

DISCUSSION PAPER SERIES

No. 6250

LEARNING AND THE GREAT INFLATION

Giacomo Carboni and Martin Ellison

INTERNATIONAL MACROECONOMICS



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP6250.asp

LEARNING AND THE GREAT INFLATION

Giacomo Carboni, University of Warwick
Martin Ellison, University of Warwick and CEPR

Discussion Paper No. 6250
April 2007

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INTERNATIONAL MACROECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Giacomo Carboni and Martin Ellison

ABSTRACT

Learning and the Great Inflation*

We respond to the challenge of explaining the Great Inflation by building a coherent framework in which both learning and uncertainty play a central role. At the heart of our story is a Federal Reserve that learns and then disregards the Phillips curve as in Sargent's *Conquest of American Inflation*, but at all times takes into account that its view of the world is subject to considerable uncertainties. Allowing Federal Reserve policy to react to these perceived uncertainties improves our ability to explain the Great Inflation with a learning model. Bayesian MCMC estimation results are encouraging and favour a model where policy reacts to uncertainty over a model where uncertainty is ignored. The posterior likelihood is higher and the internal Federal Reserve forecasts implied by the model are closer to those reported in the Greenbook.

JEL Classification: E52, E58 and E65

Keywords: great inflation, learning, monetary policy and uncertainty

Giacomo Carboni
Department of Economics
University of Warwick
Coventry
CV4 7AL
UK
Email: g.carboni@warwick.ac.uk

Martin Ellison
Department of Economics
University of Warwick
Coventry
CV4 7AL
UK
Email: M.Ellison@warwick.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=163153

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=146947

* We greatly appreciate Tao Zha making available his C++ routines for Bayesian MCMC estimation. We also thank seminar participants at Birkbeck, Cardiff, Durham, Glasgow, Humboldt and Warwick 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).

Submitted 23 March 2007

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 Cho, Williams and Sargent (2002), Cogley and Sargent (2005a, 2005b), Sargent and Williams (2005) and Primiceri (2005) 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 easily dominating a Bayesian vector autoregression in terms of its ability to match and forecast inflation dynamics. However, doubts still remain about the ability of the learning hypothesis to credibly explain the rise and fall in American inflation. The first doubt we highlight is that the good fit of Sargent, Williams and Zha (2006) is predicated on an assumption that the Federal Reserve completely ignores uncertainty when setting policy. Figure 1 shows how the Federal Reserve's view of the Phillips curve trade-off evolves over the Great Inflation period, together with the perceived two standard deviation confidence interval. The degree of uncertainty is far from negligible, with the trade-off being perceived as statistically insignificant even at the height of the Great Inflation. Our doubt crystalises as we consider the likely reaction of the Federal Reserve to such uncertainty, since Sargent, Williams and Zha (2006) simply assume that uncertainty is ignored when setting policy. Whilst this may be acceptable as a first approxima-

tion, our doubt arises because the implication is a Federal Reserve that completely disregards numerous policy implications from the vast academic literature on optimal and robust control under uncertainty. Brainard’s seminal paper on uncertainty and the effectiveness of policy was published as early as 1967, so would have been in the consciousness of the Federal Reserve throughout the Great Inflation period.

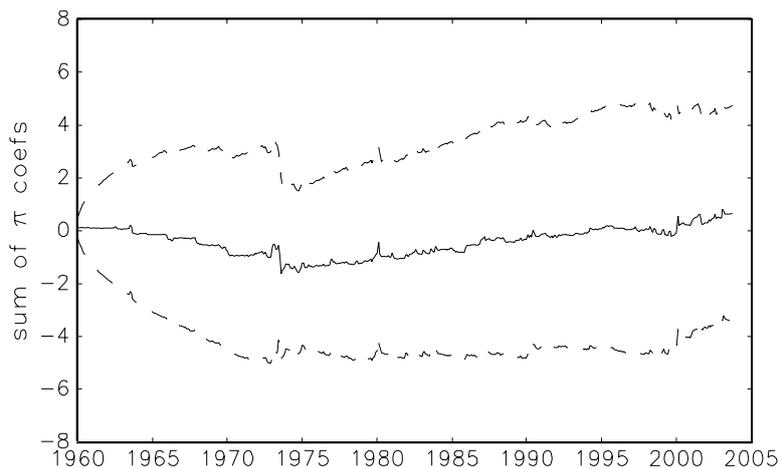


Figure 1: Sargent, Williams and Zha (2006) estimates of Phillips curve trade-off, as perceived by the Federal Reserve

The second reason to doubt learning explanations of the Great Inflation rests on the poor performance they imply for internal forecasts made by the Federal Reserve. Sargent, Williams and Zha (2006) rightly stress that a model based on the learning hypothesis embodies good internal forecasts for inflation. The prediction errors made by internal forecasts are small, being of similar magnitude to forecast errors from the Greenbook or a Bayesian vector autoregression. However, in their model the Federal Reserve has almost complete control over inflation, so good internal inflation forecasting only really implies that the Federal Reserve is able to accurately predict its own actions. We think this is important but only half the story: A good explanation of the Great Inflation should have the Federal Reserve’s internal forecasts performing well across a range of macroeconomic indicators, not just inflation. Unfortunately this is not the case in current learning explanations, as witnessed by Figure 2 which plots unemployment forecast errors attributed to the Federal Reserve by Sargent, Williams and Zha

(2006). The problem is that errors in the internal forecast of unemployment are huge, with a standard deviation of 3.1 percentage points that completely swamps the magnitude of forecast errors from the Greenbook or a Bayesian vector autoregression. We find it difficult to accept that the Federal Reserve made internal forecasting errors of this magnitude, especially since American unemployment was relatively stable at the time.

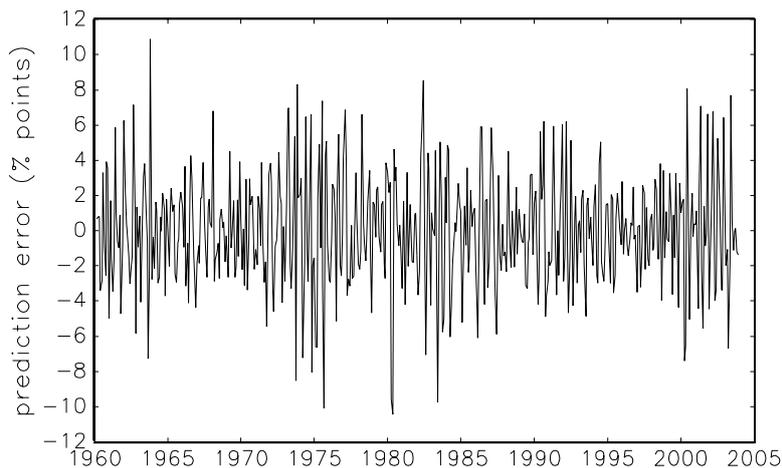


Figure 2: One-step-ahead prediction errors derived from the Federal Reserve’s internal unemployment forecast, as identified by Sargent, Williams and Zha (2006)

Our paper is motivated by doubts such as these as to whether the learning hypothesis can credibly explain the Great Inflation. We are uneasy with the assumption that the Federal Reserve completely ignores uncertainty, and are concerned that the Federal Reserve appears to make substantial errors in its own internal forecasting. To address these concerns, we embed the learning hypothesis of the Great Inflation in a model where the Federal Reserve is assumed to take uncertainty into account. Specifically, we follow Sack (2000) and allow policy to be influenced by how much parameter uncertainty the Federal Reserve perceives in its estimates of the monetary transmission mechanism, and by how much the Federal Reserve is sensitive to risk. Our model is ideal for comparing to the existing literature, since it conveniently reduces to that of Sargent, Williams and Zha (2006) if the Federal Reserve is made completely insensitive to risk. Innovating in this way directly responds to the first doubt we expressed about the learning hypothesis, namely that the Federal Reserve is assumed to

ignore uncertainty. Whether our innovation is successful depends ultimately on whether we can still explain the Great Inflation once policy is allowed to internalise uncertainty. Our second doubt is only indirectly addressed. However, our hunch is that the poor performance of internal forecasting in Sargent, Williams and Zha (2006) is intimately linked to having the Federal Reserve set policy according to estimates of the monetary transmission mechanism that are surrounded by great uncertainty. We are therefore optimistic that allowing the Federal Reserve to internalise uncertainty may also improve the apparent poor performance of its internal forecasting.

2 The model

Our model builds on that of Sargent, Williams and Zha (2006) by explicitly modelling the reaction of the monetary authority to the uncertainty it perceives in its estimation of the monetary transmission mechanism. Reacting to uncertainty has implications for the policy problem of the monetary authority, but leaves the rest of the model unchanged. The main focus of our model exposition is therefore the policy problem, but for completeness we begin by sketching out the rest of the model. The underlying structure of the economy is described by a Lucas natural-rate Phillips curve and a true inflation process:

$$u_t - u^{**} = \theta_0(\pi_t - E_{t-1}\pi_t) + \theta_1(\pi_{t-1} - E_{t-2}\pi_{t-1}) + \tau_1(u_{t-1} - u^{**}) + \sigma_1 w_{1t}, \quad (1)$$

$$\pi_t = x_{t-1} + \sigma_2 w_{2t}, \quad (2)$$

where u_t is unemployment, π_t is inflation and u^{**} is the natural rate of unemployment. Equation (1) is an expectations-augmented Phillips curve in which unemployment is an autoregressive process driven by unexpected inflation movements and an unemployment shock. Equation (2) states that the monetary authority controls inflation up to a random control error. We refer to the policy instrument x_{t-1} as intended inflation. Private agents are assumed to understand the monetary authority's actions, so $E_{t-1}\pi_t = x_{t-1}$. The shocks w_{1t} and w_{2t} are uncorrelated i.i.d. disturbances with standard normal distributions. In a step back from full rationality, the monetary authority is assumed to be unaware of the underlying structure determining unemployment in the economy. Instead, it has an approximating model

of unemployment-inflation dynamics:

$$u_t = \alpha_t' \Phi_t + \sigma_w w_t, \quad (3)$$

in which Φ_t is a vector of current inflation, lags of inflation, lags of unemployment and a constant. Compared to the true Phillips curve (1), the approximating model is misspecified in that it fails to recognise the role of inflation expectations in determining unemployment. Furthermore, the monetary authority believes (incorrectly) 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 $\mathcal{J}_t = \{u_1, \pi_1, \dots, u_t, \pi_t\}$, we have:

$$\hat{\alpha}_{t+1|t} = \hat{\alpha}_{t|t-1} + \frac{P_{t|t-1} \Phi_t (u_t - \Phi_t' \hat{\alpha}_{t|t-1})}{\sigma_w^2 + \Phi_t' P_{t|t-1} \Phi_t}, \quad (4)$$

$$P_{t+1|t} = P_{t|t-1} - \frac{P_{t|t-1} \Phi_t \Phi_t' P_{t|t-1}}{\sigma_w^2 + \Phi_t' P_{t|t-1} \Phi_t} + V. \quad (5)$$

We now turn to the policy problem faced by the monetary authority. Our contention is that policy should be based on the monetary authority'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 monetary authority's approximating model and the precision $P_{t|t-1}$ with which those coefficients are estimated. We start the mathematical derivation of optimal policy under uncertainty with a generalised objective for the monetary authority:

$$\min_{\{x_{t-1}\}_{t=0}^{\infty}} \hat{E} \sum_{j=0}^{\infty} \delta^j \{ (\pi_{t+j} - \pi^*)^2 + \lambda ((\tilde{u}_{t+j} - u^*)^2 + \sigma Var(u_{t+j})) \}. \quad (6)$$

π^* and u^* are target levels of inflation and unemployment, δ is the discount factor and λ is the relative weight given to unemployment deviations from target. The notation \tilde{u}_{t+j} indicates the expected value of u_{t+j} , so our generalisation of the objective function 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).¹ The next step in deriving optimal policy is an explanation of how the monetary authority 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 monetary authority projects forward 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 $\tilde{u}_{t+j} = \hat{\alpha}'_{t|t-1} \tilde{\Phi}_{t+j}$. For the variance term, we follow Sack (2000) and assume that the monetary authority projects forward on the basis of the precision with which the parameters in its approximating model are estimated. The monetary authority 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. Our assumption that future projections of the bias and variance terms are formed in this way means the objective function is linear-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 monetary authority faces a standard linear-quadratic control problem. The solution is a best response function:

$$x_{t-1} = h(\hat{\alpha}_{t|t-1}; P_{t|t-1})' \Phi_t. \quad (7)$$

Optimal policy under 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 monetary authority’s approximating model and the precision with which those coefficients are estimated. It is precisely here that our model differs from that used by Sargent, Williams and Zha (2006). They adopt a stricter interpretation of Kreps (1998) ‘anticipated utility’ behaviour in which only the bias term is projected forwards, so

¹Our 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.

their monetary authority ignores uncertainty and policy only depends on current coefficient estimates, not the precision with which they are estimated. Their model corresponds to a special case of ours in which the risk sensitivity parameter σ is set to zero.

3 Empirical methodology

We derive estimates of the free parameters in our model by applying the Bayesian MCMC algorithm developed in Sargent, Williams and Zha (2006). Allowing the monetary authority to react to uncertainty implies only minor changes to their methodology, so we restrict ourselves to a brief overview of the steps involved. At the heart of the algorithm is a Gibbs sampler that successively draws from three conditional distributions to generate a sample from the joint distribution of the free parameter estimates. The first and second conditional distributions are for the coefficients $\{u^{**}, \theta_0, \theta_1, \tau_1\}$ and variances $\{\sigma_1^2, \sigma_2^2\}$ in the true structure of the economy, and have normal-inverse gamma conjugate priors. The third conditional distribution is for $\{P_{1|0}, V\}$, the monetary authority's perception of initial estimation precision and the variance-covariance matrix of the drifting coefficients. There is no suitable conjugate prior for this distribution so a Metropolis algorithm is used to generate draws for the Gibbs sampler. The only change to the methodology we require is a redefinition of the third conditional distribution to allow for the reaction of policy to uncertainty. In the redefined distribution we set the risk sensitivity parameter $\sigma = 1$ when policy reacts to uncertainty and $\sigma = 0$ when uncertainty is ignored. The remaining free parameters $\{\delta, \lambda, \pi^*, u^*, \sigma_w, \hat{\alpha}_{1|0}\}$ and all priors are fixed according to the values in Sargent, Williams and Zha (2006).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 month end) change in the seasonally-adjusted PCE chain price index published by the BEA.

²The parameter σ_w is not identified when policy depends only on $\hat{\alpha}_{t|t-1}$, as in Sargent, Williams and Zha (2006). No such problems arise in our extension because policy depends on both $\hat{\alpha}_{t|t-1}$ and $P_{t|t-1}$. In estimations we retain the assumption $\sigma_w = \kappa\sigma_1$, but note that this represents calibration not normalisation once policy is allowed to react to uncertainty.

The sample period for both series is January 1960 to December 2003.

4 Results

The first set of results we present are posterior estimates of the structural parameters in the underlying Lucas natural-rate Phillips curve and the true inflation process. 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 We also report the maximum log value of the likelihood (multiplied by the prior).

Parameter	<i>Policy ignores uncertainty</i>	<i>Policy reacts to uncertainty</i>
u^{**}	6.1104 (5.2500,7.1579)	6.6758 (5.6382,7.1527)
θ_0	-0.0008 (-0.0237,0.0475)	-0.0836 (-0.1181,-0.0490)
θ_1	-0.0122 (-0.0375,0.0297)	-0.0708 (-0.1041,-0.0375)
τ_1	0.9892 (0.9852,0.9960)	0.9883 (0.9833,0.9935)
$1/\sigma_1^2$	35.6538 (28.7565,32.4947)	31.0769 (29.1923,32.9873)
$1/\sigma_2^2$	18.9767 (15.6565,18.2557)	21.1792 (19.8683,22.4902)
log-likelihood	564.92	614.46

Table 1: Posterior estimates of structural parameters

The results when policy ignores uncertainty replicate those of Sargent, Williams and Zha (2006). Against this benchmark, allowing policy to react to uncertainty has striking implications. First and foremost is the rise of the log-likelihood, with a logarithmic gain of 49.54

³68% probability intervals are given in parentheses.

implying a substantial improvement in the statistical fit of the model. Second is the greater precision of parameter estimates and the fall in the estimated variance of inflation control errors. The main conclusion of Table 1 is that allowing the Federal Reserve to react to uncertainty improves the ability of the learning hypothesis to explain the dynamics of inflation and unemployment. It appears that our first doubt about the learning hypothesis was misplaced, since it is actually easier to explain the Great Inflation when policy reacts to uncertainty than when uncertainty is ignored. The intuition behind this is that allowing policy to react to uncertainty leads to a better match between the inflation persistence implied by the model and that observed in the data. To understand why, we start by noting there are two reasons why inflation persistence changes once we allow for uncertainty. The first is a direct effect as intended inflation internalises the precision $P_{t|t-1}$ with which the coefficients $\hat{\alpha}_{t|t-1}$ in the Federal Reserve's approximating model are estimated. The second is an indirect effect due to the estimated precision matrices and coefficients themselves evolving differently when policy reacts to uncertainty.

To isolate the direct effect of uncertainty, we freeze the estimated precision matrices and coefficients at values when uncertainty is ignored and ask how policy would differ if uncertainty were to be internalised. The results are presented in the top two panels of Table 2, which show representative best response functions $h(\hat{\alpha}_{t|t-1})$ when policy ignores uncertainty and $h(\hat{\alpha}_{t|t-1}; P_{t|t-1})$ when policy reacts to uncertainty. From a theoretical perspective, the direct effect of uncertainty depends on the value of the precision matrix $P_{t|t-1}$ at the time policy is set. Elements on the leading diagonal are associated with uncertainty surrounding the impact of policy and uncertainty about transition dynamics. 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 reversed if there is sufficient uncertainty about transition dynamics.⁴ Elements lying off the leading diagonal have potentially more complex 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

⁴Craine (1979) shows that very active policy is optimal when uncertainty about transition dynamics is dominant.

high covariances between parameter estimates and a crucial role for uncertainty dynamics in policy. In Table 2 we find that uncertainty makes policy react more positively to the first lag of inflation and more negatively to the second lag of inflation. Optimal policy therefore involves an aggressive yet short-lived response to shocks, and the direct effect of uncertainty is a reduction in the degree of inflation persistence in the model.

Period	$h_{\pi_{t-1}}$	$h_{\pi_{t-2}}$	$h_{\pi_{t-3}}$	$h_{u_{t-2}}$	$h_{u_{t-3}}$
<i>Policy ignores uncertainty</i>					
1965:1	0.885	-0.237	0.288	-0.009	0.005
1975:1	1.299	-0.661	0.272	-0.012	-0.006
1985:1	1.198	-0.480	0.198	-0.010	-0.003
1995:1	1.266	-0.204	-0.142	0.021	0.005
<i>Policy reacts to uncertainty (direct effect only)</i>					
1965:1	1.350	-1.074	0.696	-0.029	0.012
1975:1	1.865	-2.639	1.649	-0.098	-0.034
1985:1	1.751	-2.288	1.457	-0.090	-0.021
1995:1	1.770	-1.661	0.086	-0.027	-0.028
<i>Policy reacts to uncertainty (direct and indirect effects)</i>					
1965:1	1.047	-0.061	0.005	0.006	-0.002
1975:1	1.641	-0.036	-0.592	0.053	0.046
1985:1	1.437	-0.251	-0.198	0.023	0.006
1995:1	1.354	-0.198	-0.170	0.018	0.004

Table 2: Representative best response functions

The indirect effect of uncertainty depends on the extent to which estimates of the initial precision matrix $P_{1|0}$ and the perceived drifting coefficients variance-covariance matrix V differ once policy reacts to uncertainty. Differing parameter estimates imply that the Federal Reserve interprets data differently and so forms a different view about the coefficients in its approximating model. This has obvious consequences for policy and the persistence of inflation. In our case, allowing policy to react to uncertainty has a large effect on the estimated value of V but not on the estimated value of $P_{1|0}$. Tables 3 and 4 give the full numerical estimation results.

$P_{1 0}$					
<i>Policy ignores uncertainty</i>					
0.1087	0.1432	0.0225	-0.2540	-0.0093	-0.1015
0.1432	0.1937	0.0296	-0.3398	-0.0119	-0.1359
0.0225	0.0296	0.0047	-0.0526	-0.0019	-0.0211
-0.2540	-0.3398	-0.0526	0.5990	0.0213	0.2396
-0.0093	-0.0119	-0.0019	0.0213	0.0008	0.0085
-0.1015	-0.1359	-0.0211	0.2396	0.0085	0.0959
<i>Policy reacts to uncertainty</i>					
0.1002	0.0932	0.0147	-0.1905	-0.0026	-0.0972
0.0932	0.0868	0.0137	-0.1772	-0.0025	-0.0903
0.0147	0.0137	0.0022	-0.0280	-0.0004	-0.0143
-0.1905	-0.1772	-0.0280	0.3621	0.0050	0.1847
-0.0026	-0.0025	-0.0004	0.0050	0.0001	0.0025
-0.0972	-0.0903	-0.0143	0.1847	0.0025	0.0943

Table 3: Estimates of initial precision matrix $P_{1|0}$

V

Policy ignores uncertainty

$$\begin{pmatrix} 0.0823 & -0.0778 & 0.0092 & 0.0498 & -0.0081 & -0.4141 \\ -0.0778 & 0.0814 & 0.0003 & -0.0509 & 0.0194 & 0.6859 \\ 0.0092 & 0.0003 & 0.0299 & 0.0012 & 0.0370 & 0.7207 \\ 0.0498 & -0.0509 & 0.0012 & 0.0320 & -0.0105 & -0.3996 \\ -0.0081 & 0.0194 & 0.0370 & -0.0105 & 0.0514 & 1.0064 \\ -0.4141 & 0.6859 & 0.7207 & -0.3996 & 1.0064 & 25.8831 \end{pmatrix}$$

Policy reacts to uncertainty

$$\begin{pmatrix} 0.2384 & -0.2305 & 0.0999 & 0.1176 & -0.0400 & -0.6626 \\ -0.2305 & 0.2301 & -0.0917 & -0.1292 & 0.0407 & 1.1812 \\ 0.0999 & -0.0917 & 0.0666 & 0.0149 & -0.0073 & 1.1595 \\ 0.1176 & -0.1292 & 0.0149 & 0.1180 & -0.0334 & -2.6817 \\ -0.0400 & 0.0407 & -0.0073 & -0.0334 & 0.0555 & 0.4893 \\ -0.6626 & 1.1812 & 1.1595 & -2.6817 & 0.4893 & 101.8092 \end{pmatrix}$$

Table 4: Estimates of perceived drifting coefficients variance-covariance matrix V

The most obvious change in Tables 3 and 4 is an increase in the magnitude of nearly all the estimated elements of the V matrix when policy reacts to uncertainty.⁵ The larger values on the leading diagonal imply that internalising uncertainty makes the Federal Reserve perceive coefficients as drifting more, although changes to the off-diagonal elements suggest subtle changes to the precise variance-covariance structure of the drift. If coefficients are perceived to drift more then it is optimal for the Federal Reserve to place more weight on recent data when estimating its approximating model. Indeed, when the precision matrices are translated

⁵The elements of V when policy ignores uncertainty are already large relative to σ_w , as stressed by Sargent, Williams and Zha (2006).

into the Kalman gain of learning equation (4) we find that the Federal Reserve generally places more weight on recent forecast errors when policy is assumed to react to uncertainty. Such a mechanism in isolation would lead to greater volatility in the estimated coefficients of the Federal Reserve’s approximating model. However, this is not the only channel through which the V matrix affects coefficient volatility in learning models. Instead, there is an additional effect because the Federal Reserve is estimating a drifting coefficient process that is misspecified vis-à-vis the true data generating process. The degree to which the misspecification matters depends directly on the value of V (relative to σ_w). At one extreme if $V = 0$ the misspecified model allows no time-variation in the coefficient estimates, whereas at the other extreme $\sigma_w = 0$ attributes all variation in inflation-unemployment dynamics to changes in the drifting coefficients. If the true data generating process is time-invariant then both these extremes are likely to track the data reasonably well. The time-invariant case naturally performs well providing the initial parameter estimates $\hat{\alpha}_{1|0}$ are sensible, but attributing all variation to drift in coefficients also does well because it effectively fits stochastic trends to the data.⁶ Intermediate values of V are likely to be more problematic. In our case the larger elements of V reduce misspecification and lead to an improvement in the ability of the approximating model to track the data. This manifests itself in the Federal Reserve making smaller forecast errors when policy is allowed to react to uncertainty, which in turn reduces the volatility in the estimated coefficients of the Federal Reserve’s approximating model.

The two channels by which V affects coefficient volatility operate in different directions. A “larger” V creates volatility by inducing the Federal Reserve to place more weight on recent forecast errors, but removes volatility by reducing the magnitude of forecast errors due to misspecification of the approximating model. In our estimations the second effect dominates and coefficient estimates become more stable once policy reacts to uncertainty. In effect, the Federal Reserve puts slightly more weight on much smaller forecast errors so overall there is greater stability. The reduced volatility is apparent in Figure 3, which shows how the Federal Reserve’s view of the Phillips curve is estimated to have evolved assuming policy either ignores

⁶It is well-known that fitting stochastic trends implies “no change” predictors that forecast well over short horizons, see Theil (1971).

or internalises uncertainty. The different evolving views translate into different best response functions in the second and third panels of Table 3, from which it is clear that the indirect effect of policy internalising uncertainty is greater persistence in intended inflation and a rise in the degree of inflation persistence in the model.

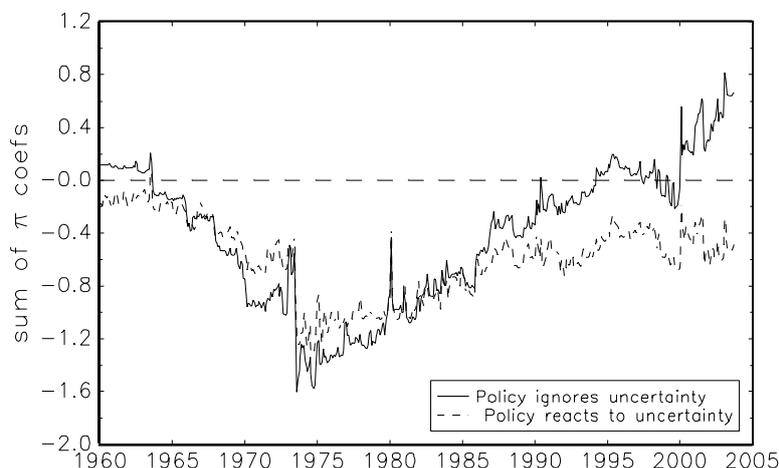


Figure 3: Estimates of Phillips curve trade-off, as perceived by the Federal Reserve

We are now in a position to explain why assuming policy reacts to uncertainty leads to a better fit of the model to the data. Our contention is that the direct and indirect effects of uncertainty combine to give a better match between the inflation persistence of the model and that observed in data. To see why this is the case consider Figure 4, which shows the inflation control errors needed to match intended and actual inflation when policy is assumed to ignore uncertainty. The problem is that here the model has so much persistence that the rapid changes in inflation observed in the 70s and 80s can only be explained by persistent control errors forcing inflation up and down. Indeed, each inflation peak in Figure 4 (marked with a vertical line) is preceded by positive control errors and followed by negative control errors. Fewer such problem arises when policy is allowed to react to uncertainty. In our estimations, the direct effect of uncertainty dominates so the inflation persistence of the model falls and there is less need for persistent control errors to explain inflation dynamics. This is reflected in inflation control errors becoming smaller and less correlated once policy reacts to uncertainty, which is the main reason behind the better fit of the model to the data. It is precisely this fall in the variance of inflation control errors that contributes most to the improvement in the

posterior likelihood in Table 1.

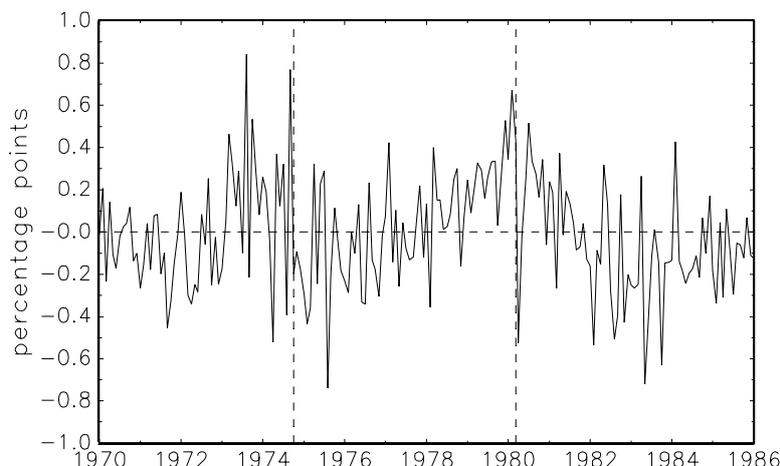


Figure 4: Estimated inflation control errors when policy ignores uncertainty

In passing, we note that the indirect effect of uncertainty partially assuages our second doubt about learning models of the Great Inflation, namely that they imply implausibly poor performance of the Federal Reserve’s internal forecasts of unemployment. Internal forecasting improves when policy reacts to uncertainty, as the Federal Reserve’s model becomes less misspecified and the estimated coefficients within it become less volatile. Our hunch that there is an intimate link between poor forecasting performance and the assumption of policy ignoring uncertainty appears to be correct. We say the doubt is only partially assuaged because the Federal Reserve still makes substantial internal forecast errors when policy reacts to uncertainty. In our estimations, the standard deviation of their one-step-ahead unemployment forecast errors only falls from 3.1 to 2.6 percentage points.

5 Conclusions

At the beginning of this paper we expressed concerns about the credibility of existing learning explanations of the Great Inflation. Our first doubt was that the good empirical fit obtained by Sargent, Williams and Zha (2006) may be contingent on their assumption that the Federal Reserve completely ignores uncertainty when setting policy. This doubt proved to be groundless, since allowing the Federal Reserve to react to uncertainty actually improves the ability

of the learning model to explain the Great Inflation. Our second doubt was that behind the good fit of Sargent, Williams and Zha (2006) appears to be a Federal Reserve that makes unrealistically large errors in its internal forecasts of unemployment. Allowing policy to react to uncertainty reduces misspecification and removes some of the excess variance in forecast errors, but fundamental tensions still exist between the structural parameter values that lead to a good statistical fit of the model and those that allow the Federal Reserve to make small forecast errors. The challenge for future research is to find additional mechanisms that relieve this tension.

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] 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
- [8] 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

- [9] Sack, B., (2000), Does the fed act gradually? A VAR analysis, *Journal of Monetary Economics*, 46(1), 229-256
- [10] Sargent, T.J., (1999), *The Conquest of American Inflation*, Princeton University Press, Princeton: New Jersey
- [11] 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
- [12] 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
- [13] Theil, H., (1971), *Principles of Econometrics*, John Wiley: New York
- [14] Whittle, P., (1990). *Risk-sensitive Optimal Control*, John Wiley & Sons, New York.