

NBER WORKING PAPER SERIES

INFLATION-GAP PERSISTENCE IN THE U.S.

Timothy Cogley
Giorgio E. Primiceri
Thomas J. Sargent

Working Paper 13749
<http://www.nber.org/papers/w13749>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2008

For comments and suggestions, we thank James Kahn, Spencer Krane, and seminar participants at the Federal Reserve Board, the Federal Reserve Bank of Chicago, the Summer 2007 meetings of the Society for Computational Economics, and the EABCN Workshop on "Changes in Inflation Dynamics and Implications for Forecasting." We are also grateful to Francisco Barillas and Christian Matthes for research assistance. Sargent thanks the National Science Foundation for research support through a grant to the National Bureau of Economic Research. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2008 by Timothy Cogley, Giorgio E. Primiceri, and Thomas J. Sargent. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inflation-Gap Persistence in the U.S.

Timothy Cogley, Giorgio E. Primiceri, and Thomas J. Sargent

NBER Working Paper No. 13749

January 2008

JEL No. C11,C15,C32,E3,E52

ABSTRACT

We use Bayesian methods to estimate two models of post WWII U.S. inflation rates with drifting stochastic volatility and drifting coefficients. One model is univariate, the other a multivariate autoregression. We define the inflation gap as the deviation of inflation from a pure random walk component of inflation and use both of our models to study changes over time in the persistence of the inflation gap measured in terms of short- to medium-term predicability. We present evidence that our measure of the inflation-gap persistence increased until Volcker brought mean inflation down in the early 1980s and that it then fell during the chairmanships of Volcker and Greenspan. Stronger evidence for movements in inflation gap persistence emerges from the VAR than from the univariate model. We interpret these changes in terms of a simple dynamic new Keynesian model that allows us to distinguish altered monetary policy rules and altered private sector parameters.

Timothy Cogley
University of California, Davis
Economics Department
One Shields Ave.
Davis, CA 95616
twcogley@ucdavis.edu

Thomas J. Sargent
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
thomas.sargent@nyu.edu

Giorgio E. Primiceri
Northwestern University
Department of Economics
2001 Sheridan Road
3218 Andersen Hall
Evanston, IL 60208-2600
and NBER
g-primiceri@northwestern.edu

Inflation-Gap Persistence in the U.S.*

Timothy Cogley,[†] Giorgio E. Primiceri,[‡] and Thomas J. Sargent[§]

Revised: December 2007

Abstract

We use Bayesian methods to estimate two models of post WWII U.S. inflation rates with drifting stochastic volatility and drifting coefficients. One model is univariate, the other a multivariate autoregression. We define the inflation gap as the deviation of inflation from a pure random walk component of inflation and use both models to study changes over time in the persistence of the inflation gap measured in terms of short- to medium-term predicability. We present evidence that our measure of the inflation-gap persistence increased until Volcker brought mean inflation down in the early 1980s and that it then fell during the chairmanships of Volcker and Greenspan. Stronger evidence for movements in inflation gap persistence emerges from the VAR than from the univariate model. We interpret these changes in terms of a simple dynamic new Keynesian model that allows us to distinguish altered monetary policy rules and altered private sector parameters.

1 Introduction

This paper studies how inflation persistence has changed since the Great Inflation. We distinguish the persistence of inflation from the persistence of a component of it called the inflation gap. Our first message is that although inflation remains highly persistent, the inflation gap became less persistent after the Volcker disinflation. Our second message is that multivariate information helps to detect changes in inflation-gap persistence. Although the univariate evidence is mixed, a clearer picture emerges

*For comments and suggestions, we thank James Kahn, Spencer Krane, and seminar participants at the Federal Reserve Board, the Federal Reserve Bank of Chicago, the Summer 2007 meetings of the Society for Computational Economics, and the EABCN Workshop on “Changes in Inflation Dynamics and Implications for Forecasting.” We are also grateful to Francisco Barillas and Christian Matthes for research assistance. Sargent thanks the National Science Foundation for research support through a grant to the National Bureau of Economic Research.

[†]University of California, Davis. Email: twcogley@ucdavis.edu.

[‡]Northwestern University. Email: g-primiceri@northwestern.edu.

[§]New York University and Hoover Institution, Stanford University. Email: ts43@nyu.edu.

from a VAR. Our third message is that the decline in inflation-gap persistence seems to be due for the most part to lower variability of changes in the Fed’s long-run inflation target.

We decompose inflation into two parts, a stochastic trend τ_t that (to a first-order approximation) evolves as a driftless random walk, and an inflation gap $g_t = \pi_t - \tau_t$ that represents temporary differences between actual and trend inflation. In general equilibrium models, trend inflation is usually pinned down by a central bank’s target, a view that associates movements in trend inflation with shifts in the Federal Reserve’s target. Because trend inflation is a driftless random walk, actual inflation has a unit autoregressive root and is highly persistent. In our view, target inflation has not stopped drifting, though its conditional variance has declined.¹

Transient movements in the inflation gap are layered on top of τ_t . (Cogley and Sargent 2001 and 2005a) reported weak evidence of a decline in inflation-gap persistence. Several authors have challenged the statistical significance of that evidence (e.g., see Sims 2001, Stock 2001, and Pivetta and Reis 2007). Here we report new evidence that is more decisive. We can now say that it is very likely that inflation-gap persistence has decreased since the Great Inflation.

We organize the discussion as follows. We begin with an unobserved components model of Stock and Watson (2007) and relate it to the drifting-parameter VARs of Cogley and Sargent (2005a) and Primiceri (2005). We use these statistical models to define trend inflation and to focus attention on the inflation gap.

Next we define a measure of persistence in terms of the predictability of the inflation gap,² in particular, as the fraction of total inflation-gap variation j quarters ahead that is due to shocks inherited from the past. We say that the inflation gap is weakly persistent when the effects of shocks decay quickly and that it is strongly persistent when they decay slowly. When the effects of past shocks die out quickly, future shocks account for most of the variation in g_{t+j} , pushing our measure close to zero. But when the effects of past shocks on g_{t+j} decay slowly, they account for a higher proportion of near-term movements, pushing our measure of persistence closer to one. Thus, a large fraction of variation over short to medium horizons that is due to past shocks signifies strong persistence and a small fraction indicates weak persistence.

Under a convenient approximation, our measure is the R^2 statistic for j -step ahead inflation-gap forecasts.³ Heuristically, a connection between predictability and

¹For evidence that the innovation variance for τ_t has declined, see Stock and Watson (2007).

²This measure is inspired by Diebold and Kilian (2001). Barsky (1987) used a closely-related measure to compare inflation persistence under the Gold Standard and after World War II.

³Strictly speaking, we should say ‘pseudo forecasts’ because we neglect complications associated with real-time forecasting. This is not a shortcut; it is intentional. Our goal is to make retrospective statements about inflation persistence. To attain as much precision as possible, we use ex post revised data and estimate parameters using data through the end of the sample.

persistence arises because past shocks give rise to forecastable movements in g_{t+j} , while future shocks contribute to the forecast error. Hence, the continuing influence of past shocks can be measured by the proportion of predictable variation in g_{t+j} .

We deduce persistence measures from the posterior distribution of a drifting-parameter VAR, then study how they have changed since the Great Inflation. A key finding is that inflation gaps were highly predictable circa 1980, but are much less so now. Furthermore, the evidence of declining persistence is statistically significant at conventional levels. Thus, the statistical results strengthen our conviction that the inflation gap has become less persistent.

Finally, we use a simple dynamic new Keynesian model to examine what caused the change in the law of motion for inflation. In our DSGE model, improved monetary policy is the single most important factor explaining the decline in inflation volatility and persistence. A key dimension is the reduction in the rate at which the Fed's target drifts. Nevertheless, nonpolicy factors are also important; in particular, we find that mark-up shocks have become less volatile and persistent, and this also contributes to better inflation outcomes. In our model, better policy and changes in the private sector both play a role.

2 Unobserved components models for inflation

Stock and Watson (2007) estimate a univariate unobserved components model for inflation. They assume that inflation π_t is the sum of a stochastic trend τ_t and a martingale-difference innovation $\varepsilon_{\pi t}$,

$$\pi_t = \tau_t + \varepsilon_{\pi t}. \quad (1)$$

The trend component evolves as a driftless random walk,

$$\tau_t = \tau_{t-1} + \varepsilon_{\tau t}. \quad (2)$$

Equation (1) is the measurement equation for a state-space representation, and equation (2) is the state equation. The innovations $\varepsilon_{\pi t}$ and $\varepsilon_{\tau t}$ are assumed to be martingale differences that are conditionally normal with variances $h_{\pi t}$ and $h_{\tau t}$, respectively. The latter are independent stochastic volatilities that evolve as geometric random walks,

$$\begin{aligned} \ln h_{\pi t} &= \ln h_{\pi t-1} + \sigma_{\pi} \eta_{\pi t}, \\ \ln h_{\tau t} &= \ln h_{\tau t-1} + \sigma_{\tau} \eta_{\tau t} \end{aligned} \quad (3)$$

where $\eta_{\pi t}$ and $\eta_{\tau t}$ are i.i.d. Gaussian shocks with means of zero that are mutually independent.

A consensus has emerged that trend inflation is well approximated by a driftless random walk. Authors who model trend inflation in this way include Cogley and Sargent (2001, 2005a), Ireland (2007), Smets and Wouters (2003), and Cogley and Sbordone (2006). There is little controversy about this feature of the data. Our main focus, however, is on the inflation gap, $g_t \equiv \pi_t - \tau_t$. We want to know how persistent g_t is and whether the degree of persistence in g_t has changed over time. Stock and Watson's (2007) model is not a suitable vehicle for investigating this issue because it imposes that $g_t \equiv \epsilon_{\pi t}$ is serially uncorrelated for all t .

Reading the literature on inflation persistence can be confusing because authors sometimes fail to state clearly what feature of the data they are trying to measure. For instance, Pivetta and Reis (2007) look for changes in inflation persistence by running rolling unit-root tests on π_t . They find that the largest autoregressive root is always close to 1 and conclude that inflation persistence is unchanged. But that finding can be viewed simply as a manifestation of shifts in target inflation: π_t has a unit root because τ_t drifts. Estimates of the largest autoregressive root in π_t would help measure inflation-gap persistence only if trend inflation were constant over time, an assumption that much of the recent literature denies.⁴

Stock and Watson's specification is a useful starting point because it highlights the role of τ_t . But it is not a good vehicle for pursuing questions about inflation-gap persistence because it assumes that $g_t = \epsilon_{\pi t}$ is a martingale difference. To address the questions about the persistence of g_t that interest us, we must modify their model.

Cogley and Sargent (2005a) estimate a closely related time-varying parameter VAR. Evidence reported there suggests that g_t is autocorrelated and that the degree of serial dependence has probably changed over time. But that model assumed no stochastic volatility in the parameter innovations, a feature that Stock and Watson say is important. In this paper, we combine and extend features of Stock and Watson's model and our earlier ones to create a new model that lets us focus on the persistence of the inflation gap.

2.1 A univariate autoregression with drifting parameters

As a first step, we introduce an autoregressive term into Stock and Watson's representation. With this addition, the measurement and state equations become

$$\pi_t = \mu_{t-1} + \rho_{t-1}\pi_{t-1} + \epsilon_{\pi t}, \quad (4)$$

$$\gamma_t = \gamma_{t-1} + \epsilon_{st}, \quad (5)$$

where $\gamma_t = [\mu_t, \rho_t]'$ and $\epsilon_{st} = [\epsilon_{\mu t}, \epsilon_{\rho t}]'$. Here the vector ϵ_{st} is the noise in a *state* vector, whose components are parameter values in the measurement equation (4). Notice that

⁴Levin and Piger (2004) pointed out this shortcoming of Pivetta and Reis. After allowing for a shift in trend inflation, Levin and Piger were able to detect a decline in inflation-gap persistence.

the constant term in the measurement equation has become an intercept rather than a local approximation to the mean.⁵ As in our earlier work, we approximate trend inflation by $\tau_t \doteq \mu_t/(1 - \rho_t)$. To a first-order approximation, this is also a driftless random walk.⁶

Equations (4) and (5) describe a univariate autoregression with drifting parameters. If the innovation variances were all constant, (4)–(5) would be a special case of the time-varying parameter model of Cogley and Sargent (2001). In Cogley and Sargent (2005a) and Primiceri (2005), the measurement-innovation variance is time-varying, but the variance of the state-innovation ε_{st} variance is constant. In contrast, Stock and Watson assume that both innovation variances are time varying. Here we follow Stock and Watson by modeling both innovation variances as stochastic volatility processes.

We retain Stock and Watson’s specification for $\text{var}(\varepsilon_{\pi t})$, and we adopt a bivariate stochastic volatility model for the state innovations ε_{st} :

$$\text{var}(\varepsilon_{st}) = Q_t = B^{-1} H_{st} B^{-1'}. \quad (6)$$

As in our earlier work, we assume that H_{st} is diagonal and that B is lower triangular,

$$H_{st} = \begin{pmatrix} h_{\mu t} & 0 \\ 0 & h_{\rho t} \end{pmatrix}, \quad (7)$$

$$B = \begin{pmatrix} 1 & 0 \\ \beta_{21} & 1 \end{pmatrix}. \quad (8)$$

The diagonal elements of H_{st} are independent, univariate stochastic volatilities that evolve as driftless, geometric random walks:

$$\ln h_{it} = \ln h_{it-1} + \sigma_i \eta_{it}, \quad (9)$$

$i = \pi, s$. The volatility innovations η_{it} are mutually independent, standard normal variates. The variance of $\Delta \ln h_{it}$ depends on the free parameter σ_i . For tractability and parsimony, we also assume that ε_{st} is uncorrelated at all leads and lags with $\varepsilon_{\pi t}$ and that the standardized state and measurement innovations are independent of the volatility innovations η_t .

This is a convenient specification for modeling recurrent persistent changes in variance. It ensures that Q_t is positive definite and allows for time-varying correlations among the elements of ε_{st} .

We constrain ρ_t to be less than one in absolute value at all dates. Having assumed that trend inflation is a driftless random walk, the stability constraint on ρ_t just rules

⁵We also adopt a slightly different dating convention. The reason for this dating convention will become clear when we discuss predictability. Nothing of substance hinges on this convention.

⁶A first-order Taylor approximation makes τ_t a linear function of γ_t , which evolves as a driftless random walk.

out a second unit or explosive root in inflation. There is an emerging consensus that the price level is best modeled as an $I(2)$ process, but few economists think that it is $I(3)$. The stability constraint just rules out an $I(3)$ representation.

The model is estimated by Bayesian methods using a Markov Chain Monte Carlo algorithm outlined in appendix A.

2.2 A vector autoregression with drifting parameters

Although a univariate autoregression is a useful first step, it is not entirely satisfactory for representing changes in the inflation process. Cogley and Sargent (2001 and 2005a) found evidence of changes in the autocorrelations of the inflation gap and also in cross-correlations with lags of other variables. Accordingly, we also consider a vector autoregression with drifting parameters. Since our definition of the persistence of g_t is based on its predictability, it is interesting to check how findings depend on the information that we use to condition predictions.

As in Cogley and Sargent (2005a), we estimate a trivariate VAR for inflation, unemployment, and a short-term nominal interest rate. The state and measurement equations for the VAR are

$$y_t = X'_{t-1}\theta_{t-1} + \varepsilon_{yt}, \quad (10)$$

$$\theta_t = \theta_{t-1} + \varepsilon_{st}. \quad (11)$$

The vector y_t contains current observations on the variables of interest, X_{t-1} includes constants plus lags of y_t , and ε_{yt} is a vector of innovations. The parameter vector θ_t evolves as a driftless random walk subject to a reflecting barrier that guarantees that the VAR has nonexplosive roots at every date.

We assume that the innovation variances follow multivariate stochastic volatility processes. The state innovation variance Q_t has the same form as in the $AR(1)$ model, but has a higher dimension to conform to the size of θ_t . We assume that the measurement innovation variance V_t also has this form, again adapting its dimensions to the size of ε_{yt} .

This model is very much like those in Cogley and Sargent (2005a) and Primiceri (2005). The main difference concerns the specification for $\text{var}(\varepsilon_{st})$. Our earlier papers assumed that the parameter innovation variance was constant; here we adopt a stochastic volatility model so that the variance is time varying. Equations (10) and (11) can also be regarded as a multivariate extension of Stock and Watson (2007). We think this model is a useful vehicle for connecting their paper to this one.

We estimate the multivariate model by a Bayesian Markov Chain Monte Carlo algorithm. Details are given in appendix A.1.

In what follows, we make frequent use of the companion form of the VAR,

$$z_{t+1} = \mu_t + A_t z_t + \varepsilon_{zt+1}. \quad (12)$$

The vector z_t includes current and lagged values of y_t , the vector μ_t contains the VAR intercepts, and the companion matrix A_t contains the autoregressive parameters. We use the companion form for multi-step forecasting. When we do that, we approximate multi-step forecasts by assuming that VAR parameters will remain constant at their current values going forward in time. This approximation is common in the literature on bounded rationality and learning, being a key element of an ‘anticipated-utility’ model (Kreps 1998). In other papers, we have found that it does a good job of approximating the mean of Bayesian predictive densities (e.g., see Cogley, Morozov, and Sargent 2005 and Cogley and Sargent 2006).

With this assumption, we can form local-to-date t approximations to the moments of z_t . For the unconditional mean, we follow Beveridge and Nelson (1981) by defining the stochastic trend in z_t as the value to which the series is expected to converge in the long run:

$$\bar{z}_t = \lim_{h \rightarrow \infty} E_t z_{t+h}. \quad (13)$$

With θ_t held constant at its current value, we approximate this as

$$\bar{z}_t \cong (I - A_t)^{-1} \mu_t. \quad (14)$$

To a first-order approximation, \bar{z}_t evolves as a driftless random walk, implying that inflation and the other variables in y_t have a unit root. As in the AR(1) model, the stability constraint on A_t just rules out an $I(2)$ representation for y_t .

After subtracting \bar{z}_t from both sides of (12) and invoking the anticipated-utility approximation, we get a forecasting model for gap variables,

$$(z_{t+1} - \bar{z}_t) = A_t(z_t - \bar{z}_t) + \varepsilon_{z,t+1}. \quad (15)$$

We approximate forecasts of gap variables j periods ahead as $A_t^j \hat{z}_t$,⁷ and we approximate the forecast-error variance by

$$\text{var}_t(\hat{z}_{t+j}) \cong \sum_{h=0}^{j-1} (A_t^h) \text{var}(\varepsilon_{z,t+1}) (A_t^h)'. \quad (16)$$

To approximate the unconditional variance of \hat{z}_{t+1} , we take the limit of the conditional variance as the forecast horizon j increases,⁸

$$\text{var}(\hat{z}_{t+1}) \cong \sum_{h=0}^{\infty} (A_t^h) \text{var}(\varepsilon_{z,t+1}) (A_t^h)'. \quad (17)$$

Under the anticipated-utility approximation, this is also the unconditional variance of \hat{z}_{t+s} for $s > 1$.

⁷By the anticipated-utility approximation, $E_t \bar{z}_{t+j} = \bar{z}_t$. This is a good approximation because \bar{z}_t is a driftless random walk to a first-order approximation.

⁸This is a second-moment analog to the Beveridge-Nelson trend.

3 Persistence and predictability

Let $\pi_t = e_\pi z_t$, where e_π is a selector vector. To measure persistence at a given date t , we calculate the fraction of the total variation in g_{t+j} that is due to shocks inherited from the past relative to those that will occur in the future. This is equivalent to 1 minus the fraction of the total variation due to future shocks. Since future shocks account for the forecast error, that fraction can be expressed as the ratio of the conditional variance to the unconditional variance,

$$\begin{aligned} R_{jt}^2 &= 1 - \frac{\text{var}_t(e_\pi \hat{z}_{t+j})}{\text{var}(e_\pi \hat{z}_{t+j})}, \\ &\cong 1 - \frac{e_\pi \left[\sum_{h=0}^{j-1} (A_t^h) \text{var}(\varepsilon_{zt+1}) (A_t^h)' \right] e_\pi'}{e_\pi \left[\sum_{h=0}^{\infty} (A_t^h) \text{var}(\varepsilon_{zt+1}) (A_t^h)' \right] e_\pi'}. \end{aligned} \quad (18)$$

We label this R_{jt}^2 because it is analogous to the R^2 statistic for j -step ahead forecasts. This fraction must lie between zero and one, and it converges to zero as the forecast horizon j lengthens.⁹ Whether it converges rapidly or slowly reflects the degree of persistence. If past shocks die out quickly, the fraction converges rapidly to zero. But if one or more shocks decay slowly, the fraction may converge only gradually to zero, possibly remaining close to one for some time. Thus, for small or medium $j \geq 1$, a small fraction signifies weak persistence and a large fraction strong persistence.

In a univariate $AR(1)$ model, things simplify because R_{jt}^2 depends on a single parameter ρ_t . In this case, the unconditional variance is $\sigma_{\varepsilon t}^2 / (1 - \rho_t^2)$, and the conditional variance is $(1 - \rho_t^{2j}) \sigma_{\varepsilon t}^2 / (1 - \rho_t^2)$. Therefore, R_{jt}^2 simplifies to ρ_t^{2j} .

Matters are more complicated if we increase the number of lags or add other variables. For a VAR, the ratio depends on all of the parameters of the companion matrix A_t . Sometimes economists summarize persistence in a VAR by focusing on the largest autoregressive root in A_t . This is problematic for at least two reasons. One is that the largest root could be associated not with inflation but with another variable in the VAR. Hence the largest root of A_t might exaggerate persistence in the inflation gap. Another problem is that two large roots could matter for inflation, in which case the largest root of A_t would understate the degree of persistence. We think it is important to retain all the information in A_t .

3.1 A caveat

Nevertheless, (18) is not entirely satisfactory because it depends on the conditional variance V_{t+1} in addition to the conditional mean parameters A_t . Changes in V_{t+1} that take the form of a scalar multiplication are not a problem because the scalar would cancel in numerator and denominator. But R_{jt}^2 is not invariant to other changes

⁹This follows from the stability constraint on A_t .

in V_{t+1} . For instance, our measure of persistence would be reduced by a change in the composition of structural shocks away from those whose impulse response functions decay slowly and toward those whose impulse response functions vanish quickly.

This problem really relates to the question of why inflation persistence has changed, not whether it has changed. For the moment, we want to focus on the latter. We think that assembling evidence about the structure of inflation persistence is a step in the right direction.

In what follows, we focus on horizons of 1, 4, and 8 quarters, those being the most relevant for monetary policy. We calculate values of R_{jt}^2 implied by a drifting-parameter VAR and study how they have changed over time.

4 Properties of inflation

Inflation is measured either as the log-difference of the GDP or PCE chain-type price index. Stock and Watson (2007) examine GDP inflation. A number of colleagues in the Federal Reserve system encouraged us to look at PCE inflation as well, saying that the Fed pays more attention to that for policy purposes.

For the VAR, we also condition on unemployment and a short-term nominal interest rate. Unemployment is measured by the civilian unemployment rate. The original monthly series was converted to a quarterly basis by sampling the middle month of each quarter. As in Cogley and Sargent (2001 and 2005a), the logit of the unemployment rate enters the VAR. The nominal interest rate is measured by the secondary market rate on three-month Treasury bills. These data are also sampled monthly, and we converted to a quarterly series by selecting the first month of each quarter in order to align the nominal interest data as well as possible with the inflation data. For the VAR, the nominal interest rate is expressed as yield to maturity.

The inflation and unemployment data are seasonally adjusted, and all the data span the period 1948.Q1 to 2004.Q4. The data were downloaded from the Federal Reserve Economic Database (FRED).¹⁰

Our priors are described in the appendices. For the most part, they follow our earlier papers. Our guiding principle was to use proper priors to ensure that the posterior is proper, but to make the priors as weakly informative as possible, so that the posterior is dominated by information in the data.¹¹

¹⁰This can be found at <http://research.stlouisfed.org/fred2/>. The series have FRED mnemonics GDPCTPI, PCECTPI, UNRATE, and TB3MS, respectively

¹¹We think this is appropriate for exploratory data analysis. However it means that we cannot compare models via Bayes factors for reasons having to do with the Lindley paradox. E.g., see Gelfand (1996).

4.1 Trend inflation and inflation volatility

A number of our findings resemble those reported elsewhere (e.g. Cogley and Sargent 2005a, Stock and Watson 2007). We briefly touch on them before moving on to novel ones.

Figure 1 portrays the posterior median and interquartile range for τ_t . The left and right-hand columns depict estimates for the $AR(1)$ and VAR, respectively, while the top and bottom rows correspond to GDP and PCE inflation. Trend inflation is estimated using data through 2004.Q4. Accordingly, the figure represents a retrospective interpretation of the data.




Figure 1: Trend Inflation

The patterns shown here are similar to those reported in earlier papers. Trend inflation was low and steady in the early 1960s, it began rising in the mid-1960s, and it attained twin peaks around the time of the 1970s oil shocks. It fell sharply during the Volcker disinflation, and then settled down to the neighborhood of 2 percent after the mid-1990s. There are some differences between the $AR(1)$ and the VAR, and those differences will influence some properties of the inflation gap. Nevertheless, the broad contour of trend inflation is similar across models.

The next two figures summarize changes in inflation volatility. Once again, we plot the posterior median and interquartile range at each date. The top row in each figure shows the standard deviation for the inflation innovation, and the bottom row plots the unconditional standard deviation, $[e_\pi V_{\hat{z}_t} e'_\pi]^{1/2}$.

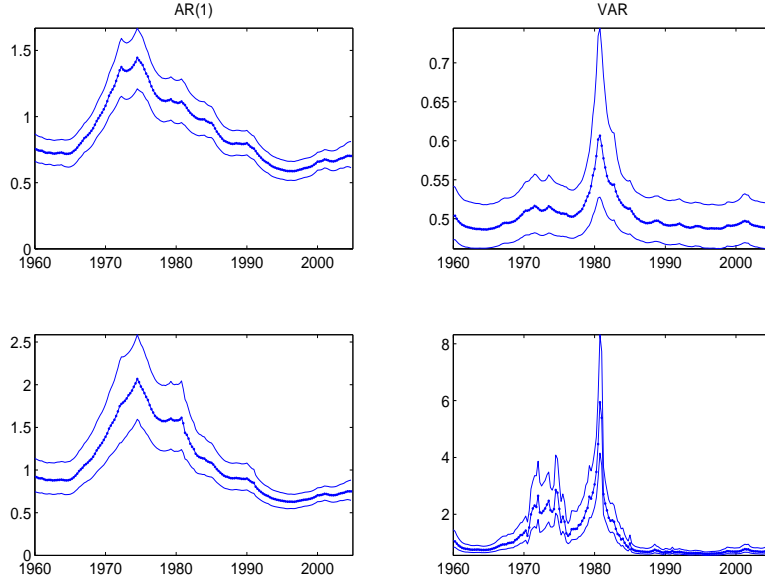


Figure 2: GDP Inflation Volatility

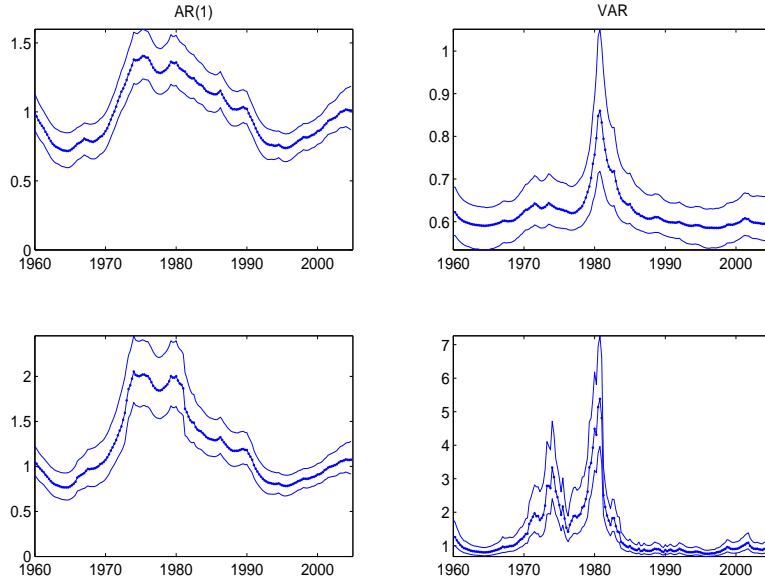


Figure 3: PCE Inflation Volatility

The patterns shown here are also familiar from earlier papers. For the univariate models, the innovation variance started rising in the mid 1960s and peaked around the time of the first oil shock. After that, the innovation variance declined gradually until the mid 1990s. The pattern for the VARs is a bit different. Instead of a gradual rise and fall, the VAR innovation variance remains roughly constant for most of

the sample, except for a spike in the late 1970s and early 1980s when the Fed was targeting monetary aggregates. That the innovation variances differ across univariate and multivariate models is not surprising because they portray different conditional variances. The VARs condition on more variables, and its innovation variance would be the same as in the univariate model only if the additional variables failed to Granger cause inflation. Since the additional variables were chosen precisely because they help forecast inflation, the VAR innovation variances are lower than the AR(1) innovation variances.

The bottom rows of figures 2 and 3 illustrate the unconditional standard deviation of inflation. For the AR(1) models, the general contour is similar to that of the innovation variance, but the magnitudes differ. The unconditional variance also rises and falls gradually, but it reaches a higher peak in the mid 1970s. For an AR(1), the unconditional variance can be expressed as $\sigma_{\pi t}^2 = \sigma_{\varepsilon t}^2 / (1 - \rho_t^2)$. If ρ_t were constant, movements in $\sigma_{\pi t}$ would mirror those in $\sigma_{\varepsilon t}$. From the patterns shown here, it follows that changes in the innovation variance account for much of the variation in the unconditional variance, but not all of it. Changes in ρ_t also matter. We say more about the contribution of ρ_t below.

Similar comments apply to the VARs, except that changes in the relative magnitudes of the two variances are even more pronounced. In the early 1980s, the standard deviation of VAR innovations rose by about 10 basis points, an increase of roughly 20 percent. At the same time, the unconditional standard deviation increased by roughly 4 percentage points, or about 200 percent. Hence for the VAR, changes in the innovation variance account for a relatively small proportion of changes in the unconditional variance. Most of the variation in the VAR unconditional variance must be due to changes in persistence.

Among other things, this means that a multivariate conditioning set is likely to be more helpful for detecting changes in inflation persistence. A univariate model may not use enough information.

4.2 Has the inflation gap become less persistent?

To focus more clearly on changes in persistence parameters, we turn to evidence on the predictability of the inflation gap. First we consider univariate evidence and then we turn to results from the VAR.

4.2.1 Univariate evidence

For the AR(1) model, everything depends on a single parameter ρ_t . Figure 4 portrays the posterior median and interquartile range for this parameter for the two inflation measures.




Figure 4: Posterior Median and Interquartile Range for ρ_t

For GDP inflation, the inflation gap is moderately persistent throughout the sample. The median estimate for ρ_t was around 0.55 in the early 1960s. It increased gradually to 0.7 by 1980, and then fell in two steps in the early 1980s and early 1990s, eventually reaching a value of 0.3 at the end of the sample. These estimates imply half-lives of 3.5, 5.8, and 1.7 months, respectively. For PCE inflation, the gap was initially less persistent, with an autocorrelation of 0.3, but otherwise movements in ρ_t are similar to those for GDP inflation. The patterns shown here are consistent with evidence reported in our earlier papers. Taken at face value, the figure suggests not only that inflation was lower on average during the Volcker-Greenspan years, but also that the inflation gap was less persistent.

The controversy about inflation persistence hinges not on the evolution of the posterior median or mean, however, but rather on whether changes in ρ_t are statistically significant. To assess this, we examine the joint posterior distribution for ρ_t across pairs of time periods. There are many possible pairs, of course, and to make the problem manageable we concentrate on two pairs, 1960-1980 and 1980-2004. The years 1960 and 2004 are the beginning and end of our sample, respectively.¹² We chose 1980.Q4 because it was the eve of the Volcker disinflation and because it splits the sample roughly in half. However, the results reported below are not particularly sensitive to this choice. Dates adjacent to 1980.Q4 tell much the same story.

Figures 5 and 6 depict results for GDP inflation. Figure 5 portrays the joint distribution for ρ_{1980} and ρ_{2004} , with values for 1980 plotted on the x -axis and those for 2004 on the y -axis. Combinations clustered near the 45 degree line represent pairs for which there was little or no change. Those below the 45 degree line represent a decrease in persistence ($\rho_{1980} > \rho_{2004}$), while those above represent increasing persistence. Similarly, figure 6 illustrates the joint distribution for ρ_{1960} and ρ_{1980} , with values for 1960 plotted on the x -axis and those for 1980 on the y -axis.

¹²Earlier data are used as a training sample for the prior.




Figure 5: Joint Distribution for ρ_{1980} and ρ_{2004} , GDP Inflation




Figure 6: Joint Distribution for ρ_{1960} and ρ_{1980} , GDP Inflation

A number of alternative perspectives can be represented on these graphs. Stock and Watson assume $\rho_t = 0$, so the point $(0,0)$ represents their model. There are some realizations in the neighborhood of the origin, but most of the probability mass lies elsewhere. The second column of table 1 reports the probability that ρ_t is close to zero in both periods, where ‘close’ is defined as $|\rho| < 0.1$. This comes out to 1.2 and 1.7 percent, respectively, for the two pairs of years. This finding motivates our extension of their model.

Table 1: Posterior Probabilities

GDP Inflation				
pair	Stock-Watson	$ \Delta\rho < 0.05$	High, No Change	Changing ρ
1980, 2004	0.012	0.122	<0.001	0.894
1960, 1980	0.017	0.384	0.027	0.758
PCE Inflation				
pair	Stock-Watson	$ \Delta\rho < 0.05$	High, No Change	Changing ρ
1980, 2004	0.006	0.056	<0.001	0.959
1960, 1980	0.008	0.066	<0.001	0.956

Sims (2001), Stock (2001), and Pivetta and Reis (2007) argue that inflation persistence is approximately unchanged. That perspective can be represented by drawing a neighborhood along the 45 degree line. As figures 5 and 6 show, the posterior attaches considerable probability mass to a ridge clustered tightly along the 45 degree line. How much probability is near that ridge depends on how a neighborhood is defined. For example, if we define ‘little change’ by the neighborhood $|\Delta\rho| < 0.05$, the posterior probability comes to 12 and 38 percent, respectively, for the two pairs of years. Obviously these probabilities would be higher if we widened the neighborhood and lower if we narrowed it, but the point is that the probability is nontrivial even for a narrowly defined interval along the 45 degree line. For the GDP deflator, the notion that univariate inflation-gap persistence is approximately constant cannot be rejected at the 10 percent level.

If we examine the ridges more closely, we notice that the scatterplots are densest along the ridge for low values of ρ and that they become sparse for high values. Thus, the notion that inflation-gap persistence is *both* unchanging *and* high has little support. For example, if we define ‘high persistence’ as a half-life of 1 year or more ($\rho \geq 0.8409$), the probability of high and unchanging persistence is less than one-tenth of 1 percent for 1980-2004 and 2.7 percent for 1960-1980. Inflation-gap persistence might have been high (especially during the Great Inflation), or it might have been unchanged, but it is unlikely that it was both. As noted above, the notion that persistence is both high and unchanging really applies to *inflation* – because of drift in τ_t – but not to the inflation *gap*.

In figure 5, the largest probability mass of points – a bit less than 90 percent – lies below the 45 degree line. For combinations in this region, $\rho_{1980} > \rho_{2004}$, so this represents the probability of declining inflation-gap persistence. We interpret this as substantial though not decisive evidence of a decline in persistence. Similarly, in figure 6, the preponderance of the combinations – approximately 75 percent – lie above the 45 degree line and are consistent with the idea that the inflation gap became more persistent between 1960 and 1980.

Thus, for GDP inflation the univariate evidence is mixed. While the preponderance of the joint distribution points to a rise and then a decline in persistence, there is enough mass along the 45 degree ridge in figures 5 and 6 to support the idea that inflation-gap persistence has not changed. This does not mean that the two interpretations stand on an equal footing; one has higher posterior probability than the other. But neither perspective overwhelms the other, and neither can be dismissed as unreasonable.

Figures 7 and 8 repeat this analysis for PCE inflation. For this measure, clear evidence emerges of a rise in persistence between 1960 and 1980 and a decline thereafter. In figures 7 and 8, the 45 degree ridges are more sparsely populated than those for GDP inflation, and the great majority of points lie below or above the line. The probability of an increase in ρ_t between 1960 and 1980 is 0.956, and the probability of a decline after 1980 is 0.959. This is significant evidence of changing inflation-gap persistence.




Figure 7: Joint Distribution for ρ_{1980} and ρ_{2004} , PCE Inflation




Figure 8: Joint Posterior for ρ_{1960} and ρ_{1980} , PCE Inflation

4.2.2 A pitfall: uncertainty at one time or across time?

Had we followed the methods of Pivetta and Reis (2007), we would not have detected these changes. Pivetta and Reis assess statistical significance by asking whether a horizontal line can be drawn through marginal confidence bands surrounding the mean or median. If it can, they conclude that the evidence for change is statistically insignificant. For both GDP and PCE inflation, marginal confidence bands for ρ_t overlap at all three dates. Hence we would have mistakenly concluded that the evidence for changing persistence is insignificant. Their procedure is difficult to interpret, however, because it confounds uncertainty about the level of ρ_t at a point in time with uncertainty about changes in ρ_t across dates. A marginal confidence band is fine for assessing level uncertainty at a point in time, but we must consult the joint distribution across dates in order to assess uncertainty about changes.¹³ For PCE inflation, the joint distribution points to significant changes in ρ_t .

4.3 Multivariate evidence

As noted above, the estimates reported in figures 2 and 3 suggest that VARs are more promising for detecting changes in inflation-gap persistence. Accordingly, we now turn to multivariate evidence. For each draw in the posterior distribution for VAR parameters, we calculate R_{jt}^2 statistics as in equation (18) and then study how they changed during and after the Great Inflation. Figure 9 portrays the posterior median and interquartile range for R_{jt}^2 for $j = 1, 4$, and 8 quarters.

The top row refers to 1-quarter ahead forecasts. In the mid 1960s, VAR pseudo forecasts accounted for approximately 50 to 55 percent of the variation of the inflation gap. During the Great Inflation, this increased to more than 90 percent and at times approached 99 percent. The inflation gap became less predictable during the Volcker disinflation, and after that R_{1t}^2 settled to the neighborhood of 50 percent. It was still around 50 percent at the end of the sample.

The second and third rows refer to 4 and 8 quarter forecasting horizons. As expected, R_{jt}^2 statistics are lower for longer horizons. For $j = 4$, VAR pseudo forecasts accounted for roughly a quarter of the inflation-gap variation in the mid 1960s, for approximately 50 to 75 percent during the Great Inflation, and for about 15 percent after the Volcker disinflation. For $j = 8$, the numbers follow a similar pattern but are lower. VAR pseudo forecasts accounted for about 10 percent of inflation-gap variation in the mid-1960s, for 20 to 35 percent during the mid 1970s and early 1980s, and for 10 percent or less after the Volcker disinflation. Thus, there was apparently a substantial decline in inflation-gap predictability after the mid 1980s.

¹³Sims and Zha (1999) make this point in the context of confidence bands for impulse response functions. Their logic applies here.

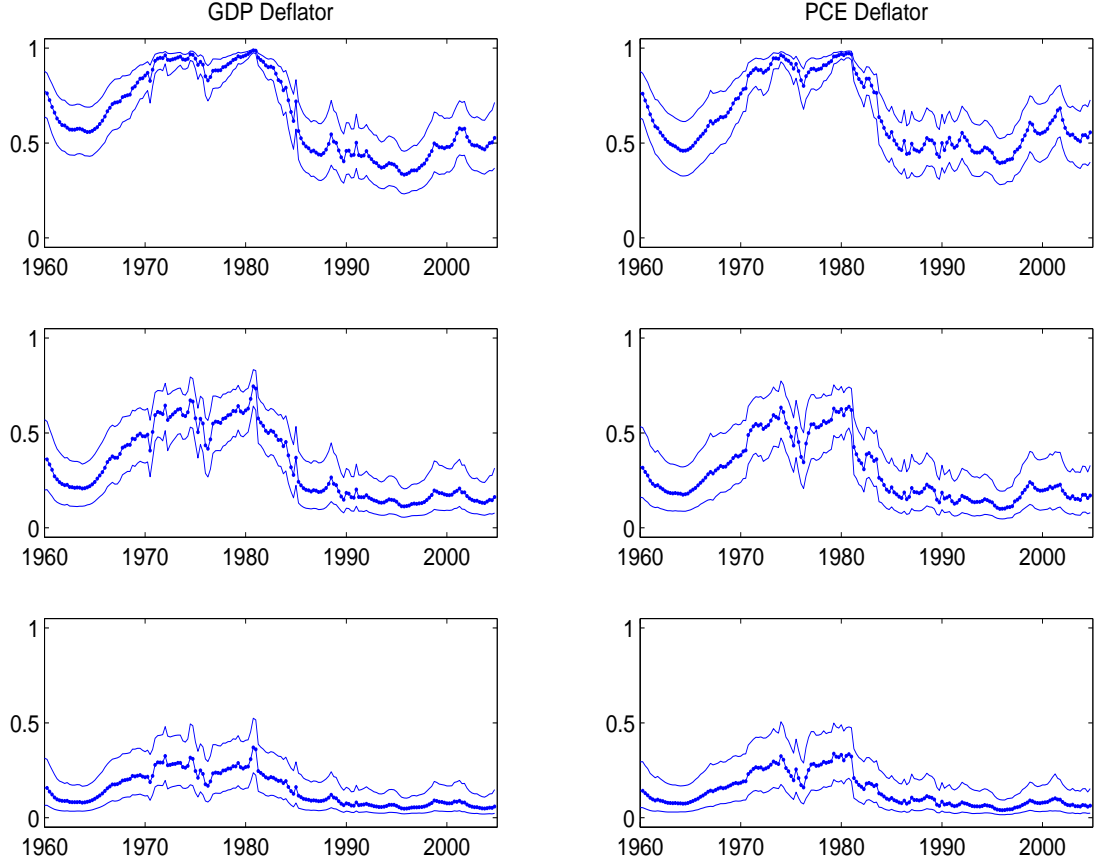


Figure 9: R_t^2 Statistics

Yet the question remains whether the changes are statistically significant. We approach this question in the same way as before, by examining the joint posterior distribution for R_{jt}^2 across pairs of years. Figures 10 and 11 plot the joint distribution for R_{1t}^2 for the years 1980 and 2004. Values for 1980 are shown on the x -axis, and those for 2004 are on the y -axis. For both measures of inflation, virtually the entire distribution lies below the 45 degree line, signifying that $R_{1,1980}^2 > R_{1,2004}^2$ with high probability. Very few points are clustered along the 45 degree line.




Figure 10: Joint Distribution for $R^2_{1,1980}$ and $R^2_{1,2004}$, GDP Inflation




Figure 11: Joint Distribution for $R^2_{1,1980}$ and $R^2_{1,2004}$, PCE Inflation

Table 2 records the fraction of posterior draws for which R^2_{jt} declined between 1980 and 2004. For 1-step ahead pseudo forecasts, the probability of a decline is 98.9 and 97.8 percent, respectively, for GDP and PCE inflation, thus confirming the visual impression conveyed by the figures. For 4- and 8-quarter ahead forecasts, the joint distributions are less tightly concentrated than those shown above, and the probabilities are a bit lower. Nevertheless, at the 4-quarter horizon, the probability of a decline in R^2_{jt} is almost 96 percent for GDP inflation and 92 percent for PCE inflation, and they are a bit less than 90 percent at the 8-quarter horizon.


Table 2: Probability of Changing R_{jt}^2

GDP Inflation			
pair	1 Quarter Ahead	4 Quarters Ahead	8 Quarters Ahead
1980, 2004	0.989	0.957	0.889
1960, 1980	0.991	0.919	0.820

PCE Inflation			
pair	1 Quarter Ahead	4 Quarters Ahead	8 Quarters Ahead
1980, 2004	0.978	0.922	0.876
1960, 1980	0.960	0.857	0.792

Figures 12 and 13 examine changes in predictability between 1960 and 1980. Here we plot $R_{1,1960}^2$ on the x -axis and $R_{1,1980}^2$ on the y -axis. Now virtually the entire distribution lies above the 45 degree line, signifying that $R_{1,1960}^2 < R_{1,1980}^2$ with high probability. Table 2 also reports the probability of an increase in $R_{j,t}^2$ between 1960 and 1980. For GDP inflation, this probability is 99.1 percent for 1-quarter ahead pseudo forecasts, 91.9 percent for 1-year ahead forecasts, and 82 percent for 2-year ahead forecasts. The probabilities are slightly lower for PCE inflation, but the results still point to a significant change in predictability at the 1-quarter horizon.

Thus, statistically significant evidence for changes in inflation persistence emerges from VARs. Estimates of R_{1t}^2 put posterior probabilities above 96 percent on the joint event of both an increase in persistence during the Great Inflation and a decline in persistence after the Volcker disinflation. The results for 4-quarter ahead forecasts also point in this direction, standing at the 90 or 95 percent levels for a fall in persistence in the second half of the sample and straddling the 90 percent level for a rise in the first half. The results for 2-year ahead forecasts hint at a change in persistence, but fall short of statistical significance at the 90 percent level.

Figure 12: Joint Distribution for $R_{1,1960}^2$ and $R_{1,1980}^2$, GDP Inflation




Figure 13: Joint Distribution for $R^2_{1,1960}$ and $R^2_{1,1980}$, PCE Inflation

5 Related research

Barsky (1987) explains an apparent violation of the Fisher equation in prewar US data in terms of changes in inflation predictability. The correlation between inflation and short-term nominal interest was negative prior to World War II but positive afterward. Barsky argues that this reflects changes in the time-series properties of inflation, not a change in the structural relation between nominal interest and expected inflation. Although inflation was highly forecastable after the mid 1960s, he documented that it was essentially unforecastable prior to World War I, and he demonstrated that this could account for the absence of a Fisher correlation in pre-war data.

Benati (2006) gathers data on inflation in a wide variety of monetary regimes and examines how inflation persistence varies across regime. Broadly speaking, he reports that high persistence occurs only in monetary regimes that lack a well-defined nominal anchor. For instance, for the modern era he contrasts countries whose central bank explicitly targets inflation with those that do not, and he finds that inflation is more autocorrelated in the latter. He also extends Barsky's work by looking at pre-WWII data from countries other than the US and confirms that inflation was close to white noise in many countries.

For the postwar US, Stock and Watson (2007) also document changes in the predictability of inflation. They find that inflation has become both easier and harder to forecast in the Volcker-Greenspan era. In an absolute sense, forecasting inflation is easier because inflation is less volatile and its innovation variance is smaller. But in a relative sense, predicting inflation has become more difficult because future inflation is less closely correlated with current inflation and other predictors. Their conclusion

agrees with ours: although the innovation variance for inflation has declined, the unconditional variance has fallen by more, implying that predictive R^2 statistics are lower.

5.1 Comparison with Atkeson-Ohanian findings

Stock and Watson also interpret a result of Atkeson and Ohanian (2001) in terms of the changing time-series properties of inflation. Atkeson and Ohanian studied the predictive power of backward-looking Phillips-curve models during the Volker-Greenspan era and found that Phillips-curve forecasts were inferior to a naive forecast that equates expected inflation over the next 12 months with the simple average of inflation over the previous year. Stock and Watson show that Phillips-curve models were more helpful during the Great Inflation, and they account for the change by pointing to two features of the data. First, like many macroeconomic variables, unemployment has become less volatile since the mid-1980s. Hence there is less variation in the predictor. Second, the coefficients linking unemployment and other activity variables to future inflation have also declined in absolute value, further muting their predictive power.

Our VARs share these characteristics. In figure 14, we illustrate how news about unemployment alters forecasts of inflation. At each date, we imagine that forecasters start with information on inflation, unemployment, and the nominal interest rate through date $t - 1$ and then see a one-sigma innovation in unemployment. They revise their inflation forecasts in light of the unemployment news. Because the VAR innovations are correlated, the forecast revision at horizon j is¹⁴

$$FR_{jt} = e_{\pi} A_t^j E(\varepsilon_{zt} | \varepsilon_{ut}) \sigma_{ut}. \quad (19)$$

Since the innovations are conditionally normal and the unemployment innovation is scaled to equal σ_{ut} , $E(\varepsilon_{zt} | \varepsilon_{ut}) = \text{cov}(\varepsilon_{zt}, \varepsilon_{ut}) / \sigma_{ut}$. The figure portrays the median and interquartile range for forecast revisions at horizons of 1, 4, and 8 quarters.

For the most part, a positive innovation in unemployment reduces expected inflation. Furthermore, in the 1970s and early 1980s, the magnitude of forecast revisions was substantial. For instance, according to the median estimates, a one-sigma innovation in unemployment would have reduced expected inflation 4 quarters ahead by close to 50 basis points in the mid-1970s and by approximately 1 to 1.5 percentage points at the time of the Volker disinflation. After the mid 1980s, however, the sensitivity of inflation forecasts to unemployment news was more muted. During the Greenspan era, a one-sigma innovation in unemployment would have had essentially no influence at all on inflation forecasts one or two years ahead.

¹⁴This follows from another anticipated-utility approximation.

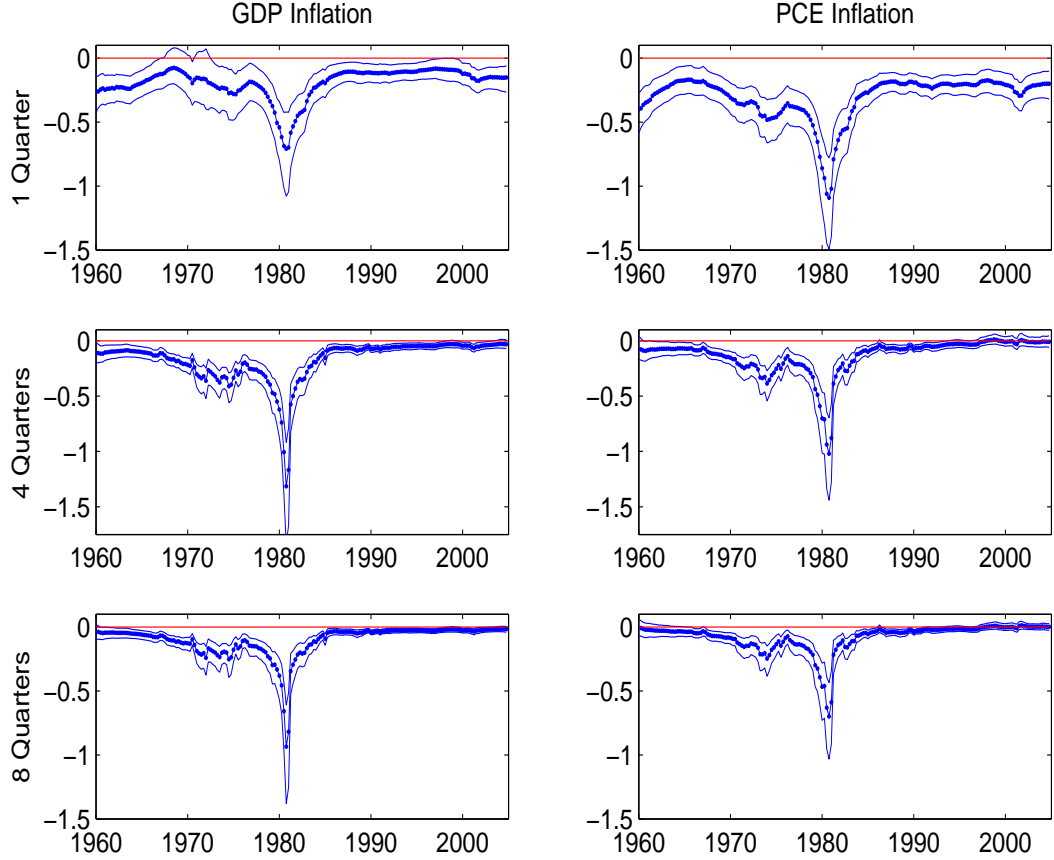


Figure 14: How Unemployment News Alters Expected Inflation

As Stock and Watson point out, these outcomes reflect both that unemployment innovations are less volatile and that inflation forecasts are less sensitive to innovations of a given size. Figure 15 depicts the posterior median and interquartile range for σ_{ut} , the standard deviation of innovations to unemployment. The magnitude of unemployment innovations was largest at the beginning of the sample and around the time of the Volcker disinflation, but it declined sharply after the mid 1980s. One reason why unemployment news has become less relevant for inflation forecasting is that there is less of it.

But this is not the whole story. Figure 16 adjusts for changes in the innovation variance by showing forecast revisions for the time-series average of the median estimate of σ_{ut} shown in figure 15. This holds the size of the hypothetical unemployment innovation constant across dates. Although less pronounced, the pattern shown here is similar to that depicted in figure 14 (the two figures are graphed on the same scale). Hence figure 14 cannot be explained solely by changes in σ_{ut} . Especially at horizons

of a 4 or 8 quarters, inflation forecasts have also become less sensitive to a given amount of unemployment news than they were during the Great Inflation.




Figure 15: Standard Deviation of Unemployment Innovations




Figure 16: Forecast Revisions with σ_u Held Constant

Ironically, the decreased predictive power of unemployment innovations for inflation coincided with a return of the Phillips correlation. Figure 17 portrays a number of conditional and unconditional correlations for inflation and unemployment.¹⁵ The unconditional correlation – shown in the bottom row – was negative prior to the 1970s, but it turned positive during the Great Inflation. A negative correlation reappeared after the Volcker disinflation and has hovered around -0.25 ever since.

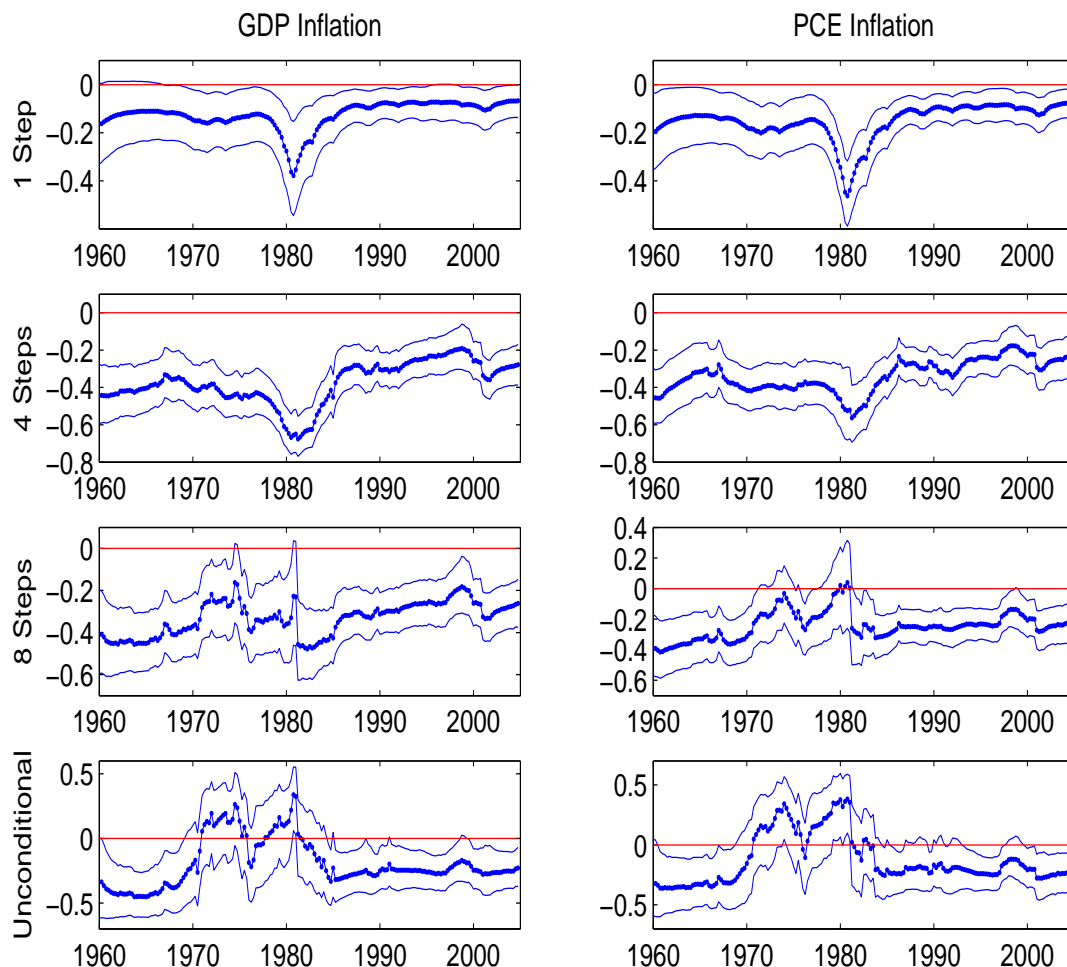


Figure 17: Conditional and Unconditional Phillips Correlations

The other rows of the figure depict conditional correlations at forecast horizons of 1, 4, and 8 quarters. The 1- and 4-quarter ahead forecasts are most relevant for

¹⁵These were calculated using the approximations in equations (16) and (17).

reconciling conventional wisdom with Atkeson and Ohanian. At these horizons, conditional correlations have indeed been negative throughout the sample, peaking in magnitude at the time of the Volcker disinflation. They are smaller now than in the past, but at the 4-quarter horizon the correlation is still around -0.25. Nevertheless, these conditional correlations are irrelevant for prediction because they summarize *unexpected* comovements in the two variables. That prediction errors in unemployment are inversely related with prediction errors in inflation tells us little about forecastable movements in the two variables. Thus, Atkeson and Ohanian's observations about predictability can coexist comfortably with conventional views about Phillips correlations.

Contrary to Atkeson and Ohanian, figures 9-11 suggest that some short-term predictability remains at the end of the sample. Two caveats should be kept in mind, however. One is that our calculations involve pseudo forecasts that depend on data and estimates through the end of the sample, while Atkeson and Ohanian look at real-time, out-of-sample forecasts. Presumably this matters only slightly at the end of the sample, but more for earlier periods.

The other caveat is that there is substantial uncertainty about R_{2004}^2 . We can state with confidence that R_{2004}^2 is smaller than R_{1980}^2 , but that is mainly because the posterior for R_{1980}^2 clusters tightly near 1. It is harder to say how predictable inflation is at the end of the sample. At the 1-quarter horizon, the probability that R_{2004}^2 exceeds 0.25 is 0.904 for GDP inflation and 0.924 for PCE inflation. Thus, although our estimates suggest more predictability than those of Atkeson and Ohanian, the fact that the posteriors portrayed in figures 10 and 11 assign non-negligible probability to values of R_{2004}^2 near zero provides at least some weak support for their point of view.

6 A More Structural Analysis

In this section we offer a structural explanation of the statistical findings presented above. We estimate a New-Keynesian model along the lines of Rotemberg and Woodford (1997) and Boivin and Giannoni (2006). This model is a simple possible framework for addressing the causes of the declines in the volatility, persistence, and predictability of inflation.

6.1 The model

The model economy is populated by a representative household, a continuum of monopolistically competitive firms, and a government. The representative household maximizes

$$E_t \sum_{s=0}^{\infty} \delta^s b_{t+s} \left[\log (C_{t+s} - hC_{t+s-1}) - \varphi \int_0^1 \frac{L_{t+s}(i)^{1+\nu}}{1+\nu} di \right], \quad (20)$$

subject to a sequence of budget constraint

$$\int_0^1 P_t(i) C_t(i) di + B_t + T_t \leq R_{t-1} B_{t-1} + \Pi_t + \int_0^1 W_t(i) L_t(i) di. \quad (21)$$

B_t represents government bonds, T_t denotes lump-sum taxes and transfers, R_t is the gross nominal interest rate, and Π_t are the profits that firms pay to the household. C_t is a Dixit-Stiglitz aggregator of differentiated consumption goods,

$$C_t = \left[\int_0^1 C_t(i)^{\frac{1}{1+\theta_t}} di \right]^{1+\theta_t}. \quad (22)$$

P_t is the associated price index, $L_t(i)$ denotes labor of type i that is used to produce differentiated good i , and $W_t(i)$ is the corresponding nominal wage. The coefficients h and ν set the degree of internal habit formation and the inverse Frisch elasticity of labor supply, respectively. Finally, b_t and θ_t are exogenous shocks that follow the stochastic processes

$$\begin{aligned} \log b_t &= \rho_b \log b_{t-1} + \varepsilon_{b,t} \\ \log \theta_t &= (1 - \rho_\theta) \log \theta + \rho_\theta \log \theta_{t-1} + \varepsilon_{\theta,t}. \end{aligned} \quad (23)$$

The random variable b_t is an inter-temporal preference shock perturbing the discount factor, and θ_t can be interpreted as a shock to the firms' desired mark-up.

Each differentiated consumption good is produced by a monopolistically competitive firm using a linear production function,

$$Y_t(i) = A_t L_t(i), \quad (24)$$

where $Y_t(i)$ denotes the production of good i , and A_t represents aggregate labor productivity. We model A_t as a unit root process with a growth rate $z_t \equiv \log(A_t/A_{t-1})$ that follows the exogenous process

$$z_t = (1 - \rho_z) \gamma + \rho_z z_{t-1} + \varepsilon_{z,t}. \quad (25)$$

As in Calvo (1983), at each point in time a fraction ξ of firms cannot re-optimize their prices and simply indexes them to the steady-state value of inflation. Subject to the usual cost-minimization condition, a re-optimizing firm chooses its price ($\tilde{P}_t(i)$) by maximizing the present value of future profits,

$$E_t \sum_{s=0}^{\infty} \xi^s \delta^s \lambda_{t+s} \left\{ \tilde{P}_t(i) \pi^s Y_{t+s}(i) - W_{t+s}(i) L_{t+s}(i) \right\}, \quad (26)$$

where π is the gross rate of inflation in steady state and λ_{t+s} is the marginal utility of consumption.

The monetary authority sets short-term nominal interest rates according to a Taylor rule,

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\left(\frac{\bar{\pi}_{4,t}}{(\pi_t^*)^4} \right)^{\frac{\phi_\pi}{4}} \left(\frac{Y_t}{Y_t^*} \right)^{\phi_Y} \right]^{1-\rho_R} e^{\varepsilon_{R,t}}. \quad (27)$$

The central bank smooths interest rates and responds to two gaps, the deviation of annual inflation ($\bar{\pi}_{4,t}$) from a time-varying inflation target and the difference between output and its flexible price level. R is the steady-state value for the gross nominal interest rate and $\varepsilon_{R,t}$ is a monetary policy shock that we assume to be *i.i.d.*

Following Ireland (2007), we model the inflation target π_t^* as an exogenous random process,

$$\log \pi_t^* = (1 - \rho_*) \log \pi + \rho_* \log \pi_t^* + \varepsilon_{*,t}. \quad (28)$$

There are many reasons that the Central Bank's inflation target might vary over time. Our preferred one is that the central bank endogenously adjusts its target as it learns about the structure of the economy. For instance, Sargent (1999), Cogley and Sargent (2005), Primiceri (2006), and Sargent, Williams, and Zha (2006) hypothesize that changing beliefs about the output-inflation tradeoff generated a pronounced low-frequency, hump-shaped pattern in inflation. We approximate outcomes of this learning process by an exogenous random variable like (28).¹⁶

6.2 Model solution and observation equation

Since the technology process A_t is assumed to have a unit root, consumption, real wages, and output evolve along a stochastic growth path. To solve the model, we first rewrite it in terms of deviations of these variables from the technology process. Then we solve the log-linear approximation of the model around the non-stochastic steady state. We specify the vector of observable variables as $[\log Y_t - \log Y_{t-1}, \pi_t, R_t]$. For estimation, we use data on per-capita GDP growth, the quarterly growth rate of the GDP deflator, and the Federal funds rate.¹⁷

6.3 Bayesian inference and priors

We use Bayesian methods to characterize the posterior distribution of the model's structural parameters.¹⁸ Table 3 reports our priors. These priors are relatively disperse and are broadly in line with those adopted in previous studies (see, for instance, Del Negro et al. 2007 or Justiniano and Primiceri 2007). But a few items deserve discussion.

¹⁶By way of analogy, technology is also endogenous, but macroeconomists often model it as an exogenous random variable.

¹⁷These variables are standard for estimating small-scale DSGE models (see, for instance, Boivin and Giannoni 2006).

¹⁸See appendix B.

Table 3: Priors for Structural Parameters

COEFFICIENT	PRIOR		
	Density	Mean	Standard Deviation
ν	<i>Calibrated</i>	2	—
$\theta - 1$	<i>Calibrated</i>	0.1	—
100γ	<i>Normal</i>	0.475	0.025
$100(\pi - 1)$	<i>Normal</i>	0.5	0.1
$100(\delta^{-1} - 1)$	<i>Gamma</i>	0.25	0.1
h	<i>Beta</i>	0.5	0.1
ξ	<i>Beta</i>	0.66	0.1
ϕ_π	<i>Normal</i>	1.7	0.3
ϕ_y	<i>Gamma</i>	0.3	0.2
ρ_R	<i>Beta</i>	0.6	0.2
ρ_z	<i>Beta</i>	0.4	0.2
ρ_θ	<i>Beta</i>	0.6	0.2
ρ_*	<i>Calibrated</i>	0.995	—
ρ_b	<i>Beta</i>	0.6	0.2
$100\sigma_R$	<i>Inverse Gamma</i>	0.15	1
$100\sigma_z$	<i>Inverse Gamma</i>	1	1
$100\sigma_\theta$	<i>Inverse Gamma</i>	0.15	1
$100\sigma_*$	<i>Uniform</i>	0.075	0.0433
$100\sigma_b$	<i>Inverse Gamma</i>	1	1

- We fix two parameters because they are not identified. In particular, we set the Frisch elasticity of labor supply ($1/\nu$) to 0.5 and the steady-state price mark-up ($\theta - 1$) to 10%.
- For all but two persistence parameters, we use a Beta prior with mean 0.6 and standard deviation 0.2. One exception concerns labor productivity, which already includes a unit root. For this reason, we center the prior for the autocorrelation of its growth rate (ρ_z) at 0.4. The other exception is the autocorrelation of the inflation target shock, which we calibrate to 0.995. In other words, we restrict π_t^* so that it captures low-frequency movements in inflation.¹⁹
- The standard deviation of the innovation to the inflation target is a crucial parameter in our analysis because it governs the rate at which π_t^* drifts. We want a weakly informative prior in order to let the data dominate the posterior. Accordingly, we adopt a uniform prior on (0,0.15). For the standard deviations of the other shocks, we follow Del Negro et al. (2007) by choosing priors that

¹⁹We do not set $\rho_* = 1$ because the DSGE model would not admit a non-stochastic steady state and the log-linearization would not be possible in that case.

are fairly disperse and that generate realistic volatilities for the endogenous variables.

- Finally, we truncate the prior at the boundary of the determinacy region.

6.4 Estimation Results

We estimate the model separately on two subsamples. The first, 1960:I - 1979:II, corresponds approximately to the period of rising inflation before the Volcker chairmanship. The second period, 1982:IV - 2006:IV, corresponds to the Volcker and Greenspan chairmanships, excluding the years of monetary targeting, for which the Taylor rule might not represent an appropriate description of systematic monetary policy (see, for instance, Sims and Zha 2006 or Hanson 2006).

Figure 18 presents the model-implied evolution of the Central Bank inflation objective. Notice that it resembles quite closely the VAR-based estimate of the permanent component of inflation plotted in figure 1.




Figure 18: The Central Bank's Inflation Target

Table 4 reports estimates of the structural parameters. While many coefficients are similar across subsamples, there are some important differences. For example, we find that the Taylor-rule coefficient for inflation (ϕ_π) increased from 1.55 in the first subsample to 1.78 in the second. While an increase is consistent with findings of Clarida, Gali, and Gertler (2000) and Lubik and Schorfheide (2004), we do not find values of ϕ_π in the pre-1980 period as low as they do. This might be due to the fact that, for simplicity, we have ruled out indeterminacy a priori. Another possibility is that the presence of a time varying inflation target reduces the differences between the reaction to inflation in the two subsamples.

Table 4: Posteriors for Structural Parameters

COEFFICIENT	1960-1979			1982-2006		
	Median	25th pct	75th pct	Median	25th pct	75th pct
100γ	0.468	0.452	0.484	0.484	0.467	0.500
$100(\pi - 1)$	0.501	0.435	0.566	0.516	0.452	0.581
$100(\delta^{-1} - 1)$	0.159	0.121	0.204	0.255	0.199	0.319
h	0.445	0.390	0.502	0.526	0.482	0.568
ξ	0.782	0.741	0.818	0.800	0.762	0.835
ϕ_π	1.557	1.372	1.746	1.784	1.598	1.974
ϕ_y	0.643	0.541	0.772	0.66	0.562	0.784
ρ_R	0.704	0.630	0.759	0.633	0.576	0.686
ρ_z	0.264	0.156	0.390	0.297	0.191	0.415
ρ_θ	0.598	0.515	0.676	0.255	0.182	0.344
ρ_b	0.699	0.632	0.758	0.876	0.850	0.898
$100\sigma_R$	0.160	0.147	0.174	0.069	0.063	0.076
$100\sigma_z$	0.641	0.527	0.797	0.493	0.426	0.562
$100\sigma_\theta$	0.118	0.097	0.139	0.126	0.114	0.137
$100\sigma_*$	0.081	0.062	0.104	0.049	0.037	0.065
$100\sigma_b$	2.533	2.226	2.889	2.429	2.146	2.785

A second notable change in monetary policy concerns the innovation variances for the two shocks, $\varepsilon_{*,t}$ and $\varepsilon_{R,t}$. According to our estimates, both declined substantially after the Volcker disinflation. The innovation variance for the shock to target inflation fell by almost 50 percent, from 0.081 to 0.049, while the variance for the funds-rate shock declined even more, from 0.16 to 0.07. The decline in σ_* should not be surprising, given the findings of Stock and Watson (2007) and our VAR statistical results. It contributes directly to the decline in inflation volatility after 1980.

Among the nonpolicy parameters, most change only slightly across the two samples. This is comforting because these parameters are supposed to be invariant to changes in monetary policy. One exception is the persistence parameter ρ_θ for the cost-push shock, which declines from 0.6 to 0.25. Thus, the cost-push shock is less persistent and has smaller unconditional variance after 1982. This decline might reflect the reduced incidence of oil-price shocks in the second half of the period. If that is correct, the estimates capture elements of good luck as well as improved policy.

Table 5 summarizes the model's implications for inflation volatility, persistence, and predictability at the posterior median of the model parameters. Column 1 reports the unconditional standard deviation of inflation, while columns 2–4 report R^2 statistics for inflation-gap predictability for forecasting horizons of 1, 4 and 8 quarters.²⁰ Notice that, in line with our statistical VAR findings, the model reproduces

²⁰The inflation gap here is defined as the difference between inflation and the central bank inflation

well the substantial decline in volatility, persistence and predictability of inflation. All decrease by roughly 50 to 70 percent.²¹

Table 5: Implications of the DSGE Model for Inflation Volatility and Predictability

	$100 \cdot std(\hat{\pi}_t)$	R_1^2	R_4^2	R_8^2	<i>Slope</i>
1960:I - 1979:II	4.702	0.631	0.433	0.409	0.132
1982:IV - 2006:IV	2.354	0.206	0.136	0.124	0.040
Percent Change	-50	-67	-69	-70	-70

Finally, column 5 addresses the results of Atkeson and Ohanian (2001). Here we report the model-implied slope β of the Phillips curve,

$$E_t(\hat{\pi}_{4,t+4} - \hat{\pi}_{4,t}) = \beta(\hat{Y}_t - \hat{Y}_t^*) + \epsilon_{t,t+4}.$$

We do not include a constant because the variables of the regression all have mean zero in the model. Except for the fact that we replace the unemployment rate with the output gap, this is the regression estimated by Atkeson and Ohanian (2001). Consistent with their results, those of Stock and Watson, and our own results reported above, our model implies a substantial decline in the predictive power of real-activity variables in conventional Phillips curve regressions after the Volcker disinflation.

6.5 Counterfactuals

In line with the statistical VAR findings, the DSGE model reproduces much of the substantial decline in volatility, persistence, and predictability of inflation. We are sufficiently encouraged by its performance to use the DSGE model to explore the structural sources of these changes.

In this subsection, we conduct some counterfactual exercises in order to understand the causes of the decline in inflation volatility, persistence, and predictability. In the first experiment, we combine the Taylor-rule coefficients ($[\phi_\pi, \phi_y, \rho_R, \sigma_R, \sigma_*]$) of the second subsample with the private-sector parameters of the first. In this way, we assess the extent to which better monetary policy would have reduced inflation volatility and persistence during the Great Inflation. In the second experiment, we combine the private-sector parameters of the second subsample with the policy parameters of the first. This scenario illustrates the contribution of nonpolicy factors to the improvement in inflation outcomes.

objective that, in the DSGE model, captures the permanent component of inflation.

²¹Since we estimate the model on two separate subsamples, the joint posterior distribution of the coefficients of the first and second subsample is not available. Therefore, we cannot report standard errors.

Table 6 reports the results. The numbers recorded there represent the proportion of the total change across subsamples accounted for by the hypothetical structural shift,

$$100 \times \frac{\text{counterfactual change}}{\text{total change}}.$$

Positive numbers signify that the counterfactual goes in the same direction as the total change, and negative numbers mean that it goes in the opposite direction.

Table 6: Counterfactual Exercises Based on the DSGE Model

COEFFICIENTS	VOLATILITY	PERSISTENCE			SLOPE (β)
		R_1^2	R_4^2	R_8^2	
Policy 2, Private 1	75	43	90	91	-94
σ_*	69	32	68	69	-46
ϕ_π	9	13	28	28	-27
Private 2, Policy 1	36	43	15	14	125
ρ_θ	7	-39	-109	-111	121

Monetary policy seems to be the most important factor behind the decline in inflation volatility. The change in policy rule accounts for 75 percent of the decline in inflation volatility. In contrast, better luck – primarily in the form of a less volatile and persistent cost-push shock – accounts for 36 percent of the decline. This is a substantial contribution, but only about half the magnitude of the effect of monetary policy.²²

The results for predictability are similar, especially at the 4 and 8 quarter horizons. At those horizons, better monetary policy accounts for approximately 90 percent of the decline, while changes in private-sector behavior account for around 15 percent. At the 1-quarter horizon, however, the two factors contribute equally to the decline in predictability, each accounting for 43 percent of the total change.

The second and third rows of the table 6 look more closely at particular aspects of monetary policy. Here we change a single Taylor-rule parameter, holding all other coefficients equal to the estimated value from subsample 1. Otherwise the experiments are the same as before.

Among monetary-policy coefficients, changes in the variability of the inflation objective (σ_*) and in the reaction to inflation (ϕ_π) have the largest impact on inflation outcomes. The more stable inflation objective is responsible for the largest portion of the decline in inflation volatility and persistence, accounting for roughly two-thirds

²²The two numbers need not sum to 100 because the model is nonlinear in the coefficients and, therefore, the total change is not the sum of the effects of the policy and nonpolicy coefficients shift.

of the total change. This is because changes in the Central Bank’s inflation target generate persistent deviations of the nominal interest rate and marginal cost from the steady state. In turn, this induces persistent deviations of inflation from the target, i.e., a persistent component in the inflation-gap. Hence, a decline in the volatility of the inflation target reduces the overall persistence of the inflation-gap by reducing the relative importance of the persistent component.

Another important contributor is stronger monetary-policy reaction to inflation. In our model, however, this is secondary to enhanced stability of the inflation target, accounting for about 10 percent of the decline in volatility and 13-28 percent of the decline in predictability. One reason that we might be finding a smaller contribution than has been found by others (e.g., Benati and Surico 2007) is that we truncate our prior on the boundary of the determinacy region. Thus, our feedback parameter rises from 1.56 to 1.78. Enhanced feedback plays a role in our model, but not the primary role.

We also look more closely at the particular aspects of private-sector behavior that have the greatest influence on changing inflation outcomes. Among nonpolicy parameters, the key change is the shift in the persistence of the mark-up shock. The final row of table 6 sheds light on its contribution. Everything else equal, the decline in persistence of the mark-up shock (ρ_θ) would have induced an *increase* in inflation-gap persistence. This might seem surprising but has a simple explanation: a reduction in ρ_θ corresponds to a decrease in the unconditional variability of the mark-up shock, which reduces the volatility of inflation due to this shock. As a consequence, the role of the inflation-target shock for inflation becomes relatively larger, and this increases persistence.

The final column of table 6 examines how changes in monetary policy and private-sector parameters contribute to the flattening of the slope in an Atkeson-Ohanian regression. Recall that the DSGE model predicts a decline in β from 0.13 to 0.04 across the two subsamples. In this case, the relative importance of better policy and better luck are reversed. Changes in private sector parameters go in the right direction and overpredict the total decline. Conditional on the mark-up shock, the output gap and future changes in the inflation rate comove positively. The drop in persistence and unconditional volatility of the mark-up shock reduces this positive comovement and results in a lower estimate of the slope coefficient. But changes in policy parameters go in the wrong direction and predict a substantial increase in β . Thus, for a complete picture of the change in inflation outcomes, both private and policy factors are needed.

7 Concluding remarks

This paper reports what autoregressions with drifting coefficients and stochastic volatility say about the persistence of the inflation gap defined as the fraction of variation of future inflation gaps that is due to past shocks. A high proportion means that past shocks retain influence for a long time, while a low proportion signifies that their influence decays quickly. Since past shocks give rise to forecastable variation in future inflation gaps, our concept of persistence is closely related to predictability. VAR estimates point to a statistically significant increase in inflation-gap predictability during the Great Inflation and to a statistically significant decline in predictability after the Volcker disinflation. Univariate estimates are mixed, with significant evidence of a rise and fall in persistence for PCE inflation and marginal evidence for GDP inflation.

We have used a new Keynesian DSGE model to examine what caused these changes. We find evidence that both better policy and better luck – in the form of less volatile and less persistent cost-push shocks – contributed to improved inflation outcomes. The enhanced stability of the Fed’s long-run inflation target stands as key improvement in policy. In our DSGE model, this is the single most important factor behind the reduction in inflation volatility and persistence.

The DSGE model treats the inflation target as an exogenous random process. Explaining why it drifts is a priority for future research. Our preferred story involves learning and changing central bank beliefs about the structure of the economy (Cogley and Sargent 2005b, Primiceri 2006, and Sargent, Williams, and Zha 2006), but more work is needed to understand this aspect of monetary policy.

A Markov chain Monte Carlo algorithm for simulating the VAR posterior

For the VAR, the posterior density is²³

$$p(\theta^T, H_y^T, H_s^T, B_y, B_s, \sigma_y, \sigma_s | Y^T). \quad (29)$$

The state and measurement innovation variances are defined as

$$\begin{aligned} Q_t &= (B_s^{-1})' H_{st} (B_s^{-1}), \\ R_t &= (B_y^{-1})' H_{yt} (B_y^{-1}), \end{aligned} \quad (30)$$

respectively, where H_{st} and H_{yt} are diagonal matrices with univariate stochastic volatilities along the main diagonal and B_s and B_y are triangular matrices with

²³The MCMC algorithm for the univariate AR is a special case of that for the VAR.

ones along the main diagonal and static covariance parameters below. The univariate stochastic volatilities are geometric random walks; the vectors σ_s and σ_y list their innovation variances. The notation x^T represents the complete history of x_t .

We use a ‘Metropolis-within-Gibbs’ algorithm to simulate the posterior. The parameters are partitioned into 7 blocks:

- $\theta^T | H_y^T, H_s^T, B_y, B_s, \sigma_y, \sigma_s, Y^T$
- $H_y^T | \theta^T, H_s^T, B_y, B_s, \sigma_y, \sigma_s, Y^T$
- $B_y | \theta^T, H_y^T, H_s^T, B_s, \sigma_y, \sigma_s, Y^T$
- $\sigma_y | \theta^T, H_y^T, H_s^T, B_y, B_s, \sigma_s, Y^T$
- $H_s^T | \theta^T, H_y^T, B_y, B_s, \sigma_y, \sigma_s, Y^T$
- $B_s | \theta^T, H_y^T, H_s^T, B_y, \sigma_y, \sigma_s, Y^T$
- $\sigma_s | \theta^T, H_y^T, H_s^T, B_y, B_s, \sigma_y, Y^T$

After substituting Q_t for Q , the samplers for the first four blocks are identical to those in Cogley and Sargent (2005a); details can be found there. Those for the last three blocks – which pertain to the state innovation variance Q_t – are isomorphic to the three blocks for the measurement innovation variance R_t . Thus, the appendices in Cogley and Sargent (2005a) cover those blocks as well.

We executed 100,000 scans of the chain and diagnosed convergence by inspecting recursive mean plots of the parameters. We discarded the first 50,000 scans to allow for burn in. The results reported in the text are based on the remaining 50,000 scans.

A.1 Priors for the VAR

The priors are similar to those in Cogley and Sargent (2005a). We assume that the hyperparameters and initial value of the drifting parameters are independent across blocks, so that the joint prior factors into a product of marginal priors. Each of the marginal priors is selected from a family of natural conjugate priors and is specified to proper yet weakly informative.

The unrestricted prior for the initial state is Gaussian,

$$f(\theta_0) \propto \mathcal{N}(\bar{\theta}, \bar{P}), \quad (31)$$

where $\bar{\theta}$ and \bar{P} are the OLS point estimate and asymptotic variance, respectively, based on a training sample covering the period 1948-58. Because the training sample is short, the asymptotic variance is large, making the prior weakly informative for θ_0 .

Priors for the blocks governing R_t are also calibrated to put considerable weight on sample information. The prior for H_{y0}^{ii} is log-normal,

$$f(\ln H_{y0}^{ii}) = \mathcal{N}(\ln R_0^{ii}, 10), \quad (32)$$

where $\ln R_0^{ii}$ is the estimate of the log of residual variance of variable i from the preliminary sample. A variance of 10 is huge on a log scale and allows a wide range of values for h_{i0} . As is the case for θ_0 , the prior mean for H_{y0} is no more than a ballpark number surrounded by considerable uncertainty.

Similarly, the prior for β_y is normal with mean zero and a large variance,

$$f(\beta) = \mathcal{N}(0, 10000 \cdot I). \quad (33)$$

Lastly, the prior for σ_{yi}^2 , the variance of the stochastic volatility innovations, is inverse-gamma

$$f(\sigma_i^2) = IG(\delta_i/2, 1/2), \quad (34)$$

with scale parameter $\delta = 0.0001$ and degree-of-freedom parameter equal to 1. This distribution is proper but has fat tails.

The priors for the blocks governing Q_t parallel those for R_t . The prior for H_{Q0}^i is also log-normal,

$$f(\ln H_{Q0}^{ii}) = \mathcal{N}(\ln Q_0^{ii}, 10), \quad (35)$$

where $Q_0 = \gamma^2 \bar{P}$ is calibrated in the same way as in Cogley and Sargent (2005a). Similarly, the priors for β_Q and σ_Q have the same form as those for β_y and σ_y . We just alter the dimensions so that they conform to H_{Qt} instead of H_{Rt} . The prior mean for H_{Q0} induces only a slight degree of time variation in θ_t , but in other respects the priors are sufficiently uninformative that they permit a wide range of outcomes for Q_t .

B Markov chain Monte Carlo algorithm for simulating the DSGE posterior

As in An and Schorfheide (2006), we use a Metropolis-Hastings algorithm to simulate the posterior distribution of the coefficients of the DSGE model. Let y^T denote the set of available data and α the vector of coefficients of the DSGE model. Moreover, let $\alpha^{(j)}$ denote the j^{th} draw from the posterior of α . The subsequent draw is obtained by drawing a candidate value, $\tilde{\alpha}$, from a Gaussian proposal distribution with mean $\alpha^{(j)}$ and variance $s \cdot V$. We then set $\alpha^{(j+1)} = \tilde{\alpha}$ with probability equal to

$$\min \left\{ 1, \frac{p(\tilde{\alpha}|y^T)}{p(\alpha^{(j)}|y^T)} \right\}. \quad (36)$$

If the proposal is not accepted, we set $\alpha^{(j+1)} = \alpha^{(j)}$.

The posterior distribution for α , $p(\alpha|y^T)$, can be computed multiplying the prior density by the likelihood function. Because the DSGE model has a linear-Gaussian state-space representation, the likelihood function can be evaluated using the prediction-error decomposition and the Kalman filter.

The algorithm is initialized around the posterior mode, found using a standard maximization algorithm. We set V to the inverse Hessian of the posterior evaluated at the mode, while s is chosen in order to achieve an acceptance rate approximately equal to 25 percent. We run two chains of 70,000 draws and discard the first 20,000 to allow convergence to the ergodic distribution.

References

- An, S. and F. Schorfheide, 2007, Bayesian Analysis of DSGE Models, *Econometric Reviews* 39, 1357-1373.
- Atkeson, A. and L.E. Ohanian, 2001, Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25, 2-11.
- Barsky, R.B., 1987, The Fisher Hypothesis and the Forecastability and Persistence of Inflation, *Journal of Monetary Economics* 19, 3-24.
- Benati, L., 2006, Investigating Inflation Persistence Across Monetary Regimes, unpublished manuscript, European Central Bank.
- Benati, L. and P. Surico, 2007, U.S. Evolving Monetary Policy and the Inflation Gap's Persistence and Predictability, mimeo, European Central Bank.
- Beveridge, S., and C.R. Nelson, 1981. A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the Business Cycle. *Journal of Monetary Economics* 7, 151-174.
- Bayer A. and R. Farmer, 2007, On the Indeterminacy of Determinacy and Indeterminacy: Comments on the Paper "Testing for Indeterminacy" by Thomas Lubik and Frank Schorfheide, *American Economic Review* 97(1), 524-529.
- Boivin, J. and M. Giannoni, 2006, Has Monetary Policy Become More Effective?, *Review of Economics and Statistics* 88(3), 445-462.
- Calvo, G., 1983, Staggered Prices in a Utility-Maximizing Framework, *Journal of Monetary Economics*, 12(3), 383-398.
- Clarida, R., J. Gali and M. Gertler, 2000, Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory, *The Quarterly Journal of Economics* 115(1), 147-180.
- Cogley, T. and T.J. Sargent, 2001, Evolving Post World War II U.S. Inflation Dynamics, *NBER Macroeconomics Annual* 16, 331-373.

- Cogley, T. and T.J. Sargent, 2005a, Drifts and Volatilities: Monetary Policies and Outcomes in the Post War U.S., *Review of Economic Dynamics* 8, 262-302.
- Cogley, T. and T.J. Sargent 2005b. The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty. *Review of Economic Dynamics* 8, 528-563.
- Cogley, T. and T.J. Sargent, 2006, Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making, forthcoming, *International Economic Review*.
- Cogley, T., S. Morozov, and T.J. Sargent, 2005, Bayesian Fan Charts for U.K. Inflation: Forecasting and Sources of Uncertainty in an Evolving Monetary System, *Journal of Economic Dynamics and Control* 29, 1893-1925.
- Cogley, T. and A.M. Sbordone, 2006, Trend Inflation and Inflation Persistence in the New Keynesian Phillips Curve, unpublished manuscript.
- Del Negro, M., F. Schorfheide, F. Smets and R. Wouters, 2007, On the Fit and Forecasting Performance of New Keynesian Models, *Journal of Business and Economic Statistics* 25(2), 123-162.
- Diebold, F.X. and L. Kilian, 2001, Measuring Predictability: Theory and Macroeconomic Applications, *Journal of Applied Econometrics* 16, 657-669.
- Gelfand, A.E. Model determination using sampling-based methods. In W.R. Gilks, S. Richardson, and D.J. Spiegelhalter, eds., *Markov Chain Monte Carlo in Practice*, London: Chapman and Hall, 1996.
- Hanson, M.S., 2006, Varying Monetary Policy Regimes: A Vector Autoregressive Investigation, *Journal of Economics and Business, Elsevier* 58(5-6), 407-427.
- Ireland, P.N., 2007, Changes in the Federal Reserve's Inflation Target: Causes and Consequences, forthcoming, *Journal of Money Credit and Banking*.
- Justiniano, A. and G.E. Primiceri, 2007, The Time-Varying Volatility of Macroeconomic Fluctuations, forthcoming, *American Economic Review*.
- Kreps, D., 1998, Anticipated Utility and Dynamic Choice, 1997 Schwartz Lecture, in *Frontiers of Research in Economic Theory*, Edited by D.P. Jacobs, E. Kalai, and M. Kamien, Cambridge University Press, Cambridge, England.
- Levin, A.T. and J.M. Piger, 2004, Is Inflation Persistence Intrinsic in Industrial Societies?, ECB working paper no. 334.
- Lubik T.A. and F. Schorfheide, 2004, Testing for Indeterminacy: An Application to U.S. Monetary Policy, *American Economic Review*, 94(1), 190-217.
- Pivetta, F. and R. Reis, 2007, The Persistence of Inflation in the United States, *Journal of Economic Dynamics and Control* 31, 1326-58.
- Primiceri, G., 2005, Time Varying Structural Vector Autoregressions and Monetary Policy, *Review of Economic Studies* 72, 821-852.

Primiceri, G., 2006, Why Inflation Rose and Fell: Policymakers' Beliefs and U.S. Postwar Stabilization Policy, *Quarterly Journal of Economics* 121, 867-901.

Rotemberg J.J. and M. Woodford, 1997, An Optimization-Based Econometric Model for the Evaluation of Monetary Policy, *NBER Macroeconomics Annual* 12, 297-346.

Sargent, T.J., N. Williams and T. Zha, 2006, Shocks and Government Beliefs: The Rise and Fall of American Inflation, *American Economic Review*, 96(4), 1193-1224.

Sims, C.A., 2001, Comment on Sargent and Cogley's 'Evolving Post World War II U.S. Inflation Dynamics,' *NBER Macroeconomics Annual* 16, 373-379.

Sims, C.A. and T. Zha, 1999, Error Bands for Impulse Responses, *Econometrica* 67, 1113-1155.

Sims, C.A. and T. Zha, 2006, Were There Regime Switches in {US} Monetary Policy?, *American Economic Review* 96(1), 54-81.

Smets, F.R. and R. Wouters, 2003, An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area, *Journal of the European Economic Association* 1, 1123-1175.

Stock, J.H., 2001, Discussion of Cogley and Sargent 'Evolving Post World War II U.S. Inflation Dynamics,' *NBER Macroeconomics Annual* 16, 379-387.

Stock, J.H. and M.W. Watson, 2007, Has Inflation Become Harder to Forecast?, *Journal of Money Credit and Banking* 39(1), 3-34.