

The Return of the Gibson Paradox^{*}

Timothy Cogley[†] Thomas J. Sargent[‡] Paolo Surico[§]

October 2011

Abstract

Before World War I, nominal interest rates were approximately uncorrelated with inflation, a fact known as Gibson's paradox. This correlation increased after World War II, however, and the paradox vanished during the Great Inflation of the 1970s. By estimating vector autoregressions with drifting parameters and stochastic volatility, we show that the statistical association between inflation and nominal interest rates declined in the U.S. in the late 1980s and that Gibson's paradox reappeared after 1995. We estimate a new Keynesian DSGE model for two subsamples – the Great Inflation and the period after 1995 – to identify structural changes that contributed to its reappearance. Counterfactual experiments point to two features: a more anti-inflationary monetary-policy rule and a decline in the extent of price indexation to past inflation. Changes in these features account for the return of the Gibson paradox.

JEL CLASSIFICATION: E4, E5, N1

KEY WORDS: Gibson paradox, inflation persistence, policy regimes

^{*}For comments and suggestions, we are grateful to seminar participants at Cornell, CREI-UPF, the IMF, Princeton, and Wharton. The graphs in this paper are best viewed in color.

[†]New York University. Email: tim.cogley@nyu.edu.

[‡]New York University and Hoover Institution. Email: ts43@nyu.edu.

[§]London Business School and CEPR. Email: psurico@london.edu.

1 Introduction

Keynes (1930) interpreted the observations that nominal interest rates were highly correlated with the aggregate price level but approximately uncorrelated with inflation as contradicting Irving Fisher's equation linking interest rates to *expected* inflation. Keynes called it the Gibson paradox in honor of A.H. Gibson (1923), who Keynes said first detected the pattern. Although those high interest rate-price level correlations long prevailed in data before World War I, they changed afterward. According to Friedman and Schwartz (1982, p. 586),

"The relation holds over neither World War I nor World War II. It is dubious whether it holds for the post-World War II period, particularly since the middle 1960s. For the period our data cover, it holds clearly and unambiguously only for the period from 1880 to 1914, and less clearly for the interwar period."

Barsky (1987) corroborates Friedman and Schwartz's findings and also demonstrates that the Gibson paradox had vanished by the early 1970s. Barsky and Summers (1988, p. 535) conclude that "Gibson's paradox is largely, or perhaps solely, a gold standard phenomenon," and they construct a commodity-money model that generates Gibson's correlation.

This paper takes up the theme of a changing Gibson correlation, making two contributions to the literature. First, by estimating a vector autoregression with time-varying parameters and stochastic volatility, we establish that the statistical association between inflation and nominal interest weakened in the U.S. in the late 1980s and that Gibson's paradox re-emerged after 1995. In the appendix, we report strong evidence for a return in U.K. data after the introduction of inflation targeting in 1992. That Gibson's paradox reappeared under fiat monetary regimes indicates that it is not solely or perhaps even largely a gold-standard phenomenon. Not only can it occur under other monetary regimes, but also its recurrence need not require reverting to a commodity-money standard a la Barsky and Summers.¹

Evidently a force transcending both commodity and fiat money regimes is at work. As Barsky (1987) emphasizes, the critical accompanying feature is the degree of inflation persistence. Gibson's paradox emerged during periods when inflation was weakly persistent, as under the gold standard, and it vanished when inflation

¹We do not question their explanation for the period before World War I. We simply note that their commodity-money mechanism is no longer operative. Of course, an inflation target regime attempts to gather all of the benefits of a gold standard with fewer of its costs.

became strongly persistent, as in the 1970s. Consistent with Barsky’s analysis, we report evidence that U.S. inflation became less persistent in the years leading up to its return.²

Sargent (1973) emphasized that inflation-nominal interest correlations are general-equilibrium outcomes that depend on all features of a macroeconomic model. He criticized then prevalent regression tests of the Fisher equation and instead analyzed the problem in the context of an IS-LM-AS model with rational expectations.³ In many respects, our paper is an updated version of Sargent (1973). We also want to understand the structural features that contribute to inflation persistence and therefore to the breakdown and revival of the Gibson paradox. To that end, we study a standard version of a dynamic new Keynesian DSGE model that includes a variety of shocks, sticky prices as well as intrinsic persistence in both inflation and consumption as exemplified, for instance, by indexation to past inflation and habit formation in households’ preferences respectively. We estimate the model over two subsamples, one for the Great Inflation, when the Gibson paradox was clearly absent, and another for the period after 1995, when it came back. We use the fitted DSGE models to conduct counterfactual experiments designed to isolate the causes of its return.

Among other things, we find that neither a decline in the variance of the shocks (“good luck”) nor a more aggressive policy response to inflation (“good policy”) completely accounts for the return of the Gibson paradox. Changes in the variances of shocks matter little in this context. Gibson’s paradox would still have reappeared in the later part of the sample had the economy been subjected to shocks like those of the 1970s, and it would still have been absent during the Great Inflation had the economy been hit by shocks like those after 1995. Similarly, changes in monetary policy rule parameters are only partially successful in explaining the return of Gibson’s paradox. The Gibson correlation would have fallen significantly in the 1970s had the Fed followed the policy of Volcker and Greenspan, but not to the level observed after 1995, and inflation persistence would have remained too high. Furthermore, the decline in the statistical link between nominal interest rate and inflation would still have reappeared after 1995, though to a lesser extent, had the Fed continued following

²See also Cogley, Primiceri, and Sargent (2010).

³In Sargent’s (1973) model, the validity of Fisher’s theory is closely linked to the proposition that real variables are invariant to the systematic component of monetary policy. In that context, the simplest way to test Fisher’s hypothesis is to test the neutrality proposition. Our model severs that linkage, however. Although the Fisher equation holds by design, systematic monetary policy affects the real interest rate.

the policy rule of the 1970s. It follows that neither changes in shock variances nor the adoption of a new monetary-policy rule fully account for the facts.

The single most important change turns out to be a decline in the backward-looking term of the new Keynesian Phillips curve, whose estimated value falls from 0.86 for the period 1968.Q1-1983.Q4 to 0.13 for the period 1995.Q1-2007.Q4.⁴ Furthermore, this single change goes a long way, though not all the way, toward accounting for both facts. Whether this represents a structural change in price-setting behavior or is itself a consequence of a more anti-inflationary policy stance is difficult to say because the relationship between NKPC parameters and monetary-policy coefficients is typically left unmodeled in the current generation of DSGE models. Our own preferred interpretation is that the backward-looking component is not structural in the sense of being invariant under alterations in monetary-policy rules and that its decline is a consequence of the change in policy. If that is so, then both NKPC and monetary-policy coefficients must be altered in order to assess the effects of a change in policy. When this is done, we are able fully to account for the return of Gibson’s paradox and the decline in inflation persistence.

2 From VAR and DSGE models to spectral densities

We characterize the Gibson paradox in terms of low-frequency comovements between inflation and nominal interest. Let y_t measure an interest rate and z_t measure an inflation rate. Let $\{y_t, z_t\}$ be a mean-zero covariance-stationary random process, and consider the infinite-order least-squares projection of y_t onto past, present, and future values of z_t ,

$$y_t = \sum_{j=-\infty}^{\infty} h_j z_{t-j} + \epsilon_t, \quad (1)$$

where ϵ_t is a random process that satisfies the population orthogonality conditions

$$E\epsilon_t z_{t-j} = 0 \quad \forall j.$$

2.1 Characterization of Gibson paradox

Lucas (1980) used unit slopes of graphs of long two-sided moving averages with geometrically declining weights to characterize the implications of the Fisher equation.

⁴Benigno and Lopez-Salido (2006) and Benati (2008) report international evidence that the backward-looking term in the New-Keynesian Phillips curve is unstable across monetary regimes.

Sargent and Surico (2011) followed Whiteman (1984), who pointed out that a unit slope in limiting versions of Lucas's graphs is equivalent to a unit sum in the regression (1):

$$\sum_{j=-\infty}^{\infty} h_j = 1. \quad (2)$$

Following Lucas, we say that the Fisher theory prevails and the Gibson paradox is absent when (2) holds. We say that a Gibson paradox emerges when $\sum_{j=-\infty}^{\infty} h_j$ is close to zero or negative.

2.2 Frequency domain characterization

Let the spectral densities of y and z be denoted $S_y(\omega)$ and $S_z(\omega)$, respectively, and let the cross-spectral density be denoted $S_{yz}(\omega)$. The Fourier transform of $\{h_j\}$ is

$$\tilde{h}(\omega) = \sum_{j=-\infty}^{\infty} h_j e^{-i\omega j} = \frac{S_{yz}(\omega)}{S_z(\omega)}. \quad (3)$$

Similarly, the sum of the distributed-lag regression coefficients is

$$\tilde{h}(0) = \sum_{j=-\infty}^{\infty} h_j = \frac{S_{yz}(0)}{S_z(0)}. \quad (4)$$

Thus, we can say that the Fisher theory prevails and that the Gibson paradox is absent when $\tilde{h}(0)$ is approximately one. We say that a Gibson paradox emerges when it is close to zero or negative.

We construct estimates of $\tilde{h}(0)$ by estimating vector autoregressions (VARs), and we interpret the results in the context of a log-linear DSGE model. Time-invariant versions of our VAR and DSGE model can both be represented in terms of the state-space system

$$\begin{aligned} X_{t+1} &= AX_t + BW_{t+1}, \\ Y_{t+1} &= CX_t + DW_{t+1}, \end{aligned} \quad (5)$$

where X_t is an $n_X \times 1$ state vector, Y_t is an $n_Y \times 1$ vector of observables, and W_{t+1} is an $n_W \times 1$ Gaussian random vector. We assume that W_{t+1} is identically and independently distributed across time with mean zero and unit covariance matrix. A, B, C, D are conformable matrices, with the absolute values of the eigenvalues of A being bounded strictly above by unity. In our DSGE model, elements of the

matrices A, B, C, D are nonlinear functions of a lower-dimensional vector of structural parameters η .

Suppose that y_t, z_t are two scalar components of Y_t . We seek a mapping from the state-space representation (5) to the sum of projection coefficients $\tilde{h}(0)$. The spectral density matrix for Y is⁵

$$S_Y(\omega) = C(I - Ae^{-i\omega})^{-1}BB'(I - A'e^{i\omega})^{-1}C' + DD'. \quad (6)$$

After extracting the appropriate elements of $S_Y(\omega)$, the sum of projection coefficients $\tilde{h}(0)$ can be computed from formula (4).

The disappearance and re-emergence of the Gibson paradox is connected with changes in inflation persistence (Barsky 1987). As a measure of inflation persistence, we use the first-order autocorrelation ($FACF_\pi$) based on VAR estimates. Similar results, not reported but available upon request, are obtained using the normalized spectrum at frequency zero to characterize inflation persistence.⁶

2.3 Smoothness across ω and inferences about $\tilde{h}(0)$

Sims (1971) shows that in models with sufficiently large numbers of parameters, even arbitrarily large data samples contain no information about features that, like $\tilde{h}(0)$, depend on $S_Y(\omega)$ at a single frequency. This is because models with infinite numbers of parameters have enough freedom to fit arbitrary spectral shapes across frequencies. Then how in this paper can we claim to make sample evidence about $\tilde{h}(0)$ a centerpiece of our story?

The answer is that both our statistical model in section 3 and our theoretical model in section 4 make the matrices A, B, C , and D in the state space representation (5), and therefore also the spectral density matrix $S_Y(\omega)$ in (6), functions of a finite-dimensional vector of parameters $\theta \in \Theta$. Therefore, the spectral density $S_Y(\omega; \theta)$

⁵The spectral density matrix is the Fourier transform of the sequence of autocovariance matrices,

$$S_Y(\omega) = \sum_{j=-\infty}^{\infty} \Gamma_j e^{-i\omega j},$$

where $\Gamma_j = \text{cov}(Y_t, Y_{t-j})$. The autocovariance matrices can be recovered from $S_Y(\omega)$ via the inversion formula

$$\Gamma_j = \frac{1}{2\pi} \int_{-\pi}^{\pi} S_Y(\omega) e^{i\omega j} d\omega.$$

⁶Cogley and Sargent (2005) normalize the spectrum for inflation by dividing by its variance. For details, see section 3.6.2 of their paper.

embodies cross-frequency, and, in the case of the putatively structural models of section 4, the cross-equation restrictions that are the hallmark of rational-expectations models. Variations in θ affect $S_Y(\omega; \theta)$ across *all* frequencies. These cross-frequency restrictions mean that inferences that we draw about single frequency features like $\tilde{h}(0)$ depend on sample information about co-movements of $\{Y_t\}$ at *all* frequencies. A way to express this is that our Bayesian prior over θ induces a prior over outcomes that imposes smoothness across frequencies $\omega \in [-\pi, \pi]$ of frequency domain features like $\tilde{h}(\omega)$.

2.4 Game plan

We proceed as follows. To summarize the patterns in the data, section 3 estimates a purely statistical model in the Sargent-Sims (1977) tradition of ‘not pretending to have an economic theory’. Section 4 then estimates models in a Hansen-Sargent (1980) rational expectations tradition that ‘pretends to have an economic theory.’ We use the section 4 model to interpret patterns revealed by the purely statistical model.

3 The return of the Gibson paradox

In this section, we use an atheoretical statistical model to establish that the Gibson paradox re-emerged in U.S. data after 1995. We fit a VAR with drifting coefficients and stochastic volatility to post-WWII quarterly data for the United States and then construct ‘temporary’ estimates of $\tilde{h}(0)$ that vary over time. A time-varying VAR is useful for summarizing the data because it allows for changes in the dynamics of inflation, money growth, the nominal interest, and output, possibly arising from changes in policy regimes and/or structural instabilities such as changes in shock variances. We want a flexible statistical model at this stage because our sample spans the Bretton Woods era, the Great Inflation, and the Great Moderation. The appendix provides details on data sources and the definitions of variables.

3.1 A time-varying VAR

The statistical model is a VAR(p) with drifting coefficients and stochastic volatility:

$$Y_t = B_{0,t} + B_{1,t}Y_{t-1} + \dots + B_{p,t}Y_{t-p} + \epsilon_t \equiv X_t'\theta_t + \epsilon_t, \quad (7)$$

where X_t' collects the first p lags of Y_t , θ_t is a matrix of time-varying parameters, ϵ_t are shocks to the systematic part of the VAR and Y_t is defined as $Y_t \equiv [\Delta m_t, \pi_t, \Delta y_t, R_t]'$. The operator Δ denotes a first log difference; m_t is the logarithm of a monetary aggregate, $M2$; π_t is the inflation rate, the first difference of the log of the GDP deflator, p_t ; and y_t is real GDP. The short-term nominal interest rate is R_t . Following Cogley and Sargent (2005), we set the lag order $p=2$. The time-varying VAR parameters, collected in the vector θ_t , are postulated to evolve as driftless random walks subject to reflecting barriers that ensure that the autoregressive roots are always nonexplosive (see Cogley and Sargent 2005). When not affected by the reflecting barrier, θ evolves as

$$\theta_t = \theta_{t-1} + \eta_t,$$

where $\eta_t \sim N(0, Q)$.

Following Primiceri (2005), the VAR innovations ϵ_t are postulated to be normally distributed with mean zero and having a time-varying covariance matrix Ω_t that is factored as

$$\text{Var}(\epsilon_t) \equiv \Omega_t = A_t^{-1} H_t (A_t^{-1})' \quad (8)$$

The time-varying matrices H_t and A_t are defined as:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad A_t \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix} \quad (9)$$

with the elements $h_{i,t}$ evolving as geometric random walks:

$$\ln h_{i,t} = \ln h_{i,t-1} + \nu_{i,t} \quad (10)$$

Again following Primiceri (2005), we postulate:

$$\alpha_t = \alpha_{t-1} + \tau_t \quad (11)$$

where $\alpha_t \equiv [\alpha_{21,t}, \alpha_{31,t}, \dots, \alpha_{43,t}]'$, and assume that the vector $[u_t', \eta_t', \tau_t', \nu_t']'$ is distributed as

$$\begin{bmatrix} u_t \\ \eta_t \\ \tau_t \\ \nu_t \end{bmatrix} \sim N(0, V), \text{ with } V = \begin{bmatrix} I_4 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & Z \end{bmatrix} \text{ and } Z = \begin{bmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_2^2 & 0 & 0 \\ 0 & 0 & \sigma_3^2 & 0 \\ 0 & 0 & 0 & \sigma_4^2 \end{bmatrix}, \quad (12)$$

where u_t is such that $\epsilon_t \equiv A_t^{-1} H_t^{\frac{1}{2}} u_t$.

The model (7)-(12) is estimated using Bayesian methods (see Kim and Nelson (2000)). The elements of S are assumed to follow an inverse-Wishart distribution centered at 10^{-3} times the prior mean(s) of the relevant element(s) of the vector α_t with the prior degrees of freedom equal to the minimum allowed. The priors for all the other hyperparameters are borrowed from Cogley and Sargent (2005). We use 80000 Gibbs sampling replications, discard the first 60000 as burn-in, and then retain every tenth one to attenuate the autocorrelation across retained draws. To calibrate the priors for the VAR coefficients, we use a training samples of ten years. Not including the training sample, we use the period 1968Q1-2007Q4 to estimate our model, with the last observation chosen to exclude effects of the financial crisis. Full descriptions of the algorithm, including the Markov-Chain Monte Carlo (MCMC) used to simulate the posterior distribution of the hyperparameters and the states conditional on the data, are provided for instance by Cogley and Sargent (2005) and Primiceri (2005).

3.2 Low-frequency comovements between inflation and the nominal interest rate

To describe the evolution of low-frequency comovements between inflation and nominal interest, we construct a local-to-date t approximation of the sum of projection coefficients,

$$\tilde{h}_{R\pi,t|T}(0) = \frac{S_{R\pi,t|T}(0)}{S_{\pi,t|T}(0)}, \quad (13)$$

using smoothed estimates of the time-varying VAR conditioned on the full sample, $[Y_1, Y_2, \dots, Y_T]$. Temporary versions of $S_{R\pi,t|T}(0)$ and $S_{\pi,t|T}(0)$ are calculated by applying formula (6) to the (t, T) versions of A, B, C, D . Ideally, we would also account for the fact that parameters drift going forward from date t , but this is computationally challenging because it requires integrating a high-dimensional predictive density across all possible paths of future parameters. Adhering to a practice in the learning literature (referred to as ‘anticipated-utility’ by Kreps, 1998), we instead update period-by-period the elements of θ_t, H_t , and A_t and then treat the updated values as if they would remain constant going forward in time.

Estimates of $\tilde{h}_{R\pi,t|T}(0)$ are reported in the top panel of figure 1. The black line portrays the median estimate at date t , and red lines depict central 68% posterior credible sets. Two results are worth emphasizing. First, median estimates vary quite a bit, increasing from values near 1 in the 1970s to more than 2 in the 1980s and then declining to values insignificantly different from zero after 1995. The bottom panel reports a local-to-date- t approximation to the first-order autocorrelation of inflation,

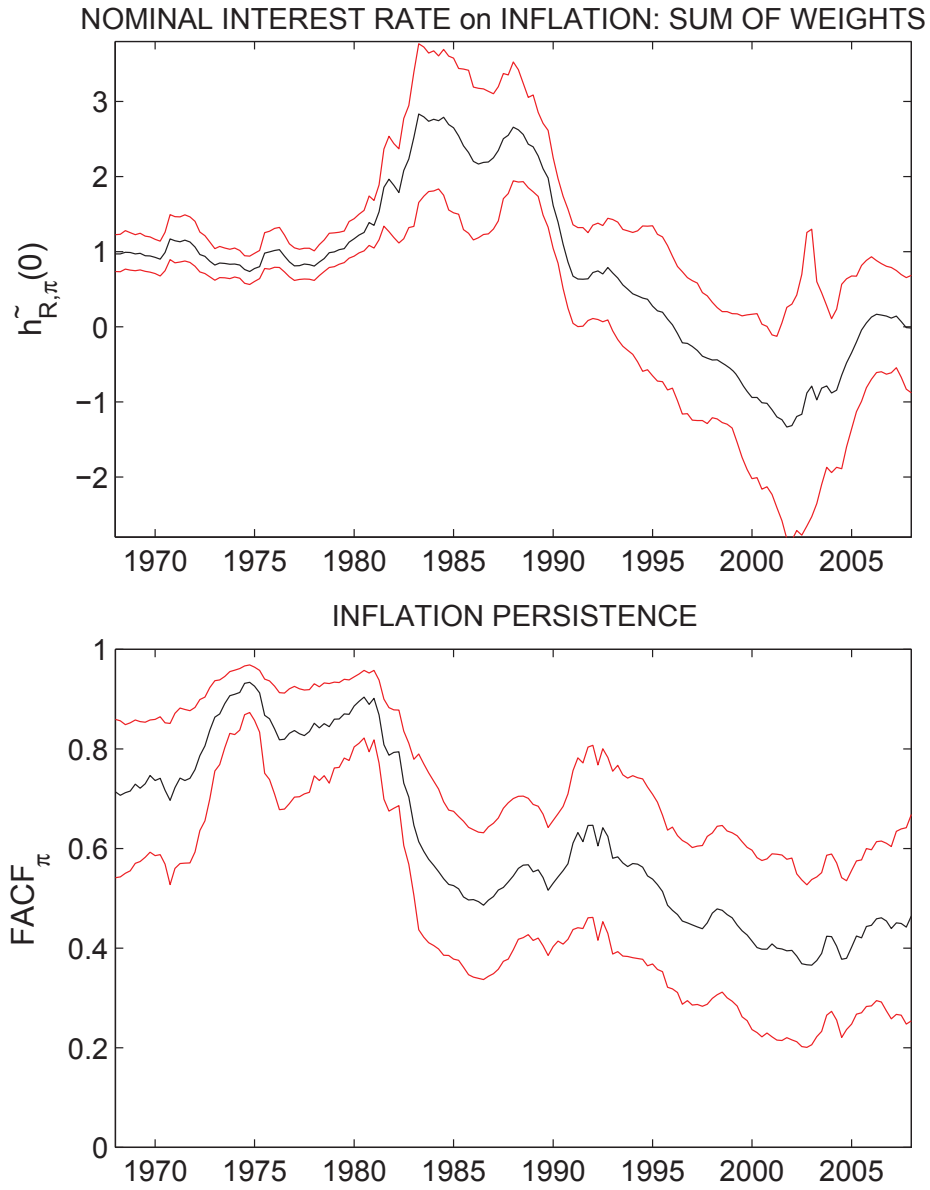


Figure 1: Median and 68% central posterior bands for $\tilde{h}_{R,\pi}(0)$ (top panel) and $FADF_{\pi}$ (bottom panel) based on a VAR with time-varying coefficient and stochastic volatility.

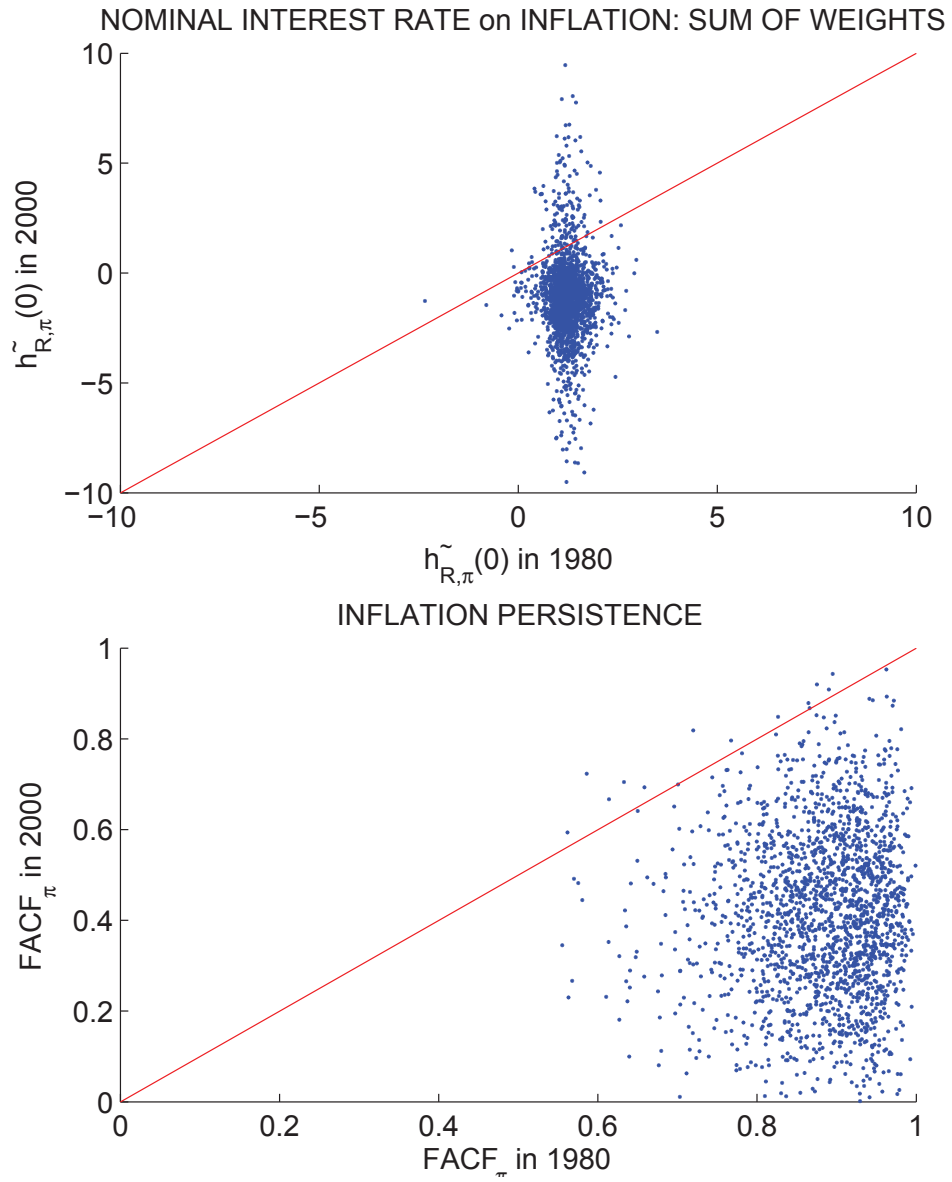


Figure 2: Joint posterior distributions for $\tilde{h}_{R,\pi}(0)$ (top panel) and $FADF_\pi$ (bottom panel), 1980-2000.

$FACF_\pi$. The median estimate of $FACF_\pi$ declines from 0.8-0.9 in the 1970s to around 0.5 after 1995. The timing of the decline in $FACF_\pi$ differs from that of $\tilde{h}_{R\pi,t|T}(0)$, however. The sharpest decline in $FACF_\pi$ occurs in the early 1980s, around the time of the Volcker disinflation, whereas $\tilde{h}_{R\pi,t|T}(0)$ actually *increases* sharply at that time. Indeed, median estimates $\tilde{h}_{R\pi,t|T}(0)$ remain above 2 for most of the 1980s, before falling in the late 1980s and early 1990s and reaching zero around 1995.

Figure 2 assesses the statistical significance of these changes by comparing joint posterior distributions for 1980 and 2000. The top panel portrays the joint distribution for $\tilde{h}_{R\pi,t|T}(0)$, with values for 1980 shown on the x -axis and those for 2000 shown on the y -axis. Points below the 45-degree line therefore represent pairs in which $\tilde{h}_{R\pi,t|T}(0)$ is lower in 2000 than in 1980, while points above the line represent draws in which $\tilde{h}_{R\pi,t|T}(0)$ was higher in 2000. The evidence of a decline is substantial although perhaps not absolutely decisive, with 92.4% of draws lying below the 45-degree line. Similarly, the bottom panel depicts the joint posterior distribution for $FACF_\pi$ in those two years. In this case, 99.3% of pairs have lower values in 2000 and 1980. While 1980 is meant to exemplify a period immediately before a policy change, the choice of this specific year is somewhat arbitrary. It should be noted, however, that the results are similar for any other year among the ten years preceding the beginning of the great moderation.

4 Interpreting the evidence

In this section, we try to understand what caused the return of Gibson’s paradox. Toward that end, we estimate a new Keynesian model over a pair of subsamples, one corresponding to the period of the great inflation (1968Q1-1983Q4) and another to the period after its reappearance (1995Q1-2007Q4). For each subsample, we calculate $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ by applying formulas (4)-(6), thereby verifying that the structural model succeeds in approximating the VAR estimates of $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ for the periods before 1980 and after 1995.⁷

To detect structural changes that can account for the outcomes, we perform a number of counterfactual exercises. For instance, in the first experiment, we consider whether changes in the variances of exogenous shocks can account for its reappearance. In a second experiment, we examine whether changes in the monetary policy

⁷The middle period, when VAR estimates of $\tilde{h}_{R,\pi}(0)$ increased to around 2, remains a bit of a mystery, possibly because the Volcker disinflation might have involved a learning transition that we do not model. In any event, we are mainly interested in the return of Gibson’s paradox, and that did not occur until after 1995.

rule can explain its return. In the literature, these are known as the good-luck and good-policy hypotheses, respectively.⁸ Somewhat to our surprise, we found that neither is enough. Exploring further, we investigate the role of structural changes in private-sector parameters other than those governing shocks. A change in the relative importance of forward- and backward-looking terms in the new Keynesian Phillips curve turns out to be critical. During the later period, when the Gibson paradox reappears, an ‘indexation parameter’ describing the influence of past inflation on current price-setting decisions apparently dropped markedly. Our structural model specifies this parameter as an object that is invariant to alterations in monetary policy. Our empirical results lead us to expect that a better model would interpret that indexation parameter as a mongrel parameter that itself depends nonlinearly on the monetary-policy rule.

4.1 The structure of the economy

Following Ireland (2004) and Rotemberg (1982), we work with a new-Keynesian DSGE model with costly price adjustment, indexation to past inflation, habit formation in households’ preferences, separability between consumption and real money balances, and a unit root in a technology shock process.⁹ After log-linearizing, the model can be represented as follows:

$$\pi_t = \beta(1 - \alpha_\pi) E_t \pi_{t+1} + \beta \alpha_\pi \pi_{t-1} + \kappa x_t - \frac{1}{\tau} e_t, \quad (14)$$

$$x_t = (1 - \alpha_x) E_t x_{t+1} + \alpha_x x_{t-1} - \sigma(R_t - E_t \pi_{t+1}) + \sigma(1 - \xi)(1 - \rho_a) a_t, \quad (15)$$

$$\Delta m_t = \pi_t + z_t + \frac{1}{\sigma\gamma} \Delta x_t - \frac{1}{\gamma} \Delta R_t + \frac{1}{\gamma} (\Delta \chi_t - \Delta a_t), \quad (16)$$

$$\tilde{y}_t = x_t + \xi a_t, \quad \Delta y_t = \tilde{y}_t - \tilde{y}_{t-1} + z_t, \quad (17)$$

where π_t , x_t , Δm_t and R_t are inflation, the output gap, nominal money growth, and the short-term interest rate, respectively. All variables are expressed as log deviations from their steady-state values. The level of detrended output is \tilde{y}_t , and Δy_t refers to output growth. The rate of technological progress is z_t . Equation (14) is an example of a new Keynesian Phillips curve, while (15) is a new Keynesian IS curve. Equation (16) is a money demand equation of a type derived by McCallum and Nelson (1999) and Ireland (2003).

⁸Contributions to this literature include, among others, Lubik and Schorfheide (2004), Primiceri (2005), Sims and Zha (2006), Canova (2009), Canova and Gambetti (2009), Benati and Surico (2009), and Liu, Waggoner and Zha (2011).

⁹See Ireland (2004) for a comprehensive discussion.

The discount factor is β , the parameter α_π may be interpreted as measuring the extent to which prices are indexed to past inflation and α_x captures the extent of habit formation. The coefficients κ and σ are the slope of the Phillips curve and the elasticity of intertemporal substitution, respectively. The parameter τ measures the cost of adjusting prices in Rotemberg's (1982) formulation, while ξ represents the inverse of the labor supply elasticity. The interest elasticity of money demand is $1/\gamma$.

The economy is exposed to four non-policy disturbances: a markup shock e_t , an aggregate demand shock a_t , a money demand shock χ_t , and a technology shock Z_t . The respective shocks evolve as

$$e_t = \rho_e e_{t-1} + \varepsilon_{et}, \text{ with } \varepsilon_{et} \sim N(0, \sigma_e^2), \quad (18)$$

$$a_t = \rho_a a_{t-1} + \varepsilon_{at}, \text{ with } \varepsilon_{at} \sim N(0, \sigma_a^2), \quad (19)$$

$$\chi_t = \rho_\chi \chi_{t-1} + \varepsilon_{\chi t}, \text{ with } \varepsilon_{\chi t} \sim N(0, \sigma_\chi^2), \quad (20)$$

$$\Delta \ln(Z_t) \equiv z_t = \varepsilon_{zt}, \text{ with } \varepsilon_{zt} \sim N(0, \sigma_z^2). \quad (21)$$

As for monetary policy, we consider two types of rules. By appealing to narrative accounts about the conduct of U.S. monetary policy, Sargent and Surico (2011) argue that during the great inflation a money-supply rule represents Federal Reserve behavior better than an interest-rate rule.¹⁰ For the first subsample, we therefore estimate a policy rule that smoothly adjusts money growth in response to movements in inflation and the output gap,

$$\Delta m_t = \rho_m \Delta m_{t-1} + (1 - \rho_m) (\phi_\pi \pi_t + \phi_x x_t) + \varepsilon_{mt}, \quad \varepsilon_{mt} \sim N(0, \sigma_m^2). \quad (22)$$

The parameters ϕ_π and ϕ_x measure long-run responses of money-growth to inflation and the output gap, respectively, while ρ_m is a partial-adjustment parameter.

For the period after 1995, we follow a conventional wisdom and adopt a Taylor rule that smoothly adjusts the short-term nominal interest rate in response to movements in inflation and the output gap,

$$R_t = \rho_r R_{t-1} + (1 - \rho_r) (\psi_\pi \pi_t + \psi_x x_t) + \varepsilon_{Rt}, \quad \varepsilon_{Rt} \sim N(0, \sigma_R^2). \quad (23)$$

The parameters ψ_π and ψ_x measure long-run responses of nominal interest to inflation and the output gap, respectively, while ρ_r is a partial-adjustment parameter.

For each subsample, we use Bayesian simulation methods to approximate the posterior distribution of structural parameters.¹¹ As in the VAR, we specify the vector of observable variables as $[\Delta m_t, \pi_t, R_t, \Delta y_t]$.

¹⁰Galí, Lopez-Salido, and Valles (2003) present empirical evidence consistent with this view.

¹¹See An and Schorfheide 2007 for details.

4.2 The great inflation

Our first subsample spans the period 1968Q1-1983Q4, with the first observation corresponding to the first data point available for VAR estimation and the last corresponding to the end of the Volcker disinflation. Because the Fed’s policy instrument is assumed to be a monetary aggregate, Volcker’s experiment with non-borrowed-reserve targeting is included in this subsample.

Our priors for the model’s structural coefficients are reported in the middle panel of table 1. The slope of the Phillips curve and the elasticity of intertemporal substitution are centered between the low point estimates of Lindé (2005) and Benati and Surico (2009) and the higher estimates of Lubik and Schorfheide (2004). For coefficients governing the degree of forward-looking behavior in the IS and Phillips curves, we adopt weakly-informative beta priors centered on 0.5, thereby putting backward- and forward-looking components on an equal footing a priori. The prior on the discount factor is tight while those on the coefficients governing the cost of price adjustment, labour supply elasticity, and interest rate semi-elasticity of money demand are quite disperse. Our priors on the reaction coefficients in the money supply rule, ϕ_π and ϕ_x , are loosely centered near the posterior means in Sargent and Surico (2011). Finally, we take an agnostic view on the relative importance of structural shocks by adopting identical weakly-informative priors on the persistence and variance of each.

Our priors on structural parameters induce priors on the sum of projection coefficients $\tilde{h}_{R,\pi}(0)$ and first-order autocorrelation $FACF_\pi$. Geweke (2005) recommends prior predictive analysis to articulate how priors on structural parameters affect priors on features of interest.¹² For our model, the latter can be found by sampling from the prior for the model’s structural parameters and calculating the implied values for $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$. The results, which are shown in the last two rows of table 1, attest that the prior on structural parameters implies weakly-informative priors for $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$, with means of 0.664 and 0.867, respectively, and centered 90 percent credible sets spanning the intervals (0.134, 0.960) for $\tilde{h}_{R,\pi}(0)$ and (0.667, 0.973) for $FACF_\pi$. Thus, our prior over the original structural parameters encodes strong views about neither the Gibson paradox nor the degree of inflation persistence.

¹²See Leeper, et al. (2011) for an example.

Table 1: Prior densities and posterior estimates - 1968Q1-1983Q4

description	coefficient	Prior			Posterior		
		density	domain	mean	[5 th ; 95 th]	mean	[5 th ; 95 th]
discount factor	β	beta	[0,1]	0.99	[0.981 ; 0.997]	0.989	[0.980 ; 0.997]
NKPC backward-looking component	α_π	beta	[0,1]	0.5	[0.171 ; 0.826]	0.864	[0.765 ; 0.972]
NKPC slope	κ	gamma	\mathbb{R}^+	0.15	[0.104 ; 0.202]	0.108	[0.070 ; 0.141]
price adjustment cost	τ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	2.706	[1.427 ; 3.987]
IS curve backward-looking component	α_x	beta	[0,1]	0.5	[0.171 ; 0.826]	0.390	[0.335 ; 0.446]
elasticity of intertemporal substitution	σ	gamma	\mathbb{R}^+	0.15	[0.104 ; 0.202]	0.100	[0.069 ; 0.131]
inverse of labour supply elasticity	ξ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	2.123	[1.282 ; 2.923]
interest elasticity of money demand	γ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	2.705	[2.124 ; 3.248]
money growth response to inflation	ϕ_π	normal	\mathbb{R}	0.5	[0.335 ; 0.665]	0.466	[0.305 ; 0.625]
money growth response to output gap	ϕ_x	normal	\mathbb{R}	-5	[-.665 ; -.335]	-5.666	[-.721 ; -.398]
money growth smoothing	ρ_m	beta	[0,1]	0.5	[0.171 ; 0.826]	0.598	[0.478 ; 0.752]
persistence of mark up shock	ρ_e	beta	[0,1]	0.5	[0.171 ; 0.826]	0.136	[0.021 ; 0.246]
persistence of demand shock	ρ_a	beta	[0,1]	0.5	[0.171 ; 0.826]	0.926	[0.872 ; 0.983]
persistence money demand shock	ρ_χ	beta	[0,1]	0.5	[0.171 ; 0.826]	0.656	[0.396 ; 0.910]
standard deviation of mark up shock	σ_e	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0098	[.0050 ; .0146]
standard deviation of demand shock	σ_a	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0045	[.0030 ; .0060]
standard deviation of money demand shock	σ_χ	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0062	[.0032 ; .0093]
standard deviation of technology shock	σ_z	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0063	[.0047 ; .0078]
standard deviation of policy shock	σ_m	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0082	[.0069 ; .0095]
nominal interest rate on inflation: sum of weights	$\tilde{h}_{R,\pi}(0)$	gamma	\mathbb{R}	0.664	[0.134 ; 0.960]	0.788	[0.621 ; 0.924]
inflation persistence	$FACF_\pi$	gamma	\mathbb{R}	0.867	[0.667 ; 0.973]	0.947	[0.911 ; 0.975]

Note: based on 1,000,000 posterior draws using the Metropolis-Hastings algorithm. Fraction of accepted draws: 22%. $FACF_\pi$: first order autocorrelation of inflation.

The last three columns of table 1 report the posterior mean for each parameter along with a centered 90 percent credible set. For many parameters, the posterior mean is not far from the prior mean. This is not surprising because the model is estimated on a relatively short sample. An exception is α_π , which governs the relative importance of forward- and backward-looking terms in the NKPC. Whereas the prior mean and standard deviation for α_π are 0.5 and 0.2, respectively, the posterior mean and standard deviation are 0.86 and 0.06. Because the estimate for α_π is not far below 1, the backward-looking term in the NKPC dominates the forward-looking term. For this subsample, the model exhibits a high degree of intrinsic inflation persistence.

Also contributing to high inflation persistence is the monetary-policy rule. The feedback parameters ϕ_π and ϕ_x are estimated to be 0.47 and -0.57, respectively, and the partial-adjustment parameter ρ_m is 0.60. Since ϕ_π is positive, money growth responds procyclically to inflation. Furthermore, that the Fed responds negatively to the output gap also contributes to this procyclicality, for almost half of the variation in output during this period was due to markup shocks (see table 4 in the appendix). Because a markup shock moves output and inflation in opposite directions, reacting negatively to the output gap entails a conditionally procyclical reaction to inflation (i.e., conditional on a markup shock). Meltzer (2009) describes how monetarists criticized the Fed throughout the 1970s for increasing money growth when inflation was high. Our estimates verify that this was indeed a systematic feature of Fed policy at that time.¹³

Figure 3 illustrates what the model implies about how the monetary policy rule affects inflation persistence and low-frequency comovements with nominal interest. The figure was constructed by freezing all parameters other than the policy-feedback coefficients ϕ_π and ϕ_x at their posterior means and then calculating $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ for various values of the feedback parameters. The figure includes a scatter plot from the posterior distribution for ϕ_π and ϕ_x . The first point to take away is that the model broadly replicates both the high persistence of inflation in the 1970s and the large low-frequency regression coefficient $\tilde{h}(0)$ of R on π (compare figures 1 and 2 with figure 3). Inflation persistence is a bit higher in the structural model than in the VAR and low-frequency comovements are weaker, but both are in the right

¹³That money growth covaries positively with inflation is not a consequence of an unfortunate mix of shocks. Money growth would have remained procyclical with respect to inflation – although less so – had the standard deviation of the demand shocks been 100 times larger than in table 1 and that of markup shocks 100 times smaller. The correlation between money growth and inflation turns negative only when the feedback parameter ϕ_π turns below -0.3, keeping the other coefficients fixed to the posterior means in table 1.

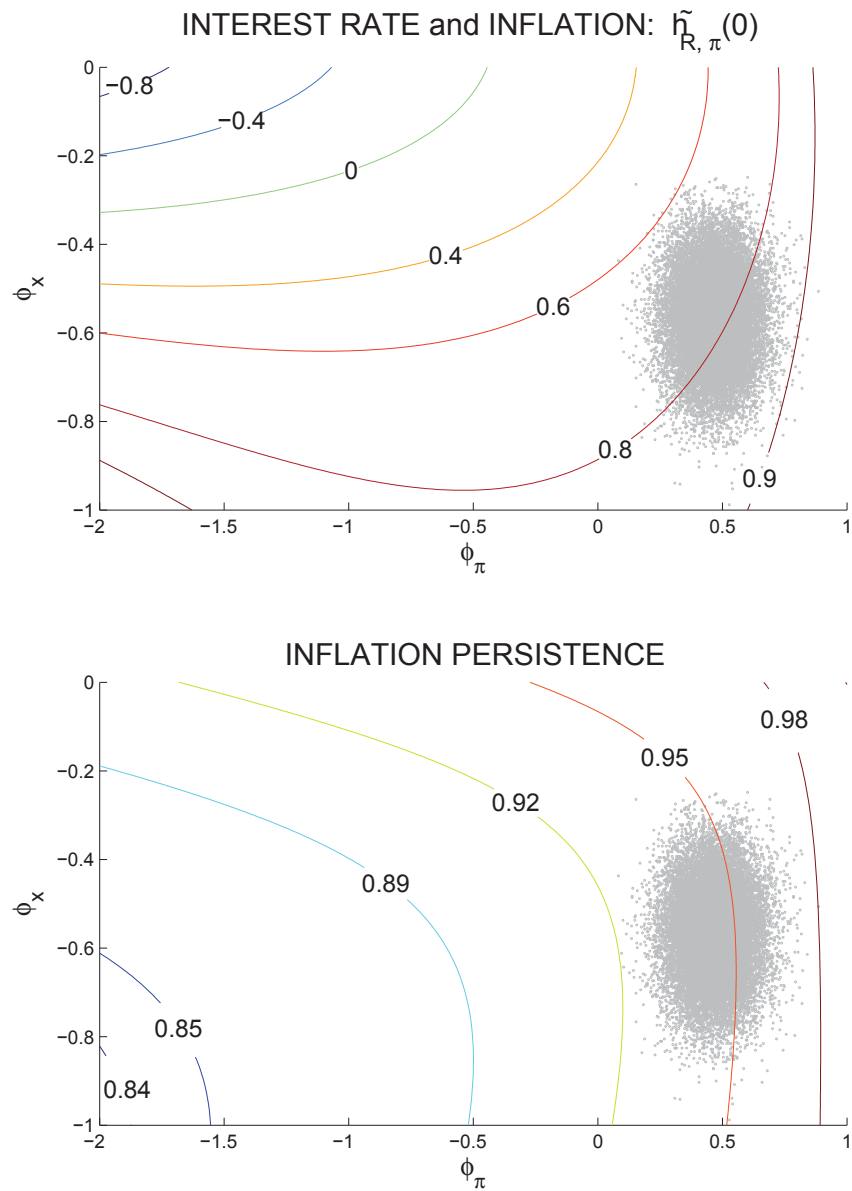


Figure 3: Sums of weights $\tilde{h}(0)$ and first order autocorrelation of inflation in the new-Keynesian model under a money growth rule. The scatter plot represents the joint posterior distribution of the policy responses to inflation and output gap estimated over the 1968Q1-1983Q4 sample. All other parameters are fixed to their posterior mean.

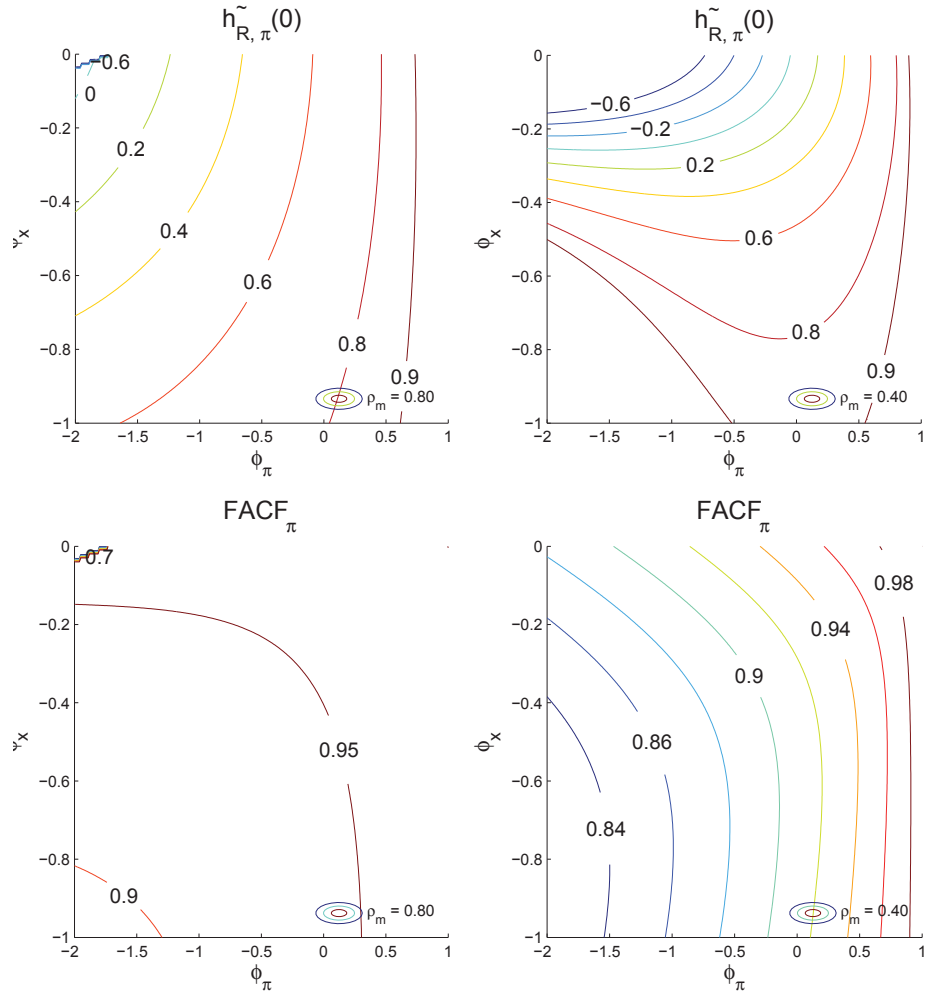


Figure 4: Sums of weights $\tilde{h}(0)$ and first order autocorrelation of inflation in the new-Keynesian model under a money growth rule. Left (right) column fixes $\rho_m=0.80$ ($\rho_m=0.40$). All other parameters are fixed to their posterior mean.

ballpark.

The figure also shows how these statistics would vary in response to changes in policy coefficients with other structural parameters held constant. In particular, the top panel shows that a Gibson paradox would re-emerge if the feedback parameters moved to the northwest of the posterior estimates. For instance, $\tilde{h}_{R\pi}(0)$ would be zero if ϕ_π were -0.75 and ϕ_x were -0.2. That a negative value of ϕ_π reduces $\tilde{h}_{R\pi}(0)$ is intuitive because this represents a more anti-inflationary policy stance. That a less negative value of ϕ_x also reduces $\tilde{h}_{R\pi}(0)$ is perhaps less obvious, but it follows from the fact that markup shocks were an important source of inflation and output-gap variation during the great inflation.

It follows that changes in monetary policy alone could in principle account for a return of the Gibson paradox. However, a second fact to be explained is the decline in inflation persistence, and changes in the coefficients of the money-growth rule matter less in this respect. The bottom panel of figure 3 shows how the first-order autocorrelation $FACF_\pi$ varies as a function of the policy responses to inflation and the output gap. Moving to the northwest of the posterior estimates reduces $FACF_\pi$, but only slightly, and none of the policy combinations shown there approach VAR estimates for the later part of our sample. Furthermore, reducing ρ_m – the partial-adjustment parameter in the money-growth rule – also helps only slightly. As shown in figure 4, $FACF_\pi$ remains about the same as ρ_m declines from 0.6 to 0.4. That more anti-inflationary policies fail significantly to reduce $FACF_\pi$ is due primarily to the high degree of intrinsic inflation persistence, with $\alpha_\pi = 0.86$. At least in this subsample, in their impacts on objects that we use to characterize the Gibson paradox, high intrinsic inflation persistence trumps more anti-inflationary policies.

4.3 After 1995

The second subsample spans the period 1995Q1-2007Q4, with the first observation coinciding with the return of Gibson’s paradox and the last chosen to exclude the financial crisis. Our priors for non-policy parameters – shown in the middle panel of table 2 – are identical to those in table 1. Synchronizing the priors for the two subsamples ensures that changes in estimates of non-policy parameters can be attributed to differences in sample information and not to discrepancies in prior information. In addition, setting the prior for the second subsample equal to that of the first and not to its posterior allows more flexibility for detecting changes in private-sector parameters, such as those governing the shocks.

Table 2: Prior densities and posterior estimates - 1995Q1-2007Q4

description	coefficient	Prior			Posterior		
		density	domain	mean	[5 th ; 95 th]	mean	[5 th ; 95 th]
discount factor	β	beta	[0,1]	0.99	[0.981 ; 0.997]	0.990	[0.982 ; 0.998]
NKPC backward-looking component	α_π	beta	[0,1]	0.5	[0.171 ; 0.826]	0.133	[0.022 ; 0.236]
NKPC slope	κ	gamma	\mathbb{R}^+	0.15	[0.104 ; 0.202]	0.138	[0.092 ; 0.185]
price adjustment cost	τ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	4.009	[2.538 ; 5.472]
IS curve backward-looking component	α_x	beta	[0,1]	0.5	[0.171 ; 0.826]	0.179	[0.068 ; 0.286]
elasticity of intertemporal substitution	σ	gamma	\mathbb{R}^+	0.15	[0.104 ; 0.202]	0.112	[0.073 ; 0.151]
inverse of labour supply elasticity	ξ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	0.972	[0.526 ; 1.393]
interest elasticity of money demand	γ	gamma	\mathbb{R}^+	3	[1.560 ; 4.811]	2.344	[1.658 ; 3.033]
interest rate response to inflation	ψ_π	truncated normal	\mathbb{R}	1.5	[1.010 ; 1.990]	1.653	[1.276 ; 2.040]
interest rate response to output gap	ψ_x	normal	\mathbb{R}	1.25	[0.000 ; 0.250]	0.117	[0.012 ; 0.222]
interest rate smoothing	ρ_r	beta	[0,1]	0.5	[0.171 ; 0.826]	0.838	[0.787 ; 0.888]
persistence of mark up shock	ρ_e	beta	[0,1]	0.5	[0.171 ; 0.826]	0.474	[0.234 ; 0.705]
persistence of demand shock	ρ_a	beta	[0,1]	0.5	[0.171 ; 0.826]	0.921	[0.868 ; 0.979]
persistence money demand shock	ρ_χ	beta	[0,1]	0.5	[0.171 ; 0.826]	0.638	[0.342 ; 0.951]
standard deviation of mark up shock	σ_e	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0050	[.0031 ; .0067]
standard deviation of demand shock	σ_a	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0044	[.0027 ; .0061]
standard deviation of money demand shock	σ_χ	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0052	[.0030 ; .0072]
standard deviation of technology shock	σ_z	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0042	[.0034 ; .0050]
standard deviation of policy shock	σ_m	inv. gamma	\mathbb{R}^+	0.01	[0.004 ; 0.022]	.0021	[.0017 ; .0025]
nominal interest rate on inflation: sum of weights	$\tilde{h}_{R,\pi}(0)$		\mathbb{R}	0.760	[-0.623 ; 1.623]	-0.278	[-1.386 ; 1.160]
inflation persistence	$FACF_\pi$		\mathbb{R}	0.792	[0.509 ; 0.960]	0.585	[0.415 ; 0.732]

Note: based on 1,000,000 posterior draws using the Metropolis-Hastings algorithm. Fraction of accepted draws: 20%. $FACF_\pi$: first order autocorrelation of inflation.

Because the Fed switched from a money-growth to an interest-rate rule, our priors on policy coefficients are entirely new. We chose informative priors on Taylor-rule parameters because we anticipated that they would be weakly identified in this subsample, as indeed they are.¹⁴ For the response coefficients on inflation and the output gap, we adopt normal priors centered on values suggested by Taylor (1993). For ψ_x , the prior variance is calibrated so that a centered 95 percent credible set ranges from 0 to 1 when nominal interest is expressed at an annual rate.¹⁵ Similarly, for ψ_π the prior variance is calibrated so that a 95 percent credible set ranges from 1 to 2. However, to enforce the Taylor principle, we truncate the prior for ψ_π at 1.¹⁶ For the interest-smoothing parameter ρ_r , we adopt a weakly informative beta prior centered on 0.5. Finally, in order to remain agnostic about the sources of fluctuations, we adopt the same prior for the policy-shock variance as for the other shocks.

The switch to an interest-rate rule and corresponding change in prior alter the implied priors for $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$. As shown in the last two rows of table 2, the prior mean for $\tilde{h}_{R,\pi}(0)$ rises from 0.66 for the first subsample to 0.76 for the second, but the prior remains uninformative, with a 90 percent prior credible set of (-0.62, 1.62). In any case the direction of change is the opposite of what we want to explain. The prior for $FACF_\pi$ moves slightly in the desired direction, however, with the mean falling from 0.87 to 0.79 and a 90 percent credible set covering the interval (0.51, 0.96). Despite that, the prior on $FACF_\pi$ remains sufficiently weakly informative that a decline in inflation persistence is not hardwired into the model.

The last three columns of table 2 report posterior means and centered 90 percent credible sets for the structural parameters. As expected, most of the shocks become less volatile. The unconditional standard deviation of markup shocks falls by 42 percent, with a decline in its innovation variance more than offsetting an increase in its persistence. The unconditional standard deviation of technology and money-demand shocks decline by 33 and 17 percent, respectively, while that of demand shocks remains about the same. In terms of relative importance, it follows that the markup shocks that bedeviled the Fed during the great inflation became less severe after 1995 and were replaced to a great extent by easier-to-manage demand, technology, and money-demand shocks. A variance decomposition, reported in the appendix, confirms that markup shocks were less important as a source of output variation

¹⁴Weak identification follows from that fact that the sample is short and that there was relatively little variation in inflation and output gaps (Mavroeidis 2010).

¹⁵Here nominal interest is expressed at a quarterly rate, hence the need to divide ψ_x by 4.

¹⁶This truncation is handled automatically within a Metropolis-Hastings algorithm because it implies an acceptance probability of zero for ψ_π proposals falling below the lower bound.

after 1995, although they remained the dominant source of inflation variation.¹⁷

A second difference relative to the great-inflation sample concerns monetary policy. Alas, the response coefficients to inflation and the output gap are weakly identified, and so their posteriors are similar to the priors. Our priors follow a conventional wisdom, however, by assuming that the policy rule satisfies the Taylor principle and responds more strongly to inflation than to the output gap. On the other hand, the interest-smoothing parameter ρ_r is well identified and precisely estimated around 0.84, implying a high degree of interest smoothing. All these features are widely considered to be desirable for improving outcomes for inflation and output in new Keynesian models.

A third difference concerns the relative importance of forward- and backward-looking components in the NKPC (see Benati 2008) and to a lesser extent in the IS curve. In particular, the posterior mean of α_π – the ‘indexation’ parameter in the NKPC – dropped from 0.86 to 0.13 and that of α_x – the ‘habit-formation’ parameter in the IS curve – fell from 0.39 to 0.18. The estimated degrees of intrinsic inflation and output persistence are therefore substantially lower than during the great inflation. The decline in α_π is especially important for the counterfactual experiments reported below.

The last two lines of table 2 record the model’s implications for $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$, and figure 5 reports additional details. Although posteriors for the period after 1995 are diffuse, especially for $\tilde{h}_{R,\pi}(0)$, plausible parameterizations can be found for which the model succeeds in approximating changes in $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$. In particular, the posterior mean for $\tilde{h}_{R,\pi}(0)$ declines from 0.79 in the first subsample to -0.28 in the second, while that of $FACF_\pi$ falls from 0.95 to 0.59. Both are roughly in line with VAR estimates from the period before 1980 and after 1995.

Figure 6 is the counterpart for the post-1995 period of calculations reported in figure 3 for the great inflation. As before, the figure portrays the model’s implications for $\tilde{h}_{R,\pi}(0)$ (top panel) and $FACF_\pi$ (bottom panel) as functions of the monetary-policy parameters ψ_π and ψ_x , with all other parameters being frozen at their posterior means. Once again, scatterplots depict the posterior sample for the Taylor-rule parameters. When the policy parameters are also set at their posterior means, the model produces values of $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ that are not far from those in figure 1 for the period after 1995. Notice, however, that as the reaction coefficients approach the boundary where the Taylor principal is violated, inflation persistence increases,

¹⁷To be precise, markup shocks accounted for a higher share of a much-reduced total inflation variance.

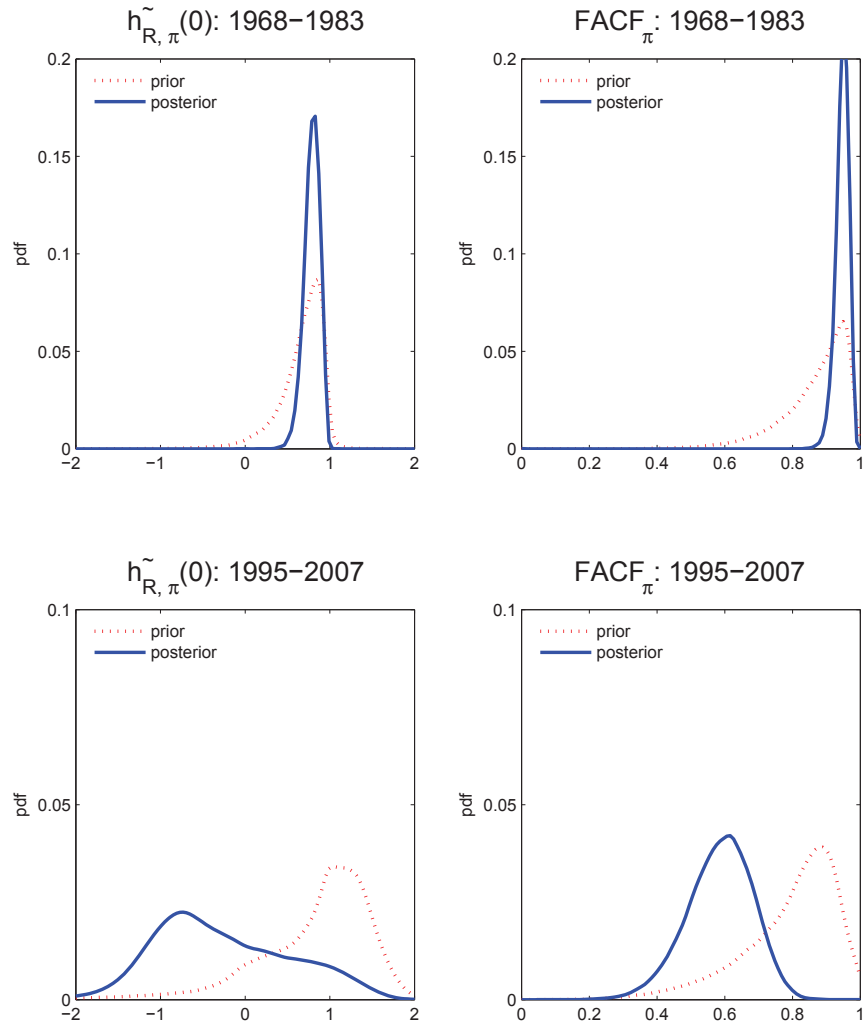


Figure 5: Prior and posterior probability density function for the long-run link between interest rate and inflation, $\tilde{h}(0)$, and first order autocorrelation of inflation, $FACF_{\pi}$, in the new-Keynesian model implied by the posterior estimates in tables 1 and 2.

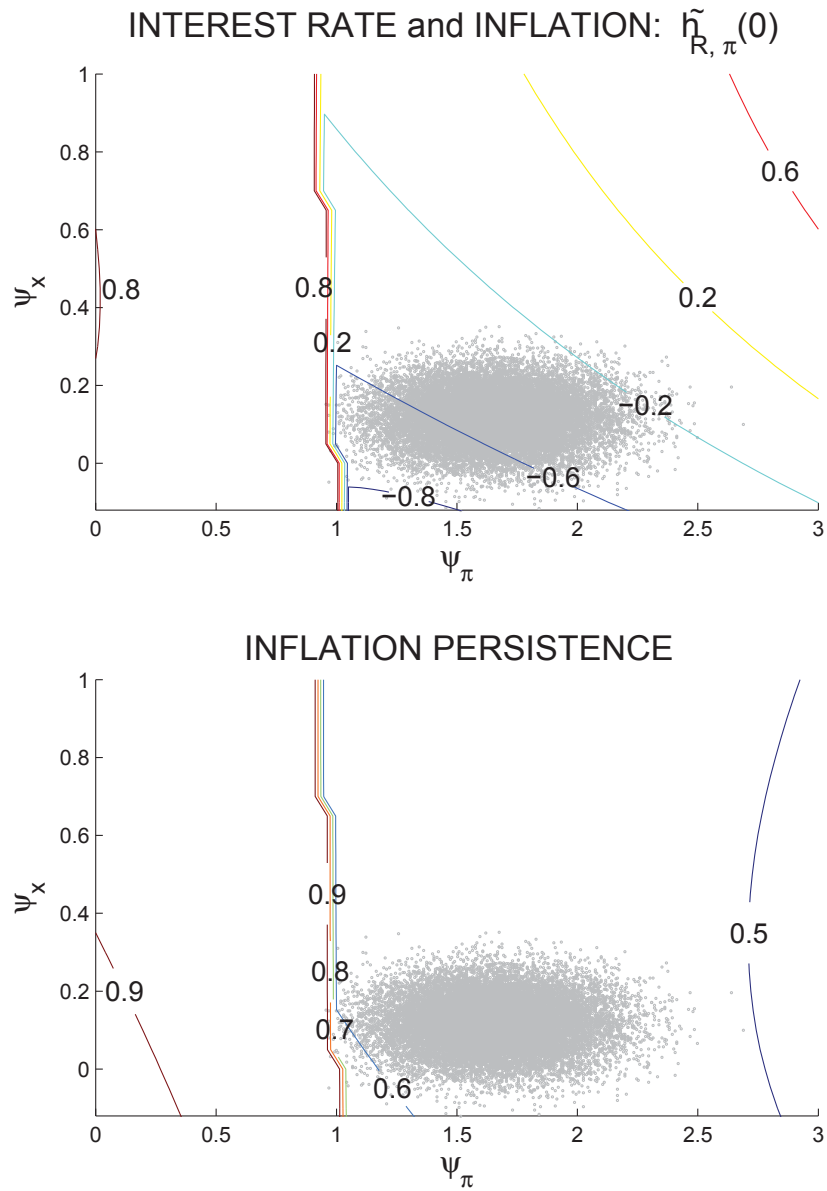


Figure 6: Sums of weights $\tilde{h}(0)$ and first order autocorrelation of inflation in the new-Keynesian model under an interest-rate rule. The scatter plot represents the joint posterior distribution of the policy responses to inflation and output gap estimated over the 1995Q1-2007Q4 sample. All other parameters are fixed to their posterior mean.

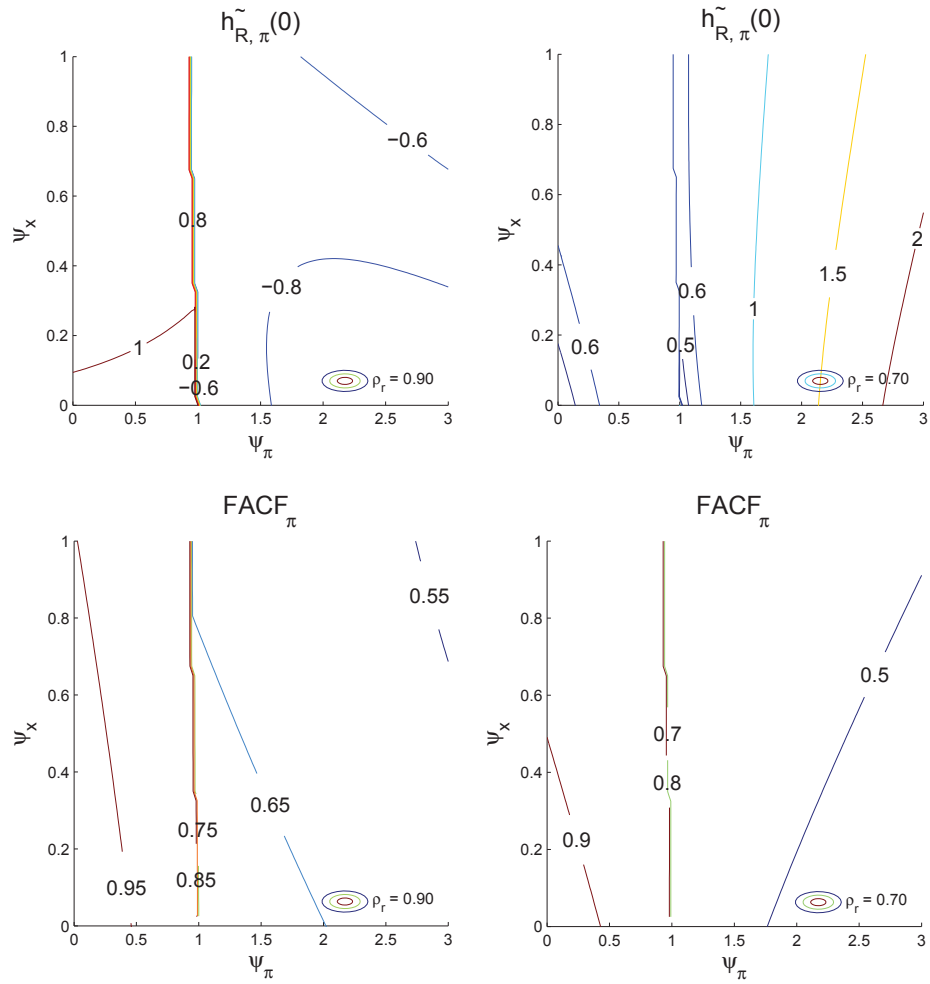


Figure 7: Sums of weights $\tilde{h}(0)$ and first order autocorrelation of inflation in the new-Keynesian model under an interest-rate rule. Left (right) column fixes $\rho_r=0.90$ ($\rho_r=0.70$). All other parameters are fixed to their posterior mean.

and the Gibson paradox vanishes.

Figure 7 portrays the same information for alternative values of the interest-smoothing parameter ρ_r . A high degree of interest-rate smoothing also contributes to the reappearance of the Gibson paradox. For instance, when $\rho_r = 0.7$, $\tilde{h}_{R,\pi}(0)$ remains above 1 for all policies for which ψ_π exceeds 1.5, and it fails to approach zero for any combination of ψ_π and ψ_x depicted there (see the top-right panel). As ρ_r increases, $\tilde{h}_{R,\pi}(0)$ declines, reaching plausible levels for $\rho_r = 0.84$ and becoming even more negative when $\rho_r = 0.9$ (see the top-left panel).

On the other hand, changes in ρ_r have little effect on inflation persistence. As ρ_r varies from 0.7 to 0.9, the autocorrelation measures in the bottom rows of figures 6 and 7 remain about the same as functions of ψ_π and ψ_x . A high degree of inflation persistence emerges only when the policy coefficients approach the boundary where the Taylor principle is violated. That the location of that boundary depends more on the long-run policy responses ψ_π and ψ_x than on the degree of interest smoothing explains why ρ_r has relatively little influence on $FACF_\pi$.

Finally, notice the long right tail that emerges in the post-1995 posterior distribution for $\tilde{h}_{R,\pi}(0)$ (see the lower left panel of figure 5). In appendix C, we summarize the configuration of structural parameters associated with right-tail draws of $\tilde{h}_{R,\pi}(0)$. The main difference between the censored posteriors shown there and the unconditional posterior distribution in table 2 concerns the persistence of markup shocks. While the unconditional posterior mean of ρ_e is 0.474, the average of draws associated with the right tail of $\tilde{h}_{R,\pi}(0)$ is 0.654. In other respects, the structural parameters in appendix C are not so different from those in table 2. It follows that Gibson's paradox would not have reappeared after 1995 had markup shocks been quite a bit more persistent than our posterior mean estimate. Notice also that the required degree of persistence is higher than in the first subsample, for which we estimate $\rho_e = 0.136$. The point is not that markup shocks like those from the earlier period would have prevented a return of Gibson's paradox, but that shocks more persistent than the means of our estimates from either subsample are needed. In other words, this is not a version of the good-luck hypotheses. Bad luck – in the form of an even larger increase in markup-shock persistence – would be needed to offset other forces contributing to a decline in $\tilde{h}_{R,\pi}(0)$.

4.4 Counterfactual experiments

In this section, we examine a number of counterfactual scenarios in order to pinpoint what caused the return of the Gibson paradox and decline in inflation persistence. The structural estimates in tables 1 and 2 differ in several respects, and we want to know which of these differences contribute most to the reemergence of Gibson’s paradox. The first row of table 3 reports the statistics for the baseline model, which are obtained by fixing the parameters of the model to the posterior means in tables 1 and 2.¹⁸

We begin by assessing a version of the good-luck hypothesis, viz. that the return of the Gibson paradox and decline in inflation persistence are due to changes in the innovation variances of shocks. We examine this hypothesis from two angles, first by asking what would have happened during the great inflation had the economy been driven by the low-variance shocks of the post-1995 period, and secondly by turning the question around and asking what would have happened after 1995 had the economy been driven by the high-variance shocks of the great inflation. Just to be clear, for each subsample we change only the shock variances. All other parameters are frozen at the sample-specific posterior means reported in tables 1 and 2.

The results are recorded in the second row of table 3. Replacing the shock variances in the great-inflation model with those of the post-1995 period alters the two statistics only slightly, with $\tilde{h}_{R,\pi}(0)$ rising from 0.79 in the baseline model to 0.83 in the counterfactual model and $FACF_\pi$ falling from 0.94 to 0.93. Similarly, replacing the shock variances in the post-1995 model with those of the great inflation makes only a slight difference. The persistence measure $FACF_\pi$ rises, but only slightly, from 0.55 to 0.63. The comovement statistic $\tilde{h}_{R,\pi}(0)$ moves in the wrong direction, falling from -0.51 in the baseline model to -1.31 in the counterfactual model. Thus, for both settings of shock variances, a Gibson paradox emerges in the second subsample and not in the first, while inflation is strongly autocorrelated in the first subsample and not in the second. It follows that changes in the shock variances explain neither the reappearance of Gibson’s paradox nor the decline in inflation persistence.¹⁹

¹⁸The entries of the baseline model for the post-1995 sample in table 3 differ slightly from the corresponding posterior means in table 2. The reason for this is that the latter are computed using the distributions of $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ obtained by drawing the parameters of the model from their posterior distributions. As shown in figure 5, the posterior distributions of the long-run statistics are not symmetric over the post-1995 sample.

¹⁹We also examine an alternative version of the ‘good-luck’ hypothesis in which we swap both the innovation variances and autoregressive coefficients for the shocks. For this version, $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ would have been 0.93 and 0.96, respectively, for the Great inflation and -1.78 and 0.64 for the post-1995 period. Hence there is still little support for the notion that changes in shocks can

Table 3: Counterfactual Scenarios

	<i>Great Inflation</i>		<i>Post-1995</i>	
	$\tilde{h}_{R,\pi}(0)$	$FACF_{\pi}$	$\tilde{h}_{R,\pi}(0)$	$FACF_{\pi}$
Baseline Model	0.79	0.94	-0.51	0.55
Shock Variances	0.83	0.93	-1.31	0.63
Policy	0.05	0.92	0.24	0.63
NKPC	0.38	0.74	1.01	0.89
NKPC plus Policy	-0.67	0.68	0.81	0.94

Note: The great-inflation and post-1995 samples span the periods 1968Q1-1983Q4 and 1995Q1-2007Q4, respectively. The baseline models for each subsample are calibrated at the respective posterior means reported in tables 1 and 2. For the shock variances hypothesis, the counterfactual models replace the shock variances in the baseline model with those from the other subsample, holding all other parameters constant. For the policy hypothesis, the counterfactual models replace the monetary-policy rule in the baseline model with that of the other subsample, holding all other parameters constant. For the NKPC hypothesis, the counterfactual models replace the NKPC parameters in the baseline model with those of the other subsample, holding all other parameters constant. For the row labeled 'NKPC plus Policy', the counterfactual models replace both the NKPC parameters and the monetary-policy rule in the baseline model with those of the other subsample, holding all other parameters constant.

Next we examine the role of changes in monetary policy. The third row of table 3 reports the results of counterfactual calculations in which the interest-rate rule from the post-1995 model replaces the money-growth rule in the great-inflation model and vice versa, with all non-policy parameters frozen at the levels in the respective baseline models. The results show that changes in monetary policy alone also fail fully to account for the re-emergence of the Gibson paradox. For instance, when we substitute the post-1995 policy into the great-inflation model, $\tilde{h}_{R,\pi}(0)$ falls from 0.79 to 0.05, and $FACF_{\pi}$ remains about the same. This hypothesis is partially successful: a Gibson paradox now emerges in the first subsample, but not quite to the same extent as that which emerged later in the data, and inflation persistence remains too high. Similarly, when we substitute the great-inflation policy into the post-1995 model, $\tilde{h}_{R,\pi}(0)$ rises from -0.51 to 0.24, and $FACF_{\pi}$ increases from 0.55 to 0.63. A partial Gibson paradox would still have reappeared in the second subsample and inflation would have been only slightly more persistent even if monetary policy had not changed.

Since neither the swap of shock variances nor the swap of monetary policy rules explain the return of the Gibson paradox.

fully account for the changes, it must be the case that changes in private-sector parameters other than the shocks also matter. After rounding up the usual suspects, we found that changes in NKPC parameters are especially important. In particular, recall that α_π – the backward-looking term in the NKPC – declined from 0.86 in the first subsample to 0.13 in the second, implying that the Fed faced less intrinsic inflation persistence after 1995. This seems to be the single most important change, and it goes a long way toward accounting for the changes in $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$.

In the fourth row of table 3, we report the results of counterfactuals that swap NKPC parameters across subsamples, holding all other parameters constant.²⁰ For instance, replacing the NKPC parameters in the post-1995 model with those of the great-inflation period causes the Gibson paradox to disappear and makes inflation highly autocorrelated, with $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ increasing from -0.51 and 0.55 to 1.01 and 0.89, respectively. Similarly, replacing the NKPC parameters in the great-inflation model with those from after 1995 causes a Gibson paradox partially to emerge in the first subsample and reduces inflation persistence, with $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ falling from 0.79 and 0.94 to 0.38 and 0.74, respectively. Thus, irrespective of monetary policy or the shocks, a Gibson paradox would have emerged at least partially and inflation would have been moderately persistent when price setters were more forward-looking, and not otherwise.

A critical question, therefore, concerns why the parameter α_π declined. At this level of modeling it is impossible to say because α_π is treated as a primitive. One respectable interpretation, however, is that α_π declined because of the change in policy. Indeed, the significant decline in estimates of α_π after the Volcker disinflation might be taken as *prima facie* evidence that it is not structural in the sense of being invariant to altered government policy functions (see Lucas, 1976). If the decline in α_π is in fact a consequence of a more anti-inflationary policy stance, then both sets of parameters must be swapped in order properly to assess the effects of a change in monetary policy. According to this interpretation, a change in monetary policy operates through two channels, the first being a direct effect coming from changes in the policy rule itself and the second an indirect channel working through changes in the extent of indexation to past inflation.

The results of this joint counterfactual experiment are reported in the fifth row of table 3. This combination explains changes in $\tilde{h}_{R,\pi}(0)$ and $FACF_\pi$ quite well. For example, when we substitute the great-inflation policy and NKPC parameters into

²⁰The results are similar when α_π is the only parameter that changes and all other NKPC parameters are also frozen.

the post-1995 model, $\tilde{h}_{R,\pi}(0)$ rises from -0.51 to 0.81, and $FACF_\pi$ increases from 0.55 to 0.94, matching quite closely the values implied for the baseline great-inflation model. Similarly, when we substitute the post-1995 policy and NKPC parameters into the great-inflation model, $\tilde{h}_{R,\pi}(0)$ falls from 0.79 to -0.67, and $FACF_\pi$ declines from 0.94 to 0.68, well approximating outcomes for the baseline great-moderation model. It follows that the other changes across subsamples are secondary for understanding the return of the Gibson paradox.

5 Conclusion

Our counterfactuals point to a change in monetary policy as being the origin of the return of Gibson’s paradox. To say this, we push the New Keynesian econometric model beyond its usual limits. To make our hypothesis work, we posit unmodeled nonlinearities linking New Keynesian Phillips curve parameters such as the degree of indexation and cost of price adjustment to parameters of the policy rule. Although we think this is economically defensible, we are slightly uncomfortable about manipulating the model in this way because the content of the New Keynesian model is that those parameters really are structural. As such, they are critical for determining the properties of inflation, both directly and indirectly through their influence on the design of monetary policy. From that perspective, those NKPC parameters really *are* the nominal anchor in this model.²¹ Tampering with someone’s nominal anchor is risky.

Here our focus is on historical data analysis, which perhaps places lighter demands on a model because its parameters can simply be re-estimated across policy regimes. But to the extent that key parameters fail to be invariant, we should worry about a model’s reliability for predicting the consequences of policies unseen in the samples used for estimation.

²¹Calling these parameters the nominal anchor is not original with us. Guillermo Calvo said this to Sargent.

References

- Barsky, R.B., 1987, The Fisher Hypothesis and the Forecastability and Persistence of Inflation, *Journal of Monetary Economics* 19, 3-24.
- Barsky, R.B. and L.H. Summers, 1988, Gibson's Paradox and the Gold Standard, *Journal of Political Economy* 96, 528-550.
- Benati, L., 2008, Investigating Inflation Persistence Across Monetary Regimes, *Quarterly Journal of Economics* 123, 1005-1060.
- Benati, L. and P. Surico, 2009, VAR Analysis and the Great Moderation, *American Economic Review* 99, 1636-52.
- Benigno, P. and D. Lopez-Salido, 2006, Inflation Persistence and Optimal Monetary Policy in the Euro Area, *Journal of Money Credit and Banking* 38, 587-614.
- Canova, F., 2009, What Explains the Great Moderation in the U.S.? A Structural Analysis, *Journal of the European Economic Association* 7, 697-721.
- Canova F. and L. Gambetti, 2009, Structural Changes in the U.S. Economy: Is There a Role for Monetary Policy?, *Journal of Economic Dynamics and Control* 33, 477-490.
- Cogley, T., and T. J. Sargent, 2005, Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII US, *Review of Economic Dynamics* 8, 262-302.
- Cogley, T., G.E. Primiceri, and T.J. Sargent, 2010, Inflation-Gap Persistence in the U.S., *American Economic Journal: Macroeconomics* 2, 43-69.
- Friedman, M. and A.J. Schwartz, 1982, *Monetary Trends in the United States and the United Kingdom: Their Relation to Income, Prices, and Interest Rates, 1867-1975*, University of Chicago Press: Chicago IL.
- Galí, J., D. Lopez-Salido, and J. Valles, 2003, Technology Shocks and Monetary Policy: Assessing the Fed's Performance, *Journal of Monetary Economics* 50, 723-743.
- Geweke, J., 2005, *Contemporary Bayesian Econometrics and Statistics*, Wiley-Interscience.
- Gibson, A.H., 1923, The Future Course of High Class Investment Values, *Banker's Magazine* (London) 115, 15-34.

- Hansen, L. P. and Sargent, T. J., 1980, Formulating and estimating dynamic linear rational expectations models, *Journal of Economic Dynamics and Control*, 2(1), 7-46.
- Ireland, P., 2004, Technology Shocks in the New-Keynesian Model, *Review of Economics and Statistics* 86, 923-936.
- Keynes, J.M., 1930, *A Treatise on Money*, Macmillan: London.
- Kim, C.J. and C.R. Nelson, 2000, *State-Space Models with Regime Switching*, MIT Press: Cambridge MA.
- Kreps, D., 1998. Anticipated utility and dynamic choice. In: Jacobs, D.P., Kalai, E., Kamien, M. (Eds.), *Frontiers of Research in Economic Theory*. Cambridge University Press.
- Leeper, E., N. Traum, and T.B. Walker, 2011, The Fiscal Multiplier Morass: A Bayesian Perspective, unpublished working paper.
- Lindé, J., 2005, Estimating New-Keynesian Phillips curves: A full information maximum likelihood approach, *Journal of Monetary Economics* 52, 1135-1149.
- Liu, Z., D.F. Waggoner and T. Zha, 2011, Sources of Macroeconomic Fluctuations: a Regime-Switching DSGE Approach, *Quantitative Economics*, forthcoming.
- Lubik, T., and F. Schorfheide, 2004, Testing for Indeterminacy: An Application to US Monetary Policy, *American Economic Review*, 94(1), 190-217.
- Lucas, R.E., Jr, 1976, Econometric policy evaluation: A critique, *Carnegie-Rochester Conference Series on Public Policy* 1(1), 19-46.
- Lucas, R. E., Jr, 1980. Two Illustrations of the Quantity Theory of Money, *American Economic Review* 70(5), pp. 1005-14.
- Mavroeidis, S., 2010, Monetary Policy Rules and Macroeconomic Stability: Some New Evidence, *American Economic Review* 100, 491-503.
- McCallum, B. and E. Nelson, 1999, The Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model, in J.B. Taylor ed, *Monetary Policy Rules*, University of Chicago Press, 15-45.

- Meltzer, A.H., 2009, *A History of the Federal Reserve, Volume 2, Book 2, 1970-1986*, University of Chicago Press.
- Primiceri, G.E., 2005, Time-Varying Vector Autoregressions and Monetary Policy, *Review of Economic Studies*, 72, 821-852.
- Rotemberg, J., 1982, Sticky Prices in the United States, *Journal of Political Economy* 90, 1187-1211.
- Sargent, T.J., 1973, Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment, *Brookings Papers on Economic Activity* 2, 429-472.
- Sargent, T.J. and C. Sims, 1977, Business cycle modeling without pretending to have too much a priori economic theory in *New methods in business cycle research*, edited by C. Sims, Federal Reserve Bank of Minneapolis.
- Sargent, T.J. and P. Surico, 2011, Two Illustrations of the Quantity Theory of Money: Breakdowns and Revivals, *American Economic Review* 101, 113-132.
- Sims, C., 1971, Distributed Lag Estimation when the Parameter Space is Explicitly Infinite Dimensional, *Annals of Mathematical Statistics*, Vol. 42, No. 5, 1622-1636.
- Sims, C. and T. Zha, 2006, Were There Regime Switches in U.S. Monetary Policy?, *American Economic Review* 96, 54-81.
- Whiteman, C.H., 1984. Lucas on the Quantity Theory: Hypothesis Testing without Theory, *American Economic Review* 74(4), 742-49.

A The data

Inflation π is measured by the first difference of the logarithm of GDP deflator, output growth Δy is the first difference of the logarithm of real GDP, the short-term nominal interest rate R is the three month Treasury bill rate and money growth Δm is the first difference of the logarithm of M2. All data are quarterly and are available from the B.E.A. and Federal Reserve Bank of St. Louis (FRED). The series for M2 is available from FRED since 1959Q1, which is therefore the starting date for our analysis. The data are displayed in figure 8

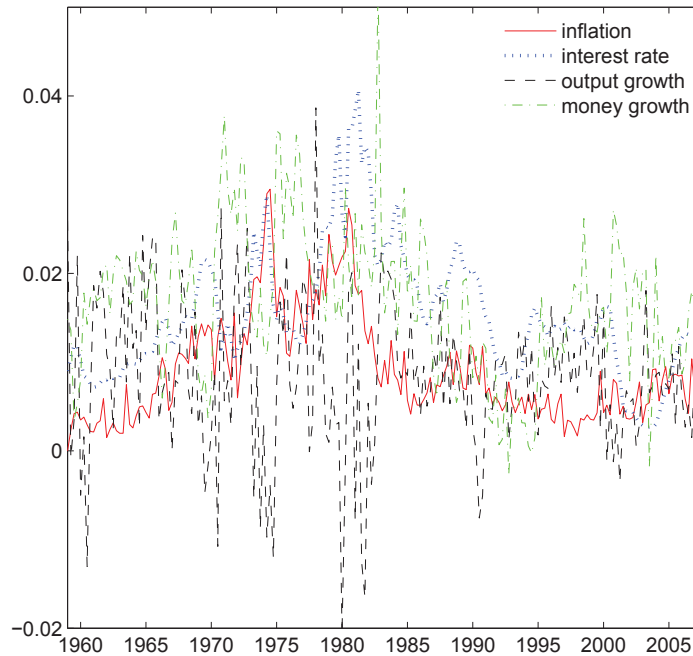


Figure 8: quarterly data for the United States- 1959Q1-2007Q4.

B Variance decompositions

Tables 4 and 5 report variance decompositions for the great-inflation and great-moderation samples, respectively. For the great-inflation subsample, mark-up shocks were the major driver of fluctuations in all variables but output growth, for which demand shocks played a predominant role. In line with Ireland (2004), technology

shocks account for only a small fraction of aggregate fluctuations. The last column displays the values (multiplied by 100) implied by the mean estimates in table 1 for the standard deviation of the four observables and the output gap.

Table 5 makes clear that the composition of shocks changed in the great-moderation sample. The contribution of mark-up shocks to fluctuations in inflation (output gap) increased (decreased). The contribution of policy shocks shifted towards the output gap, possibly reflecting a larger policy response to inflation during the great-moderation period. The large share of interest rate variance explained by policy shocks confirms the difficulty of estimating a policy rule over a sample of stable inflation and output gap, whose standard deviations in the last column are six times smaller than the standard deviations in table 4.

Table 4: Variance decomposition - 1968Q1-1983Q4

<i>series</i>	<i>Shocks</i>					<i>std. dev.</i>
	mark-up	demand	money ^d	technology	policy	
inflation	60.94	0.11	0.04	3.52	35.39	1.35
output gap	47.00	0.30	0.29	5.94	46.47	1.12
money growth	39.43	0.06	0.04	1.39	59.08	1.32
output growth	1.26	74.38	0.15	19.32	4.89	1.16
interest rate	58.86	1.97	2.83	12.30	24.04	1.23

Note: results are reported in percent and they are based on the mean estimates in table 1. Shares may not add up to 100% due to rounding. *money^d* stands for money demand, *std.dev.* stands for standard deviation.

Table 5: Variance decomposition - 1995Q1-2007Q4

<i>series</i>	<i>Shocks</i>					<i>std. dev.</i>
	mark-up	demand	money ^d	technology	policy	
inflation	74.60	0.00	0.00	0.00	25.40	0.28
output gap	28.25	0.00	0.00	0.00	71.75	0.25
money growth	1.24	5.05	8.61	25.27	59.84	0.84
output growth	0.80	49.24	0.00	45.44	4.53	0.62
interest rate	23.58	0.01	0.00	0.00	76.42	0.30

Note: results are reported in percent and they are based on the mean estimates in table 2. Shares may not add up to 100% due to rounding. *money^d* stands for money demand, *std.dev.* stands for standard deviation.

C The tail of the post-1995 posterior distribution for $\tilde{h}_{R,\pi}(0)$

The following table summarizes a pair of conditional posterior distributions for the structural parameters. These distributions were found by isolating right-tail draws from the posterior distribution of $\tilde{h}_{R,\pi}(0)$ and examining the associated structural parameters. Our objective is simply to understand what accounts for the long right tail in the post-1995 posterior distribution for $\tilde{h}_{R,\pi}(0)$. By comparing the unconditional posterior distribution in table 2 with the conditional posteriors summarized here, we see that the main difference concerns the persistence of markup shocks. While the unconditional posterior mean is 0.474, values associated with right-tail draws of $\tilde{h}_{R,\pi}(0)$ have a mean of 0.614 and 0.654, respectively, depending on how the tail is defined.

Table 6: Posterior estimates conditional to the tail of the $\tilde{h}_{R,\pi}(0)$ distribution - 1995Q1-2007Q4

description	coefficient	Conditional to $\tilde{h}_{R,\pi}(0) \geq 0.5$			Conditional to $\tilde{h}_{R,\pi}(0) \geq 1$		
		mean	[5 th ; 95 th]	95 th	mean	[5 th ; 95 th]	95 th
discount factor	β	0.9898	[0.9800 ; 0.9967]	0.9967	0.990	[0.982 ; 0.998]	0.998
NKPC backward-looking component	α_π	0.1444	[0.0417 ; 0.2847]	0.2847	0.1442	[0.0417 ; 0.2866]	0.2866
NKPC slope	κ	0.1328	[0.0901 ; 0.1817]	0.1817	0.1302	[0.0893 ; 0.1778]	0.1778
price adjustment cost	τ	3.6352	[2.4608 ; 5.1514]	5.1514	3.5444	[2.3832 ; 4.9962]	4.9962
IS curve backward-looking component	α_x	0.1670	[0.0632 ; 0.2744]	0.2744	0.1603	[0.0580 ; 0.2652]	0.2652
elasticity of intertemporal substitution	σ	0.1075	[0.0729 ; 0.1489]	0.1489	0.1053	[0.0716 ; 0.1456]	0.1456
inverse of labour supply elasticity	ξ	0.9850	[0.5634 ; 1.5215]	1.5215	0.972	[0.526 ; 1.393]	1.393
interest elasticity of money demand	γ	2.3529	[1.7141 ; 3.1143]	3.1143	2.3551	[1.7341 ; 3.1258]	3.1258
interest rate response to inflation	ψ_π	1.7034	[1.3249 ; 2.0867]	2.0867	1.7660	[1.4293 ; 2.1297]	2.1297
interest rate response to output gap	ψ_x	.1233	[0.0195 ; 0.2282]	0.2282	0.1209	[0.0160 ; 0.2282]	0.2282
interest rate smoothing	ρ_r	0.8174	[0.7585 ; 0.826]	0.826	0.8107	[0.7497 ; 0.8622]	0.8622
persistence of mark up shock	ρ_e	0.6142	[0.4183 ; 0.7813]	0.7813	0.6535	[0.4718 ; 0.8081]	0.8081
persistence of demand shock	ρ_a	0.9238	[0.8597 ; 0.9749]	0.9749	0.9236	[0.8611 ; 0.9749]	0.9749
persistence money demand shock	ρ_χ	0.6384	[0.2827 ; 0.9212]	0.9212	0.6389	[0.2859 ; 0.9212]	0.9212
standard deviation of mark up shock	σ_e	0.0055	[0.0037 ; 0.0078]	0.0078	0.0056	[0.0038 ; 0.0081]	0.0081
standard deviation of demand shock	σ_a	0.0045	[0.0030 ; 0.0067]	0.0067	0.0045	[0.0029 ; 0.0067]	0.0067
standard deviation of money demand shock	σ_χ	0.0052	[0.0033 ; 0.0079]	0.0079	0.0052	[0.0033 ; 0.0079]	0.0079
standard deviation of technology shock	σ_z	0.0042	[0.0035 ; 0.0051]	0.0051	0.0042	[0.0035 ; 0.0051]	0.0051
standard deviation of policy shock	σ_m	0.0021	[0.0018 ; 0.0025]	0.0025	0.0021	[0.0017 ; 0.0025]	0.0025

D The return of the Gibson paradox in UK data

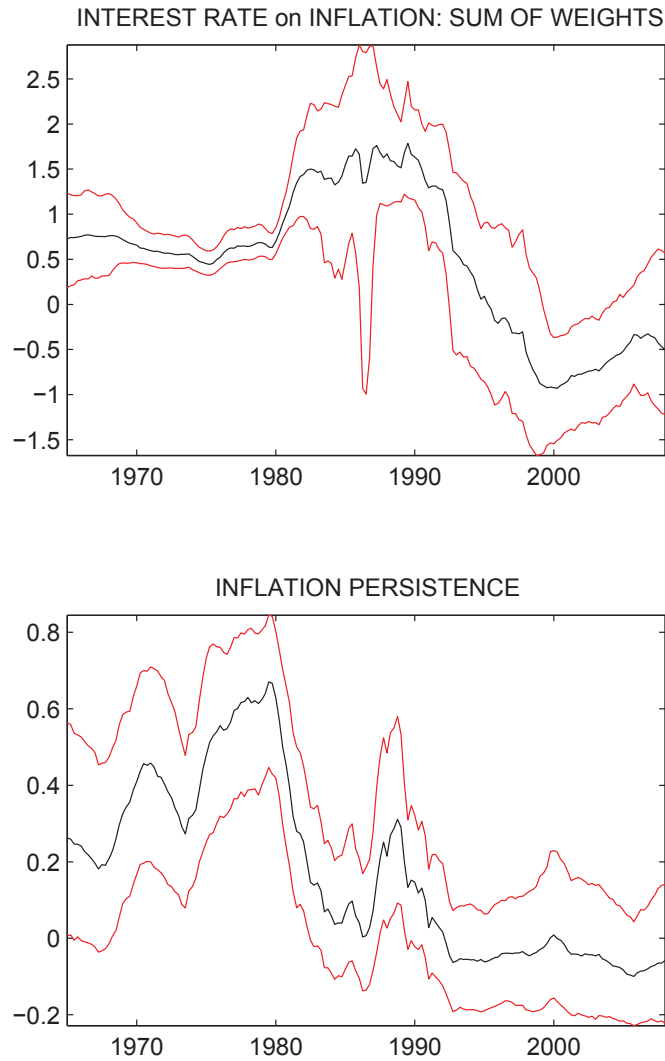


Figure 9: Median and 68% central posterior bands for $\tilde{h}_{R,\pi}(0)$ (top panel) and $FACF_{\pi}$ (bottom panel) based on a time-varying VAR(2) in the three month nominal interest rate and the first log difference of M4, GDP deflator and real GDP.