

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Benati, Luca

Working Paper Evolving Phillips trade-off

ECB Working Paper, No. 1176

Provided in Cooperation with: European Central Bank (ECB)

Suggested Citation: Benati, Luca (2010) : Evolving Phillips trade-off, ECB Working Paper, No. 1176, European Central Bank (ECB), Frankfurt a. M.

This Version is available at: https://hdl.handle.net/10419/153610

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



WORKING PAPER SERIES NO 1176 / APRIL 2010

EVOLVING PHILLIPS TRADE-OFF

by Luca Benati











NO 1176 / APRIL 2010

EVOLVING PHILLIPS TRADE-OFF¹

by Luca Benati²

NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the author and do not necessarily reflect those of the ECB.



In 2010 all ECB publications feature a motif taken from the €500 banknote.



This paper can be downloaded without charge from http://www.ecb.europa.eu or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=1585204.





© European Central Bank, 2010

Address Kaiserstrasse 29 60311 Frankfurt am Main, Germany

Postfach 16 03 19 60066 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Internet http://www.ecb.europa.eu

Fax +49 69 1344 6000

All rights reserved.

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the author.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, http://www. ecb.europa.eu/pub/scientific/wps/date/ html/index.en.html

ISSN 1725-2806 (online)

CONTENTS

Abstract	4		
Non-technical summary			
1 Introduction	6		
1.1 Main results	7		
1.2 Related literature	8		
2 Methodology	9		
2.1 A bayesian time-varying parameters			
VAR with stochastic volatility	9		
2.2 Identification	П		
3 Reduced-form evidence	12		
3.1 Time-series properties of the estimated			
cyclical components of inflation			
and unemployment	13		
3.2 The unconditional correlation between			
inflation and unemployment	14		
3.3 The evolution of forecast errors			
and the reduced-form Phillips loop	15		
3.4 The lead of unemployment over inflation	17		
4 Structural Analysis	18		
4.1 The unconditional correlation between			
inflation and unemployment	8		
4.2 The conditional Phillips trade-off	81		
5 Globalisation and the evolution of the Phillips			
trade-off	20		
6 Conclusions	21		
References	22		
Appendices	25		

Abstract

We characterise the evolution of the U.S. unemployment-inflation tradeoff since the late XIX century era *via* a Bayesian time-varying parameters structural VAR.

The Great Inflation episode appears as *historically unique* along several dimensions. In particular, the shape of the 'Phillips loop'—which is defined in terms of the impulse-response functions of inflation and unemployment's deviations from equilibrium—was, during those years, clearly out of line with respect to the rest of the sample period for all structural innovations except money demand shocks.

During the Great Depression, on the other hand, the Phillips trade-off did not exhibit any peculiar *qualitative* feature, so that, when seen through these lenses, the 1930s only stand out because of the sheer size of the macroeconomic fluctuation.

The historical evolution of the Phillips trade-off exhibits virtually *no* connection with the evolution of the extent of trade openness of the U.S. economy. Although, by itself, this does not rule out a possible impact of globalisation on the slope of the trade-off in recent years, it clearly suggests that, historically, the evolution of the trade-off has been dominated by factors other than trade openness.

Keywords: Phillips trade-off; Lucas critique; Bayesian VARs; time-varying parameters; stochastic volatility; identified VARs; Great Inflation; Great Depression; globalisation.

Non Technical Summary

Since the publication of A.W. Phillips' classic 1958 paper, the unemployment-inflation trade-off has been one of the most intensely investigated relationships in macroeconomics, playing a key role—in particular, during the 1970s—in shaping the evolution of both macroeconomic thought and policymaking. This paper uses a Bayesian time-varying parameters structural VAR to characterise the evolution of the U.S. unemployment-inflation trade-off since the Gold Standard era. Four structural shocks (money demand and supply, and non-monetary demand and supply), are identified by imposing sign restrictions upon the estimated reduced-form VAR on a quarter-by-quarter basis.

The Great Inflation episode clearly appears as *historically unique* along several dimensions. In particular, the shape of the 'Phillips loop'—which is defined in terms of the impulse-response functions of inflation and unemployment's deviations from equilibrium—was, during those years, clearly out of line with respect to the rest of the sample period for all structural innovations except money demand shocks. During the Great Depression, on the other hand, the Phillips trade-off did not exhibit any peculiar *qualitative* feature, so that, when seen through these lenses, the 1930s only stand out because of the sheer size of the macroeconomic fluctuation. The historical evolution of the Phillips trade-off exhibits virtually *no* connection with the evolution of the extent of trade openness of the U.S. economy. Although, by itself, this does not rule out a possible impact of globalisation on the slope of the trade-off has been dominated by factors other than trade openness.

Most economists of my generation have made a career of analysing so-called 'trade-offs' between inflation and unemployment, between external and domestic stability, between the long and the short run. But that theorising has been rooted in certain assumptions— assumptions that are now suspect—about the stability of expectations. When expectations of future inflation are so strong and potentially volatile as they have become, the 'trade-offs' disappear, or they appear in a much different light. (emphasis added)

—Paul Volcker (1979)

In earlier periods before roughly 1965, the monetary regime guaranteed some long-run stability in monetary growth, and therefore in long-term inflation, which in turn restricted the effects of shifting inflationary expectations [...]. The international economy has been moving gradually away from this type of monetary setup since World War I, and especially since the 1930s, although some remnants of the Gold Standard and fixed exchange rates in the form of the post-WWII Bretton Woods arrangements were in operation as recently as 1971. [...] Although there were earlier periods when the United States did not adhere to a gold or silver standard, these episodes typically occurred in times of war and could reasonably be perceived as temporary. The period since 1971 seems to be the first time that we have completely severed, both currently and prospectively, the link between our money and a commodity base.

-Robert Barro (1982)

[I]f we fail to maintain a situation which is conducive to price stability, we could find ourselves caught up very quickly in an inflationary spiral.

—William McChesney Martin (1965)

1 Introduction

Since the publication of A.W. Phillips' paper,¹ the unemployment-inflation tradeoff has been one of the most intensely investigated relationships in macroeconomics, playing a key role—in particular, during the 1970s—in shaping the evolution of both macroeconomic thought² and policymaking.

In recent years, a vast literature has produced evidence of a flattening of the Phillips trade-off—defined in terms of either the unemployment rate, or of proxies for the output gap—over the last two decades,³ and has explored alternative possible

¹See Phillips (1958). The existence of a negative relationship between wage inflation and unemployment had been previously documented by several researchers, most notably Fisher (1926) and Tinbergen (1937)—see the discussion in Bacon (1973).

²See in particular Lucas (1972a), Lucas (1972b), and Lucas (1973).

³See in particular the discussion in Mishkin (2007). To mention just two papers within a vast literature, see Brainard and Perry (2000), who estimate a time-varying parameters Robert Gordon-type Phillips curve for the United States, and Benati (2007), who uses complex demodulation techniques to explore changes in the correlation between inflation and measures of economic activity at the business-cycle frequencies for a panel of countries.

explanations for this phenomenon. The two most commonly advanced explanations are (i) the better anchoring of inflation expectations which characterises the current period compared, first and foremost, with the Great Inflation years,⁴ and (ii) the impact of globalisation.⁵

In this paper we use a Bayesian time-varying parameters structural VAR to characterise the evolution of the U.S. unemployment-inflation trade-off since the Gold Standard era. We identify four structural shocks (money demand and supply, and non-monetary demand and supply), by imposing sign restrictions upon the estimated reduced-form VAR on a quarter-by-quarter basis.

1.1 Main results

Our main findings can be summarised as follows.

First, when seen from a long-term perspective, the Great Inflation episode clearly appears as *historically unique* along several dimensions. In particular, the shape of the 'Phillips loop'—which is defined in terms of the impulse-response functions of inflation and unemployment's deviations from equilibrium—was, during those years, clearly out of line with respect to the rest of the sample period for all structural innovations except money demand shocks. During the Great Depression, on the other hand, the Phillips trade-off did not exhibit any peculiar *qualitative* feature, so that, when seen through these lenses, the 1930s only stand out because of the sheer size of the macroeconomic fluctuation.

Second, overall, the long-term evolution of the Phillips trade-off broadly reflects the evolution of U.S. monetary regimes since the end of the XIX century, and their impact on the anchoring—or de-anchoring—of inflation expectations. This clearly emerges, for example, from an analysis of the evolution of the slope of the structural Phillips trade-off in response to demand non-policy shocks. Whereas during the Gold Standard inflation was followed by deflation, so that the last portion of the Phillips loop was clearly in negative territory, during the Great Inflation episode (i) the negative portion of the Phillips loop entirely disappeared, and (ii) the trade-off exhibited the steepest slope in history.

Third, the historical evolution of the Phillips trade-off exhibits virtually no connection with the evolution of the extent of trade openness of the U.S. economy. In particular, (i) the transition from the period leading up to World War I—which was characterised by the first period of globalisation—to the trade restrictions of the interwar period, and the resulting collapse in world trade, was not accompanied by any perceptible change in either the reduced-form or the structural Philips trade-off; and (ii) the evolution of the trade-off over the post-WWII era does not exhibit any obvious connection with the evolution of the extent of trade openness of the U.S.

 $^{^{4}}$ See FED Chairman Bernanke's speech at the 2007 NBER Summer Institute (Bernanke (2007)), and the literature cited therein.

⁵See e.g. Borio and Filardo (2007).

economy. Although, by itself, this does not rule out a possible impact of globalisation on the slope of the trade-off in recent years, it clearly suggests that, historically, the evolution of the trade-off has been dominated by factors other than the extent of openness in trade (and financial) flows.

1.2 Related literature

In spite of the unemployment-inflation trade-off having been so intensely investigated over the last several decades, to the very best of our knowledge King and Watson (1994) is the only study ever to have been based on structural VAR methods.⁶ Different from the present work, however, King and Watson's (1994) exclusive focus was on estimating the *long-run* trade-off between the two series, and, in particular, on assessing the impact of alternative identifying assumptions—which were derived from three contrasting theoretical approaches (traditional Keynesian, rational-expectations monetarist, and real business-cycle)—on the estimated trade-off. In contrast, the focus of the present work is on the *short-to-medium run* trade-off—which we define in terms of the deviations of the two series from their time-varying equilibria—whereas, by construction, in the long-run both inflation and unemployment return to such equilibria. So we completely eschew the issue of the long-run trade-off, which within the present context cannot be meaningfully addressed.

This paper has been largely inspired by the analysis of the evolution, over the post-1960 period, of U.S. inflation's forecast errors in response to reduced-form unemployment innovations, which was contained in the NBER working paper version of Cogley, Primiceri, and Sargent $(2010)^7$ (those results were not retained in the final published version). The present work paper extends Cogley *et al.*'s (2008) analysis back in time to the Classical Gold Standard era, and it supplements it with a structural analysis based on sign restrictions.

The paper is organised as follows. The next section discusses the reduced-form specification for the time-varying parameters VAR with stochastic volatility we will use throughout the paper, and the identification strategy, which is based on sign

⁶Roberts (1993) estimated a fixed-coefficients structural VAR for M2 velocity and the first differences of inflation and the unemployment rate for the United States for the period 1962:Q3-1988:Q4. He identified both natural rate shocks, and shocks to the FED's 'inflation target' (which might be interpreted as trend, or equilibrium inflation), but his analysis completely eschewed the Phillips trade-off, and was instead almost exclusively focused on determining the portion of U.S. businesscycle fluctuations which could be attributed to the various shocks.

By the same token, King and Morley (2007) used a fixed-coefficients structural VAR to estimate the natural rate of unemployment for the post-WWII United States. In Section 6, pp. 560-562, they then computed some simple correlations between unemployment's deviation from the estimated natural rate and *raw* inflation (as opposed to inflation's deviation from equilibrium, as it is done in the present work), but their work was almost exclusively focused on estimating the natural rate, rather than analysing the Phillips trade-off.

⁷See Cogley, Primiceri, and Sargent (2008), Section 5.1.

restrictions, whereas (standard) technical aspects of the Bayesian inference—in particular, our choices for the priors, and the Markov chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data—are relegated to an appendix. Section 3 presents and discusses the reduced-form evidence, in particular, the time-series properties of the estimated cyclical components of inflation and unemployment (computed as the differences between the actual series and their VAR-implied trends); the evolution of the unconditional correlation between inflation and unemployment; the evolution of the reduced-form conditional loop, defined in terms of the ratio between the responses of the two series to a reduced-form shock to unemployment; changes in the relative persistence of inflation and unemployment; and the evolution of the lead of unemployment over inflation. Section 4 turns to a structural analysis, by discussing the evolution of the conditional Phillips loop, defined as the ratio between the impulse-response functions of inflation and unemployment to each of the four structural shocks. Section 5 discusses the implications of our findings for the two previously mentioned explanations of the recent flattening of the Phillips trade-off, that is, a better anchoring of inflation expectations and the impact of globalisation. Section 6 concludes, and outlines possible directions for future research.

2 Methodology

2.1 A Bayesian Time-Varying Parameters VAR with Stochastic Volatility

Following Cogley, Primiceri, and Sargent (2010), in what follows we will work with the time-varying parameters VAR(p) model

$$Y_{t} = B_{0,t} + B_{1,t}Y_{t-1} + \dots + B_{p,t}Y_{t-p} + \epsilon_{t} \equiv X_{t}'\theta_{t} + \epsilon_{t}$$
(1)

where the notation is obvious, with $Y_t \equiv [\pi_t, u_t, r_t, m_t]'$, with π_t and m_t being inflation and the rate of growth of M2, r_t being the logarithm of the commercial paper rate, and u_t being the unemployment rate. Appendix A describes the data, whereas Appendix B discusses the methodology, along the lines of Bernanke, Gertler, and Watson (1997), which we use to produce a quarterly interpolated series for the rate of unemployment for the period 1890:1-1929:2.⁸

Consistent with the vast majority of the papers in the literature, and mostly for reasons of computational feasibility, the lag order is set to p=2. Following, e.g., Cogley and Sargent (2002), Cogley and Sargent (2005), and Primiceri (2005) the VARs' time-varying parameters, collected in the vector θ_t , are postulated to evolve

⁸Monthly data for the U.S. unemployment rate are available starting from April 1929, whereas for previous years data are only available at the annual frequency.

according to

$$p(\theta_t \mid \theta_{t-1}, Q_t) = I(\theta_t) \ f(\theta_t \mid \theta_{t-1}, Q_t)$$
(2)

with $I(\theta_t)$ being an indicator function rejecting unstable draws—thus enforcing a stationarity constraint on the VAR⁹—and with $f(\theta_t \mid \theta_{t-1}, Q_t)$ given by

$$\theta_t = \theta_{t-1} + \eta_t \tag{3}$$

with $\eta_t \equiv [\eta_{1,t}, \eta_{2,t}, ..., \eta_{N \cdot (1+Np),t}]'$. $\eta_t \sim N(0, Q_t)$. We postulate a stochastic volatility specification for the evolution of the covariance matrix of the innovations to the VAR's random-walk coefficients, Q_t . Specifically, we assume that Q_t is given by¹⁰

$$Q_t \equiv \begin{bmatrix} q_{1,t} & 0 & \dots & 0 \\ 0 & q_{2,t} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & q_{N \cdot (1+Np),t} \end{bmatrix}$$
(4)

with the $q_{i,t}$'s evolving as geometric random walks,

$$\ln q_{i,t} = \ln q_{i,t-1} + \omega_{i,t} \tag{5}$$

For future reference, we define $q_t \equiv [q_{1,t}, q_{2,t}, ..., q_{N \cdot (1+Np),t}]'$.

The VAR's reduced-form innovations in (1) are postulated to be zero-mean normally distributed, with $Var(\epsilon_t) \equiv \Omega_t$ with time-varying covariance matrix Ω_t which, following established practice, we factor as

$$\Omega_t = A_t^{-1} H_t (A_t^{-1})' \tag{6}$$

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

$$H_{t} \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0\\ 0 & h_{2,t} & 0 & 0\\ 0 & 0 & h_{3,t} & 0\\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \qquad A_{t} \equiv \begin{bmatrix} 1 & 0 & 0 & 0\\ \alpha_{21,t} & 1 & 0 & 0\\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0\\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix}$$
(7)

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

$$\ln h_{i,t} = \ln h_{i,t-1} + \nu_{i,t} \tag{8}$$

⁹It is important here to be precise about the meaning of such stationarity constrainty. Although, due to the time-varying parameters specification (1), inflation contains a stochastic trend, the constraint (2) implies that its fluctuations *around* such trend cannot be explosive.

¹⁰This specification is simpler than the one used by Cogley, Primiceri, and Sargent (2010), who factor the covariance matrix of the innovations to the VAR's random-walk parameters as $Q_t = (B_s^{-1})'H_{s,t}B_s^{-1}$, where $H_{s,t}$ has exactly the same specification which is postulated herein for Q_t , and B_s is a triangular matrix with ones along the main diagonal and static covariance parameters below. (To put it differently, our specification is obtained from Cogley *et al.*'s by setting B_s equal to the identity matrix.)

For future reference, we define $h_t \equiv [h_{1,t}, h_{2,t}, h_{3,t}, h_{4,t}]'$. Following Primiceri (2005), we postulate the non-zero and non-one elements of the matrix A_t —which we collect in the vector $\alpha_t \equiv [\alpha_{21,t}, ..., \alpha_{43,t}]'$ —to evolve as driftless random walks,

$$\alpha_t = \alpha_{t-1} + \tau_t , \qquad (9)$$

and we assume the vector $[u_t', \tau_t', \nu_t', \omega_t']'$ to be distributed as N(0, V), with

$$V = \begin{bmatrix} I_4 & 0 & 0 & 0 \\ 0 & S & 0 & 0 \\ 0 & 0 & Z_{\nu} & 0 \\ 0 & 0 & 0 & Z_{\omega} \end{bmatrix}, \ Z_{\nu} = \begin{bmatrix} \sigma_{\nu,1}^2 & 0 & 0 & 0 \\ 0 & \sigma_{\nu,2}^2 & 0 & 0 \\ 0 & 0 & \sigma_{\nu,3}^2 & 0 \\ 0 & 0 & 0 & \sigma_{\nu,4}^2 \end{bmatrix} \text{ and}$$
$$Z_{\omega} = \begin{bmatrix} \sigma_{\omega,1}^2 & 0 & \dots & 0 \\ 0 & \sigma_{\omega,2}^2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \sigma_{\omega,N\cdot(1+Np)}^2 \end{bmatrix}$$
(10)

where u_t is such that $\epsilon_t \equiv A_t^{-1} H_t^{\frac{1}{2}} u_t$. Finally, following Primiceri (2005) we adopt the additional simplifying assumption of postulating a block-diagonal structure for S, too—namely

$$S \equiv \operatorname{Var}(\tau_t) = \operatorname{Var}(\tau_t) = \begin{bmatrix} S_1 & 0_{1\times 2} & 0_{1\times 3} \\ 0_{2\times 1} & S_2 & 0_{1\times 3} \\ 0_{3\times 1} & 0_{3\times 2} & S_3 \end{bmatrix}$$
(11)

with $S_1 \equiv \text{Var}(\tau_{21,t})$, $S_2 \equiv \text{Var}([\tau_{31,t}, \tau_{32,t}]')$, and $S_3 \equiv \text{Var}([\tau_{41,t}, \tau_{42,t}, \tau_{43,t}]')$, thus implying that the non-zero and non-one elements of A_t belonging to different rows evolve independently. As discussed in Primiceri (2005, Appendix A.2), this assumption drastically simplifies inference, as it allows to do Gibbs sampling on the non-zero and non-one elements of A_t equation by equation.

We estimate (1)-(11) via Bayesian methods. Appendix B discusses our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data.

2.2 Identification

We identify four structural shocks by imposing sign restrictions upon the estimated reduced-form VAR on a quarter-by-quarter basis. Two shocks pertain to the demand and supply sides of the money market—we label them as ϵ_t^{MD} and ϵ_t^{MS} respectively whereas the remaining two are non-monetary demand and supply shocks, which we label as ϵ_t^D and ϵ_t^S , respectively. The following table summarises the sign restrictions we impose upon the estimated VAR. The restrictions are the same we used in Benati (2008), with the obvious difference that, since the real activity indicator used therein was output growth, instead of the unemployment rate, the signs reported in the second row in the table below are the opposite of the corresponding signs for output growth in Benati (2008). The sign restrictions are imposed both on impact, and for two subsequent quarters. A '+' and a '-' mean 'greater than or equal to zero' and 'smaller than or equal to zero', respectively, whereas a '?' means that the sign of this specific impact has been left unconstrained.

	Shock:			
Variable:	ϵ_t^{MS}	ϵ_t^{MD}	ϵ_t^D	ϵ_t^S
inflation	+	—	+	_
unemployment	—	+		—
commercial paper rate	_	+	+	?
M2 growth	+	+	+	+

We compute the time-varying structural impact matrix, $A_{0,t}$, via the procedure proposed by Rubio-Ramirez, Waggoner, and Zha (2005). Specifically, let $\Omega_t = P_t D_t P'_t$ be the eigenvalue-eigenvector decomposition of the VAR's time-varying covariance matrix Ω_t , and let $\tilde{A}_{0,t} \equiv P_t D_t^{\frac{1}{2}}$. We draw an $N \times N$ matrix, K, from the N(0, 1) distribution, we take the QR decomposition of K—that is, we compute matrices Q and R such that $K=Q \cdot R$ —and we compute the time-varying structural impact matrix as $A_{0,t}=\tilde{A}_{0,t} \cdot Q'$. We then check whether the impact matrix satisfies the sign restrictions. If it does, we keep the draw, otherwise we discard it and we keep drawing until all the restrictions are satisfied. If, after drawing 10,000 rotation matrices, we still fail to find one such that the corresponding impact matrix satisfies the sign restrictions, we move on to the next draw from the ergodic distribution.

3 Reduced-Form Evidence

Since our investigation of changes over time in the U.S. unemployment-inflation tradeoff will be based on the cyclical components of the two series produced by the timevarying VAR, a crucial preliminary step, before delving into the analysis, involves exploring their time-series properties. How do these cyclical components look like, and what do they resemble? Are they 'reasonable', in the specific sense of closely mimicking cyclical components produced by widely-accepted detrending methods? In the next sub-section we explore this issue, by comparing the cyclical components of the two series generated by the VAR with those produced by the Hodrick-Prescott filter.

3.1 Time-series properties of the estimated cyclical components of inflation and unemployment

Figure 1 shows, in the top panel, the median estimate of the 'inflation gap'—which, in line with Cogley, Primiceri, and Sargent (2010), we define as the difference between inflation and its VAR-implied time-varying equilibrium—together with HP-filtered inflation, whereas the bottom panel shows the corresponding objects for the rate of unemployment. The visual impression suggests a close resemblance between the unemployment gap and the HP-filtered unemployment rate over the entire sample period, with the exception of the Great Depression, for which the HP filter produces a significantly smaller cyclical component. Although, quite obviously, we have no way of determining which of the two proxies for the cyclical component of the rate of unemployment should be regarded as the more reliable, two things ought to be noticed. First, our estimate of the unemployment gap, which suggest that unemployment returned to equilibrium only towards the end of 1941, is exactly in line with the results of Romer (1992), who, in discussing the evolution of U.S. real GNP relative to her estimate of potential during the 1930s, stated¹¹ that '....' GNP was about 38 per cent below its trend level in 1935 and 26 per cent below it in 1937. Only in 1942 did GNP return to trend.' ¹² She then went on to stress that¹³ '...] full employment was not reached again until 1942.' Second, it is to be noticed that linear filtering methods such as the HP filter split a series into a trend and a cyclical components uniquely based on frequency-domain logic,¹⁴ so that large, prolonged, and very highly persistent cyclical fluctuations—such as the one that might be thought to have been associated with the Great Depression—are automatically interpreted, by such methods, as partly reflecting fluctuations in the trend. This is why the HP-filtered cyclical components of unemployment for these two episodes are significantly smaller than the ones produced by the VAR, which, different from the HP filter, estimates the time-varying properties of the economy at each point in time. So, although we do not want to take a strong stand on which of the two cyclical proxies should be regarded as the more reliable, we are inclined to believe that precisely because the method used herein estimates the stochastic properties of the economy via multivariate methods, it might be thought as more reliable than simple linear filtering methods. As for the inflation gap, on the other hand, it appears as remarkably close to HP-filtered inflation, with the partial and minor exception of the Great Depression.

Overall these results are reassuring, in the sense that (i) the VAR-implied cyclical components exhibit reasonable properties; (ii) they closely co-move with—and they

¹¹See Romer (1992, p. 760).

 $^{^{12}\}mathrm{On}$ this, see in particular her Figure 2, p. 761.

 $^{^{13}}$ See Romer (1992, p. 761).

¹⁴The frequency-domain logic at the root of linear filtering is explicit in the case of band-pass filters—see e.g. Baxter and King (1999) and Christiano and Fitzgerald (2003)—and is instead only implicit in the case of the Hodrick-Prescott filter. For an analysis of the HP filter from a frequency-domain perspective, see Baxter and King (1999) and Ravn and Uhlig (2002).

are often indistinguishable from—the cyclical components generated by the Hodrick-Prescott filter; and (*iii*) the significant departure of the unemployment gap from HP-filtered unemployment around the time of the Great Depression can be easily rationalised, and given the superiority, in principle, of the method used herein, it should not be regarded as problematic. Let's now turn to a reduced-form analysis of the evolution of the correlation between the cyclical components of inflation and unemployment.

Unless otherwise stated, from now on, by 'inflation' and 'unemployment' we will mean 'the VAR-implied cyclical component of inflation' and 'the VAR-implied cyclical component of unemployment', respectively—that is, the two gaps.

3.2 The unconditional correlation between inflation and unemployment

In recent years higher inflation has often been accompanied by higher, not lower unemployment, especially for periods of several years in length. A simple statistical Phillips curve for such periods seems to be positively sloped, not vertical.

-Milton Friedman (1977)

In this section, and in the next one, we extend the results of Section 5.1 of Cogley et al. (2008) backwards in time to the Classical Gold Standard era. The left-hand side panel of Figure 2 shows the median and the 16th and 84th percentiles of the posterior distribution of the unconditional correlation between inflation and unemployment,¹⁵ whereas the right-hand side panel shows the fraction of the draws from the ergodic distribution for which the correlation is negative. In line with Cogley et al. (2008), the most interesting finding emerging from the figure is the change in the sign of the correlation—from negative to *positive*—around the time of the Great Inflation episode, with the fraction of draws for which the correlation is estimated to have been negative reaching a historical low of 28.9 per cent in 1980:2. During the rest of the sample period, on the other hand, the correlation is estimated to have been uniformly negative, with the fraction of draws shown in the right-hand side panel fluctuating between about 54 per cent, in the immediate aftermath of World War II, and more than 97 per cent towards the end of the 1930s. After the Great Inflation, the correlation has clearly returned negative, with the fraction of draws shown in the right-hand side panel fluctuating between about 65 and 85 per cent.

The dramatic change in the sign of the reduced-form unconditional correlation between unemployment and inflation documented in Figure 2 illustrates a first important dimension along which the Great Inflation episode appears as *unique*, when

¹⁵For each quarter, we compute the unconditional correlation between the inflation and unemployment gaps based on the unconditional variance-covariance matrix of the time-varying VAR, which we compute as in Cogley *et al.* (2008).

seen from a very long-term perspective. In the following pages we will document several other dimensions along which the Great Inflation clearly stands out from the rest of recorded macroeconomic history (at least, since the end of the XIX century ...). An important point to stress is that the extent of the dislocation in the Phillips trade-off caused by the Great Inflation—compared with the historical norm—was significantly greater than that caused by World War I.¹⁶ As the right-hand side panel of Figure 2 shows, around the time of World War I the fraction of draws for which the correlation was negative dropped below 60 per cent only for a few quarters. On the other hand, during the Great Inflation episode it dropped below 40 per cent, and it stayed there for several years. The fact that the dislocation caused by the Great Inflation in such a fundamental macroeconomic relationship such as the unemployment-inflation trade-off was so significantly greater than that caused by a major war brings home, in the starkest possible way, the truly unique nature of that episode.

3.3 The evolution of forecast errors and the reduced-form Phillips loop

Figure 3 shows, for selected quarters since 1897Q4, the median and the 16th and 84th percentiles of the distributions of the reduced-form 'Phillips loop', where such loop is defined as the ratio between the forecast errors of inflation and unemployment at the various horizons in response to a one per cent¹⁷ negative reduced-form innovation to unemployment at each point in time. Since reduced-form innovations are correlated (as encoded in the VAR's covariance matrix), a reduced-form innovation to unemployment at time zero is associated to reduced-form innovations to the other three variables, and the entire exercise is therefore performed following the steps detailed in Cogley *et al.*'s (2008) Section 5.1.¹⁸ For each quarter *t*, and for each draw *k* from the ergodic distribution generated by the Gibbs sampler, we shock the unemployment rate by minus one per cent, computing the corresponding innovation for the other three variables based on the VAR's covariance matrix. The Phillips loop is then nothing but the scatterplot of the forecast errors of inflation and unemployment

¹⁶Due to the widespread use of price controls during World War II, on the other hand, a proper comparison with this conflict is much more fraught with difficulties.

¹⁷Since we are focusing on the *ratio* between the IRFs of the two variables to a reduced-form innovation to unemployment, the specific *size* of such innovation is obviously irrelevant. For example, instead of considering, at each point in time, a one per cent shock to the unemployment rate, we could have considered a one-standard deviation innovation, as Cogley *et al.* (2008) do.

 $^{^{18}}$ In a previous version of the paper we presented results from a different exercise, in which the unemployment rate is decreased by one percentage point at time t, whereas all other variables are postulated to remain unaffected on impact (and only on impact). Results from this exercise were broadly in line with those shown in Figure 4 (these results are available upon request).

at the various horizons.¹⁹ Several things ought to be stressed about this figure.²⁰ In particular,

- evidence for the Gold Standard accords remarkably well with what we would expect *ex ante* based on our knowledge of the workings of metallic standards. Since, as it is well known,²¹ a key feature of such regimes was (broad) stabilisation of the price level, we would expect the inflationary outburst caused by a decrease in the unemployment rate to be subsequently followed by a deflationary spell. As the first two panels on the left show, this is indeed what our evidence suggests, with inflation during the first quarters after the shock being followed by deflation along the path that led the economy back to equilibrium. Further, as the panels clearly show, the slope of the loop switches sign as time goes by, going from negative to positive.
- Intriguingly, the interwar period is not associated with any significant change compared with the Gold Standard era. Most importantly, the Great Depression does not appear to have had *any* discernible impact on the reduced-form Phillips loop.²²
- The post-WWII period is associated with the emergence of a different pattern, which is especially apparent during the Great Inflation episode. Specifically, focusing on median estimates, inflation was still followed by deflation at the very beginning of the post-WWII era, and started instead being followed by more and more inflation during subsequent years. This is especially clear for the Great Inflation of the 1970s, which, based on median estimates, exhibited the steepest trade-off since the late XIX century.

 20 The scatterplots reported in the figure only refer to the quarters *after* the impact.

 21 See e.g. Barro (1979).

 22 A *caveat* to these results is that they obviously crucially depend on the correct measure of the unemployment gap being the one generated by our VAR, rather than something closer to the one generated by the HP filter. As we previously mentioned, there are strong reasons why we should prefer our measure, but this *caveat* should nonetheless be kept in mind.

¹⁹A subtle issue here is the following. In principle, scatterplots such those reported in Figure 4 can be constructed in two alternative ways. *First*, you can sort inflation and unemployment's IRFs, and then you can plot the k-th percentile of the distribution of inflation's IRFs against the k-th percentile of the distribution of unemployment's IRFs. The key problem with this approach is that the sorting gets performed along a dimension which, for the present purposes, is not the relevant one. *Second*, you can sort the *slopes* of the trade-off—defined, for each stochastic simulation, as the ratio between inflation's and unemployment's IRFs—and then compute the percentiles of the distribution of the Phillips loop based on the percentiles of the distribution of the slopes. Since what we are ultimately interested in for the present purposes is the slope of the trade-off, the percentiles of the distributions reported in Figure 4—as well as the percentiles of the distributions reported in subsequent figures—are based on the second approach. To be clear, in order to recover the loops reported in figure 4, we have applied the percentiles of the distributions of the slopes to the median of the distribution of the IRF of unemployment.

Overall, the simplest way to interpret the evidence reported in Figure 3 is therefore in terms of the evolution of the U.S. monetary regime since the end of the XIX century, with the transition from a regime which aimed at stabilising the price level to a regime characterised instead by a continuous, and sometimes accelerated drift in the price level. Another way of putting this is that the evolution of the U.S. reduced-form tradeoff since the Gold Standard era clearly reflects the evolution of the monetary regime, and is therefore nothing but a simple manifestation of the Lucas (1976) critique.

Focusing on the post-WWII era, one simple interpretation of the evidence reported in the bottom row is in terms of changes in the *relative* persistence of inflation and unemployment. Figure 4 reports the median and the 16th and 84th percentiles of the distributions of the normalised spectrum of the inflation and unemployment gaps at the frequency $\omega=0$, together with the same objects for the distribution of the ratio between the normalised spectra of the inflation and unemployment gaps at $\omega=0$. As the figure clearly shows,

(i) the persistence of the unemployment gap exhibits, overall, very weak evidence of time-variation over the sample period.

(*ii*) Inflation persistence, on the other hand, exhibits a clear hump-shaped pattern over the sample period, with a peak around the time of the Great Inflation episode, and a decline thereafter.²³

(*iii*) As a consequence of (*i*) and (*ii*), the time-pattern of the *relative* persistence of inflation and unemployment closely mimics the corresponding time-pattern of inflation persistence, with the result that, around the time of the Great Inflation, a reduced-form innovation to unemployment is associated with a comparatively more prolonged and drawn out fluctuation in inflation than either before or after the Great Inflation, thus resulting in the previously mentioned steepening of the slope.

3.4 The lead of unemployment over inflation

Figure 5 reports evidence on changes over time in the average lead of unemployment over inflation at the business-cycle frequencies. Specifically, for each quarter t, and for each draw k from the ergodic distribution generated by the Gibbs sampler, we simulate the economy 200 quarters (that is, 50 years) into the future. Then, for each of these simulated paths, we extract from the generated, artificial unemployment and inflation series the business-cycle components *via* the Christiano and Fitzgerald (2003) band-pass filter,²⁴ and we compute the cross-correlations between such components at leads and lags. For each simulated path, we store the lead corresponding to the maximum of the cross-correlation between unemployment and inflation.

²³Needless to say, this is very much in line with the work, first and foremost, of Cogley and Sargent (2002), Cogley and Sargent (2005), and Cogley, Primiceri, and Sargent (2010).

²⁴Following established conventions in business-cycle analysis—see e.g. Baxter and King (1999), Stock and Watson (1999), and Christiano and Fitzgerald (2003)—the business-cycle frequency band is defined as the one associated with frequencies of oscillation between 6 quarters and 8 years.

The left-hand side panel of Figure 5 reports the median, the mean, and the 16th and 84th percentiles of the distributions of the leads corresponding to the maximum of the cross-correlations between unemployment and inflation, whereas the right-hand side panel reports the fraction of the simulated paths for which the lead corresponding to the maximum of the cross-correlations in quarter t is greater than in the quarter of reference. As the results reported in the figure clearly show, evidence of changes over time in the lead-lag relationship between unemployment and inflation is, overall, extremely weak, and only partly suggestive of a shift from a lead, overall, of unemployment over inflation before WWII to the opposite pattern over the post-WWII era.

Let's now turn to a structural analysis.

4 Structural Analysis

4.1 The unconditional correlation between inflation and unemployment

The top row of Figure 6 shows the medians and the 16th and 84th percentiles of the posterior distributions of the unconditional correlations between inflation and unemployment²⁵ generated by each of the four structural shocks,²⁶ whereas the bottom row shows the fractions of the draws from the ergodic distribution for which the correlations are negative. As the figure clearly shows, the result discussed in Section 3.2—the unconditional reduced-form correlation turning from negative to positive around the time of the Great Inflation—is largely driven by the money supply shock, and to a significantly lower extent by the demand and money demand shocks, whereas the supply shock does not play any discernible role.

4.2 The conditional Phillips trade-off

An important limitation of the analysis of Section 3.3 is that weak evidence of changes in the reduced-form Phillips loop may conceal significant changes in the structural loops—that is, the inflation-unemployment loops generated by structural innovations—which may partially cancel out in the reduced-form. As we will now

 $^{^{25}}$ For each quarter, we compute the unconditional correlation between the inflation and unemployment gaps as in Section 3.2.

²⁶To be clear, the results reported in each of the columns of Figure 7 have been generated by setting to zero the variances of all of the structural shocks *except* the one of interest. This is trivially done as follows. Since $\Omega_t = A_{0,t}A'_{0,t}$ —with the structural shocks being unit-variance by construction the VAR's variance-covariance matrix which is obtained by setting to zero the variances of all the structural shocks except the first one is given by $\Omega_{1,t} = A_{0,t}\Xi_1A'_{0,t}$, where $\Xi_1 = \text{diag}([1 \ 0 \ 0 \ 0]')$. The variance-covariance matrices which are obtained by setting to zero each of the other structural shocks are computed in the same way. Then, based on the $\Omega_{j,t}$'s, we compute the unconditional correlations between the inflation and unemployment gaps as in Section 3.2.

see, this appears indeed to be case here, with the evidence of changes over time being non-existent for one shock, money demand, and very strong for the other three.

Figures 8-11 report, for the same selected quarters shown in Figure 3, the medians and the 16th and 84th percentiles of the distributions of the structural Phillips loops, where such loops are defined in terms of the generalised IRFs of inflation and unemployment to a positive structural shock. The IRFs to each of the shocks have been normalised so that the absolute value of the impact at zero on the unemployment rate is equal to one per cent. Generalised IRFs have been computed *via* the Monte Carlo integration procedure proposed by Koop, Pesaran, and Potter (1996), which allows to properly take into account of future variation in the state of the economy.²⁷ Figures 8-11 show the loops for k = 1, 2, 3, ..., 20 quarters after the impact.

Several things are readily apparent from Figures 7-10. In particular,

- the structural loop in response to a money demand shock does not exhibit any significant change since the late XIX century, and, in particular, it does not appear to have been affected at all by the Great Inflation episode.
- Evidence for the other three shocks, on the other hand, points towards significant changes over the sample period, with all the three structural loops clearly exhibiting a highly idiosyncratic behaviour—and a remarkable steepness—around the time of the Great Inflation, and much less variation both before and after those years. This is especially clear for both the supply non-monetary shock whose steepness dramatically increases during the 1970s—and for the money supply shock, which during those years exhibits, after the very first quarters,²⁸ a *negative* slope completely out of line with the rest of the sample period.
- Crucially, *all* the three shocks for which significant evidence of time-variation in the Phillips loop has been detected exhibit one important similarity over the post-WWII period. For either of them the immediate aftermath of World War II appears as very similar not only to the interwar era, but also to the period which is closest to us, whereas the intervening years appear as clearly 'off kilter'. This is especially clear from a comparison—for either non-monetary demand and supply shocks, or money supply shocks—between the structural loops in 1946:4 and in 2006:4. This 'hump-shaped' pattern in the evolution of three structural loops out of four over the post-WWII era will play a crucial role, in Section 5, in an assessment of the role played by globalisation in shaping the evolution of the Phillips trade-off.

Summing up, an analysis of the structural Phillips loops therefore highlights—in a way which is significantly starker than for the reduced-form loops—the uniqueness

²⁷The methodology is briefly described in Benati (2008).

²⁸It is important to remember that we are here imposing the sign restrictions both on impact and for the subsequent two quarters, so that the negative slope in response to a money supply shock necessarily pertains to quarters beyond this horizon.

of the Great Inflation episode, and the extent to which it dislocated the Phillips trade-off compared even with just the rest of the post-WWII era.

Let's now turn to the issue of the (supposed) impact of globalisation on the slope of the Phillips trade-off.

5 Globalisation and the Evolution of the Phillips Trade-Off

As mentioned in the introduction, in recent years several papers have argued that the flattening of the Phillips trade-off which has taken place over the last two decades largely reflects the impact of globalisation. The key idea here is that an increase in trade openness weakens the link between inflation and *domestic* economic activity, thus causing a flattening of the trade-off.

If such conjecture were correct, we should expect to find a strong link between the previously documented evolution of the slope of the Phillips trade-off and measures of the extent of openness of the U.S. economy, such as the ratio between the sum of exports and imports and GDP. Figure 11 shows, for the period following World War I^{29} the evolution of (i) the sum of nominal exports and imports over nominal GNP, based on the data found in Appendix B of Balke and Gordon (1986), and (ii) the sum of nominal exports over nominal GDP, based on Table 1.1.5 of the *NIPA*. Overall, empirical evidence provides essentially *no* support to the notion that the evolution of the extent of openness of the U.S. economy may have played more than a secondary role in shaping the evolution of the Phillips trade-off. In particular,

- whereas the interwar period was characterised by a significant decrease in the extent of trade openness—first in the immediate aftermath of World War I, and then following the Smoot–Hawley Tariff Act of June 1930—as we previously mentioned both the reduced-form and the structural Phillips trade-offs exhibited, during those years, very little change. If the 'globalisation conjecture' were correct, a collapse in trade openness from about 20 per cent to below 10 per cent *should* be expected to have some impact on the Phillips trade-off. Evidence from the interwar period, however, appears to clearly reject such a notion.
- Further, as for the post-WWII period there appears to be, once again, little connection between the extent of trade openness and the slope of the Phillips trade-off. In particular, *first*, the significant increase in trade openness during the 1970s should have led—according to this view of the world—to a *decrease*, rather than an increase, in the slope of the trade-off; *second*, whereas the slope of the trade-off exhibits a hump-shaped pattern over the post-WWII period, trade

 $^{^{29}}$ We were not able to find data on exports and imports for the pre-1919 period.

openness displays a very broad upward trend, although with some significant temporary oscillations.

To be sure, our evidence does not disprove that globalisation and the extent of trade openness may have exterted some impact on the slope of the Phillips trade-off. What it does clearly suggest, however, is that this effect has been at most of second order, and that the historical evolution of the Phillips trade-off has been dominated by factors other than the extent of trade openness.

6 Conclusions

In this paper we have characterised the evolution of both the reduced-form and the structural U.S. unemployment-inflation trade-off since the Gold Standard era via a Bayesian time-varying parameters (structural) VAR. The Great Inflation episode appears as historically unique along several dimensions. In particular, the shape of the 'Phillips loop'—which is defined in terms of the impulse-response functions of inflation and unemployment's deviations from equilibrium—in response to both reduced-form unemployment shocks, and especially all structural innovations except money demand shocks, appears to have been clearly out of line with the rest of the sample period. We have related changes in the loop in response to reducedform unemployment innovations to changes in the persistence of inflation's deviations from equilibrium, and therefore, ultimately, to the anchoring of inflation expectations. During the Great Depression, on the other hand, the Phillips trade-off did not exhibit any peculiar qualitative feature, so that, when seen through these lenses, the 1930s only stand out because of the sheer size of the macroeconomic fluctuation. Overall, the long-term evolution of the Phillips trade-off broadly reflects the evolution of U.S. monetary regimes since the end of the XIX century, and their impact on the anchoring (or de-anchoring) of inflation expectations. Finally, although our results do not rule out a possible impact of globalisation on the slope of the Phillips trade-off in recent years, they clearly suggest that the historical evolution of the Phillips trade-off has been dominated by factors other than the extent of trade openness.

References

- BACON, R. (1973): "The Phillips Curve: Another Forerunner," *Economica*, 40(159), 314–15.
- BALKE, N., AND R. GORDON (1986): "Appendix B: Historical Data," in Gordon, R.J., ed. (1986), The American Business Cycle: Continuity and Changes, The University of Chicago Press.
- BARRO, R. (1982): "United States Inflation and the Choice of a Monetary Standard," in Hall, R.E., ed., 'Inflation: Causes and Effects', University of Chicago Press.
- BARRO, R. J. (1979): "Money and the Price Level Under the Gold Standard," Economic Journal, 89(353), 13–33.
- BAXTER, M., AND R. KING (1999): "Approximate Band-Pass Filters for Economic Time Series: Theory and Applications," *Review of Economics and Statistics*, 81(4), 575–593.
- BENATI, L. (2007): "The Time-Varying Phillips Correlation," Journal of Money, Credit and Banking, 39(5), 1275–1283.
- (2008): "The Great Moderation in the United Kingdom," Journal of Money, Credit, and Banking, 39(1), 121–147.
- BERNANKE, B., M. GERTLER, AND M. WATSON (1997): "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 1997(1), 91–157.
- BERNANKE, B. S. (2007): "Inflation Expectations and Inflation Forecasting," Speech given at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts July 10, 2007.
- BORIO, C., AND A. FILARDO (2007): "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation," *BIS Working Papers No 227.*
- BRAINARD, W., AND G. PERRY (2000): "Making Policy in a Changing World," in Economic Events, Ideas, and Policies: The 1960s and After, Perry, G.L., and Tobin, J., eds., Brookings Institution Press.
- CARTER, C. K., AND R. P. KOHN (2004): "On Gibbs Sampling for State Space Models," *Biometrika*, 81, 541–553.
- CHRISTIANO, L., AND T. FITZGERALD (2003): "The Band-Pass Filter," International Economic Review, 44(2), 435–465.

COGLEY, T., AND T. J. SARGENT (2002): "Evolving Post-WWII U.S. Inflation Dynamics," in B. Bernanke and K. Rogoff, eds. (2002), NBER Macroeconomics Annuals 2001.

(2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.," *Review of Economic Dynamics*, 8(April), 262–302.

- COGLEY, T. W., G. E. PRIMICERI, AND T. J. SARGENT (2008): "Inflation-Gap Persistence in the U.S.," *NBER Working Paper 13749, January 2008.*
- (2010): "Inflation-Gap Persistence in the U.S.," *American Economic Jour*nal: Macroeconomics, forthcoming.
- CORANA, A., M. MARCHESI, C. MARTINI, AND S. RIDELLA (1987): "Minimizing Multimodal Functions of Continuous Variables with the Simulated Annealing Algorithm," *ACM Transactions on Mathematical Software*, 13, 262–80.
- FISHER, I. (1926): "A Statistical Relation between Unemployment and Price Changes," *International Labor Review*, June, 785–92.
- FRIEDMAN, M. (1977): "Nobel Lecture: Inflation and Unemployment," Journal of Political Economy, 85(3), 451–472.
- GOFFE, W. L., G. FERRIER, AND J. ROGERS (1994): "Global Optimization of Statistical Functions with Simulated Annealing," *Journal of Econometrics*, 60,(1/2), 65–99.
- HARVEY, A. (1989): Forecasting, Structural Time Series Models, and the Kalman Filter. Cambridge University Press.
- JACQUIER, E., N. G. POLSON, AND P. ROSSI (1994): "Bayesian Analysis of Stochastic Volatility Models," Journal of Business and Economic Statistic, 12, 371–418.
- KIM, C. J., AND C. NELSON (2000): *State-Space Models with Regime Switching*. Cambridge, Mass., The MIT Press.
- KING, R., AND M. WATSON (1994): "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," Carnegie-Rochester Conference Series on Public Policy.
- KING, T., AND J. MORLEY (2007): "In Search of the Natural Rate of Unemployment," *Journal of Monetary Economics*, 54(March), 550–564.
- KOOP, G., M. H. PESARAN, AND S. POTTER (1996): "Impulse response analysis in nonlinear multivariate models," *Journal of Econometrics*, 74(1), 119–147.

LEBERGOTT, S. (1964): Manpower In Economic Growth. New York, McGraw Hill.

- LUCAS, R. E. (1972a): "Econometric Testing of the Natural Rate Hypothesis," in Eckstein, O., ed. The Econometrics of Price Determination, Washington, D.C., Board of Governors of the Federal Reserve System, pp. 50–59.
 - (1972b): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4(April 1972), 103–24.
 - (1973): "Some International Evidence on Output-Inflation Trade-Offs," *American Economic Review*, 63(3), 326–334.
- LUCAS, R. E. (1976): "Econometric Policy Evaluation: A Critique," Carnegie-Rochester Conference Series on Public Policy, 1, 19–46.
- MARTIN, W. M. (1965): "Statement Before the Joint Economic Committee," *February 26, 1965.*
- MISHKIN, F. (2007): "Inflation Dynamics," NBER Working Paper n. 13147, June 2007.
- PHILLIPS, A. (1958): "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, 25(100), 283–99.
- PRIMICERI, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *The Review of Economic Studies*, 72, 821–852.
- RAVN, M. O., AND H. UHLIG (2002): "On Adjusting the HP-Filter for the Frequency of Observations," *Review of Economics and Statistics*, 84(1), 371–380.
- ROBERTS, J. M. (1993): "The Sources of Business Cycles: A Monetarist Interpretation," International Economic Review, 34(4), 923–934.
- ROMER, C. (1992): "What Ended the Great Depression?," Journal of Economic History, 52(4), 757–784.
- RUBIO-RAMIREZ, J., D. WAGGONER, AND T. ZHA (2005): "Regime Changes in the Euro Area," Federal Reserve Bank of Atlanta, mimeo.
- STOCK, J., AND M. WATSON (1999): "Business Cycle Fluctuations in U.S. Macroeconomic Time Series," in Taylor, J. B., and Woodford, M., eds., Handbook of Macroeconomics, Amsterdam, North Holland.
- TINBERGEN, J. (1937): An Econometric Approach to Business Cycle Problems. Amsterdam.
- VOLCKER, P. A. (1979): "A Time of Testing," Remarks by P.A. Volcker, Chairman, Board of Governors of the Federal Reserve System, before the American Bankers Association in New Orleans, Louisiana, 9 October 1979.

A The Data

Quarterly seasonally adjusted series for the GNP deflator and real GNP have been constructed by linking the GNP deflator and real GNP series from Balke and Gordon (1986), appendix B, Table 2, which are available from 1875:1 to 1983:4, to GNPDEF ('Gross National Product: Implicit Price Deflator, Seasonally Adjusted, Quarterly, Index 2000=100') and GNPC96 ('Real Gross National Product, Seasonally Adjusted Annual Rate, Quarterly, Billions of Chained 2000 Dollars'), respectively, which are both available since 1947:1, and are both from the U.S. Department of Commerce, Bureau of Economic Analysis.

A monthly seasonally unadjusted series for the commercial paper rate has been constructed by linking the commercial paper rate series from the *NBER Historical Database* ('NBER series: 13002; U.S. Commercial Paper Rates, New York City'), available for the period January 1857-December 1971; CP3M ('3-Month Commercial Paper Rate, Monthly, Percent, Average of offering rates on commercial paper placed by several leading dealers for firms whose bond rating is AA or equivalent, quoted on a discount basis. Averages of daily figures.') from the Board of Governors of the Federal Reserve System, which is available for the period April 1971-August 1997; and CPF3M ('3-Month AA Financial Commercial Paper Rate, Averages of Business Days, Discount Basis') from the Board of Governors of the Federal Reserve System, which is available since January 1997. Specifically, the linked series consists of the NBER series up to March 1971, of CP3M from April 1971 until December 1996, and of CPF3M after that. Finally, a quarterly series was obtained by taking averages within the quarter of the original monthly series.

A quarterly seasonally adjusted series for M2 has been constructed by linking the quarterly M2 series from Balke and Gordon (1986), appendix B, Table 2, which is available from 1875:1 to 1983:4, to M2SL ('M2 Money Stock, Board of Governors of the Federal Reserve System, H.6 Money Stock Measures, Seasonally Adjusted, Monthly, Billions of Dollars'), which is available since January 1959, and has been converted to the quarterly frequency by taking averages within the quarter.

A monthly seasonally adjusted series for the unemployment rate for the period April 1929-April 2009 has been constructed by linking the seasonally adjusted unemployment rate series from the *NBER Historical Database* ('NBER series: 08292; unemployment rate, seasonally adjusted, area covered: U.S. '), available for the period April 1929-June 1942; a seasonally unadjusted unemployment rate series computed as the ratio between series for the unemployment level and the level of the labor force (both in millions) from the *NBER Historical Database* ('NBER series: 08084 and 08325, respectively'), available for the period March 1940-December 1957 and March 1940-June 1968, respectively, which has then been seasonally adjusted *via* the ARIMA X-12 procedure as implemented in *EViews*; and UNRATE ('Civilian unemployment rate, persons 16 years of age and older, seasonally adjusted, monthly, percent') from the *U.S. Department of Labor, Bureau of Labor Statistics*. The resulting monthly easonally adjusted series has been converted to the quarterly frequency by taking averages within the quarter, and has been linked to a quarterly series for the period 1890:1-1929:2 which has been interpolated a described in Appendix B.

B A Quarterly Interpolated Series for the U.S. Unemployment Rate, 1890:1-1929:2

As mentioned in Appendix A, monthly data for the U.S. unemployment rate are only available starting from April 1929. To the very best of our knowledge, for the years before 1929 the only available data is a single annual series for the rate of unemployment starting in 1890, which can be found (e.g.) in Lebergott (1964). We interpolated this series to the quarterly frequency, and we then linked it to monthly linked unemployment rate series described in Appendix A (which we converted to the quarterly frequency by taking averages within the quarter), via the following procedure along the lines of Bernanke, Gertler, and Watson (1997).

Let U_{τ}^{A} and U_{t}^{Q} be the annual unemployment rate for year τ and the quarterly unemployment rate for quarter t, respectively, with

$$U_{\tau}^{A} \equiv \frac{1}{4} \left[U_{t}^{Q} + U_{t-1}^{Q} + U_{t-2}^{Q} + U_{t-3}^{Q} \right]$$
(B1)

where t, in (B1), is the last quarter of year τ . Let U_t^Q be defined as

$$U_t^Q \equiv U_{N,t}^Q + U_{C,t}^Q \tag{B2}$$

where $U_{N,t}^Q$ and $U_{C,t}^Q$ are the 'natural' and cyclical components of the unemployment rate, respectively. In line with, e.g., Kim and Nelson (2000), we postulate $U_{N,t}^Q$ and $U_{C,t}^Q$ to evolve according to a random-walk specification,

$$U_{N,t}^Q \equiv U_{N,t-1}^Q + u_t \qquad u_t \sim N(0, \sigma_u^2)$$
(B3)

and a 'generalised Okun's law',

$$U_{C,t}^{Q} = \sum_{k=-L}^{L} \alpha_{k} y_{C,t-k}^{Q} + v_{t}$$
(B4)

respectively, where $y_{C,t}^Q$ is the output gap, and

$$v_t = \rho_v v_{t-1} + \tilde{v}_t \qquad \qquad \tilde{v}_t \sim N(0, \sigma_v^2) \tag{B5}$$

As a reasonable proxy for the output gap, in what follows we use the HP-filtered logarithm of real GNP. An important point to stress is that, given the presence of

the autocorrelated disturbance v_t in (B5), for the interpolation procedure to work well we do not need to have a *perfect* measure of the output gap (which, in practice, would also be impossible to have), but only a a reasonably good proxy, and under this respect the HP-filtered logarithm of real GNP clearly foots the bill.

Model (B1)-(B5) can be trivially cast in state-space form, with (B1) as observation equation and (B2)-(B5) rearranged into a single transition equation in vector and matrix form. The model can then be estimated *via* maximum likelihood, by computing the likelihood function *via* the prediction-error decomposition formula found (e.g.) in Harvey (1989) and Kim and Nelson (2000), and maximising it numerically with respect to σ_u^2 , σ_v^2 , ρ_v , and the α_k 's—specifically, numerical maximisation is performed *via* simulated annealing.³⁰

The obvious problem with this approach is that σ_u^2 and σ_v^2 cannot be separately identified. We therefore consider a grid of five values for σ_u^2 , from 0.1 to 0.5, and conditional on each value in the grid we estimate the remaining parameters and we compute the quarterly interpolated unemployment rate series as one of the unobserved states in the state-space representation of the model. Estimation is performed for the period 1890-1929. In estimation we impose the constraint that the interpolated quarterly unemployment series for the first quarter for which actual unemployment data are available, 1929Q2, and and for the subsequent two quarters, be equal to the actual quarterly unemployment rate.

Figure 12 reports the results for the five values of $\sigma_u^{2,31}$ On the other hand, we do not report the maximum likelihood estimates of the model's parameters as, for the present purposes, they are essentially irrelevant, but they are available from the author upon request.³² As the figure clearly shows, although different values of σ_u^2 produce (quite obviously) different estimates of the natural rate—with an increase in σ_u^2 being associated with a more volatile natural rate—the interpolated quarterly unemployment rate series conditional on different values of σ_u^2 are remarkably close

 31 Specifically, both the estimates of the natural rate of unemployment shown in the first row of Figure 15, and the interpolated quarterly unemployment rate series shown in the second row, are the *one-sided* estimates generated by the Kalman filter conditional on the maximum likelihood estimates of the model's parameters.

³²The only feature of the parameters' estimates which might deserve to be mentioned is that, for each single value of σ_u^2 , the α_k 's MLE estimates clearly suggest the output gap to lead the cyclical component of the unemployment rate by several quarters. (To put it differently, estimates of the α_k 's are significantly negative for k>0 and essentially zero for k<0.)

³⁰Following Goffe, Ferrier, and Rogers (1994) we implement simulated annealing via the algorithm proposed by Corana, Marchesi, Martini, and Ridella (1987), setting the key parameters to $T_0=100,000, r_T=0.9, N_t=5, N_s=20, \epsilon=10^{-6}, N_{\epsilon}=4$, where T_0 is the initial temperature, r_T is the temperature reduction factor, N_t is the number of times the algorithm goes through the N_s loops before the temperature starts being reduced, N_s is the number of times the algorithm goes through the function before adjusting the stepsize, ϵ is the convergence (tolerance) criterion, and N_{ϵ} is number of times convergence is achieved before the algorithm stops. Finally, initial conditions were chosen stochastically by the algorithm itself, while the maximum number of functions evaluations, set to 1,000,000, was never achieved.

to one another. In our work we consider the estimate conditional on $\sigma_u^2=0.3$. An interesting feature of the quarterly interpolated series which is not apparent from the original annual series is that, closely mirroring fluctuations in the HP-filtered log of real GNP, it exhibits several comparatively high-frequency fluctuations towards the very end of the XIX century, which on the contrary were smoothed out, and therefore hidden, in the annual data.

C Details of the Markov-Chain Monte Carlo Procedure

We estimate (1)-(11) via Bayesian methods. The next two subsections describe our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data, while the third section discusses how we check for convergence of the Markov chain to the ergodic distribution.

C.1 Priors

For the sake of simplicity, the prior distributions for the initial values of the states— θ_0 , α_0 , h_0 , and q_0 —which we postulate all to be normal, are assumed to be independent both from one another, and from the distribution of the hyperparameters. In order to calibrate the prior distributions for θ_0 , α_0 , h_0 , and q_0 we estimate a time-invariant version of (1) based on the first 10 years of data, and we set

$$\theta_0 \sim N\left[\hat{\theta}_{OLS}, 4 \cdot \hat{V}(\hat{\theta}_{OLS})\right]$$
 (C1)

As for α_0 and h_0 we proceed as follows. Let $\hat{\Sigma}_{OLS}$ be the estimated covariance matrix of ϵ_t from the time-invariant VAR, and let C be the lower-triangular Choleski factor of $\hat{\Sigma}_{OLS}$ —i.e., $CC' = \hat{\Sigma}_{OLS}$. We set

$$\ln h_0 \sim N(\ln \mu_0, 10 \times I_3) \tag{C2}$$

where μ_0 is a vector collecting the squared elements on the diagonal of C. We then divide each column of C by the corresponding element on the diagonal—let's call the matrix we thus obtain \tilde{C} —and we set

$$\alpha_0 \sim N[\tilde{\alpha}_0, \tilde{V}(\tilde{\alpha}_0)] \tag{C3}$$

where $\tilde{\alpha}_0$ —which, for future reference, we define as $\tilde{\alpha}_0 \equiv [\tilde{\alpha}_{11,0}, \tilde{\alpha}_{21,0}, \tilde{\alpha}_{32,0}]'$ —is a vector collecting all the non-zero and non-one elements of \tilde{C}^{-1} (i.e, the elements below the diagonal), and its covariance matrix, $\tilde{V}(\tilde{\alpha}_0)$, is postulated to be diagonal, with each individual (j,j) element equal to 10 times the absolute value of the corresponding

j-th element of $\tilde{\alpha}_0$. Such a choice for the covariance matrix of α_0 is clearly arbitrary, but is motivated by our goal to scale the variance of each individual element of α_0 in such a way as to take into account of the element's magnitude.

As for q_0 we proceed as follows. Let Q_0 be the prior matrix for the extent of random-walk drift of the VAR's parameters (that is, the random walks collected in the vector θ_t) that we would use if we were working with a traditional Bayesian time-varying parameters VAR with a *constant* extent of random-walk drift over the sample. We set $Q_0 = \gamma \times \hat{\Sigma}_{OLS}$, with $\gamma = 1.0 \times 10^{-4}$, the same value used in Primiceri (2005), and a relatively 'conservative' prior for the extent of drift compared (e.g.) to the 3.5×10^{-4} used by Cogley and Sargent (2005). We set

$$\ln q_0 \sim N(10^{-2} \times \ln \bar{q}_0, 10 \times I_{N \cdot (1+Np)}) \tag{C4}$$

where \bar{q}_0 is a vector collecting the elements on the diagonal of Q_0 .

Turning to the hyperparameters, we postulate independence between the parameters corresponding to the matrices S and Z—an assumption we adopt uniquely for reasons of convenience—and we make the following, standard assumptions. The two blocks of S are assumed to follow inverted Wishart distributions, with prior degrees of freedom set equal to the minimum allowed, respectively, 2 and 3:

$$S_1 \sim IW\left(\bar{S}_1^{-1}, 2\right) \tag{C5}$$

$$S_2 \sim IW\left(\bar{S}_2^{-1}, 3\right) \tag{C6}$$

As for \bar{S}_1 and \bar{S}_2 , we calibrate them based on $\tilde{\alpha}_0$ in (C3) as $\bar{S}_1=10^{-3} \times |\tilde{\alpha}_{0,11}|$, $\bar{S}_2=10^{-3}\times \text{diag}([|\tilde{\alpha}_{0,21}|, |\tilde{\alpha}_{0,31}|]')$. Such a calibration is consistent with the one we adopted for Q, as it is equivalent to setting \bar{S}_1 and \bar{S}_2 equal to 10^{-4} times the relevant diagonal block of $\tilde{V}(\tilde{\alpha}_0)$ in (C3). As for the variances of the innovations to the stochastic volatilities for the VAR's reduced-form shocks, we follow Cogley and Sargent (2002, 2005) and we postulate an inverse-Gamma distribution for the elements of Z_{ν} ,

$$\sigma_{\nu,i}^2 \sim IG\left(\frac{10^{-4}}{2}, \frac{1}{2}\right) \tag{C7}$$

Finally, as for the variances of the innovations to the stochastic volatilities for the VAR's random-walk parameters' innovations, we postulate an inverse-Gamma distribution for the elements of Z_{ω} ,

$$\sigma_{\omega,i}^2 \sim IG\left(\frac{10^{-4}}{2}, \frac{5}{2}\right) \tag{C8}$$

(C8) implies that the prior for $\sigma_{\omega,i}^2$ has the same mean as in Cogley *et al.* (2009), but it has a smaller variance.

C.2 Simulating the posterior distribution

We simulate the posterior distribution of the hyperparameters and the states conditional on the data via the following MCMC algorithm, combining elements of Primiceri (2005) and Cogley and Sargent (2002, 2005). In what follows, x^t denotes the entire history of the vector x up to time t—i.e. $x^t \equiv [x'_1, x'_2, ..., x'_t]'$ —while T is the sample length.

(a) Drawing the elements of θ_t Conditional on Y^T , α^T , and H^T , the observation equation (1) is linear, with Gaussian innovations and a known covariance matrix. Following Carter and Kohn (2004), the density $p(\theta^T | Y^T, \alpha^T, H^T, V)$ can be factored as

$$p(\theta^{T}|Y^{T}, \alpha^{T}, H^{T}, V) = p(\theta_{T}|Y^{T}, \alpha^{T}, H^{T}, V) \prod_{t=1}^{T-1} p(\theta_{t}|\theta_{t+1}, Y^{T}, \alpha^{T}, H^{T}, V)$$
(C9)

Conditional on α^T , H^T , and V, the standard Kalman filter recursions nail down the first element on the right hand side of (C9), $p(\theta_T|Y^T, \alpha^T, H^T, V) = N(\theta_T, P_T)$, with P_T being the precision matrix of θ_T produced by the Kalman filter. The remaining elements in the factorization can then be computed via the backward recursion algorithm found, e.g., in Kim and Nelson (2000), or Cogley and Sargent (2005, appendix C.2.1). Given the conditional normality of θ_t , we have

$$\theta_{t|t+1} = \theta_{t|t} + P_{t|t} P_{t+1|t}^{-1} \left(\theta_{t+1} - \theta_t \right)$$
(C10)

$$P_{t|t+1} = P_{t|t} - P_{t|t}P_{t+1|t}^{-1}P_{t|t}$$
(C11)

which provides, for each t from T-1 to 1, the remaining elements in (1), $p(\theta_t|\theta_{t+1}, Y^T, \alpha^T, H^T, V) = N(\theta_{t|t+1}, P_{t|t+1})$. Specifically, the backward recursion starts with a draw from $N(\theta_T, P_T)$, call it $\tilde{\theta}_T$ Conditional on $\tilde{\theta}_T$, (C10)-(C11) give us $\theta_{T-1|T}$ and $P_{T-1|T}$, thus allowing us to draw $\tilde{\theta}_{T-1}$ from $N(\theta_{T-1|T}, P_{T-1|T})$, and so on until t=1.

(b) Drawing the elements of α_t Conditional on Y^T , θ^T , and H^T , following Primiceri (2005), we draw the elements of α_t as follows. Equation (1) can be rewritten as $A_t \tilde{Y}_t \equiv A_t (Y_t - X'_t \theta_t) = A_t \epsilon_t \equiv u_t$, with $\operatorname{Var}(u_t) = H_t$, namely

$$\tilde{Y}_{2,t} = -\alpha_{21,t}\tilde{Y}_{1,t} + u_{2,t} \tag{C12}$$

$$\tilde{Y}_{3,t} = -\alpha_{31,t}\tilde{Y}_{1,t} - \alpha_{32,t}\tilde{Y}_{2,t} + u_{3,t}$$
(C13)

—plus the identity $\tilde{Y}_{1,t} = u_{1,t}$ —where $[\tilde{Y}_{1,t}, \tilde{Y}_{2,t}, \tilde{Y}_{3,t}]' \equiv \tilde{Y}_t$. Based on the observation equations (C12)-(C13), and the transition equation (9), the elements of α_t can then be drawn by applying the same algorithm we described in the previous paragraph separately to (C12) and (C13). The assumption that S has the block-diagonal structure (11) is in this respect crucial, although, as stressed by Primiceri (2005, Appendix D), it could in principle be relaxed.

(c) Drawing the elements of H_t Conditional on Y^T , θ^T , and α^T , the orthogonalised innovations $u_t \equiv A_t(Y_t - X'_t \theta_t)$, with $\operatorname{Var}(u_t) = H_t$, are observable. Following Cogley and Sargent (2002), we then sample the $h_{i,t}$'s by applying the univariate algorithm of Jacquier, Polson, and Rossi (1994) element by element.³³

(d) Drawing the elements of Q_t Conditional on θ^T , the innovations $\eta_t = \theta_t - \theta_{t-1}$, with $\operatorname{Var}(\eta_t) = Q_t$, are observable, and, along the lines of point (c), we therefore sample the $q_{j,t}$'s by applying the univariate algorithm of Jacquier, Polson, and Rossi (1994) element by element.

(e) Drawing the hyperparameters Finally, conditional on Y^T , θ^T , H^T , and α^T , the innovations to θ_t , α_t , the $h_{i,t}$'s and the $q_{i,t}$'s are observable, which allows us to draw the hyperparameters—the elements of S_1 , S_2 , and the $\sigma^2_{\nu,i}$ and the $\sigma^2_{\omega,i}$ —from their respective distributions.

Summing up, the MCMC algorithm simulates the posterior distribution of the states and the hyperparameters, conditional on the data, by iterating on (a)-(e). In what follows we use a burn-in period of 50,000 iterations to converge to the ergodic distribution, and after that we run 10,000 more iterations sampling every 10th draw in order to reduce the autocorrelation across draws.³⁴

³³For details, see Cogley and Sargent (2005, Appendix B.2.5).

³⁴In this we follow Cogley and Sargent (2005). As stressed by Cogley and Sargent (2005), however, this has the drawback of 'increasing the variance of ensemble averages from the simulation'.













35



























