The Great Inflation and the Greenbook*

Giacomo Carboni
European Central Bank

Martin Ellison
University of Warwick and CEPR

October 2007

Abstract

We assess whether recent empirical evidence that Federal Reserve learning caused the Great Inflation is consistent with forecasts published in the Greenbook. If the rise and fall in inflation really was caused by the Federal Reserve learning the Phillips curve then that should be fully reflected in Greenbook forecasts. It is not. The difficulty is that empirical evidence is predicated on the Federal Reserve making forecasts that are much more volatile than those in the Greenbooks. If consistency with Greenbook forecasts is required then evidence that Federal Reserve learning caused the Great Inflation is much weaker. Our results suggest a larger role for other causes than previously thought.

*We greatly appreciate Tao Zha making available his C++ routines for Bayesian MCMC estimation. We also thank Knut Anton Mork, Antti Ripatti, and seminar participants at Bank of England, Birkbeck, Cardiff, Durham, Glasgow, Humboldt, Oslo (2nd Workshop on Monetary Policy), Queen's Belfast, Warwick and York for helpful comments and suggestions. Martin Ellison acknowledges support from an ESRC Research Fellowship, "Improving Monetary Policy for the 21st Century" (RES-000-27-0126). The views expressed in this paper do not necessarily reflect those of the ECB.

1 Introduction

One of the key challenges faced by macroeconomics today is to understand the rise and fall of American inflation in the latter half of the twentieth century. Aside from obvious academic interest, determining the causes of the Great Inflation is also crucial for government monetary authorities as they plan future macroeconomic policy and attempt to avoid repeating any mistakes they may have made in the past. In recent years there has been growing interest in explaining the Great Inflation as resulting from changes in the conduct of monetary policy itself, which occurred as the monetary authority learned and revised its view of the monetary transmission mechanism. At the forefront of this research is Sargent (1999), whose *The Conquest of American Inflation* treatise puts forward the hypothesis that American inflation dynamics can be explained by the Federal Reserve discovering and subsequently abandoning the Phillips curve. Important contributions by Ireland (1999), Cho, Williams and Sargent (2002), Cogley and Sargent (2005a, 2005b), Sargent and Williams (2005), Primiceri (2005) and Tetlow and Ironside (2007) have given further momentum to this research agenda.

The most thorough empirical assessment of the learning hypothesis to date is Sargent, Williams and Zha (2006), who operationalise the Sargent (1999) model and estimate its parameters using a Bayesian MCMC algorithm. Their results show the learning hypothesis receives remarkable support from real-world data, with the learning model dominating a Bayesian vector autoregression in terms of its ability to match and forecast inflation dynamics. However, it can be argued that matching the dynamics of the Great Inflation should only be interpreted as weak evidence in favour of the learning hypothesis. The problem is that matching dynamics tests the ability of the learning hypothesis to explain *what* the Federal Reserve did, but does not test *why* the Federal Reserve acted in this way. In this paper, we use information published in the Greenbook for each FOMC meeting to implicitly identify why the Federal Reserve acted as it did. Specifically, we extract unemployment forecasts from historical Greenbooks as a proxy for real-time estimates of how the Federal Reserve believed its policy would impact on the economy. We then assess the learning hypothesis against both the dynamics of the Great Inflation and the Greenbook unemployment forecasts, requiring a match to both before concluding there is strong evidence in its favour.

The central contribution of the paper is to show that matching just the dynamics of the Great Inflation should indeed be taken only as weak evidence in favour of the learning hypothesis. We arrive at this conclusion by estimating a learning model very similar to that in Sargent, Williams and Zha (2006), and finding that the unemployment forecasts associated with the learning model are inconsistent with those reported in the Greenbooks. If some consistency is imposed then the estimated match to Great Inflation dynamics deteriorates and evidence supporting the learning hypothesis is less convincing. In effect, the learning model we propose incorporates a cross-equation restriction that requires what the Federal Reserve did to be consistent with why it did it. Our use of Greenbook data allows us to explicitly impose the cross-equation restriction, and perform an 'Irrational Expectations Econometrics' exercise of the type advocated by Ireland (2003).[1] The weakness of supporting evidence is also found to be robust to relaxation of some of the simplifying assumptions of the learning model.

The paper is organised as follows. In Section 2 we estimate a simple learning model without Greenbook data to show a good match to inflation dynamics but inconsistency between unemployment forecasts from the model and the Greenbooks. Section 3 explicitly incorporates Greenbook data into estimation and finds the fit to inflation dynamics worsens considerably once consistency is imposed. The robustness of the result is examined in Section 4, which shows that evidence for the learning hypothesis remains weak even if simplifying assumptions that the Federal Reserve ignores parameter uncertainty and has no explicit policy smoothing objective are relaxed. Section 5 embeds the learning model in a structural model of the economy to enable direct comparison of our results with those of Sargent, Williams and Zha (2006). A final Section 6 concludes.

2 Estimation without Greenbook data

2.1 A simple learning model of optimal policy

To highlight the pitfalls of only matching the dynamics of the Great Inflation, we use a learning model very similar to that in Sargent, Williams and Zha (2006) using only inflation

[1] Ireland (2003) suggests deriving cross-equation restrictions from learning models in the same way as cross-equation restrictions are derived in the 'Rational Expectations Econometrics' of Hansen and Sargent (1980). To the best of our knowledge, our paper is one of the first to seriously take up the suggestion.

and unemployment data. At the heart of the model is a step back from full rationality, in that the Federal Reserve is assumed to be unaware of the underlying structure determining unemployment in the economy. Instead, it has an approximating model of unemployment-inflation dynamics:

$$u_t = \alpha_t' \Phi_t + \sigma_w w_t, \tag{1}$$

in which Φ_t is a vector of current inflation, lags of inflation, lags of unemployment and a constant. Furthermore, the Federal Reserve believes that the coefficients in the approximating model follow a simple drifting process $\alpha_t = \alpha_{t-1} + \Lambda_t$, where the innovation term Λ_t is i.i.d. Gaussian with mean zero and variance-covariance matrix V. Λ_t is perceived as independent of w_t. Given the simplicity of the perceived drifting process, the monetary authority obtains current estimates of the coefficients in its approximating model from a standard Kalman filter recursion. Defining $\hat{\alpha}_{t|t-1} \equiv E(\alpha_t \,|\mathcal{J}_{t-1})$, $P_{t|t-1} \equiv Var(\alpha_t \,|\mathcal{J}_{t-1})$ and the time t dataset as $\mathcal{J}_t = \{u_1, \pi_1, \ldots, u_t, \pi_t\}$, we have:

$$\hat{\alpha}_{t+1|t} = \hat{\alpha}_{t|t-1} + \frac{P_{t|t-1}\Phi_t \left(u_t - \Phi_t' \hat{\alpha}_{t|t-1}\right)}{\sigma_w^2 + \Phi_t' P_{t|t-1}\Phi_t}, \tag{2}$$

$$P_{t+1|t} = P_{t|t-1} - \frac{P_{t|t-1}\Phi_t \Phi_t' P_{t|t-1}}{\sigma_w^2 + \Phi_t' P_{t|t-1}\Phi_t} + V. \tag{3}$$

The objective for the Federal Reserve is set inflation π_t to minimise deviations in inflation and unemployment from their target levels π^* and u^*. δ is the discount factor and λ is the relative weight given to unemployment deviations from target:

$$\min_{\{\pi_t\}_{t=0}^{\infty}} \hat{E} \sum_{j=0}^{\infty} \delta^j \left\{ (\pi_{t+j} - \pi^*)^2 + \lambda \left(\tilde{u}_{t+j} - u^*\right)^2 \right\}. \tag{4}$$

To improve tractability, we follow Kreps (1998) and Sargent (1999) and assume the Federal Reserve forms forward-looking expectations using its approximating model of unemployment-inflation dynamics, but with coefficients fixed at their current estimates. Mathematically, such 'anticipated utility' behaviour implies that expected future values of unemployment are defined by the linear recursion $\tilde{u}_{t+j} = \hat{\alpha}_{t|t-1}' \tilde{\Phi}_{t+j}$, where the notation $\tilde{u}_{t+j}$ indicates the expected value of u_{t+j}. The assumption that forward-looking expectations are formed in this way means the objective function is quadratic in the vector of expected values $\tilde{\Phi}_{t+j}$. This is convenient as it simplifies the derivation of optimal policy considerably. With an objective that is quadratic in expected values and expected values themselves defined by a simple linear recursion, the

Federal Reserve faces a standard linear-quadratic control problem. The solution is a best response function $\pi_t = h(\hat{\alpha}_{t|t-1})'\phi_t$, where ϕ_t is a subset of Φ_t containing a constant and lagged values of inflation and unemployment. Optimal policy has inflation reacting linearly to the current state of the economy, with the strength of the reaction depending on the estimates of the coefficients in the Federal Reserve's approximating model.

2.2 Estimation, priors and data

We derive estimates of the free parameters in the model by acknowledging that the best response function provides only an approximate representation of Federal Reserve policy. The empirical specification therefore allows an i.i.d. Gaussian residual w_{2t} to explain discrepancies between the model and the data:

$$\pi_t = h(\hat{\alpha}_{t|t-1})'\phi_t + \sigma_2 w_{2t}. \tag{5}$$

The model is estimated by applying the Bayesian MCMC algorithm developed in Sargent, Williams and Zha (2006). Our estimation involves only minor changes to their methodology, so we restrict ourselves to a brief overview of the steps involved. At the centre of the algorithm is a Gibbs sampler that successively draws from two conditional distributions to generate a sample from the joint distribution of the free parameter estimates. The first conditional distribution is for the variance σ_2^2 of the residual w_{2t}, and has an inverse gamma conjugate prior. The second conditional distribution is for $\{P_{1|0}, V\}$, the Federal Reserve's perception of initial estimation precision and the variance-covariance matrix of the drifting coefficients. There is no suitable conjugate prior for this so a Metropolis algorithm is used to generate draws for the conditional posterior distribution. The remaining free parameters $\{\delta, \lambda, \pi^*, u^*, \hat{\alpha}_{1|0}\}$ and all priors are set to the values in Sargent, Williams and Zha (2006), reproduced for completeness in Appendix A.[2]

Our data series and sample period are chosen to match Sargent, Williams and Zha (2006). As the empirical counterpart of unemployment we use the civilian unemployment rate, 16 years and older, seasonally adjusted from the BLS. Inflation is measured by the annual (12

[2] There is a small error in the C++ codes used in Sargent, Williams and Zha (2006), meaning their results are derived under a more diffuse prior on V than that stated in their paper. We adopt a similar diffuse prior on V to maintain comparability with the earlier results.

month end) change in the seasonally-adjusted PCE chain price index published by the BEA. The sample period for both series is January 1960 to December 2003.

2.3 Results

The first set of results we present are posterior estimates of the structural parameters in the model. The estimates in Table 1 are based on 40000 draws of the Gibbs sampler, taken after a sufficiently long burn-in period to ensure that the Markov chain has converged to its ergodic distribution.[3]

Parameter	*Posterior mean*
$P_{1\|0}$	$\begin{pmatrix} 0.1535 & 0.1710 & 0.0289 & -0.3292 & -0.0092 & -0.1423 \\ 0.1710 & 0.2060 & 0.0335 & -0.3801 & -0.0114 & -0.1643 \\ 0.0289 & 0.0335 & 0.0055 & -0.0631 & -0.0018 & -0.0273 \\ -0.3292 & -0.3801 & -0.0631 & 0.7175 & 0.0208 & 0.3102 \\ -0.0092 & -0.0114 & -0.0018 & 0.0208 & 0.0006 & 0.0090 \\ -0.1423 & -0.1643 & -0.0273 & 0.3102 & 0.0090 & 0.1341 \end{pmatrix}$
V	$\begin{pmatrix} 0.1274 & -0.1198 & 0.0227 & 0.0824 & 0.0020 & -0.5246 \\ -0.1198 & 0.1254 & -0.0029 & -0.0811 & 0.0222 & 0.8687 \\ 0.0227 & -0.0029 & 0.0671 & 0.0129 & 0.0822 & 1.0614 \\ 0.0824 & -0.0811 & 0.0129 & 0.0548 & -0.0012 & -0.3992 \\ 0.0020 & 0.0222 & 0.0822 & -0.0012 & 0.1070 & 1.4659 \\ -0.5246 & 0.8687 & 1.0614 & -0.3992 & 1.4659 & 26.2302 \end{pmatrix}$
$1/\sigma_2^2$	$\underset{(\cdot,\cdot)}{19.2741}$

Table 1: Posterior estimates without Greenbook data

[3] 68% probability intervals are given in parentheses.

The structural parameter estimates are close to those of Sargent, Williams and Zha (2006). We therefore replicate their finding that the learning hypothesis receives remarkable support from real-world data when matched to the dynamics of the Great Inflation. This is apparent in Figure 1 in the small variance of the residual w_{2t} and the closeness of actual and fitted values of inflation.

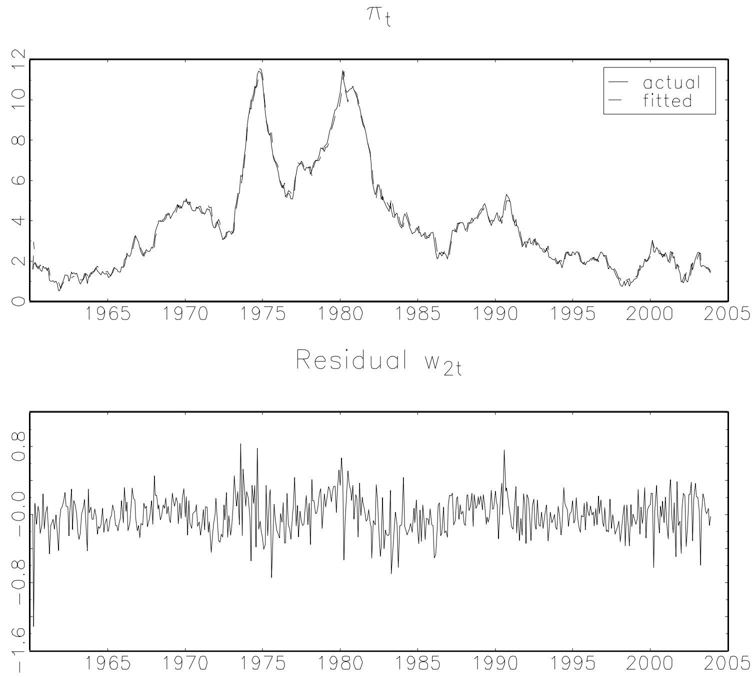

Figure 1: Actual inflation and inflation fitted without Greenbook data

Further support for the learning hypothesis is provided by the evolving Federal Reserve beliefs that underpin the good statistical fit of the model. Figure 2 shows the evolution of the Federal Reserve's view of the Phillips curve trade-off, as measured by the sum of the coefficients on current and lagged inflation in its approximating model of unemployment-inflation dynamics. Starting with almost no perceived trade-off between unemployment and inflation in the 1960s, there is clear identification of a discovery and subsequent abandonment of the Phillips curve during the sample period. The trade-off was perceived at its strongest in the mid 1970s, and was almost completely abandoned by the mid 1990s.[4]

[4] Whilst the evolution of beliefs does provide support for the learning hypothesis and Sargent's view that there has been a conquest of American inflation, it is worth noting that the Federal Reserve is supposed to believe in a positively-sloped Phillips curve towards the end of the sample period. We return to this potentially anomalous situation at the end of the next section.

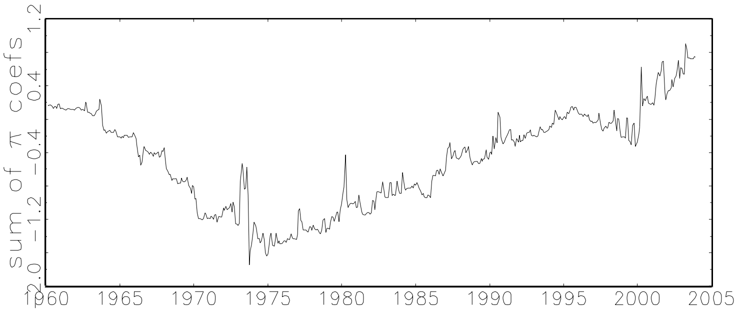

Figure 2: Perceived Phillips curve trade-off estimated without Greenbook data

2.4 Implications for Greenbook forecasts

Matching the dynamics of the Great Inflation in a model where the Federal Reserve discovers and then abandons the Phillips curve is impressive evidence in favour of the learning hypothesis. However, for the story to be totally convincing it should be that the Federal Reserve's approximating model of unemployment-inflation dynamics makes forecasts of unemployment that are consistent with those published in the Greenbook. The unemployment forecasts of the approximating model are defined by $\hat{E}(u_t) = \hat{\alpha}'_{t|t-1}\Phi_t$ from equation (1) and $\pi_t = h(\hat{\alpha}_{t|t-1})'\phi_t$ from the best response function. The contemporaneous unemployment forecast is therefore a linear function $\hat{E}(u_t) = g(\hat{\alpha}_{t|t-1})'\phi_t$ of a constant and lagged values of inflation and unemployment, with coefficients determined by the current estimates of parameters in the approximating model.

The unemployment forecasts implied by the approximating model are shown in Figure 3, together with actual forecasts extracted from historical Greenbooks.[5] The upper panel shows $\hat{E}(u_t)$, the contemporaneous forecast of unemployment given the current level of inflation set by policy. It is immediately apparent that a large discrepancy exists between forecasts from the approximating model and those from the Greenbooks. The problem is the huge volatility

[5]The forecasts in the Greenbook are 'judgemental' in exactly the same way as forecasts derived from the Federal Reserve's approximating model in our analysis. As summarised by Kozicki and Tinsley (2006), 'The multiperiod forecasts in a Greenbook provide repeated observations of predictions by the implicit forecast model of that Greenbook. Importantly, Greenbook forecasts provide measures of real-time central bank *perceptions* that are not evident in real-time data.' (original italics of authors). The details of how we constructed unemployment forecasts from the Greenbooks is described in Appendix B.

in forecasts from the approximating model, with a standard deviation of 3.7 percentage points that completely swamps the volatility in Greenbook forecasts. The correlation coefficient between forecasts from the approximating model and the Greenbooks is significant at 0.40, but the high frequency volatility in forecasts from the approximating model makes it difficult to argue that the two forecasts are consistent.

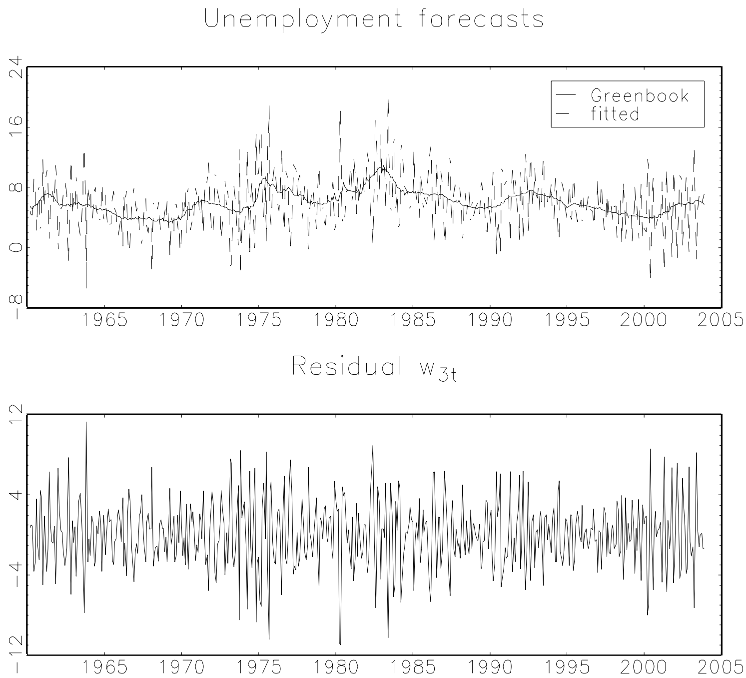

Figure 3: Unemployment forecasts without Greenbook data

The forecasting performance of the approximating model is an even greater cause for concern once we examine its ability to forecast the change in unemployment $\hat{E}(u_t - u_{t-1})$. In this case we find absolutely no correlation between the change in unemployment forecast by the approximating model and that forecast in the Greenbooks. The correlation coefficient between the two change forecasts is only 0.001, with problems of excess volatility in forecasts from the approximating model. Such evidence is problematic for the learning hypothesis of the Great Inflation. Whilst the hypothesis can explain inflation dynamics well, it appears that the motivation behind Federal Reserve policy identified in the estimation is not consistent with that implicitly defined in the Greenbooks. In other words, the learning model estimated without Greenbook data explains well *what* the Federal Reserve did, but offers a poor explanation as to *why* the Federal Reserve acted in the way it did.

3 Estimation with Greenbook data

3.1 Incorporating Greenbook forecasts

A natural response to the tensions in results without Greenbook data is to explicitly incorporate information from the Greenbook when estimating the model. This is easily achieved by continuing to acknowledge that the model is only an approximate representation of Federal Reserve behaviour, in which case the empirical specification permits an i.i.d. Gaussian residual w_{3t} to explain any discrepancies between Greenbook forecasts $E^{GB}(u_t)$ and the forecasts $g(\hat{\alpha}_{t|t-1})'\phi_t$ derived from the Federal Reserve's approximating model:

$$E^{GB}(u_t) = g(\hat{\alpha}_{t|t-1})'\phi_t + \sigma_3 w_{3t}. \tag{6}$$

The addition of a second measurement equation has only very minor consequences for the Bayesian MCMC algorithm used to estimate the model. There are now two variance parameters $\{\sigma_2^2, \sigma_3^2\}$ in the first conditional distribution of the Gibbs sampler, but the second conditional distribution is unchanged. The new residual w_{3t} has the same inverse gamma conjugate loose prior as the w_{2t} residual. All remaining free parameters and priors are kept at values used in the previous section when estimating without Greenbook data.

3.2 Results

Table 2 shows posterior estimates of the structural parameters in the model when estimation incorporates forecast information from the Greenbook. Compared to the results in Table 1 when Greenbook data was ignored, a rise in σ_2^2 signals a deterioration in the model's ability to match the dynamics of the Great Inflation. The worsening fit is further reflected in Figure 4, where larger and more persistent residuals are needed to reconcile fitted and observed inflation. In particular, large residuals are needed around the turning points of inflation and the model struggles to explain inflation patterns over the last decade of the sample. The evidence in Figure 4 suggests a more prominent role for shocks in the rise and fall of American inflation than previously identified.

Parameter	*Posterior mean*
$P_{1\|0}$	$\begin{pmatrix} 0.0420 & 0.0016 & -0.0286 & -0.0061 & 0.0142 & 0.0054 \\ 0.0016 & 0.0141 & 0.0007 & 0.0010 & -0.0039 & -0.0115 \\ -0.0286 & 0.0007 & 0.0197 & 0.0043 & -0.0103 & -0.0050 \\ -0.0061 & 0.0010 & 0.0043 & 0.0901 & -0.0213 & -0.0239 \\ 0.0142 & -0.0039 & -0.0103 & -0.0213 & 0.0103 & 0.0100 \\ 0.0054 & -0.0115 & -0.0050 & -0.0239 & 0.0100 & 0.0165 \end{pmatrix}$
V	$\begin{pmatrix} 0.0159 & -0.0138 & 0.0279 & -0.0017 & -0.0272 & -0.0607 \\ -0.0138 & 0.0184 & -0.0362 & -0.0041 & 0.0317 & 0.1021 \\ 0.0279 & -0.0362 & 0.0725 & 0.0072 & -0.0637 & -0.2003 \\ -0.0017 & -0.0041 & 0.0072 & 0.0050 & -0.0040 & -0.0366 \\ -0.0272 & 0.0317 & -0.0637 & -0.0040 & 0.0576 & 0.1656 \\ -0.0607 & 0.1021 & -0.2003 & -0.0366 & 0.1656 & 0.6304 \end{pmatrix}$
$1/\sigma_2^2$	$\underset{(\cdot,\cdot)}{3.6956}$
$1/\sigma_3^2$	$\underset{(\cdot,\cdot)}{10.7410}$

Table 2: Posterior estimates with Greenbook data

The estimates of $P_{1|0}$ and V reported in Table 2 are generally smaller than the corresponding estimates without Greenbook data. The implication is that incorporating Greenbook data in the estimation changes our view of how the Federal Reserve perceived the coefficients in its approximating model to be drifting. The smaller elements of the initial precision matrix $P_{1|0}$ mean that the Federal Reserve had more confidence than previously thought in its model at the start of the sample period. Similarly, smaller elements in the variance-covariance matrix V of the perceived drift process imply that the Federal Reserve should be viewed as placing less weight than previously thought on coefficient drift as a source of forecast errors. Both these factors contribute to the Federal Reserve being less inclined to change the estimates of the coefficients in its model.

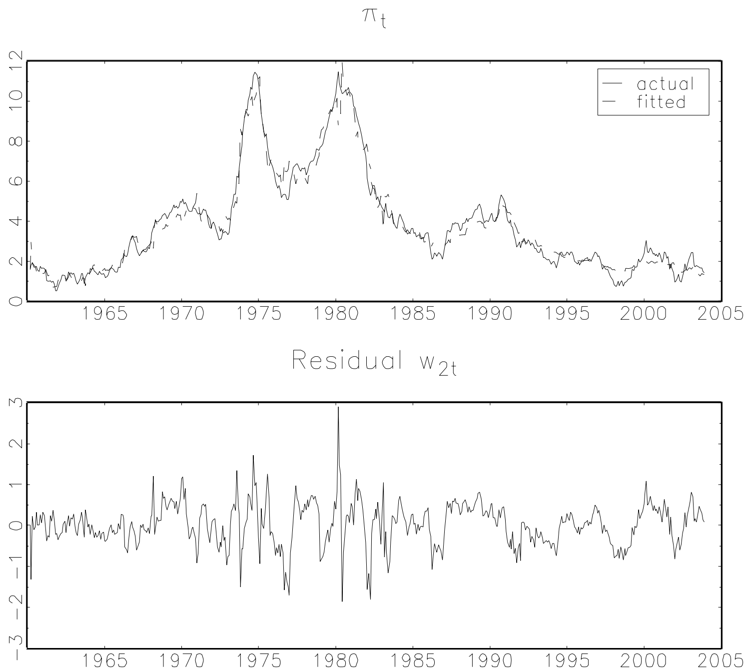

Figure 4: Actual inflation and inflation fitted with Greenbook data

The reluctance of the Federal Reserve to entertain coefficient drift as a source of forecast error translates into a less dramatic discovery and abandonment of the Phillips curve during the Great Inflation. In Figure 5 there is still a shadow of the conquest story, but the magnitude of the perceived trade-off peaks at about one sixth of that estimated without Greenbook data.[6] The evidence in favour of the learning hypothesis is weak if Federal Reserve beliefs are relatively stable in this way. With shocks also playing a more prominent role in inflation dynamics, estimation with Greenbook data appears to call for a reassessment of the relationship between the learning hypothesis and the Great Inflation.

[6] Interestingly, incorporating Greenbook data removes the Federal Reserve's anomalous belief in a positively-sloped Phillips curve at the end of the sample period. However, the fit of the model is not good so something other than learning must be driving inflation dynamics at that time.

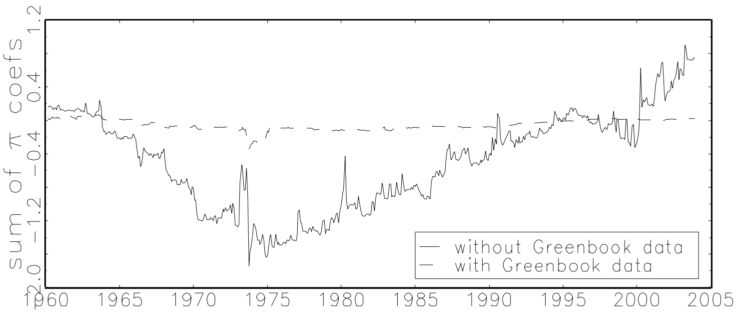

Figure 5: Perceived Phillips curve trade-off estimated with Greenbook data

3.3 Fit to Greenbook forecasts

The unemployment forecasts from the Federal Reserve's approximating model by definition match those from the Greenbook better once Greenbook data is incorporated in estimation. Figure 6 shows that much of the high frequency volatility in model forecasts has been removed. The correlation between the two forecasts is 0.98 and even the change forecasts have a highly-significant correlation of 0.19.

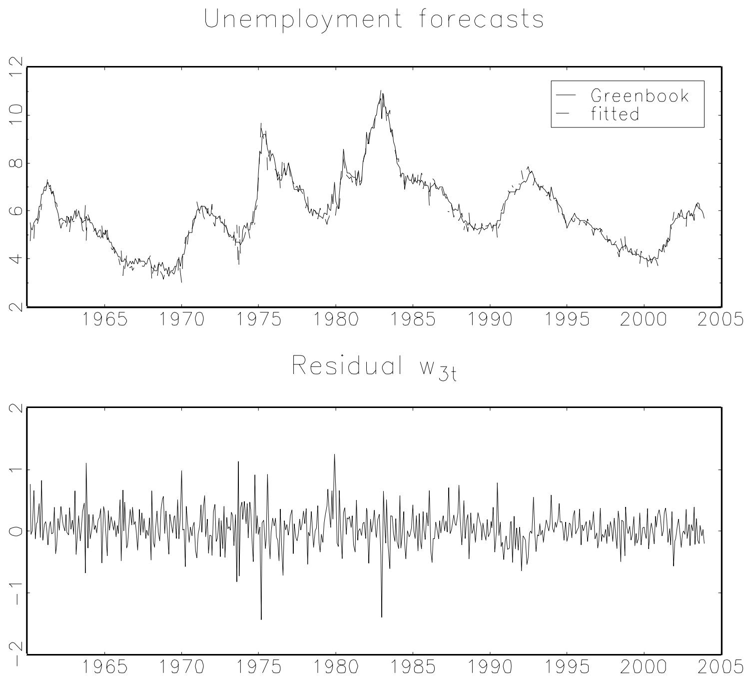

Figure 6: Unemployment forecasts with Greenbook data

The conclusion from estimating the model with Greenbook data is that it is difficult to simultaneously explain the dynamics of the Great Inflation and the unemployment forecasts published in the Greenbooks with the learning hypothesis. The Bayesian MCMC algorithm is based on identical priors across residuals in the two measurement equations, so effectively places equal weight on matching inflation dynamics and unemployment forecasts. Gibbs sampling then shades the estimation in favour of matching unemployment forecasts over matching inflation dynamics, with the greater residuals in the inflation equation then opening the door to alternative explanations of the Great Inflation.

4 Two robustness exercises

The results in the previous sections of the paper have been obtained under quite strict simplifying assumptions. A good robustness check is therefore whether the results continue to hold when some of the simplifying assumptions are relaxed. The first simplification we relax is the assumption that the Federal Reserve completely ignores uncertainty when setting policy. Whilst this may be acceptable as a first approximation, it does beg the question of why the Federal Reserve would completely disregard the numerous policy implications in the vast academic literature on optimal and robust control under uncertainty. Brainard's paper on uncertainty and the effectiveness of policy was published as early as 1967, so should have been in the consciousness of the Federal Reserve throughout the Great Inflation period. The second simplification to relax is the assumption that the Federal Reserve has no explicit incentive to smooth its policy. Whilst it is difficult to find solid microfoundations for a smoothing term in the objective function, many empirical studies suggest that policy inertia is pervasive in the economy. For example, Sack and Wieland (2000) discuss strong empirical evidence that interest rates are smoothed by the Federal Reserve.

4.1 Parameter uncertainty

The contention in this section is that policy should be based on the Federal Reserves's current view of the monetary transmission mechanism, but needs to explicitly take estimated parameter uncertainty into account. In other words, policy should respond to the current estimated coefficients $\hat{\alpha}_{t|t-1}$ of the Federal Reserve's approximating model and the precision $P_{t|t-1}$ with which those coefficients are estimated. This is potentially important because our estimation results so far suggest that the Federal Reserve perceives a large degree of uncertainty at all times. Figure 7 shows perceived two standard deviation confidence intervals around the evolution of the Federal Reserve's view of the Phillips curve, as identified in the estimation with Greenbook data. The degree of uncertainty is far from negligible, with the trade-off being perceived as almost statistically insignificant even at the height of the Great Inflation.

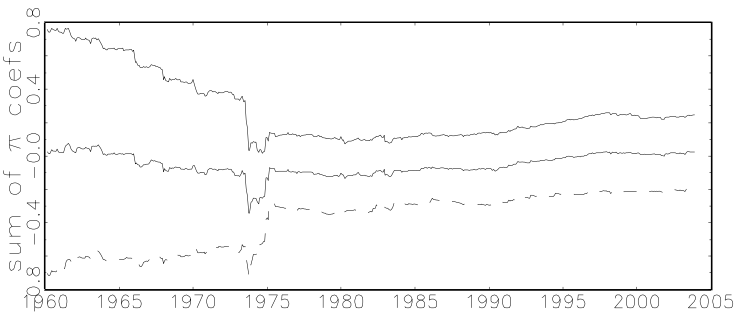

Figure 7: Perceived uncertainty in the perceived Phillips curve trade-off, estimated with Greenbook data

We start the mathematical derivation of optimal policy under uncertainty by generalising the Federal Reserve's objective function (4) to:

$$\min_{\{\pi_t\}_{t=0}^{\infty}} \hat{E} \sum_{j=0}^{\infty} \delta^j \left\{ (\pi_{t+j} - \pi^*)^2 + \lambda((\tilde{u}_{t+j} - u^*)^2 + \sigma Var(u_{t+j})) \right\}. \tag{7}$$

The notation $\tilde{u}_{t+j}$ again indicates the expected value of u_{t+j}, so our generalisation is akin to the bias-variance decomposition familiar in econometric forecasting. Indeed, increasing σ makes the monetary authority place less weight on expected unemployment being close to target (the bias term), and more weight on unemployment being certain (the variance term).[7] The next step is to explain how the Federal Reserve forms projections of the future bias and variance terms in its objective function. For the bias term, we follow Kreps (1998) and assume that the Federal Reserve projects forward using its approximating model of unemployment-inflation dynamics, but with coefficients fixed at their current estimates. This 'anticipated utility' behaviour implies that expected future values of unemployment are

[7] Our choice of σ to characterise the monetary authority's attitude to uncertainty is not coincidental. There is a direct analogy to Whittle's (1990) specification of risk-sensitive preferences, since $-2\sigma^{-1} \log E \exp(-0.5\sigma(u_t - u^*)^2) \approx (\hat{u}_t - u^*)^2 + \sigma Var(u_t)$. Our measures of risk sensitivity are therefore equivalent up to an approximation error.

given by $\tilde{u}_{t+j} = \hat{\alpha}'_{t|t-1}\tilde{\Phi}_{t+j}$. For the variance term, we follow Sack (2000) and assume that the Federal Reserve projects forward on the basis of the precision with which the parameters in its approximating model are estimated. The Federal Reserve therefore approximates future uncertainty by $Var(u_{t+j}) \approx \tilde{\Phi}'_{t+j}P_{t|t-1}\tilde{\Phi}_{t+j}$, where the timing indicates that future projections are based on the current estimate of the precision matrix. The assumed form of future projections keeps the objective function linear-quadratic in the vector of expected values $\tilde{\Phi}_{t+j}$, and the Federal Reserve continues to face a standard linear-quadratic control problem. The solution is a best response function:

$$\pi_t = h(\hat{\alpha}_{t|t-1}; P_{t|t-1})'\phi_t. \tag{8}$$

Optimal policy under parameter uncertainty has intended inflation reacting linearly to the current state of the economy, with the strength of the reaction depending on both the estimates of the coefficients in the Federal Reserve's approximating model and the precision with which those coefficients are estimated. It is precisely here that policy differs from that in the previous sections. There we adopted a stricter interpretation of Kreps (1998) 'anticipated utility' behaviour in which only the bias term was projected forwards, so the Federal Reserve ignored uncertainty and policy only depended on current coefficient estimates, not the precision with which they are estimated. The results in the previous sections and Sargent, Williams and Zha (2006) correspond to a special case of the generalised objective function where the risk sensitivity parameter σ is set to zero.

The only change in the Bayesian MCMC estimation algorithm we require is a redefinition of the conditional distribution for $\{P_{1|0}, V\}$ to allow for the reaction of policy to uncertainty. In the redefined distribution we set the risk parameter σ to unity to balance the incentives for policy to minimise the bias and variance terms. Table 3 presents our estimation results and compares them to the baseline where policy ignores parameter uncertainty. We also report the maximum log value of the likelihood (multiplied by the prior).

Parameter	*Baseline model*	*Parameter uncertainty*
$1/\sigma_2^2$	$\underset{(\cdot,\cdot)}{3.6956}$	$\underset{(\cdot,\cdot)}{3.0259}$
$1/\sigma_3^2$	$\underset{(\cdot,\cdot)}{10.7410}$	$\underset{(\cdot,\cdot)}{15.3524}$
log-likelihood	-258.1	-208.3

Table 3: Posterior estimates for parameter uncertainty with Greenbook data

The results when policy ignores uncertainty replicate those of Section 3. Against this benchmark, allowing policy to react to uncertainty has striking implications. First and foremost is the rise in the log-likelihood, with a logarithmic gain of 49.8 implying a substantial improvement in the statistical fit of the model.[8] Second is the improved matching of Greenbook forecasts, achieved at the cost of a worse match to inflation dynamics. However, the improvement in statistical fit does not translate into significant changes in the economic fit of the model. The estimates of $\{P_{1|0}, V\}$ are not sufficiently different to their baseline values to restore confidence in the learning hypothesis of the Great Inflation.[9] This is clear in Figure 8, where the estimated evolution of the Federal Reserve's view of the Phillips curve is robust to whether or not policy reacts to uncertainty. The changes in the variances of the residuals are similarly not large enough to overturn our view that the learning hypothesis is not completely convincing.

[8] Formally, given equal prior weights the posterior odds ratio almost completely favours the model where policy reacts to uncertainty. The likelihoods are directly comparable because both models have the same free parameters.

[9] Full estimation results are available from the authors on request.

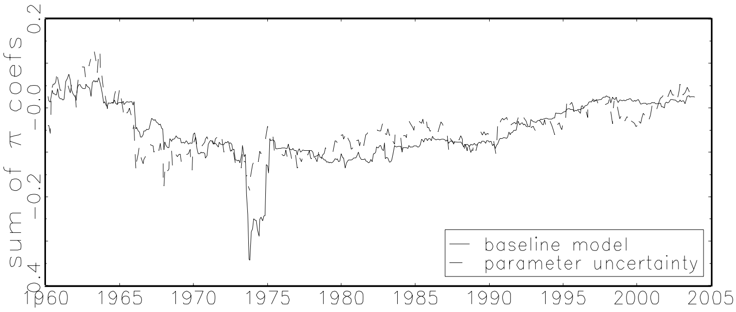

Figure 8: Perceived Phillips curve trade-off with parameter uncertainty, estimated with Greenbook data

The differences in estimation results are difficult to interpret because parameter uncertainty has a complex and potentially ambiguous effect on optimal policy. If uncertainty about the impact of policy dominates then the seminal result of Brainard (1967) applies and policy tends to be cautious, but this result can be overturned if there is sufficient uncertainty about transition dynamics.[10] In addition, elements lying off the leading diagonal of the precision matrices have potentially ambiguous effects because they give incentives for optimal policy to exploit the dynamic structure of uncertainty, as discussed in Chow (1977). Our precision matrices are dominated by the off-diagonal elements, implying high covariance between parameter estimates and a complex role for uncertainty dynamics in policy

4.2 Policy smoothing

The argument in this section is that the Federal Reserve has an incentive to smooth policy for reasons that are not explicitly articulated in its approximating model of unemployment-inflation dynamics. Perhaps the most compelling idea is the observation of Goodfriend (1987) that central banks smooth interest rates to maintain 'orderly money markets'. The Federal Reserve's approximating model abstracts from the impact of policy on financial stability, in which case the costs of volatile policy may be understated. Our framework does not permit direct modelling of the risks of interest rate volatility, but we can investigate the effects of

[10] Craine (1979) shows that very active policy is optimal when uncertainty about transition dynamics is dominant.

policy smoothing in general by expanding the Federal Reserve's objective function to include a term in the change in the policy instrument, inflation:

$$\min_{\{\pi_t\}_{t=0}^{\infty}} \hat{E}\sum_{j=0}^{\infty}\delta^j \left\{(\pi_{t+j}-\pi^*)^2+\lambda(\tilde{u}_{t+j}-u^*)^2+\omega\left(\Delta\pi_{t+j}\right)^2\right\}. \tag{9}$$

The strength of the incentive to smooth policy is measured by the parameter ω. The additional term in the objective function has similar implications for Bayesian MCMC estimation as in the previous exercise where the objective was generalised to allow for parameter uncertainty. The conditional distribution for $\{P_{1|0}, V\}$ has to be redefined but otherwise the algorithm remains unchanged. We set $\omega = 0.5$ to allow for an incentive to smooth policy that does not jeopardise the fundamental focus of policy on stabilising inflation and unemployment around their target values π^* and u^*. The estimation results are given in Table 4.

Parameter	*Baseline model*	*Policy smoothing*
$1/\sigma_2^2$	3.6956 $(\cdot,\cdot)$	4.1678 $(\cdot,\cdot)$
$1/\sigma_3^2$	10.7410 $(\cdot,\cdot)$	14.0912 $(\cdot,\cdot)$
log-likelihood	-258.1	-152.2

Table 4: Posterior estimates for policy smoothing with Greenbook data

Introducing policy smoothing unambiguously improves the statistical fit of the model. The log-likelihood rises by 105.9 so a formal Bayesian odds ratio test overwhelmingly supports the policy smoothing model.[11] The match to inflation dynamics and unemployment forecasts also improves. Unfortunately for the learning hypothesis, the huge improvement in statistical fit is not reflected in substantially changed estimates of the structural parameters $P_{1|0}$ and V. The economic fit of the model is consequently unaffected by the introduction of policy smoothing, and the evidence in favour of the learning hypothesis remains weak. Figure 9 shows only slight differences between estimates of the Federal Reserve's view of the Phillips curve obtained with

[11] The log-likelihood also rises by 56.1 relative to the parameter uncertainty model, suggesting that the policy smoothing model also dominates the parameter uncertainty model.

and without policy smoothing. Our fundamental message that learning plays a lesser role (and shocks a greater role) in the Great Inflation than previously thought appears robust to the introduction of policy smoothing.

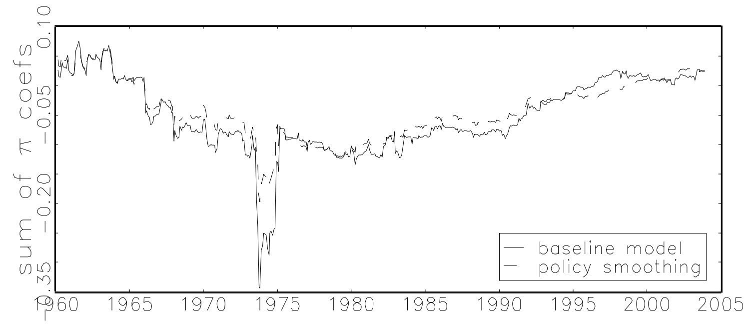

Figure 9: Perceived Phillips curve trade-off with policy smoothing, estimated with Greenbook data

The differences that do exist between estimation results can be explained by inertia induced by the smoothing term in the objective function. Policy inertia removes some of the need for persistent shocks to explain low frequency fluctuations in inflation, and helps dampen the volatile unemployment forecasts produced by the Federal Reserve's approximating model.

5 A structural model

In this section we follow Sargent, Williams and Zha (2006) and embed the learning hypothesis in a structural model of the economy. The advantages in doing so are threefold. Firstly, unemployment becomes endogenous so it is possible to test whether the learning hypothesis can also explain unemployment dynamics. Secondly, the ability to distinguish between anticipated and unanticipated changes in unemployment helps identify whether corresponding changes in inflation were themselves anticipated or unanticipated. Thirdly, it can be shown that the structural model converges to a well-defined self confirming equilibrium.[12] The underlying structure of the economy adopted by Sargent, Williams and Zha is described by a Lucas

[12] Sargent and Williams (2005) use techniques from stochastic approximation theory to characterise the possible outcomes of the Federal Reserve's learning process.

natural-rate Phillips curve and a true inflation process:

$$
\begin{aligned}
u_t - u^{**} &= \theta_0(\pi_t - E_{t-1}\pi_t) + \theta_1(\pi_{t-1} - E_{t-2}\pi_{t-1}) + \tau_1(u_{t-1} - u^{**}) + \sigma_1 w_{1t}, && (10) \\
\pi_t &= x_{t-1} + \sigma_2 w_{2t}, && (11)
\end{aligned}
$$

where u^{**} is the natural rate of unemployment. Equation (10) is an expectations-augmented Phillips curve in which unemployment is driven by unexpected inflation movements and an unemployment shock. Equation (11) states that the Federal Reserve controls inflation up to a random control error. We refer to the policy instrument x_{t-1} as intended inflation. If we assume that private agents form expectations based on the empirical specification (6) then the random control error w_{2t} corresponds exactly to the residual w_{2t} in the previous sections. In this case expected inflation is given by $E_{t-1}\pi_t = h(\hat{\alpha}_{t|t-1})'\phi_t = x_{t-1}$ and unexpected inflation is simply w_{2t}.

The natural-rate Phillips curve has five structural parameters that need to be estimated alongside the parameters of the Federal Reserve's learning model. The first four $\{u^{**}, \theta_0, \theta_1, \tau_1\}$ are assigned a normal conjugate prior and are drawn in a third step of the Gibbs sampler. The fifth σ_1^2 has an inverse gamma conjugate prior and is drawn alongside σ_2^2 and σ_3^2 in the second step of the Gibbs sampler. The precise specification of these additional prior distributions follows the details of Sargent, Williams and Zha (2006) summarised in Appendix A. The estimation results when embedding the learning hypothesis in a structural model are so close to the results in Sections 3 and 4 that we only report estimates of the new parameters in Table 5. The values of $P_{1|0}$ and V do not change much, so the evolving views of the Federal Reserve identified with and without Greenbook data are very similar to those in Figure 5.

Sargent, Williams and Zha (2006) claim that estimation of the full structural model provides substantive evidence in support of the learning hypothesis of the Great Inflation. We are more sceptical because embedding the learning hypothesis in a structural model does little to solve the problems identified in previous sections of this paper. When estimation is performed without Greenbook data there is still a problem with excess volatility in the forecasts produced by the Federal Reserve's approximating model. Similarly, estimations with Greenbook data still lead to considerable deterioration in the fit to inflation dynamics, and open the door to other explanations of the Great Inflation.

Parameter	*Structural model estimated without Greenbook data*	*Structural model estimated with Greenbook data*
u^{**}	6.1104 (5.2500,7.1579)	6.0584 $(\cdot,\cdot)$
θ_0	-0.0008 $(-0.0237,0.0475)$	0.0095 $(\cdot,\cdot)$
θ_1	-0.0122 $(-0.0375,0.0297)$	0.0328 $(\cdot,\cdot)$
τ_1	0.9892 (0.9852,0.9960)	0.9910 $(\cdot,\cdot)$
$1/\sigma_1^2$	35.6538 (28.7565,32.4947)	30.9804 $(\cdot,\cdot)$
$1/\sigma_2^2$	18.9767 (15.6565,18.2557)	2.5315 $(\cdot,\cdot)$
$1/\sigma_3^2$	–	14.4114 $(\cdot,\cdot)$

Table 5: Posterior estimates of structural parameters

The estimate of σ_1^2 without Greenbook data suggests that the structural model is successful in explaining unemployment dynamics. However, the success is partially illusionary because the estimated values of θ_0 and θ_1 have such small magnitude that unexpected inflation only plays a very minor role in the determination of unemployment.[13] The small coefficients on unexpected inflation also mean that decomposing unemployment into anticipated and unanticipated changes is not very helpful when deciding whether changes in inflation are anticipated (from $h(\hat{\alpha}_{t|t-1})'\phi_t$) or unanticipated (from w_{2t}). With unemployment effectively rendered exogenous by the estimation process, it is no great surprise that results from previous sections are robust to embedding the learning hypothesis in the structural model. The same broad intuition applies to estimation results with Greenbook data, although in this case unexpected inflation is sufficiently correlated with changes in unemployment to make the estimates of θ_0 and θ_1 positive. We interpret the results in this section as suggesting that there is only limited value-added in embedding the learning hypothesis in a structural model based on the Lucas natural-rate Phillips curve.

[13] Variance decomposition analysis of the results without Greenbook data attribute only 0.005% of the total variance in unemployment to unexpected inflation effects.

6 Conclusions

The main conclusion in this paper is that empirical evidence on what caused the Great Inflation is less conclusive than previously thought. The current state-of-the-art empirical assessment is Sargent, Williams and Zha (2006), who claim there is strong empirical support for a learning hypothesis of the Great Inflation in which inflation rose and fell as the Federal Reserve discovered and then abandoned the Phillips curve. We argue that their claim is overstated. The problem is that they offer a very good explanation of *what* the Federal Reserve did but a very poor explanation of *why* the Federal Reserve did it. In particular, if they are correct then the unemployment forecasts published by the Federal Reserve should have been extremely volatile. Our analysis of historic Greenbooks shows this is not the case. Require the learning hypothesis to explain both the 'what' and 'why' of Federal Reserve policy produces supporting empirical evidence that is much less convincing than before.

We see our conclusion as a positive step in understanding the factors behind the Great Inflation. The problem in the existing literature is that the learning hypothesis performs very well and there is little scope for alternative explanations. Our decision to incorporate Greenbook forecast data in estimations tempers the good performance of the learning hypothesis and so re-opens the door to other possible explanations. It remains to be seen whether alternative ideas based on monetary financing of fiscal expansion or oil shocks are able to bridge the gap we have uncovered. Such an exercise is beyond the scope of this paper at present, in part because much of the literature on the Great Inflation draws heavily on narrative evidence and is not immediately amenable to formal analysis.[14]

[14] The debate between Meltzer (2005) and Romer (2005) is an excellent summary of current thinking on the causes of the Great Inflation.

References

[1] Brainard, W.C., (1967), Uncertainty and the Effectiveness of Policy, *American Economic Review*, 57(2), 411-425

[2] Cho, I-K., Williams, N. and T.J. Sargent, (2002), Escaping Nash Inflation, *Review of Economic Studies*, 69(1), 1-40

[3] Chow, G., (1977), *Analysis and control of dynamic economic systems*, Wiley

[4] Cogley, T. and T.J. Sargent, (2005a), The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty, *Review of Economic Dynamics*, 8(2), 528-563

[5] Cogley, T. and T.J. Sargent, (2005b), Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S., *Review of Economic Dynamics*, 8(2), 262-302

[6] Craine, R., (1979), Optimal monetary policy with uncertainty, *Journal of Economic Dynamics and Control*, 1(1), 59-83

[7] Goodfriend, M., (1987), Interest rate smoothing and price level trend stationarity, *Journal of Monetary Economics* 19(3), 335–348.

[8] Hansen, L.P. and T.J. Sargent, (1990), Formulating and Estimating Dynamic Linear Rational Expectations Models, *Journal of Economic Dynamics and Control*, 2(1), 7-46

[9] Ireland, P.N., (1999), Does the Time-Consistency Problem Explain the Behavior of Inflation in the United States?, *Journal of Monetary Economics*, 44(2), 279-91

[10] Ireland, P.N., (2003), Irrational Expectations and Econometric Practice: Discussion of Orphanides and Williams, "Inflation Scares and Forecast-Based Monetary Policy", Federal Reserve Bank of Atlanta Working Paper 2003-22

[11] Kozicki, S. and P.A. Tinsley, (2006), Minding the Gap: Central Bank Estimates of the Unemployment Natural Rate, *Computational Economics*, 27(2-3), 295-327.

[12] Kreps, D.M., (1998), Anticipated utility and dynamic choice, in *Frontiers of research in economic theory: the Nancy L. Schwartz memorial lectures, 1983-1997*, Cambridge University Press

[13] Meltzer, A.H., (2005), Origins of the Great Inflation, Federal Reserve Bank of St. Louis *Review*, 87(2, Part 2), 145-175

[14] Primiceri, G.E., (2005), Why Inflation Rose and Fell: Policymakers' Beliefs and US Post-war Stabilization Policy, National Bureau of Economic Research Working Paper 11147

[15] Romer, C.D., Commentary on 'Origins of the Great Inflation', Federal Reserve Bank of St. Louis *Review*, 87(2, Part 2), 177-185

[16] Sack, B., (2000), Does the fed act gradually? A VAR analysis, *Journal of Monetary Economics*, 46(1), 229-256

[17] Sack, W. and V. Wieland, (2000), Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence, *Journal of Economics and Business*, 52(1-2), 205-228

[18] Sargent, T.J., (1999), The Conquest of American Inflation, *Princeton University Press*, Princeton: New Jersey

[19] Sargent, T.J. and N. Williams, (2005), Impact of Priors on Convergence and Escape from Nash Inflation, *Review of Economic Dynamics*, 8(2), 360-391

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

[21] Tetlow, R and B. Ironside, (2007), Real-Time Model Uncertainty in the United States: The Fed, 1996–2003, *Journal of Money, Credit and Banking*, 39(7), 1533-1561

[22] Whittle, P., (1990). *Risk-Sensitive Optimal Control*, John Wiley & Sons, New York.

A Calibration and priors

The calibrated values in Table A.1 are taken directly from Sargent, Williams and Zha (2006). As they explain, the parameter σ_w (the perceived standard deviation of shocks in the Federal Reserve's approximating model) is unidentified when policy depends only on $\hat{\alpha}_{t|t-1}$ as in Sections 2, 3 and 4.2. We follow their lead and normalise σ_w to their estimate of the standard deviation of unemployment shocks in a structural model. No such problems arise with parameter uncertainty in Section 4.1 because policy reacts to both $\hat{\alpha}_{t|t-1}$ and $P_{t|t-1}$. In this case we retain the same value of σ_w, but note that here it represents a calibration and not normalisation. The values for $\hat{\alpha}'_{1|0}$ are derived from estimating the Federal Reserve's approximating model on presample data from January 1948 to December 1959.

Parameter	*Value*
δ	0.9936
λ	1
π^*	2
u^*	1
σ_w	0.175
$\hat{\alpha}'_{1\|0}$	$\begin{pmatrix} -0.1324 & 0.1419 & 1.0928 & -0.0216 & -0.1338 & 0.2190 \end{pmatrix}$

Table A.1: Calibrated parameter values

The priors in Table A.2 are also based on Sargent, Williams and Zha (2006). The matrices $\mathbf{C}_P$ and $\mathbf{C}_V$ are upper triangular Choleski decompositions of $P_{1|0}$ and V such that $P_{1|0} = \mathbf{C}'_P\mathbf{C}_P$ and $V = \mathbf{C}'_V\mathbf{C}_V$. The scaling factor of 400 in the prior distribution of V corrects for the small C++ coding error highlighted in footnote 2. The stated prior in Sargent, Williams and Zha (2006) has a scaling factor of 0.5, but in their C++ codes the prior distribution is completely flat. To ensure comparability of our results we use a high scaling factor to create a very diffuse prior distribution for V.

Parameter	*Distribution*	*Mean*	*Standard error*
$\mathbf{C}_P$	Normal	0	0.5×5 on diagonals, 0.5×2.5 on off-diagonals
$\mathbf{C}_V$	Normal	0	400×5 on diagonals, 400×2.5 on off-diagonals
$u^{**}(1-\tau_1)$	Normal	0.12	0.06
θ_0	Normal	-0.20	0.10
θ_1	Normal	-0.16	0.08
τ_1	Normal	0.98	0.01
$1/\sigma_1^2$	gamma	50	25
$1/\sigma_2^2$	gamma	50	25
$1/\sigma_3^2$	gamma	50	25

Table A.2: Prior distributions

B Construction of Greenbook forecasts

The Federal Reserve Bank of Philadelphia currently publishes historical projections from the Greenbooks of July 1966 to December 2001. Each Greenbook typically provides projections for the current quarter and a few quarters ahead, and we convert the Greenbook data into monthly unemployment forecasts by selecting the most appropriate quarterly projection for each month. January and February forecasts are Q1 projections from the same year; March, April and May forecasts are Q2 projections from the same year; June, July and August forecasts are Q3 projections from the same year; September, October and November forecasts are Q4 projections from the same year. The December forecast is the Q1 projection from the following year. There are no publicly-available Greenbook projections at the beginning and end of our sample period. For January 1960 to June 1966 we adopt a simple 'no change' forecast that unemployment will stay at its current level. For January 2002 to December 2003 we use the two year ahead projection published in the Greenbook of December 2001. Alternative methods for constructing the missing Greenbook forecasts were investigated, but found to have only minor implications our results.