain in the later part of the sample, instead, are more sensitive to its choice.

Since agents are unsure about the model of the economy and whether this is changing over time, one might argue that agents may be better off always using constant-gain learning, rather than reverting to a decreasing gain when their forecasting performance is satisfactory. I follow this argument here and assume that agents always adopt constant-gain learning, only switching from a 'low' to a 'high' gain when the conditions in (2.7) are met (the 'low' gain

is fixed at 0.02, whereas the ‘high’ gain is estimated). The estimates are reported in table 5. The switches to the higher gain occur in similar periods to those found under the baseline case (see Figure 9); the estimated gains equal 0.096 for inflation and 0.042 for the output gap.

5. SIMULATION

I repeat the simulation of the model, but now using the estimated parameter values (shown in Table 2) from the previous section and fixing the agents’ learning to resemble the one estimated from US data (i.e, assuming that the endogenous gain switches as in Figure 7). I simulate an economy with 185 periods (the same length as the estimated New Keynesian model) for 10,000 times. The shocks that hit the economy are drawn from distributions with constant variance over the whole sample.²⁰ We have previously seen that learning can imply time-varying volatility in the variables about which agents are forming expectations. But suppose that learning is neglected in an empirical exercise. Let’s consider the following experiment. Suppose that an econometrician would estimate inflation and output equations on the simulated data, but without taking learning into account. Would the econometrician find evidence of time-varying volatility, even if the true data-generating process had shocks with *constant* variance through the whole sample?

To answer this question, I regress artificially-generated inflation and output gap series on a constant and their first lag (a similar regression on actual data gives the plot for conditional standard deviations in Figure 1); then I take the implied residuals and perform a test on the existence of

²⁰I use a ‘projection facility’ in the simulation to ensure that the economy does not become unstable. As in Orphanides and Williams (2005, 2007), in fact, I assume that agents recognize that the economy is stable and every time the matrix of autoregressive coefficients in their VAR has an eigenvalue larger than 1 in absolute value, they do not update their estimates, keeping $\hat{\phi}_t = \hat{\phi}_{t-1}$ and $R_t = R_{t-1}$. If this is not enough to guarantee non-explosive dynamics, I reject the specific draw.

ARCH/GARCH effects (at the 5% significance level).²¹ Table 6 reports the percentages of rejections of the null hypothesis of no ARCH/GARCH effects from simulated data. In the case that the data derive from an economy with no learning (i.e., imposing $g_{t,y} = 0$ at all times), the test rejects the null of no ARCH effects only about 5% of the times. In the case with learning, even though the variances of the shocks are constant by construction, the test concludes that ARCH/GARCH effects are a feature of the data in 52% of the cases for inflation and 78% for output gap (see Table 6 for more results under different cases).²²

These results are suggestive that estimations that abstract from agents' learning can significantly overestimate the time variation in the volatility of exogenous shocks.²³

5.1. Time-Varying Volatility. The model can generate time-varying volatility in macroeconomic variables even if the exogenous disturbances have maintained constant variance over the whole sample. I try to verify whether the model can generate a pattern of volatility that roughly resembles what is observed in actual data. That is, I aim to verify whether the outcomes from the simulations imply volatility series that increase in the first part of the sample (when the endogenous gain switches to constant) and decline in the second part (when the endogenous gain reverts to decreasing). In

²¹To test for evidence of ARCH(q) effects against the hypothesis of no ARCH effects, the squared residuals are regressed on a constant and q lagged values and the statistic $\chi^2 = TR^2$ is computed. Under the null hypothesis of no ARCH effects this statistic has a limiting chi-square distribution with q degrees of freedom. The test for GARCH(p, q) is equivalent to a test for ARCH($p + q$).

²²A more sophisticated version of the same experiment would imply estimating the full DSGE model under Rational Expectations, hence disregarding learning dynamics, to test for the existence of spurious ARCH effects or stochastic volatility in the exogenous shocks. This is, however, beyond the scope of this paper.

²³A similar point, albeit in a largely different contexts, is reached by Bullard and Singh (2007), who find that learning about different regimes may have had a role for the "Great Moderation": they conclude that 30% of the decline in variance may be due to learning, rather than pure good luck.

each simulation, I take the residuals from the inflation and output equations and I look at the point in the sample in correspondence of which the maximum rolling standard deviation is obtained (using a rolling window of 20 periods). I then estimate the Kernel density of the maxima across all simulations. Figure 11 displays the results. The standard deviations often increase near those observations that would correspond to the late 1960s and 1970s, and become typically lower in the second part of the sample. If the economy was simulated without learning, instead, one would find that the maxima of the volatility series are uniformly spread across all observations in the sample.

Therefore, as in the previous section, neglecting learning in empirical work, if learning is a feature of the data-generating process of the economy, may lead researchers to spuriously find time variation in the volatility of residuals from estimated univariate equations (or possibly of exogenous shocks in DSGE models under RE). The pattern of volatility is in the ballpark of that estimated on actual data and reported in Figure 1. One would, in fact, conclude that the volatility of shocks has increased in periods in which agents' forecast errors are generally large and in which learning introduces more noise in the economy.

The estimated model, however, is still unlikely to match the magnitude and exact timing of the changes that we observe in Figure 1. Taking more seriously the model to the data to judge what fraction of the decline in volatility can be explained by learning (that is, what fraction of the decline in exogenous volatility is removed if we allow for endogenous stochastic volatility from learning) is an important topic that I leave for future work.

5.2. The Great Moderation. The model estimation has shown that the gain coefficients were typically larger in the 1960s and 1970s than in the following decades, and this has affected the degree of volatility in the economy. The volatility of output gap and inflation has fallen in the post-1984 sample.

Is the model with learning able to generate the Great Moderation? For each simulation, I compute the ratio of the standard deviations of inflation and output gap in the second part of the sample (corresponding to 1985-2006 observations) versus the first part (which correspond to 1960-1984). Taking the median from the simulations, for the baseline model, the ratios between post-1984 and pre-1984 standard deviations equal 0.39 for inflation (versus 0.35 on actual data) and 0.42 for output gap (versus 0.50 on actual data). The model with learning, with parameters fixed at their posterior mean estimates, is, therefore, in principle capable of endogenously generating a reduction in the volatility of the main macroeconomic variables of a magnitude comparable to the Great Moderation (Table 7 shows that the results remain similar under the other cases).

5.3. The Effects of Policy. The literature that studies the main sources of the Great Moderation has often focused on testing explanations based on changes in monetary policy versus explanations based on reductions in the volatility of the exogenous shocks that hit the economy. A decline in the estimated volatility of shocks is usually taken as evidence in favor of the bad luck-good luck hypothesis versus the alternative hypothesis of transition from bad to good policy.

Changes in volatility, however, may not be unrelated to the monetary policy regime. Chairman Bernanke, in a speech about the Great Moderation in February 2004, for example, argued

[I am not convinced that the decline in macroeconomic volatility of the past two decades was primarily the result of good luck.]

[...changes in monetary policy could conceivably affect the size and frequency of shocks hitting the economy, at least as an econometrician would measure those shocks]

He continues:

[changes in inflation expectations, which are ultimately the product of the monetary policy regime, can also be confused with truly exogenous shocks in conventional econometric analyses.]

Therefore,

[some of the effects of improved monetary policies may have been misidentified as exogenous changes in economic structure or in the distribution of economic shocks.]

This paper makes time-varying volatility endogenous. In this way, it is easier to study how the volatility of the shocks (as they are measured by the econometrician) may itself depend on policy. This section can, therefore, provide an initial evaluation of the extent to which changes in the volatility of inferred economic shocks may, as in Bernanke's claims, simply reflect better monetary policies.

To examine the interaction between volatility and monetary policy, I simulate the economy with an endogenous gain as in section 3, but now under a wide range of policy feedback coefficients to inflation (using a grid from 0 to 5 in 0.5 increments). Figure 12 shows the result under the model parametrization obtained in table 2. A more aggressive monetary policy reduces agents' forecast errors and, therefore, it affects the frequency of switches in their gain coefficients (from almost 70% of the times when $\chi_\pi = 0$ to less than 30% when $\chi_\pi = 5$, with an average gain in the sample that goes from above 0.07 to below 0.05), and, through this channel, it affects also changes in volatility in the economy.²⁴ The more aggressive monetary policy, the less often the econometrician would spuriously find time-varying volatility in the reduced-form residuals (from more than 80 to 55% of the times). Changes in the volatility of estimated shocks, therefore, may not in principle be a matter of luck after all, but an implication of better policy (notice, however, that under the estimated coefficients in this paper, large changes in volatility due to policy changes are unlikely).²⁵

This result echoes, although in a different setting and with a different focus, the argument in a recent paper by Benati and Surico (2007). They

²⁴The importance of monetary policy in reducing agents' forecast errors is also discussed in Orphanides and Williams (2003).

²⁵A more detailed analysis of this channel, however, is certainly needed.

artificially generate data from a New Keynesian model assuming a policy change from a ‘passive’ to an ‘active’ policy rule and they ask whether a common VAR estimation would be able to recover the change in policy. They show that the estimated VAR would lead researchers to inaccurately conclude that the variances of the shocks have changed, but not the policy coefficients.

This paper’s results similarly suggest caution: simply finding that the variances have changed from reduced-form regressions may not necessarily imply changes in luck, but it might be an effect of policy changes or, as in this paper, of the interaction between changing policies and private agents’ learning.

6. CONCLUSIONS AND FUTURE DIRECTIONS

The paper has presented a New Keynesian model in which agents’ learning with a switching gain coefficient endogenously generates time-varying volatility in the economy. The estimation of the model has shown that there is evidence of large changes in the gain over the post-war US sample. The changes in the gain can imply important changes in macroeconomic volatility, which can roughly match the magnitude of the Great Moderation. An econometrician that would abstract from such learning dynamics, however, would be lead to overestimate the importance of changes in the volatility of exogenous shocks. Moreover, time variation in volatility may not be simply a matter of luck, but it may itself be affected by changes in monetary policy and, in particular, it can stem from the interaction between policy and learning by private agents.

A more ambitious scope for future research will be to test whether extensions of the model would be able to generate endogenous stochastic volatility series able to match those estimated in DGSE models by Justiniano and Primiceri (2006) and Fernandez-Villaverde and Rubio-Ramirez (2007).

More generally, this would allow to test what fraction of the changes in volatility may be due to learning, rather than to the decline in the volatility of exogenous disturbances.

Finally, the model can similarly be used to study the formation of expectations and learning in emerging market economies. Their bigger instability is likely to imply large forecast errors by private agents and probably learning with a higher constant gain than the one estimated here, which may help to explain part of the large changes in volatility over time in these economies.

REFERENCES

- [1] Adam, K., (2005). "Learning to Forecast and Cyclical Behavior of Output and Inflation", *Macroeconomic Dynamics*, Vol. 9(1), 1-27.
- [2] Benati, L., and P. Surico, (2006). "VAR Analysis and the Great Moderation", mimeo, Bank of England.
- [3] Best, G., and F. Milani, (2007). "Central Bank Transparency and Imperfect Credibility in a DSGE Model with Private Sector Learning", mimeo, University of California, Irvine.
- [4] Branch, W., and G.W. Evans, (2006). "A Simple Recursive Forecasting Model", *Economics Letters*, vol. 91(2), 158-166.
- [5] Branch, W., and G.W. Evans, (2007). "Model Uncertainty and Endogenous Volatility", *Review of Economic Dynamics*, vol. 10(2), 207-237.
- [6] Bullard, J., and K. Mitra, (2007). "Determinacy, Learnability, and Monetary Policy Inertia", forthcoming, *Journal of Money, Credit, and Banking*.
- [7] Bullard, J., and A. Singh, (2007). "Learning and the Great Moderation", mimeo.
- [8] Colucci, D., and V. Valori, (2004). "Adaptive Learning in the Cobweb Model with an Endogenous Gain Sequence", mimeo, Universita' di Firenze.
- [9] Colucci, D., and V. Valori, (2005). "Error Learning Behaviour and Stability Revisited", *Journal of Economic Dynamics and Control*, Vol. 29/03, 371-388.
- [10] Evans, G. W., and S. Honkapohja (2001). *Learning and Expectations in Economics*. Princeton, Princeton University Press.
- [11] Evans, G. W., and G. Ramey (2006). "Adaptive Expectations, Underparameterization and the Lucas Critique", *Journal of Monetary Economics*, Vol. 53, 249-264.
- [12] Fernandez-Villaverde, J., and J. Rubio-Ramirez, (2007). "Estimating Macroeconomic Models: A Likelihood Approach", *Review of Economic Studies*, forthcoming.
- [13] Honkapohja, S., and K. Mitra, (2003). "Learning with Bounded Memory in Stochastic Models", *Journal of Economic Dynamics and Control*, vol. 27(8), 1437-1457.
- [14] Justiniano A. and G.E. Primiceri (2006). The Time Varying Volatility of Macroeconomic Fluctuations, mimeo.
- [15] Kim, C.J., and C.R. Nelson, (1999). "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," *The Review of Economics and Statistics*, vol. 81(4), pages 608-616.
- [16] Laforte, J.P. (2005). "DSGE Models and Heteroskedasticity: a Markov-Switching Approach", mimeo, Board of Governors of the Federal Reserve System.
- [17] Lubik, T., and F. Schorfheide, (2004). "Testing for Indeterminacy: an Application to U.S. Monetary Policy" , *American Economic Review*, 94(1), pp. 190-217.

- [18] Marcet, A., and J.P. Nicolini, (2003). "Recurrent Hyperinflations and Learning," *American Economic Review*, vol. 93(5), pages 1476-1498.
- [19] McConnell, M.M., and G. Perez-Quiros (2000). Output Fluctuations in the United States: What Has Changed Since the Early 1980s? *American Economic Review*, 90, 1464-1476.
- [20] Milani, F., (2006). "Learning, Monetary Policy Rules, and Macroeconomic Stability", mimeo, University of California, Irvine.
- [21] Milani, F., (2007). "Expectations, Learning and Macroeconomic Persistence", forthcoming, *Journal of Monetary Economics*.
- [22] Murray, J., (2007). "Empirical Significance of Learning in a New Keynesian Model with Firm-Specific Capital", mimeo, Indiana University.
- [23] Orphanides, A., and J. Williams, (2003). "Imperfect Knowledge, Inflation Expectations and Monetary Policy", in Ben Bernanke and Michael Woodford, eds., *Inflation Targeting*. Chicago: University of Chicago Press.
- [24] Orphanides, A., and J. Williams, (2005a). "Inflation Scares and Forecast-Based Monetary Policy", *Review of Economic Dynamics*, Vol. 8 (2), pages 498-527.
- [25] Orphanides, A., and J. Williams, (2005b). "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations", *Journal of Economic Dynamics and Control*, vol. 29(11), pages 1927-1950.
- [26] Orphanides, A., and J. Williams, (2006). "Monetary Policy with Imperfect Knowledge", *Journal of the European Economic Association*, 4(2-3), 366-375.
- [27] Orphanides, A., and J. Williams, (2007). "Inflation Targeting and Imperfect Knowledge", *Federal Reserve of San Francisco Economic Review*.
- [28] Owyang, M.T., (2001). "Persistence, Excess Volatility, and Volatility Clusters in Inflation", *Federal Reserve Bank of St. Louis Review*, 83(6), pp. 41-52.
- [29] Primiceri, G.E., (2005). Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies*, 72, 821-852.
- [30] Smets, F., and R. Wouters, (2003). "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, vol. 1(5), pages 1123-1175.
- [31] Smets, F., and R. Wouters, (2007). "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach", forthcoming, *American Economic Review*.
- [32] Stock, J.H., and M.W. Watson, (2002). Has the Business Cycle Changed, and Why?," *NBER Macroeconomics Annual*, 17, 159-218.
- [33] Woodford, M., (2003). *Interest and Prices*. Princeton University Press.

Description	Param.	Prior Distribution			
		Range	Distr.	Mean	95% Int.
Inverse IES	σ^{-1}	\mathbb{R}^+	G	1	[.12, 2.78]
Slope PC	κ	\mathbb{R}^+	G	.25	[.03, .7]
Discount Rate	β	.99	—	.99	—
Interest-Rate Smooth	ρ_{pre79}	[0, 1]	B	.8	[.46, .99]
Feedback to Infl.	$\chi_{\pi,pre79}$	\mathbb{R}	N	1.5	[.51, 2.48]
Feedback to Output	$\chi_{x,pre79}$	\mathbb{R}	N	.5	[.01, .99]
Interest-Rate Smooth	ρ_{post79}	[0, 1]	B	.8	[.46, .99]
Feedback to Infl.	$\chi_{\pi,post79}$	\mathbb{R}	N	1.5	[.51, 2.48]
Feedback to Output	$\chi_{x,post79}$	\mathbb{R}	N	.5	[.01, .99]
Std. MP shock	σ_{ε}	\mathbb{R}^+	IG	1	[.34, 2.81]
Std. g_t	σ_g	\mathbb{R}^+	IG	1	[.34, 2.81]
Std. u_t	σ_u	\mathbb{R}^+	IG	1	[.34, 2.81]
Constant Gain infl.	$\bar{\mathbf{g}}_{\pi}$	[0, 0.3]	U	.15	[.007, .294]
Constant Gain gap	$\bar{\mathbf{g}}_x$	[0, 0.3]	U	.15	[.007, .294]
Constant Gain FFR	$\bar{\mathbf{g}}_i$	[0, 0.3]	U	.15	[.007, .294]

Table 1 - Prior Distributions.

Description	Parameter	Posterior Distribution	
		Mean	95% Post. Prob. Int.
Inverse IES	σ^{-1}	5.92	[4.23-8.34]
Slope PC	κ	0.022	[0.003-0.06]
Discount Factor	β	0.99	-
IRS pre-79	ρ_{pre79}	0.938	[0.85-0.99]
Feedback Infl. pre79	$\chi_{\pi,pre-79}$	1.37	[0.87-1.89]
Feedback Gap pre79	$\chi_{x,pre-79}$	0.58	[0.18-1.02]
IRS post-79	ρ_{post79}	0.93	[0.88-0.97]
Feedback Infl. post79	$\chi_{\pi,post-79}$	1.53	[1.05-2.05]
Feedback Gap post79	$\chi_{x,post-79}$	0.48	[0.04-0.92]
Autoregr. Cost-push shock	ρ_u	0.40	[0.28-0.52]
Autoregr. Demand shock	ρ_g	0.84	[0.75-0.92]
Std. Cost-push shock	σ_u	0.89	[0.8-1.00]
Std. Demand shock	σ_g	0.65	[0.58-0.72]
Std. MP shock	σ_ε	0.97	[0.87-1.07]
Constant gain (Infl.)	$\bar{\mathbf{g}}_\pi$	0.082	[0.07-0.09]
Decreasing gain (Infl.)	t^{-1}	-	-
Constant gain (Gap)	$\bar{\mathbf{g}}_x$	0.073	[0.06-0.083]
Decreasing gain (Gap)	t^{-1}	-	-
Constant gain (FFR)	$\bar{\mathbf{g}}_i$	0.001	[0,0.01]
Decreasing gain (FFR)	t^{-1}	-	-

Table 2 - Posterior Distributions: baseline case with $J = 4$.

Description	Parameter	Posterior Distribution	
		Mean	95% Post. Prob. Int.
Inverse IES	σ^{-1}	6.89	[4.62,10]
Slope PC	κ	0.02	[0.003,0.06]
Discount Factor	β	0.99	-
IRS pre-79	ρ_{pre79}	0.936	[0.86-0.99]
Feedback Infl. pre79	$\chi_{\pi,pre-79}$	1.31	[0.91-1.76]
Feedback Gap pre79	$\chi_{x,pre-79}$	0.61	[0.13-0.98]
IRS post-79	ρ_{post79}	0.914	[0.86-0.96]
Feedback Infl. post79	$\chi_{\pi,post-79}$	1.54	[1.09-1.93]
Feedback Gap post79	$\chi_{x,post-79}$	0.41	[0.006-0.87]
Autoregr. Cost-push shock	ρ_u	0.39	[0.26-0.5]
Autoregr. Demand shock	ρ_g	0.83	[0.74-0.91]
Std. Cost-push shock	σ_u	0.97	[0.87-1.07]
Std. Demand shock	σ_g	0.67	[0.61-0.75]
Std. MP shock	σ_ε	0.96	[0.87-1.07]
Constant gain (Infl.)	$\bar{\mathbf{g}}_\pi$	0.065	[0.055-0.075]
Decreasing gain (Infl.)	t^{-1}	-	-
Constant gain (Gap)	$\bar{\mathbf{g}}_x$	0.064	[0.049-0.072]
Decreasing gain (Gap)	t^{-1}	-	-
Constant gain (FFR)	$\bar{\mathbf{g}}_i$	0.005	[0,0.032]
Decreasing gain (FFR)	t^{-1}	-	-

Table 3 - Posterior Distributions: Alternative case with J ($J = 20$).

Description	Parameter	Posterior Distribution	
		Mean	95% Post. Prob. Int.
Inverse IES	σ^{-1}	6.37	[4.43,9.13]
Slope PC	κ	0.02	[0.003,0.054]
Discount Factor	β	0.99	-
IRS pre-79	ρ_{pre79}	0.94	[0.87-0.99]
Feedback Infl. pre79	$\chi_{\pi,pre-79}$	1.37	[0.91-1.82]
Feedback Gap pre79	$\chi_{x,pre-79}$	0.63	[0.26-1.11]
IRS post-79	ρ_{post79}	0.93	[0.88-0.97]
Feedback Infl. post79	$\chi_{\pi,post-79}$	1.57	[1.04-2.04]
Feedback Gap post79	$\chi_{x,post-79}$	0.49	[0.04-0.88]
Autoregr. Cost-push shock	ρ_u	0.402	[0.28-0.52]
Autoregr. Demand shock	ρ_g	0.83	[0.73-0.92]
Std. Cost-push shock	σ_u	0.90	[0.81-0.99]
Std. Demand shock	σ_g	0.64	[0.58-0.71]
Std. MP shock	σ_ε	0.97	[0.87-1.07]
Constant gain (Infl.)	$\bar{\mathbf{g}}_\pi$	0.078	[0.062-0.097]
Decreasing gain (Infl.)	t^{-1}	-	-
Constant gain (Gap)	$\bar{\mathbf{g}}_x$	0.073	[0.06-0.084]
Decreasing gain (Gap)	t^{-1}	-	-
Constant gain (FFR)	$\bar{\mathbf{g}}_i$	0.001	[0,0.01]
Decreasing gain (FFR)	t^{-1}	-	-
Forecast Errors Window	J	4	[1,6]

Table 4 - Posterior Distributions: Alternative case with estimated \hat{J} .

Description	Parameter	Posterior Distribution	
		Mean	95% Post. Prob. Int.
Inverse IES	σ^{-1}	7.44	[5.22-10.34]
Slope PC	κ	0.025	[0.003-0.07]
Discount Factor	β	0.99	-
IRS pre-79	ρ_{pre79}	0.932	[0.856-0.99]
Feedback Infl. pre79	$\chi_{\pi,pre-79}$	1.33	[0.86-1.82]
Feedback Gap pre79	$\chi_{x,pre-79}$	0.62	[0.22-1]
IRS post-79	ρ_{post79}	0.91	[0.86-0.96]
Feedback Infl. post79	$\chi_{\pi,post-79}$	1.56	[1.1-2.02]
Feedback Gap post79	$\chi_{x,post-79}$	0.44	[-0.04-0.85]
Autoregr. Cost-push shock	ρ_u	0.31	[0.18-0.46]
Autoregr. Demand shock	ρ_g	0.79	[0.7-0.87]
Std. Cost-push shock	σ_u	0.91	[0.82-1.01]
Std. Demand shock	σ_g	0.74	[0.66-0.82]
Std. MP shock	σ_ε	0.96	[0.87-1.07]
High Constant gain (Infl.)	$\bar{\mathbf{g}}_\pi^H$	0.096	[0.09-0.104]
Low Constant gain (Infl.)	$\bar{\mathbf{g}}_\pi^L$	0.02	-
High Constant gain (Gap)	$\bar{\mathbf{g}}_x^H$	0.042	[0.036-0.052]
Low Constant gain (Gap)	$\bar{\mathbf{g}}_x^L$	0.02	-
Constant gain (FFR)	$\bar{\mathbf{g}}_i$	0.01	-

Table 5 - Posterior Distributions: constant-gain learning ('low' and 'high' constant gains).

	Endogenous TV Gain				No Learning	
	$J = 4$		$J = 20$		ARCH(1)	GARCH(1,1)
	ARCH(1)	GARCH(1,1)	ARCH(1)	GARCH(1,1)		
Inflation	0.517	0.61	0.48	0.56	0.05	0.06
Output Gap	0.785	0.89	0.85	0.90	0.045	0.05

Table 6 - Test for the existence of ARCH/GARCH effects (5% significance): proportion of rejections of the null hypothesis of no ARCH/GARCH effects.

	Endogenous TV Gain			No Learning	Data
	Baseline	$J = 20$	CG		
Ratio $\frac{Std. Infl. 1985-2006}{Std. Infl. 1960-1984}$	0.39	0.42	0.43	1.00	0.35
Ratio $\frac{(Std. OutputGap 1985-2006)}{(Std. Output Gap 1960-1984)}$	0.42	0.52	0.54	1.00	0.50

Table 7 - The Great Moderation: ratio of standard deviations for inflation and output gap in the second versus the first part of the simulated samples (median across simulations).

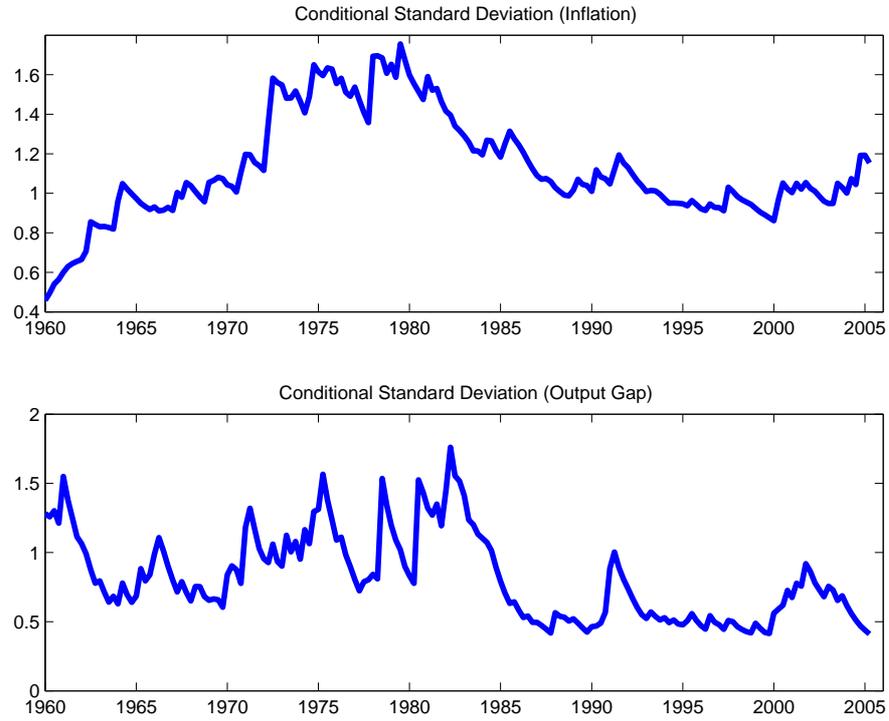


FIGURE 1. Conditional Standard Deviation series for Inflation and Output Gap

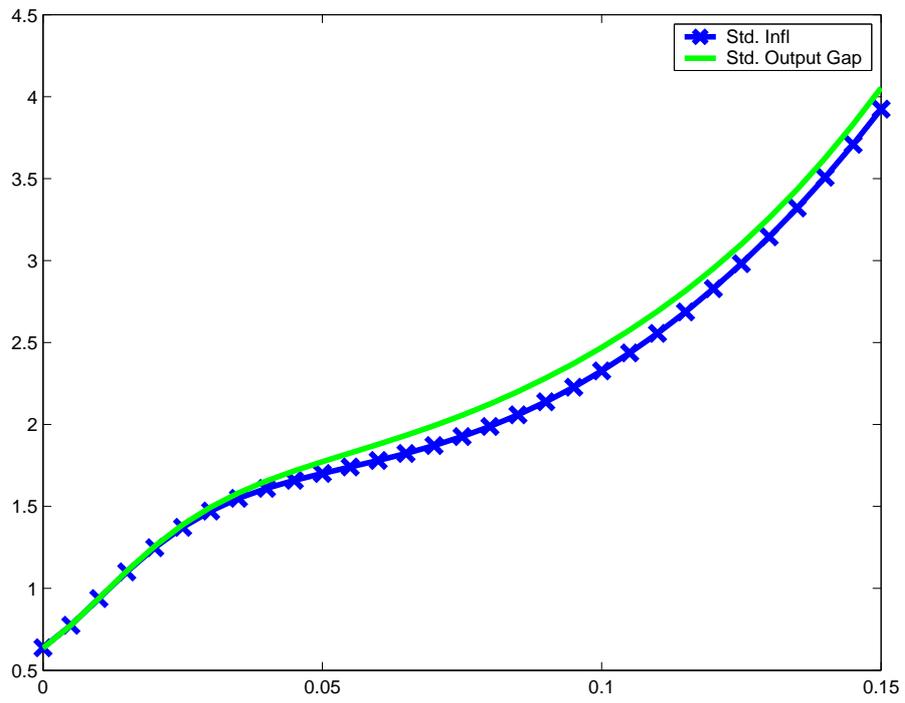


FIGURE 2. Volatility of simulated Inflation and Output Gap as a function of the constant gain coefficient.

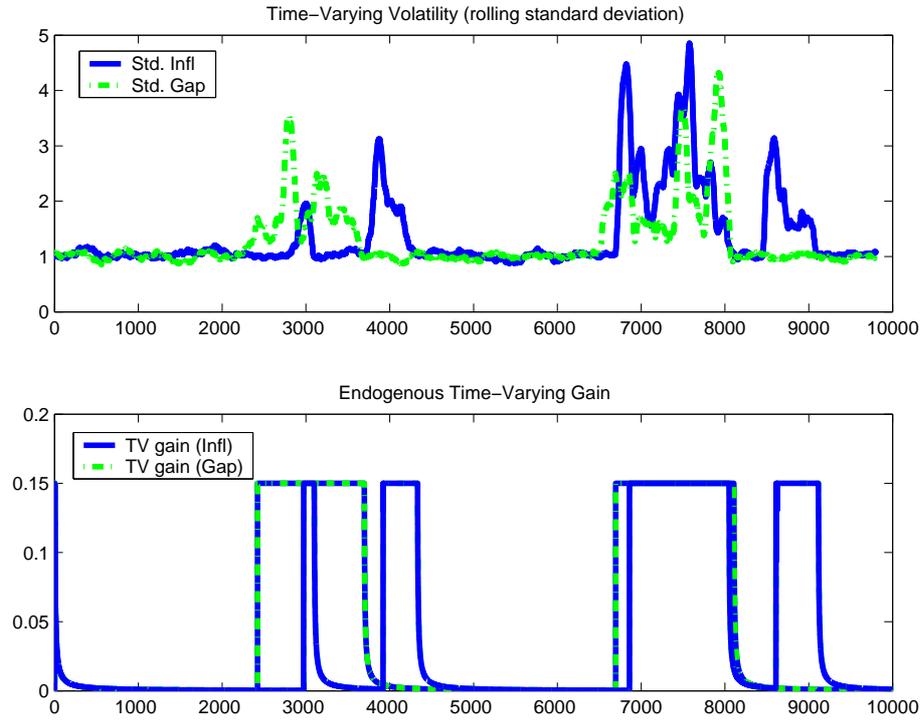


FIGURE 3. Time-Varying Volatility with Time-Varying Endogenous Gain Coefficient.

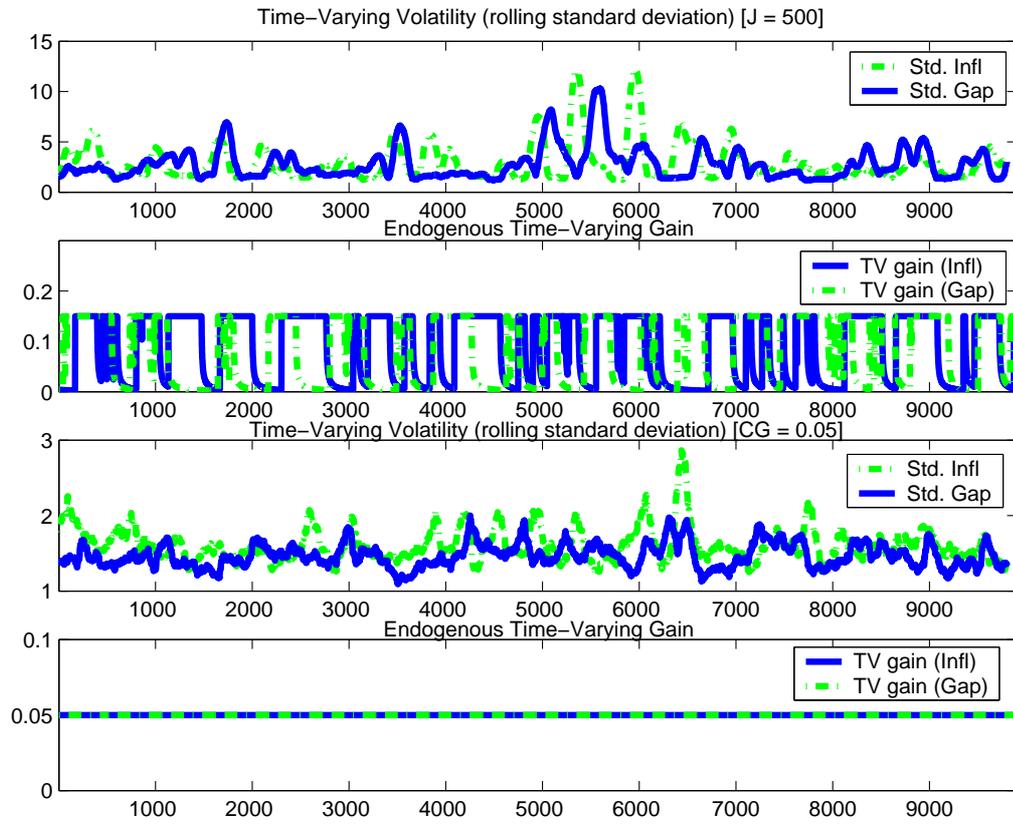


FIGURE 4. Time-Varying Volatility: Additional Cases.

Upper Plots: Time-Varying Endogenous Gain with $J = 500$ and $\bar{g} = 0.15$.
 Lower Plots: Constant Gain with $\bar{g} = 0.05$.

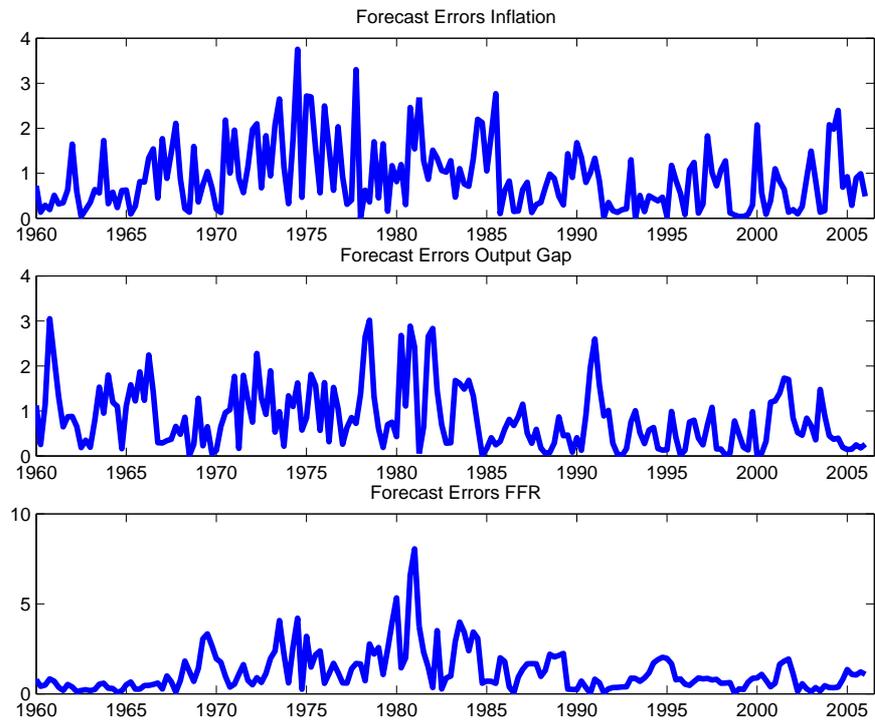


FIGURE 5. Forecast errors for inflation, output gap, and federal funds rate (absolute values).

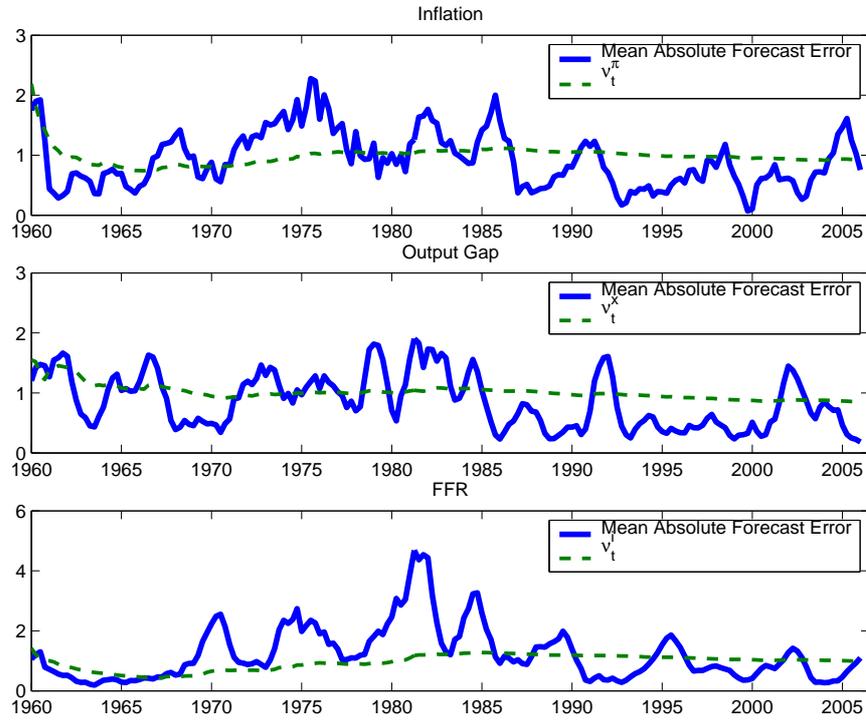


FIGURE 6. Rolling Mean Absolute Forecast errors vs. Updated ν_t for inflation, output gap, and federal funds rate series.

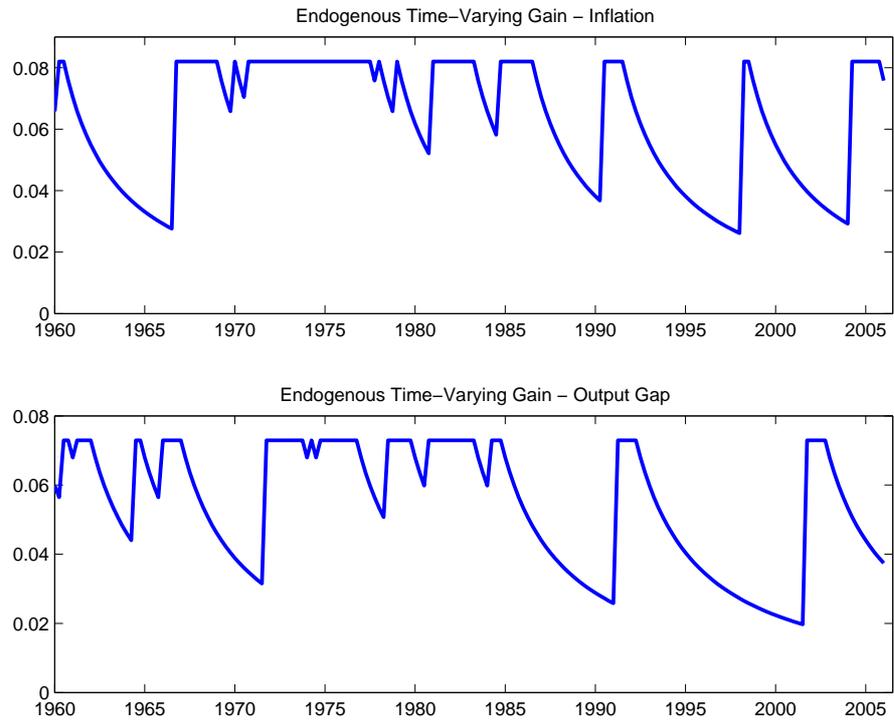


FIGURE 7. Endogenous Time-Varying Gain Coefficients (estimated constant gain). Baseline Case

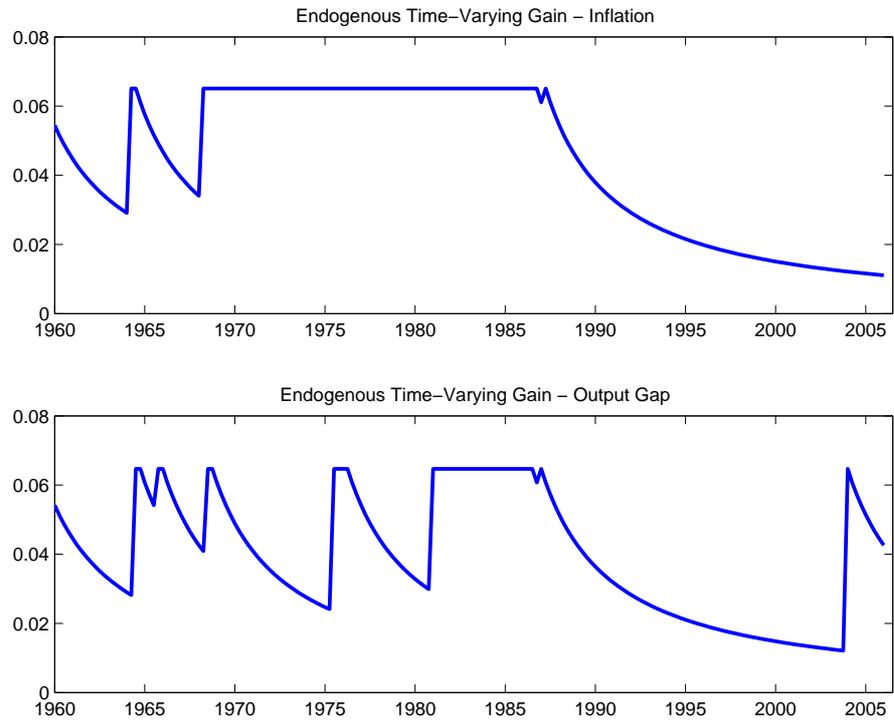


FIGURE 8. Endogenous Time-Varying Gain Coefficients (estimated constant gain). Case with $J = 20$

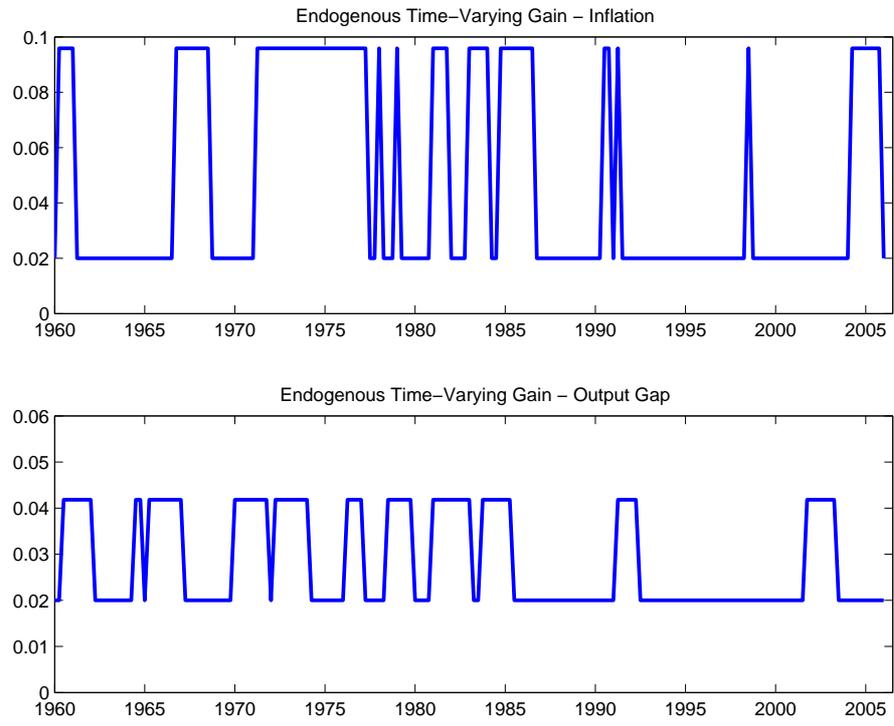


FIGURE 9. Endogenous Time-Varying Gain Coefficients
(Case with low and high constant gain coefficients only).

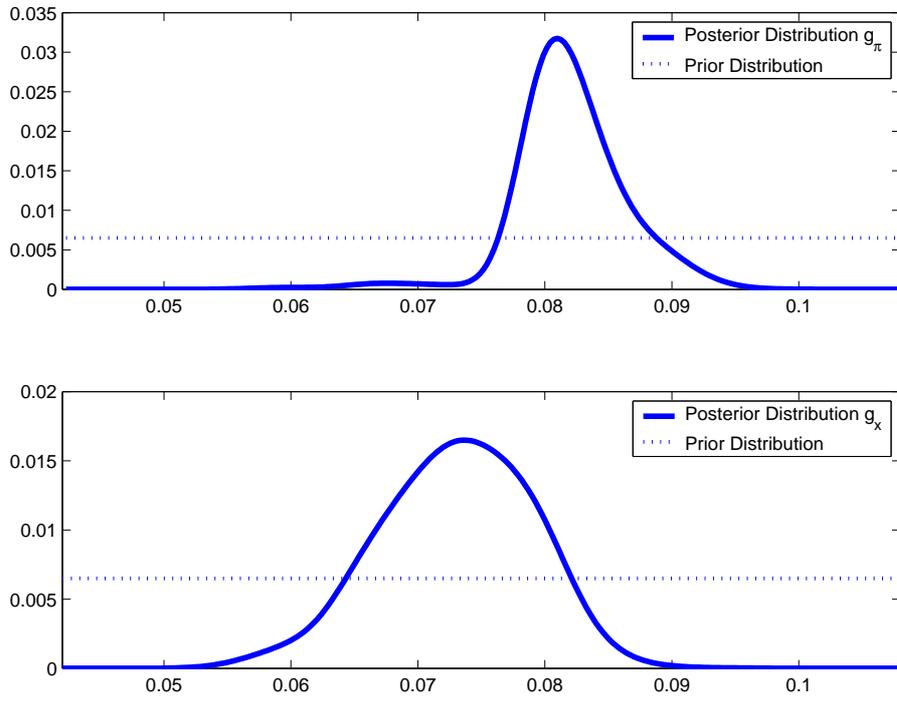


FIGURE 10. Constant Gain Coefficients: Prior and Posterior Distributions.

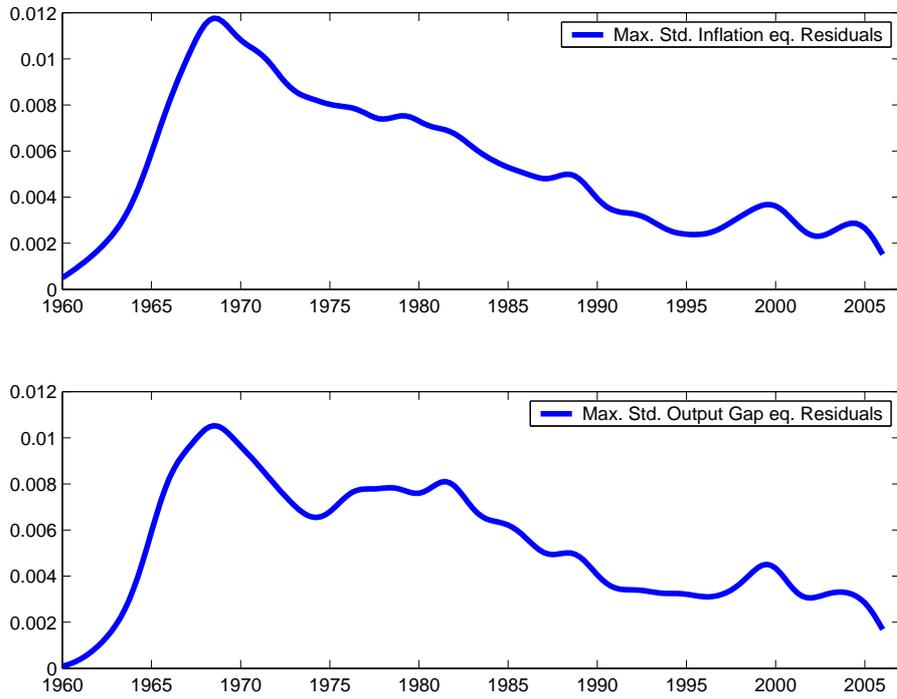


FIGURE 11. Kernel Density Estimation: sample observation in correspondence of which the Maximum Rolling Standard Deviation of residuals in the sample is obtained, across simulations.

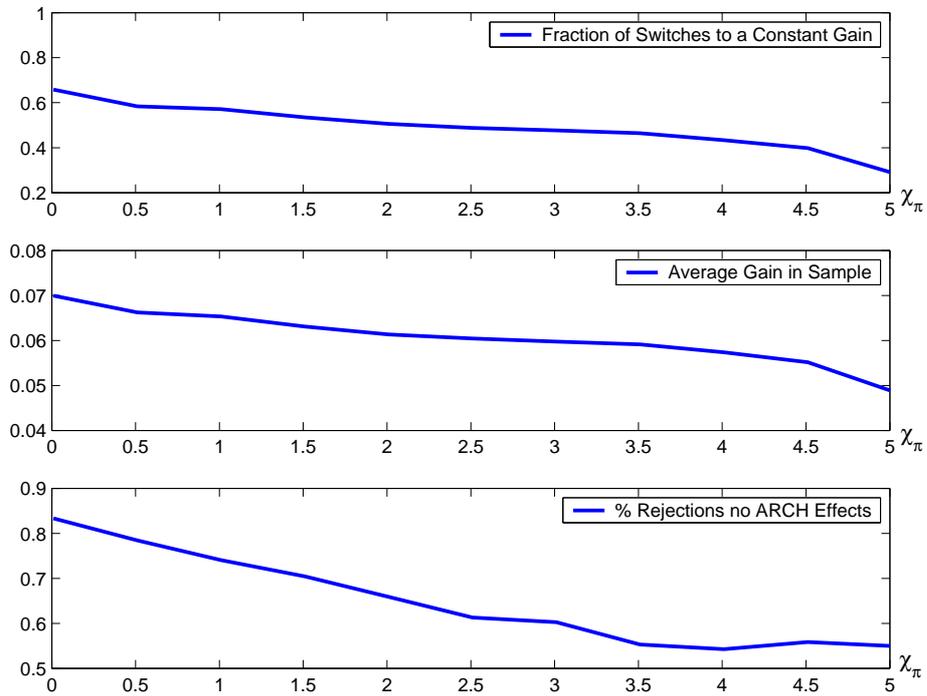


FIGURE 12. Effects of Monetary Policy on Volatility (Simulation under estimated parameters).

Graph 1: Fraction of sample periods in which agents' learning switches to constant gain as function of policy feedback to inflation χ_π ; Graph 2: Average gain in the sample as function of policy feedback to inflation χ_π ; Graph 3: Percentage of rejections of the null of no ARCH effects on the residuals as function of policy feedback to inflation χ_π .