

On the Sources of the Great Moderation*

Jordi Galí[†] Luca Gambetti[‡]

June 2008
(first draft: September 2006)

Abstract

The remarkable decline in macroeconomic volatility experienced by the U.S. economy since the mid-80s (the so-called Great Moderation) has been accompanied by large changes in the patterns of comovements among output, hours and labor productivity. Those changes are reflected in both conditional and unconditional second moments as well as in the impulse responses to identified shocks. Among other changes, our findings point to (i) an increase in the volatility of hours relative to output, (ii) a shrinking contribution of non-technology shocks to output volatility, and (iii) a change in the cyclical response of labor productivity to those shocks. That evidence suggests a more complex picture than that associated with "good luck" explanations of the Great Moderation.

JEL Code: E32

Keywords: Great Moderation, structural VARs, technology shocks, structural change, monetary policy rules, labor hoarding.

*We are grateful for comments and suggestions to Régis Barnichon, Olivier Blanchard, Steve Davis, Davide Debortoli, Luigi Guiso, Sharon Kozicki, Steve Nickell, Gabriel Pérez-Quirós, Valerie Ramey, Thijs van Rens, Todd Walker, Mark Watson, two anonymous referees, and participants in numerous seminars and conferences. The authors acknowledge the financial support from Ministerio de Educación y Ciencia (grants SEJ 2005-01124 and SEJ-2004-21682-E, respectively), the Barcelona GSE Research Network, and the Generalitat de Catalunya.

[†]CREI and Universitat Pompeu Fabra

[‡]Universitat Autònoma de Barcelona (UAB)

1 Introduction

A large body of empirical research has provided evidence of a substantial decline in the volatility of most U.S. macroeconomic time series over the postwar period. That phenomenon, which has also been experienced by other industrialized economies, has come to be known as "the Great Moderation."¹

Table 1 reminds us of the magnitude of the volatility decline associated with the Great Moderation. It shows the standard deviation for two indicators of economic activity, (log) GDP and (log) non-farm business output, before and after 1984, a date which is generally viewed as the starting point of the period of enhanced stability in the U.S. economy. We use quarterly data covering the period 1948:I-2005:IV. Both variables are normalized by the size of the working age population.² We report evidence for both the first-differenced and band-pass filtered transformations of each variable.³ As shown in the Table, and for the two variables and transformations considered, the standard deviation for the post-84 period is less than half that corresponding to the pre-84 period. Tests of equality of the variance across sub-periods reject that null hypothesis in all cases with a minuscule p -value.

While there is widespread consensus among macroeconomists on the existence and rough timing of the Great Moderation, its interpretation is still controversial. The various hypotheses put forward in the literature can be thought of as falling under two broad categories. The first view, often referred to as the "good luck" hypothesis, suggests that the greater macroeconomic stability of the past twenty years is largely the result of smaller shocks impinging on the economy, with structural changes having played at most a secondary role.⁴ A

¹Early papers on the Great Moderation include those of Kim and Nelson (1999), McConell and Pérez-Quirós (2000), and Blanchard and Simon (2001). A survey of the literature, as well as a discussion of alternative interpretations, can be found in Stock and Watson (2002). Stock and Watson (2005) and Cecchetti, Flores-Lagunes and Krause (2006) present and discuss some international evidence.

²Below we provide a detailed description of the data and its sources.

³We use the approximate band-pass filter of Baxter and King (1999). Following widespread practice, we identify the cyclical component of fluctuations as that corresponding to an interval between 6 and 32 quarters.

⁴See, e.g., Justiniano and Primiceri (2006) and Arias, Hansen and Ohanian (2006) for examples of papers making a case for smaller shocks as an explanation for the volatility

second view attributes instead the reduction in aggregate volatility to changes in the economy's structure and/or in the way policy has been conducted.⁵

In the present paper we provide evidence on some of the changes experienced by the U.S. economy over the postwar period and, in particular, around the time of the volatility break associated with the Great Moderation. Our evidence is based on (i) the observed comovements among output, hours and productivity, (ii) the identification of the sources of those comovements, and (iii) the study of their changes over time. The focus on those three variables is motivated by their central role in existing theories of the business cycle and the frequent use of their comovements in efforts to sort out among competing theories.⁶ We believe that such evidence can be useful in assessing the merits of alternative explanations for the Great Moderation, including the two broad hypotheses mentioned above.

Much of the evidence reported below is based on an estimated structural vector autoregression (SVAR) with time-varying coefficients and stochastic volatility, applied to (log) labor productivity and (log) hours. Following Galí (1999) we interpret variations in those variables –as well as in (log) output, which is given by their sum– as the result of two types of shocks impinging on the economy: technology and non-technology shocks. Technology shocks are assumed to be the source of the unit root in labor productivity; accordingly, they are identified as the only shocks that may have a permanent effect on that variable. Following Cogley and Sargent (2005), Primiceri (2005) and Benati and Mumtaz (2007), our estimated model allows for time-varying coefficients. The latter feature makes it possible to uncover, in a flexible way, changes over time in unconditional and conditional comovements, in the responses of different variables to each type of shock, as well as the contribution of the different shocks to the decline in volatility. Furthermore, as emphasized in Gambetti (2006), the use of

decline of the past two decades.

⁵Such explanations include better monetary policy (e.g. Clarida, Galí and Gertler (2000)), improvements in inventory management (e.g. Kahn, McConnell and Perez-Quirós (2002)), financial innovation and better risk sharing (e.g. Dynan, Elmendorf and Sichel (2006)), and the optimal response of production and inventories policies to a decline in the persistence of automobile sales (Ramey and Vines (2006)).

⁶Christiano and Eichenbaum (1992), Hansen and Wright (1992), and Galí (1999) are examples of work in that tradition.

time-varying coefficients overcomes the potential bias caused by the presence of significant low frequency comovements between productivity growth and hours in postwar U.S. data, a problem first diagnosed by Fernald (2008).⁷

In a way consistent with the literature, we uncover a large, (seemingly) permanent, decline in the volatility of output around the mid-80s. But the analysis of other statistics point to a more complex picture, as implied by the following findings:

- While the volatility of hours and labor productivity has also declined in absolute terms, it has risen considerably *relative* to the volatility of output. Furthermore, the timing and pattern of decline in the volatility of those three variables display considerable differences.
- Several correlations display remarkable changes. In particular, the correlation of hours with labor productivity has experienced a large decline, shifting from values close to zero in the early postwar period to large negative values in more recent times. Interestingly, and as stressed in Stiroh (2008), much of that decline appears to be concentrated in the 80s, and tracks to a large extent the fall in output volatility. Similarly, when BP-filtered data are used, the correlation of output with labor productivity shows a substantial decline, from positive values to values close to zero.⁸ The size of that change is weaker (though still statistically significant) when a first-difference transformation of the two series is used instead.
- According to our time-varying SVAR, the Great Moderation can be largely explained by a sharp fall in the contribution of non-technology shocks to the variance of output, both in absolute and relative terms. By contrast, the contribution of technology shocks to output volatility appears to have

⁷Fernald (2008) makes a forceful case for the important role played by the positive low frequency comovement between labor productivity growth and (log) hours per capita in accounting for the conflicting evidence in Galí (1999) and Christiano, Eichenbaum and Vigfusson (2003).

⁸Barnichon (2006), in work conducted independently, stresses the change in the correlation between unemployment and labor productivity, as well as the decline in the procyclicality of the latter variable.

remained largely stable in absolute terms (and has thus increased in relative terms).

- Several *conditional* correlations also display large changes over the postwar period. Most remarkably, the correlation of labor productivity with both output and hours conditional on non-technology shocks shows a rapid decline starting in the early 1980s and accelerating in the 1990s. Such a decline reflects the sizable changes over time in the pattern of the response of labor productivity to non-technology shocks, as well as the smaller relative importance of those shocks. On the other hand, the correlation of hours with both output and labor productivity conditional on technology shocks displays sizable medium-run fluctuations, often shifting sign during particular episodes. Thus, for instance, it rises considerably during the second half of the 1970s (the oil shocks period) and the second half of the 1990s (the dotcom era). Those changes mirror to a large extent the pattern of the response of hours to technology shocks.
- Most of the key findings above are robust, at least qualitatively, to using an augmented specification of our time-varying SVAR based on Fisher (2006), which distinguishes between neutral and investment-specific technology shocks.

While our analysis, by its very nature, does not allow one to uncover the deep structural sources behind the Great Moderation and other changes experienced by the postwar U.S. economy, we believe it can still be helpful at ruling out some hypotheses and shedding light on the relative merits of alternative explanations for the Great Moderation, while imposing a minimal structure.

Thus, for instance, many of the findings listed above are clearly inconsistent with a "strong" version of the good luck hypothesis that attributes the Great Moderation to a (roughly) proportional decline in the variance of *all* relevant shocks, for that hypothesis would imply a counterfactual stability of relative standard deviations and unconditional correlations among macro variables.

Our evidence is also inconsistent with a weaker version of the same hypothesis, namely, one that attributes the decline in aggregate volatility to a reduction in the variance of *a subset* of the relevant shocks, since that explanation cannot account, by itself, for the changes over time in *conditional* second moments and the patterns of impulse responses.⁹ On the other hand, the observed variation in conditional second moments points to the existence of at least some structural changes influencing the joint dynamics of output, hours and productivity over the postwar period. The fact that the timing of some of those changes coincides with the onset of the Great Moderation is, at the very least, suggestive of some connection between the two.

In that regard, and as discussed in more detail below, our evidence is consistent with either a decline in the size of non-technology shocks as well as more effective countercyclical policies in response to those shocks. The hypothesis of a change in policy is reinforced when the variations in the responses to technology and non-technology shocks are considered jointly: some key features of those changes can in principle be explained by the adoption since the early 1980s of a monetary policy that focuses on the stabilization of inflation, for that policy would also tend to stabilize output in response to a variety of demand shocks, while accommodating the changes in potential output resulting from technology shocks. Furthermore, the gradual change in the response of labor productivity to non-technology shocks (with an eventual change in the sign of that response) is consistent with a declining importance of labor hoarding by firms, possibly as a consequence of better labor input management practices or more flexible labor markets (that make it less costly to hire and fire workers in response to changes in demand).

The remainder of the paper is organized as follows. Section 2 reports estimates of the standard deviations and correlations of output, hours and labor productivity and their changes over time. Section 3 introduces the time-varying

⁹Of course, under a view of the business cycle in which the latter is largely driven by a single shock—a view held by proponents of early RBC models—the distinction between the two versions of the good luck hypothesis is meaningless.

VAR approach used to estimate changes over time in conditional second moments and impulse responses, and presents the associated evidence. Section 4 presents the main empirical findings. Section 5 show the evidence based on the augmented SVAR model. Section 6 discusses possible interpretations and concludes.

2 The Labor Market and the Great Moderation: Basic Evidence

2.1 Changes in Volatilities

Table 2 summarizes the evidence on volatility changes in output, hours and labor productivity by showing their respective standard deviations for the pre-84 and post-84 periods, as well as the ratio between the two. On the right hand panel we also report the corresponding standard deviation relative to output, and the ratio of relative standard deviations between the two sub-periods. We use quarterly data covering the sample period 1948:I-2005:IV. All variables refer to the nonfarm business sector.¹⁰ Again, we report estimates for both first-differenced and BP-filtered data, after taking natural logarithms.

Turning to the main findings, we see that independently of the transformation used, all three variables considered have experienced a large (and highly significant) reduction in their volatility in the post-84 period. The size of that decline is, however, not proportional. Thus, the percent decline in the standard deviations of hours and labor productivity is not as large as that experienced by output, as reflected in the increase in their relative standard deviations shown in the last three columns of the table. That increase in the relative volatility of hours and productivity is our first piece of evidence pointing to the presence of changes beyond those that would result from a mere proportional scaling down

¹⁰We obtained our raw data from the USECON data base. The time series used include output in the nonfarm business sector (LXBO) and hours of all persons in nonfarm business (LXBH). Both variables were normalized by the civilian non-institutional population of 16 years and over (LNN). Labor productivity was computed as a the ratio between the output and hours measures mentioned above. The GDP measure used in table 1 was drawn from the same database, with the mnemonic GDPH.

of volatility in all variables.

2.2 Changes in Comovements

Next we turn to the examination of the comovements among labor market variables and their changes over time. For each pair of variables considered, Table 3 reports their estimated correlation in the pre-84 and post-84 sample periods, as well as the difference between the two. As above, evidence is reported for two different transformations of the data, the first-differenced and BP-filtered logarithms of the original variables.

As the statistics shown in Table 3 make clear, many of the estimated changes in comovements are large and highly significant. In particular, the cyclical behavior of labor productivity, measured by its comovement with either output or hours, has experienced a considerable decline. Thus, when we use output as the cyclical indicator of reference and the BP-filter as a detrending method, labor productivity becomes an (essentially) acyclical variable in the post-84 period. That result is considerably weaker, however, when we use first-differenced data, though the decline is still statistically significant. That finding is of substantial interest since the strong procyclicality of productivity was one of the empirical cornerstones of the technology-driven view of the business cycle endorsed by RBC theory.

When we take hours as a reference cyclical indicator, the change in the cyclical behavior of labor productivity is even more dramatic: we see that the behavior of labor productivity switches from being largely acyclical to being countercyclical, with the change in correlations being highly significant, independently of the transformation used. As emphasized by Stiroh (2008), that decline in the covariance between labor productivity and hours can explain, from an accounting point of view, a substantial fraction of the decline in output volatility.

Overall we view that variation in the pattern of correlations and relative standard deviations across sample periods as evidence against a *strong* version

of the good luck hypothesis, and reflecting instead changes in either the composition of shocks or in the structure and transmission mechanisms operating in the U.S. economy. In the remainder of the paper we try to enrich the evidence presented above along two dimensions. First, we use a flexible econometric framework that allows for continuous variations in the joint dynamics of labor market variables. This allows us to contrast the timing of changes in those dynamics with that of the Great Moderation. Secondly, we identify the role played by shocks of different nature as a source of those changes.

3 A VAR Model with Time-Varying Coefficients and Stochastic Volatility

The present section describes our baseline empirical model, which consists of an SVAR with time-varying coefficients. Though focusing on different variables, the specification of the reduced form time-varying VAR follows closely that in Primiceri (2005). Our identification of the structural shocks follows that in Galí (1999).

Let y_t and n_t denote, respectively, (log) output and (log) hours, both in per capita terms. We define $x_t \equiv [\Delta(y_t - n_t), n_t]$, and assume that the joint process for (log) labor productivity and (log, per capita) hours admits a time-varying VAR representation given by

$$x_t = A_{0,t} + A_{1,t} x_{t-1} + A_{2,t} x_{t-2} + \dots + A_{p,t} x_{t-p} + u_t \quad (1)$$

where $A_{0,t}$ is a vector of time-varying intercepts, and $A_{i,t}$, $i = 1, \dots, p$, are matrices of time-varying coefficients.¹¹ We assume that all the roots of the VAR polynomial lie outside the unit circle for all t ; i.e. the process is "locally stationary." The sequence of innovations $\{u_t\}$ follows a Gaussian white noise process with zero mean and time-varying covariance matrix Σ_t , and uncorrelated with all lags of x_t . Letting $A_t = [A_{0,t}, A_{1,t}, \dots, A_{p,t}]$, we define $\theta_t = \text{vec}(A_t')$

¹¹As stressed in Gambetti (2006), the presence of a time-varying intercept in the VAR absorbs the low frequency comovement between $\Delta(y_t - n_t)$ and n_t , thus overcoming the potential distortions in the estimates pointed out by Fernald (2008).

where $vec(\cdot)$ is the column stacking operator. Conditional on the roots of the associated VAR polynomial being outside the unit circle for all t , we assume θ_t evolves over time according to the process

$$\theta_t = \theta_{t-1} + \omega_t \quad (2)$$

where ω_t is a Gaussian white noise process with zero mean and constant covariance Ω , and independent of u_t at all leads and lags.

We model the time variation for Σ_t as follows. Let $\Sigma_t \equiv F_t D_t F_t'$ where F_t is lower triangular with ones in the main diagonal and D_t a diagonal matrix.¹² Let γ_t be a vector containing all the elements of F_t^{-1} below the diagonal, stacked by rows, and σ_t the vector of diagonal elements of D_t .

$$\gamma_t = \gamma_{t-1} + \zeta_t \quad (3)$$

$$\log \sigma_t = \log \sigma_{t-1} + \xi_t \quad (4)$$

where ζ_t and ξ_t are Gaussian white noise processes with zero mean and (constant) covariance matrices Ψ and Ξ , respectively. We assume that Ψ has a block diagonal structure, i.e. all the covariances between coefficients belonging to different equations are zero, and that Ξ is diagonal. Finally we assume that ξ_t, ζ_t , and ω_t are all mutually independent.

We assume that the vector of VAR innovations u_t is a (time-varying) linear transformation of the vector of underlying "structural" shocks $\varepsilon_t \equiv [\varepsilon_t^a, \varepsilon_t^d]'$, satisfying $E\{\varepsilon_t \varepsilon_t'\} = I$ for all t , where ε_t^a represents a technology shock and ε_t^d is a non-technology shock (which occasionally refer to for convenience as a "demand" shock). Thus we assume $u_t = K_t \varepsilon_t$ for all t for some non-singular matrix K_t satisfying $K_t K_t' = \Sigma_t$. Note that, given our normalization, changes in the contribution of different structural shocks to the volatility of innovations in output, hours or productivity will be captured by changes in K_t .

Our identification of structural shocks follows Galí (1999), by assuming that only technology shocks may affect labor productivity in the long-run. As we

¹²Cogley and Sargent (2005) adopt a more restrictive specification of the time-varying VAR, characterized by a constant matrix F . That assumption imposes some restrictions on the evolution of Σ_t that are absent here.

will see next, that assumption imposes some restrictions that allow us to recover matrix K_t from our estimated reduced form model (1).

Before we proceed it is convenient to rewrite (1) in companion form:

$$\mathbf{x}_t = \boldsymbol{\mu}_t + \mathbf{A}_t \mathbf{x}_{t-1} + \mathbf{u}_t$$

where $\mathbf{x}_t \equiv [x'_t, x'_{t-1}, \dots, x'_{t-p+1}]'$, $\mathbf{u}_t \equiv [u'_t, 0, \dots, 0]'$, $\boldsymbol{\mu}_t \equiv [A'_{0,t}, 0, \dots, 0]'$ and \mathbf{A}_t is the corresponding companion matrix. We use a *local approximation* of the implied response at $t+k$ of (log) labor productivity growth and (log) hours to a realization of the innovation vector in period t . Formally, that local response is given by

$$\frac{\partial x_{t+k}}{\partial u'_t} = \mathcal{E}_{2,2} \mathbf{A}_t^k \equiv B_{t,k}$$

for $k = 1, 2, \dots$ where $\mathcal{E}_{2,2}(M)$ is a function which selects the first 2 rows and 2 columns of any matrix M , and where $B_{t,0} \equiv I$. Thus, the k -period horizon impulse responses of labor productivity growth and hours to structural shocks hitting the economy at time t are given by

$$\begin{aligned} \frac{\partial x_{t+k}}{\partial \varepsilon'_t} &= \frac{\partial x_{t+k}}{\partial u'_t} \frac{\partial u_t}{\partial \varepsilon'_t} \\ &= B_{t,k} K_t \equiv C_{t,k} \end{aligned}$$

for $k = 0, 1, 2, \dots$ Notice that in contrast with the fixed-coefficient model, the impulse response of a variable to a shock at any given horizon may vary over time.

Let $\tilde{B}_{t,k} \equiv \sum_{j=0}^k B_{t,j}$ and $\tilde{C}_{t,k} \equiv \sum_{j=0}^k C_{t,j}$. The assumed absence of a long run effect of non-technology shocks on the level of labor productivity implies that the matrix of long-run cumulative multipliers $\tilde{C}_{t,\infty} \equiv \tilde{B}_{t,\infty} K_t$ is lower triangular. This, combined with the fact that $K_t K'_t = \Sigma_t$, yields

$$\tilde{C}_{t,\infty} \tilde{C}'_{t,\infty} = \tilde{B}'_{t,\infty} \Sigma_t \tilde{B}_{t,\infty}$$

which in turn allows us to determine (up to column sign) $\tilde{C}_{t,\infty}$ as the Cholesky factor of $\tilde{B}'_{t,\infty} \Sigma_t \tilde{B}_{t,\infty}$. Given $\tilde{C}_{t,\infty}$, the structural impulse responses of shocks occurring at time t can be obtained using

$$\frac{\partial x_{t+k}}{\partial \varepsilon'_t} = B_{t,k} \tilde{B}_{t,\infty}^{-1} \tilde{C}_{t,\infty}$$

for $k = 0, 1, 2, \dots$ which is a function of parameters describing the reduced form time-varying VAR (1) only. We refer the reader to Appendix 1 for a detailed description of the method used to estimate that model, which follows Primiceri (2005).

Our analysis below focuses on the second moments (conditional and unconditional) of the growth rates of output (Δy_t), labor productivity ($\Delta(y_t - n_t) \equiv \Delta q_t$), and hours (Δn_t). Our model allows us to write each of those variables as a time-varying distributed lag of the two structural disturbances. Thus, letting $x_{i,t}$ represent one of those variables we have

$$x_{i,t} = \mu_t^i + \sum_{k=0}^{\infty} C_{t,k}^{ia} \varepsilon_{t-k}^a + \sum_{k=0}^{\infty} C_{t,k}^{id} \varepsilon_{t-k}^d$$

Given estimates of the coefficients of such distributed lags, we can construct time-varying measures of unconditional and conditional second moments of the three variables under consideration. Thus, for instance, the unconditional variance at time t of variable $x_{i,t}$ is given by

$$var(x_{i,t}) = \sum_{k=0}^{\infty} (C_{t,k}^{ia})^2 + \sum_{k=0}^{\infty} (C_{t,k}^{id})^2$$

where the two terms on the right hand side represent the contribution of each of the shocks to that variance (or, equivalently, the variances conditional on each of the shocks).

Similarly, the covariance at time t between $x_{i,t}$ and $x_{j,t}$ is given by

$$cov(x_{i,t}, x_{j,t}) = \sum_{k=0}^{\infty} C_{t,k}^{ia} C_{t,k}^{ja} + \sum_{k=0}^{\infty} C_{t,k}^{id} C_{t,k}^{jd}$$

with each of the terms on the right hand side representing the covariances at time t conditional on technology and non-technology shocks, respectively. Time-varying conditional and unconditional correlations can then be computed in a straightforward way, using the above information.

In the next section we report estimates for a number of such time-varying second moments and analyze the timing of their changes, relative to that of the Great Moderation.

4 Changing Labor Market Dynamics and the Great Moderation

4.1 Unconditional Second Moments

Next we report some unconditional second moments implied by our estimated time-varying VAR. Figure 1a displays the evolution over time of the unconditional standard deviation of output, hours and labor productivity (all in log first-differences).¹³ The observed pattern for output volatility is consistent with the existing evidence on the Great Moderation: its standard deviation experiences a remarkable decline between 1980 and 1986, stabilizing after that date at a level below that of the 1960s. Before that transition the estimated volatility is far from constant, experiencing instead a substantial increase in the mid and late 1970s.¹⁴ A similar pattern, at least qualitatively, is observed for the standard deviation of hours, though for the latter variable the hump in the 1970s is relatively more pronounced than the overall decline in volatility. Finally, and by way of contrast, we see that the volatility of labor productivity declines very gradually over the postwar period, without showing any abrupt changes around the onset of the Great Moderation.

Figure 1b complements the previous evidence by showing the evolution of the relative standard deviations of hours and labor productivity, taking the volatility of output as a benchmark. In a way consistent with the evidence in Table 2 discussed above, we observe an upward trend in both measures of relative volatility. In the case of labor productivity, the observed pattern is the mirror image of that seen in the standard deviation of output, thus showing a large increase in the early 1980s, coinciding with the onset of the Great Moderation. On the other hand, the (smaller) fluctuations around an upward trend in the relative standard deviation of hours do not display any obvious pattern that one could relate to the Great Moderation or any other event.

¹³Here and in subsequent figures we report statistics starting in 1962:I, since the earlier sample is needed for the purpose of calibration of priors' parameters. Unless noted otherwise the value reported corresponds to the median of the posterior distribution of the statistic of interest, at each point in time.

¹⁴A similar observation is made in Blanchard and Simon (2001).

Figure 2 displays the evolution of the unconditional (pairwise) correlations among output, hours and labor productivity, measured by the left-hand scale. As a reference, the figure also shows the time-varying standard deviation of output (measured by the right-hand scale). The figure confirms the decline (and change of sign) in the hours-labor productivity correlation (dash-dotted line) already uncovered in Table 3, now making clear that the bulk of that decline takes place in the early 1980s, thus coinciding in its timing with the onset of the Great Moderation. Before that turning point, the correlation shows a gradual increase.¹⁵ A similar pattern, though less pronounced, can be observed in the hours-output correlation.

We view the findings above as *prima facie* evidence against a strong version of the good luck hypothesis for, as argued in the introduction, the latter would predict a scaling down of fluctuations in all variables without a corresponding change in their correlations. The evidence so far, however, does not allow us to determine whether those changes reflect a mere composition effect (resulting from variations in the relative importance of different types of shocks) or whether, instead, there has been a genuine change in the economy's response to each kind of shock. In order to address that question we turn to the analysis of the estimated conditional moments.

4.2 Conditional Volatilities: What Shocks are Responsible for the Great Moderation?

We start by examining the sources of the changes in the standard deviation of output, hours and labor productivity over time (all in log first-differences). Figures 3a through 3c plot the estimates of the (time-varying) standard deviations of each of those variables *conditional* on technology (dashed line) and non-technology shocks (dotted line), as implied by our estimated SVAR. In each case, and as reference, we also plot the unconditional standard deviation (solid line).

¹⁵That observation confirms a key finding in Stiroh (2008), even though our statistical approaches are different (we use a time-varying VAR vs rolling correlations in Stiroh (2008)).

The pattern that emerges in Figure 3a is unambiguous: the Great Moderation can be largely accounted for by the decline in the contribution of non-technology shocks to the variance of output. In particular, the timing and magnitude of the fall in the conditional standard deviation of output, between 1980 and 1985, matches well that of its unconditional standard deviation. On the other hand, the contribution of technology shocks to output volatility, appears to have been much more stable over the postwar period, with a small decline in the early 1980s followed by an (equally small) increase over the past two decades. It is interesting to note that, starting from a dominant role of non-technology shocks in the early 60s, the different trends in the conditional volatilities mentioned above have implied a gradual convergence in the contribution of both shocks, with their weights being essentially the same at the end of the sample.

Figure 3b reports analogous evidence for hours. As in the case of output, changes in the contribution of non-technology shocks explain the bulk of the pattern in the standard deviation of hours, including its rise in the 1970s and the subsequent fall in the 1980s. The contribution of technology shocks is much smaller, and appears to display a slight downward trend.

The previous two figures have shown that technology shocks have had, at least until recently, a relatively small role as a source of fluctuations in U.S. output. In the case of hours a similar finding holds for the entire postwar period. Figure 3c makes clear that this is not the case for labor productivity: fluctuations in the latter are largely accounted for by technology shocks. Yet, the figure also makes clear that non-technology shocks are responsible for the secular decline over the postwar period in the volatility of labor productivity. Interestingly, the decline in the contribution of non-technology shocks to that volatility is seen to start in the mid 1970s, well before the onset of the Great Moderation period.

Tables 4 allows us to examine the sources of the observed changes in volatilities from a different perspective. It reports the (conditional) standard deviations

of the estimated technology and non-technology components of output, hours and labor productivity, for both the pre-84 and post-84 sample periods. In contrast with the evidence reported in Figures 3a-3c, the statistics reported in Tables 4 depend not only on the estimated moving average coefficients (the $C_{t,k}^{ij}$'s of section 3) but also *on the specific realizations of the structural shocks* in each sample period. As we did for the original data (see Table 2), we report statistics for both the first-differenced and BP-filtered transformations of each of those components and test for the significance of the estimated changes across the two subsamples.¹⁶ The statistics in Table 4 point to the following findings uncovered by our analysis. First, non-technology shocks appear to be the main source of the decline in the volatility of output and labor productivity. Second, although both shocks contribute to the drop in the volatility of hours, the larger share of that decline (and the only one significant at the 5 percent level) is that associated with technology shocks.

An important caveat must be raised at this point: our analysis so far cannot identify whether the changes in conditional volatilities are the result of changes in the variance of the underlying structural shocks ("good luck") or, alternatively, of a different impact of a shock of a given size on the variable considered, which could be the result of a change in the systematic policy response to that shock or of other structural changes. Thus, for instance, the lower contribution of non-technology shocks in the more recent period could be due either to smaller demand disturbances or to a stronger countercyclical policy in response to those shocks (or both, of course). The evidence on conditional correlations provided below, however, is inconsistent with an explanation based exclusively on changes in the variance of some of the underlying structural shocks.

The previous caveat notwithstanding, the evidence shown in Figures 3a-3c is clearly at odds with the hypothesis of a declining contribution of technology shocks to output variability put forward in Arias, Hansen and Ohanian (2006;

¹⁶We should note that the tests reported in Tables 4 and 5 treat the estimates of the $C_{t,k}^{ij}$ coefficients as the "true" coefficients, i.e. they do not take into account the sampling error associated with the estimation. Thus, they should just be viewed as a quantitative summary of the estimated changes in conditional second moments.

AHO henceforth), and which is claimed by the latter authors to fully account for the decline in the cyclical volatility of output. To be more specific, those authors show that the standard deviation of measured total factor productivity (TFP) has declined by a factor of about 1/2 between the pre-84 and post-84 periods. As shown by AHO, when two alternative calibrations of the technology process consistent with that observation are considered, an RBC model predicts a decline in the volatilities of output and its components similar to those observed in the data. The empirical evidence presented here shows no sign of a decline in the contribution of technology shocks to output volatility that could account for the Great Moderation, and hence calls into question the conclusions of AHO's analysis.

4.3 Conditional Correlations and Structural Change

In Figures 4a through 4c we display the evolution of the conditional correlations between output and hours (Figure 4a), labor productivity and hours (Figure 4b), and labor productivity and output (Figure 4c). Correlations conditional on technology (non-technology) shocks are represented by the dashed (dotted) line, while the solid line represents the unconditional correlation. In order to interpret the subsequent evidence it is worth noting the relationship linking the unconditional and conditional correlations between two generic variables x and z :

$$corr(x_t, z_t) = \lambda_a corr_a(x_t, z_t) + \lambda_d corr_d(x_t, z_t)$$

where $\lambda_i \equiv \frac{\sigma_i(x_t)}{\sigma(x_t)} \frac{\sigma_i(z_t)}{\sigma(z_t)}$ and where $corr_i(x_t, z_t)$ and $\sigma_i(z_t)$ denote, respectively, the correlation and standard deviation conditional on i -shocks, for $i = a, d$. Note that the weight given to each conditional correlation in the above expression is proportional to the geometric average of the shares of the corresponding conditional variances in the unconditional variance of each variable. As a result, that weight will be small if the associated shock accounts for a small fraction of the variance of one of the two variables, even if it plays a large role in accounting for the volatility of the other variable.

As seen in Figure 4a, the strong positive correlation between output and hours masks a more complex underlying reality: the coexistence of a stable near-unity correlation generated by non-technology shocks (dotted line) with a correlation that fluctuates between positive and (slightly) negative values as a result of technology shocks (dashed line). The weak correlation between output and hours conditional on technology shocks is consistent with much of the evidence uncovered by the recent literature on the macroeconomic effects of technology shocks.¹⁷ Our approach here allows us to uncover a novel result: the changing pattern of the output-hours correlation conditional on technology shocks. In particular, it is worth noting the increases in that correlation in the 1970s and in the second half of the 1990s, when it takes non-negligible positive values (above 0.5), before returning to negative territory. Note, however, that the two surges in the conditional correlations are hardly reflected in the corresponding unconditional correlation, given the relatively small weight of technology shocks in accounting for the total variance of hours during those episodes (see Figure 4b).

Figure 4b reports conditional and unconditional correlations between labor productivity and hours. The figure confirms the large decline in their correlation conditional on non-technology shocks (dotted line), which falls from a value of about 0.6 in the 1960s to somewhere between -0.6 and -0.8 in more recent years. Note, however, that the bulk of that decline occurs in the 1990s, once the Great Moderation is well underway and after the large decline in the unconditional correlation. On the other hand we see that the hours-productivity correlation conditional on technology shocks (dashed line) hovers around a value close to -0.8 with the exception of two spikes: one around 1980, and a larger spike in the second half of the 1990s. The previous findings, combined with those in Figures 3b and 3c suggest that the large decline in the unconditional correlation in the early 1980s is the result of a variety of factors, including a decline in both conditional correlations and an increase in the relative importance of technology shocks, given that the latter induce a negative correlation

¹⁷See Galí and Rabanal (2004) for a survey of that literature.

between hours and labor productivity,

Finally, we show in Figure 4c the evolution of the conditional and unconditional productivity-output correlations. Note that the correlation conditional on technology shocks (dashed line) is close to unity during much of the sample period. This fact, combined with the dominant role of those shocks as a source of labor productivity fluctuations (see Figure 3c), explains the relative stability of the unconditional productivity-output correlation around a high positive value. By way of contrast, the correlation conditional on non-technology shocks (dotted line) follows a rapidly declining pattern that roughly mirrors that observed for the corresponding correlation between productivity and hours in Figure 4b.

Table 5 quantifies the (pairwise) conditional correlations among output, hours and labor productivity in the pre-84 and post-84 periods. As in Table 4, we report statistics for both the first-differenced and BP-filtered transformations of each of those components and test for the significance of the estimated changes across the two subsamples. The results of that exercise confirm that non-technology shocks are largely responsible for the significant decline in the correlation between labor productivity and hours on the one hand, and labor productivity and output on the other.¹⁸

The evidence provided above suggests that at least two of the observed changes in unconditional correlations (those involving labor productivity) described in section 2 and earlier in the present section can be attributed to a (large) change in conditional correlations, the ones associated with non-technology shocks. Furthermore, and as discussed above, the timing of some of those changes matches pretty well that of the Great Moderation. That finding provides some evidence that the latter episode cannot be characterized *exclusively* in terms of a decline in the volatility of one or more shock, hinting instead (though admittedly without proving it) at a potential role for structural change.

¹⁸Note that the latter decline is (partly) offset by a small, but significant, increase in the correlation between labor productivity and output resulting from technology shocks.

4.4 Impulse Responses

Conditional volatilities and correlations summarize some dimensions of the impulse responses to different shocks. Accordingly, the changes experienced over the postwar period in those conditional second moments must be reflecting parallel changes in the underlying impulse responses. Next we present and briefly discuss the evolution over time of the impulse responses that can account for three of the most significant findings uncovered above, namely, (i) the decline in output volatility resulting from a smaller contribution of non-technology shocks, (ii) the sign and changes over time in the conditional correlations between labor productivity and hours.

As discussed above, the decline in output volatility initiated in the 1980s is the result of a smaller contribution of non-technology shocks. Figure 5a displays the evolution over time of the dynamic response of output to a non-technology shock. More specifically, the figure shows the response corresponding to the first quarter of each calendar year to a unit innovation in ε_t^d . Given our normalization, that size corresponds to a one standard deviation. Throughout the sample period the response of output to a non-technology shock shows a characteristic hump shape, and displays substantial persistence. But, as is clearly captured by the figure, the scale of the response goes down dramatically in the early 1980s, and remains subdued from then on. The magnitude of that change is reflected more clearly in figure 5b, which displays, side by side, the average impulse responses in the pre-84 and post-84 periods. Figure 5c shows the difference between those two impulse responses, together with a 68% (dashed) and 95% (dotted) confidence bands implied by the posterior distribution. Perhaps not surprisingly given the nature of our empirical approach, the uncertainty associated with the estimated impulse responses is large (as reflected in the size of the confidence bands). Yet, the posterior distribution strongly rejects the hypothesis of no differential response over the 6 quarters subsequent to the shock at a 5 percent significance level.

A second key finding emphasized above is the decline in the cyclicity of

labor productivity conditional on non-technology shocks. Figure 6a uncovers the source of that change, by showing the evolution over the postwar period of the dynamic response of labor productivity to a unit innovation in ε_t^d (i.e. the same pattern of shocks responsible for the output responses shown in Figure 6a). Thus, we see that an expansionary non-technology shock has a large and persistent *positive* effect on labor productivity in the early part of the sample, an observation consistent with the evidence of so-called "short-run increasing returns to labor" (SRIRL) uncovered by a number of economists.¹⁹ Starting in the early 80s, however, the SRIRL phenomenon vanishes gradually: the response of labor productivity keeps getting smaller over time until eventually switches its sign and becomes persistently negative, as would be implied by a technology displaying decreasing returns to labor. As shown in Figure 6b, the average impulse responses of labor productivity over the pre-84 and post-84 periods differ considerably, with the gap between the two at the time of the shock being significant at the 5 percent level (see Figure 6c).

Finally, we turn our attention to the response of hours to a technology shock and its evolution over the postwar period, which is shown in Figure 7a. For much of the sample period considered, hours display a persistent decline in response to a positive technology shock, i.e. one that increases labor productivity permanently (responses not shown here). That finding is consistent with the evidence in Galí (1999), Basu, Fernald and Kimball (2005), and Francis and Ramey (2005), and accounts for the negative *conditional* correlation between hours and labor productivity estimated for much of the sample period (see Figure 4b). Our time-varying estimates allow us to go beyond the existing evidence and examine the changes over time in the size and pattern of response. In that respect we note that, some fluctuations notwithstanding, the size of the negative response of hours appears to have gone down over time (in absolute value).²⁰

¹⁹See Gordon (1990) for a review of that literature.

²⁰The previous finding accords with the evidence, reported in Galí, López-Salido, and Vallés (2003), of large and significant contractionary effects of aggregate technological improvements on employment in the pre-Volcker period, in contrast with the small and largely insignificant short term effects over the Volcker-Greenspan period.

This is reflected in the gap between the "average" impulse responses for the pre- and post-84 periods shown in Figure 7b, though the gradual change combined with the large confidence bands associated to our time-varying impulse responses cannot reject equality between the two average responses for any horizon at any reasonable significance level (see Figure 7c).

Perhaps most interestingly, we note how the negative response of hours is more muted in the late 1970s and in the second half of the 1990s (in the latter period it even becomes positive). Those observations would seem to account for the spikes in the pattern of hours-labor productivity correlations conditional on technology shocks shown in Figure 4b.

4.5 Robustness: The Role of Investment-Specific Technology Shocks

In this section we extend our empirical analysis along the lines of Fisher (2006), thus allowing for both neutral technology shocks (henceforth, N-shocks) and investment-specific technology shocks (I-shocks), in order to check the robustness of our main findings. This extension is of particular interest in light of the findings in Justiniano and Primiceri (2006) based on time-varying estimates of a DSGE model, and which point to the smaller size of I-shocks as the main explanation for the decline in output growth volatility.

Following Fisher (2006), we identify I-shocks as the only source of the unit root in the relative price of investment, i.e. we restrict N-shocks and non-technology shocks not to have a permanent effect on that variable. On the other hand, we allow both N-shocks and I-shocks to have a long run effect on labor productivity. Following Justiniano and Primiceri (2006) we construct a series for the (log) real price of investment as a weighted average of the (log) deflators of nondurables and services consumption minus the weighted average of the (log) deflators for investment and durable consumption, with the weights given by the relative (nominal) shares of each spending category.²¹

²¹The data used to construct the relative price of investment series were drawn from the FRED-II database of the St. Louis Fed. The deflators are constructed as the ratios of nominal to real expenditure in each category, using the following formulas:

Given space limitations we focus our discussion on two key aspects of the evidence presented above: the contribution of the different shocks to the decline in output volatility and their role in accounting for the change in the labor productivity-hours correlation.

Figure 8 plots estimates of the (time-varying) standard deviation of output growth conditional on the three types of shocks, as well as the corresponding unconditional standard deviation. First, note that N-shocks (dashed line) have a relatively small and stable contribution to the volatility of output throughout the sample period, with the exception of a transitory increase around 1980. Secondly, both I-shocks (dashed-dotted line) and non-technology shocks (dotted line) play an important role in the Great Moderation. Interestingly, however, the patterns of their contribution differ substantially. Roughly speaking, while non-technology shocks account for the downward trend in volatility, I-shocks (and to a lesser extent, N-shocks) appear to be responsible for the hump observed during the second half of the 1970s.

Our augmented model thus points to an important role of I-shocks as a source of both the extraordinary increase in volatility of the 1970s and the subsequent decline in the mid-1980s. On the other hand our previous finding of an important contribution of non-technology shocks to the decline in output volatility appears to be robust to the alternative specification considered here, though the (previously dominant) role of non-technology shocks in the abrupt volatility decline of the early 80s is now shared to some degree with technology shocks.

Figure 9 displays the conditional and unconditional correlations between labor productivity and hours based on the time-varying estimates of our augmented VAR. We note that a key finding of our bivariate model, namely, the decline in the hours-labor productivity correlation conditional on non-technology shocks re-emerges here, though it appears to be less abrupt than in our bivariate model. In fact, while that correlation declines from a value close to 0.7 to one

PCDG/PCDGCC96 (durables), PCND/PCNDGC96 (nondurables), PCESV/PCESVC96 (services) and FPI/FPIC1 (investment).

about 0.1, it remains positive over the whole sample period. On the other hand both types of technology shocks generate a correlation between the same two variables that displays no strong downward trend over time, but instead shows a hump centered around 1980, though somewhat less pronounced than the one obtained in the bivariate model.

5 Tentative Interpretations and Caveats

The remarkable decline in macroeconomic volatility experienced by the U.S. economy since the mid-80s (the so-called Great Moderation) has involved more than a mere scaling down of the size of fluctuations. In particular, and as the evidence provided in the present paper makes clear, that volatility decline has been accompanied by large changes in the patterns of comovements among output, hours and labor productivity. Those changes are reflected in both conditional and unconditional second moments as well as in the impulse responses to identified shocks.

Two of our findings appear particularly relevant and worthy of further discussion. First, the decline in output volatility appears to be the result of a smaller contribution of non-technology shocks. Secondly, the Great Moderation period has witnessed a dramatic fall (with sign switch included) in the correlation between hours and labor productivity generated by non-technology shocks.

The shrinking contribution of non-technology shocks to output volatility can be due, in principle, to two non mutually exclusive developments. First, the average size of the underlying shocks may have become smaller. Secondly, the response of output may have become more muted over time, even when controlling for shock size, as a result of some structural change in the mechanisms propagating the effects of the shock (e.g. a change in the systematic policy response to those shocks).

Given our identification scheme, a variety of structural disturbances fall under the broad heading of non-technology shocks, including exogenous monetary

and fiscal policy shocks or preference shocks, among others.²² A number of authors have provided independent evidence pointing to a smaller volatility of those shocks in the post-84 period, relative to the earlier period.²³ That evidence is consistent with our finding of a smaller contribution of non-technology shocks. Yet, and at least in the case of policy shocks, it can hardly be interpreted as being consistent with the "good luck" hypothesis, at least to the extent that the decline in the volatility of those shocks is viewed as the result of a better understanding of the destabilizing effects of "erratic" policies. The key role of non-technology shocks in accounting for the Great Moderation is also consistent with the empirical literature on interest rate rules, which points to an increase in the weight attached by the Fed to inflation stabilization during the Volcker-Greenspan years relative to the pre-Volcker period.²⁴ To the extent that the non-technology shocks identified by our VAR largely lead to changes in aggregate demand with limited impact on potential output, a stronger anti-inflationary stance by the Fed should bring about greater output stability as a by-product, in a way consistent with our evidence. Furthermore, and as discussed in Galí, López-Salido, and Vallés (2003), the Fed's greater focus on inflation stabilization should automatically lead to a greater accommodation of changes in potential output resulting from technology shocks. That mechanism could thus account for the stability in the contribution of technology shocks to output volatility suggested by our estimates, even in the face of a likely reduction in the size of the underlying shocks.²⁵ It is also consistent with conventional

²²Of course, that diversity combined with changes in the relative importance of each of the shock-types and the possible differences in their respective joint responses of output, hours, and labor productivity could be a spurious source of some of the changes we detect. Unfortunately there is little we can do to assess the quantitative relevance of that hypothesis without imposing additional (and likely controversial) identifying assumptions. A further limitation of our approach results from the underlying linear structure assumed, that implies that small and large shocks generate the same conditional comovements and relative volatilities among the variables of interest. Thus, some of the estimated changes in correlations could be *in principle* caused by non-linearities combined with differences in the size of shocks across periods. Unfortunately, and due to the reasons pointed out in the text, our identification approach does not allow us to separately identify the size of the shocks and its changes over time.

²³See, in particular, section 5.4 in Stock and Watson (2002) and section 5.D in Smets and Wouters (2008).

²⁴See, e.g., Taylor (1999), Clarida, Galí, and Gertler (2000) and Boivin and Giannoni (2006).

²⁵Evidence of smaller technology shocks in the post-84 period can be found in Stock and

accounts of the role played by the Greenspan Fed in accommodating the output and employment boom during the second half of the 1990s, generally attributed to the high productivity growth brought about by the IT revolution

How can one explain our second main finding, i.e. the large decline in the hours-labor productivity correlation conditional on non-technology shocks? One way to approach this question is to consider what may have caused the high and positive conditional correlation in the early postwar period. A common explanation found in the literature is the presence of labor hoarding, understood as firms' desire to smooth employment and/or hours hired in the face of fluctuations in demand and output, possibly as a result of a variety of costs associated with the adjustment of labor. In that environment, *measured* hours will fluctuate less than their *effective* counterpart, since firms will elicit procyclical variations in (unobservable) effort.²⁶ To formalize this idea let $n_t^* = n_t + e_t$ where n_t^* and n_t denote, respectively, effective and measured (log) labor input, and e_t represents (log) effort. Suppose that, in the face of shocks that call for an adjustment of effective labor input, firms make use of both margins (hours and effort) to a greater or lesser degree. For simplicity, let us assume that $e_t = \gamma n_t^*$, where $\gamma \in [0, 1]$ measures the extent to which changes in effective labor input are achieved without adjusting measured hours (i.e. the extent of labor hoarding) and ξ_t is an i.i.d. disturbance uncorrelated with n_t^* . Assuming, for the sake of illustration, a simple a production function (in logs) of the form

$$y_t = a_t + (1 - \alpha) n_t^* + \xi_t$$

where ξ_t represents variations in non-labor inputs.²⁷ Combining the previous assumptions we obtain

$$y_t = a_t + \left(\frac{1 - \alpha}{1 - \gamma} \right) n_t + \xi_t$$

$$y_t - n_t = a_t + \left(\frac{\gamma - \alpha}{1 - \gamma} \right) n_t + \xi_t$$

Watson (2002) and Smets and Wouters (2008), among others.

²⁶See Sbordone (1996), Galí (1999), and Barnichon (2006) for examples of structural models generating such SRIRL as a result of variable effort.

²⁷For simplicity we assume the latter to be independent of the degree of labor hoarding.

In the above setup, a reduction in the degree of labor hoarding γ could potentially account for three of our findings: (i) the increase in the volatility of hours relative to output, (ii) the decline in the response of labor productivity to an expansionary non-technology shocks, with an eventual switch in the sign of that response (if γ becomes smaller than α), and (iii) the shrinking correlation between hours and labor productivity conditional on non-technology shocks.

Given the nature of our empirical analysis, the previous explanations can only be viewed as speculative. Establishing their relevance will require more direct evidence (e.g. of a decline in labor hoarding practices in response to more flexible labor markets) or the estimation of full fledged DSGE models with time-varying parameters (but at the cost of having a less flexible framework relative to the VAR).

An additional important limitation of our analysis is worth emphasizing: We have not attempted to establish a *causal* relationship between some of our findings regarding patterns of second moments and the Great Moderation. In particular, we have only pointed to a rough coincidence in time between the decline in both output volatility and in the comovement of labor productivity with hours, and have shown that those changes in second moments are largely associated to changes in the economy's response to non-technology shocks and/or in the relative importance of the latter's contribution to fluctuations. Determining whether both phenomena have a common underlying explanation, perhaps related to the evolution of the labor market structure, is a challenging task that remains beyond the scope of the present paper.

Those caveats notwithstanding, we believe that many of the findings reported in the present paper may provide a useful reference for the evaluation of alternative explanations of the Great Moderation. At the very least, our findings should convey a clear message, namely, that changes in the macroeconomic performance of the U.S. economy since the early 1980s, including the Great Moderation, are far more complex than implied by some stylized versions of the "good luck" hypothesis.

Appendix

The present appendix describes the method used to estimate the time-varying SVAR. Our approach follows closely Cogley and Sargent (2005), Primiceri (2005), and Benati and Mumtaz (2007).

A. Priors

Let z^T denote a sequence of z 's up to time T . We assume that the conditional prior density of θ^T is given by:

$$p(\theta^T | \gamma^T, \sigma^T, \Psi, \Xi, \Omega) \propto I(\theta^T) f(\theta^T | \gamma^T, \sigma^T, \Psi, \Xi, \Omega) \quad (5)$$

where $I(\theta^T) = \prod_{t=0}^T I(\theta_t)$,

$$f(\theta^T | \gamma^T, \sigma^T, \Psi, \Xi, \Omega) = f(\theta_0) \prod_{t=1}^T f(\theta_t | \theta_{t-1}, \gamma^T, \sigma^T, \Psi, \Xi, \Omega) \quad (6)$$

and $f(\theta_t | \theta_{t-1}, \gamma^T, \sigma^T, \Psi, \Xi, \Omega)$ is consistent with (2). The function $I(\theta_t)$ takes a unit value if all the roots of the VAR polynomial associated to θ_t are larger than one in modulus and 0 otherwise. To calibrate the prior densities of the coefficients we estimate a time invariant VAR using data up to 1961:IV. Following Benati and Mumtaz (2007) and Primiceri (2005) we make the following assumptions about prior densities and parameters:

$$\begin{aligned} p(\theta_0) &\propto I(\theta_0) N(\hat{\theta}_{OLS}, \hat{V}(\hat{\theta}_{OLS})) \\ p(\log \sigma_0) &= N(\log \hat{\sigma}_{OLS}, 10 \times I) \\ p(\gamma_0) &= N(\hat{\gamma}_{OLS}, |\hat{\gamma}_{OLS}|) \\ p(\Omega) &= IW(\bar{\Omega}^{-1}, T_0) \\ p(\Psi) &= IW(\bar{\Psi}^{-1}, 2) \\ p(\Xi_{i,i}) &= IG\left(\frac{0.0001}{2}, \frac{1}{2}\right) \end{aligned}$$

where $\hat{\theta}_{OLS}$ is the vector of OLS estimates of the VAR coefficients and $\hat{V}(\hat{\theta}_{OLS})$ the estimates of their covariance matrix using the initial sample, $\hat{\sigma}_{OLS}$ is a vector containing the elements of the diagonal matrix \hat{D} and $\hat{\gamma}_{OLS}$ is the element

(2,1) of the lower triangular matrix \hat{F}^{-1} , where $\hat{F}\hat{D}\hat{F}' = \hat{\Sigma}_{OLS}$, and $\bar{\Omega} = 0.005 \times V(\hat{\theta}_{OLS})$, T_0 is the number of observations in the initial sample, and $\bar{\Psi} = 0.001 * |\hat{\gamma}_{OLS}|$.

B. Estimation

To draw realizations from the posterior density we use an MCMC algorithm which works in an iterative way. Each iteration is done in six steps and consists in drawing a subset of coefficients conditional on a particular realization of the remaining coefficients and then using such a realization in the conditional densities of the remaining coefficients. Under regularity conditions and after a burn-in period, iterations on these four steps produce draws from the joint density.

- Step 1: $p(\theta^T | x^T, \gamma^T, \sigma^T, \Psi, \Xi, \Omega)$

Conditional on $x^T, \gamma^T, \sigma^T, \Psi, \Xi, \Omega$, the unrestricted posterior of the states is normal. To draw from the conditional posterior we employ the algorithm of Carter and Kohn (1994). The conditional mean and variance of the terminal state θ_T is computed using standard Kalman filter recursions while for all the other states the following backward recursions are employed

$$\theta_{t|t+1} = \theta_{t|t} + P_{t|t}P_{t|t+1}^{-1}(\theta_{t+1} - \theta_{t|t}) \quad (7)$$

$$P_{t|t+1} = P_{t|t} - P_{t|t}P_{t+1|t}^{-1}P_{t|t} \quad (8)$$

where $p(\theta_t | \theta_{t+1}, x^T, \gamma^T, \sigma^T, \phi) \sim N(\theta_{t|t+1}, P_{t|t+1})$.

- Step 2: $p(\gamma^T | x^T, \theta^T, \sigma^T, \Psi, \Xi, \Omega)$

This is done following the same procedure described in Primiceri (2005). Conditional on θ^T $\hat{y}_t = x_t - A_{0,t} - A_{1,t} x_{t-1} - \dots - A_{p,t} x_{t-p}$ is observable. We can rewrite our system of equations as $F_t^{-1}\hat{y}_t = D_t\nu_t$ where $\nu_t \sim N(0, I)$. Conditional on σ^T we use the algorithm of Carter and Kohn to obtain a draw for γ_t taking the above system as observational equations and (3) as unobserved states equations. Given that the γ_t 's and the ν_t 's are independent across equations the algorithm can be applied equation by equation. However notice that in the bivariate case we have one observable equation and one state.

- Step 3: $p(\sigma^T|x^T, \theta^T, \gamma^T, \Psi, \Xi, \Omega)$

This is done using the univariate algorithm by Jacquier, Polson and Rossi (2004) used in Cogley and Sargent (2005) (see Appendix B.2.5 of the latter for details).

- Step 4: $p(\Psi|x^T, \theta^T, \gamma^T, \sigma^T, \Xi, \Omega), p(\Xi_{:,i}|x^T, \theta^T, \gamma^T, \sigma^T, \Psi, \Omega), p(\Omega|x^T, \theta^T, \gamma^T, \sigma^T, \Xi, \Psi)$

Conditional on $x^T, \theta^T, \gamma^T, \sigma^T$ all the remaining hyperparameters, under conjugate priors, can be sampled in a standard way from Inverted Wishart and Inverted Gamma densities (see Gellman et al. (2001)).

We perform 30,000 repetitions, we discard the first 10,000 draws and we keep one for every 20 of the remaining 20,000 draws to break the autocorrelations of the draws. The densities for the parameters are typically well behaved. We made many robustness checks for prior specifications and the length of the chain with the main results not being affected significantly.

References

- Arias, Andres, Gary D. Hansen, and Lee E. Ohanian (2006): "Why Have Business Cycle Fluctuations Become Less Volatile?" NBER WP 12079.
- Barnichon, Régis (2006): "Productivity, Aggregate Demand and Unemployment Fluctuations," London School of Economics, unpublished manuscript.
- Baxter, Marianne and Robert G. King (1999): "Measuring Business cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 575-593.
- Benati, Luca and Haroon Mumtaz (2007): "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation," ECB Working paper series #746.
- Blanchard, Olivier J. and Simon (2001): "The Long and Large Decline in Output Volatility," *Brookings Papers on Economic Activity*, issue 1, 135-164.
- Boivin, Jean and Marc Giannoni (2006): "Has Monetary Policy Become More Effective?" *Review of Economics and Statistics*, 88(3), 445-462.
- Box, George E.P., and Gwilym M. Jenkins (1976): *Time Series Analysis: Forecasting and Control*, Holden-Day, Oakland, CA.
- Carter, Chris K. and Kohn, Robert. "On Gibbs Sampling for State Space Models," *Biometrika*, August 1994, 81(3), pp. 541-53.
- Cecchetti, Stephen G., Alfonso Flores-Lagunes, and Stefan Krause (2006): "Assessing the Sources of Changes in the Volatility of Real Growth," NBER WP# 11946.
- Christiano, Lawrence J., and Martin Eichenbaum (1992): "Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations," *American Economic Review* 82, 430-450.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson (2003): "What Happens after a Technology Shock?" NBER Working Paper #9819.
- Clarida, Richard, Jordi Galí, and Mark Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, vol. 105, issue 1, 147-180.
- Cogley, Timothy, and Thomas J. Sargent (2005): "Drifts and Volatilities: Monetary Policy and Outcomes in the Post WWII U.S.," *Review of Economic*

Dynamics, 8, 262-302

Dynan, Karen E., Douglas W. Elmendorf, and Daniel E. Sichel (2006): "Can financial Innovation Help to Explain the Reduced Volatility of Economic Activity?," *Journal of Monetary Economics*, vol. 53, issue 1, 123-150.

Fernald, John G. (2008): "Trend Breaks, Long Run Restrictions and the Contractionary Technology Improvements," *Journal of Monetary Economics* 54 (8), 2467-2485.

Fisher, Jonas D.M. (2006): "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy* 114(3), 413-451

Francis, Neville and Valerie Ramey (2005): "Is the Technology-driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," *Journal of Monetary Economics*, vol. 52, issue 8, 1379-1400.

Galí, Jordi (1999): "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review* 89, 249-271.

Galí, Jordi, J. David López-Salido, and Javier Vallés (2003): "Technology Shocks and Monetary Policy: Assessing the Fed's Performance," *Journal of Monetary Economics* 50, 723-743.

Gambetti, Luca (2006): "Technology Shocks and the Response of Hours Worked: Time-Varying Dynamics Matter," unpublished manuscript.

Gelman Andrew, John B. Carlin, Hal S. Stern and Donald B. Rubin (1995), *Bayesian Data Analysis*, Chapman and Hall, London.

Gordon, Robert J. (1990): "Are Procyclical Productivity Fluctuations a Figment of Measurement Error?," Northwestern University, mimeo.

Hansen, Gary D., and Randall Wright (1992): "The Labor Market in Real Business Cycle Theory," *Federal Reserve Bank of Minneapolis Quarterly Review*, Spring issue, 2-12.

Jaquier, E., N.G. Polson and P. Rossi (2004) "Bayesian Analysis of Stochastic Volatility Models", *Journal of Business and Economics Statistics*, 12, 371-418.

Justiniano, Alejandro and Giorgio E. Primiceri (2006): "The Time-Varying

Volatility of Macroeconomic Fluctuations," NBER WP # 12022.

Kahn, James A., Margaret M. McConnell, and Gabriel Perez-Quirós (2002): "On the Causes of the Increased Stability of the U.S. Economy," *Economic Policy Review of the Federal Reserve Bank of New York*, vol. 8, no. 1, 183-206.

Kim, Chang-Jin and Charles R. Nelson (1999): "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov Switching Model of the Business Cycle," *The Review of Economics and Statistics*, 81 (4) 608-616.

McConnell, Margaret M. and Gabriel Perez-Quiros (2000): "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?," *American Economic Review*, vol. 90, no. 5, 1464-1476

Priestley, M.B. (1981): *Spectral Analysis and Time Series*, Academic Press, San Diego, CA.

Primiceri, Giorgio E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, pp. 821-852.

Ramey, Valerie A. and Daniel J. Vine (2006): "Declining Volatility in the U.S. Automobile Industry," *American Economic Review* vol. 96, no. 5, 1876-1889.

Sbordone, Argia (1996): "Cyclical Productivity in a Model of Labor Hoarding," *Journal of Monetary Economics*, vol. 38, 331-361.

Smets, Frank, and Rafael Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, vol 97, no. 3, 586-606.

Stirol, Kevin J. (2006): "Volatility Accounting: A Production Perspective on Increased Economic Stability," unpublished manuscript, Federal Reserve Bank of New York.

Stock, James, and Mark W. Watson (2002): "Has the Business Cycle Changed and Why?," *NBER Macroeconomics Annual 2002*, MIT Press.

Stock, James, and Mark W. Watson (2005): "Understanding Changes in International Business Cycle Dynamics," *Journal of the European Economic Association*, vol. 3, issue 5, 968-1006.

Taylor, John B. (1999): "An Historical Analysis of Monetary Policy Rules,"
in J.B. Taylor ed., *Monetary Policy Rules*, University of Chicago Press.

Table 1. The Great Moderation

	<i>Standard Deviation</i>			p-value
	Pre-84	Post-84	$\frac{\text{Post-84}}{\text{Pre-84}}$	
First-Difference				
<i>GDP</i>	1.21	0.54	0.44	<0.01
<i>Nonfarm Business Output</i>	1.57	0.68	0.43	<0.01
BP-Filter				
<i>GDP</i>	2.00	0.94	0.47	<0.01
<i>Nonfarm Business Output</i>	2.59	1.23	0.47	<0.01

Note: All variables transformed by taking the natural logarithm and applying the transformation indicated in the table (first difference or band-pass filter). P-values correspond to a test of equality of variances across the two subsamples based on the asymptotic standard errors of variance estimates computed using an 8-lag window.(see, Priestley (1991), p. 327).

Table 2. Changes in Volatility

	<i>Standard Deviation</i>				<i>Relative Standard Deviation</i>		
	Pre-84	Post-84	$\frac{\text{Post-84}}{\text{Pre-84}}$	p-value	Pre-84	Post-84	$\frac{\text{Post-84}}{\text{Pre-84}}$
First-Difference							
<i>Output</i>	1.57	0.68	0.43	0.00	1.00	1.00	1.00
<i>Hours</i>	1.05	0.65	0.61	0.00	0.66	0.95	1.41
<i>Productivity</i>	1.00	0.61	0.62	0.00	0.63	0.89	1.44
BP-Filter							
<i>Output</i>	2.59	1.23	0.47	0.00	1.00	1.00	1.00
<i>Hours</i>	2.08	1.39	0.67	0.06	0.79	1.10	1.40
<i>Productivity</i>	1.18	0.68	0.57	0.01	0.45	0.55	1.21

Note: P-values correspond to a test of equality of variances across the two subsamples based on the asymptotic standard errors of variance estimates computed using an 8-lag window.(see, Priestley (1991), p. 327)

Table 3. Changes in Cross-Correlations

First-Difference	<i>pre-84</i>	<i>post-84</i>	<i>change</i>
<i>Output, Hours</i>	0.78	0.57	-0.20** (0.08)
<i>Hours, Productivity</i>	0.18	-0.41	-0.59** (0.10)
<i>Output, Productivity</i>	0.75	0.50	-0.24** (0.11)
BP-Filter	<i>pre-84</i>	<i>post-84</i>	<i>change</i>
<i>Output, Hours</i>	0.89	0.86	-0.02 (0.09)
<i>Hours, Productivity</i>	0.18	-0.46	-0.65** (0.15)
<i>Output, Productivity</i>	0.61	0.03	-0.58** (0.19)

Note: Test of equality of correlations across the two subsamples based on the asymptotic standard errors of estimated correlations computed using an 8-lag window.(see, e.g., Box and Jenkins (1976), p. 376). One asterisk denotes significance at the 10 percent level. Two asterisks indicate significance at the 5 percent level.

Table 4. Changes in Conditional Volatility

	<i>Non-Technology Shocks</i>				<i>Technology Shocks</i>			
	Pre-84	Post-84	$\frac{\text{Post-84}}{\text{Pre-84}}$	p-value	Pre-84	Post-84	$\frac{\text{Post-84}}{\text{Pre-84}}$	p-value
First-Difference								
<i>Output</i>	1.18	0.56	0.48	0.00	0.46	0.47	1.01	0.91
<i>Hours</i>	0.78	0.59	0.76	0.12	0.39	0.22	0.56	0.00
<i>Productivity</i>	0.53	0.17	0.33	0.00	0.68	0.93	0.88	0.51
BP-Filter								
<i>Output</i>	1.97	1.17	0.59	0.07	0.59	0.49	0.83	0.43
<i>Hours</i>	1.60	1.29	0.81	0.40	0.57	0.32	0.56	0.04
<i>Productivity</i>	0.50	0.25	0.51	0.03	0.80	0.67	0.84	0.40

Note: P-values correspond to a test of equality of variances across the two subsamples based on the asymptotic standard errors of variance estimates computed using an 8-lag window.(see, Priestley (1991), p. 327)

Table 5. Changes in Conditional Correlations

	<i>Non-Technology Shocks</i>			<i>Technology Shocks</i>		
	pre-84	post-84	change	pre-84	post-84	change
First-Difference						
<i>Output, Hours</i>	0.93	0.95	-0.02 (0.10)	-0.27	-0.65	-0.37** (0.11)
<i>Hours, Productivity</i>	0.61	-0.30	-0.93** (0.15)	-0.75	-0.83	-0.07 (0.06)
<i>Output, Productivity</i>	0.85	-0.01	-0.87** (0.17)	0.83	0.96	0.12** (0.02)
BP-Filter						
<i>Output, Hours</i>	0.98	0.98	-0.01 (0.12)	0.06	-0.33	-0.40* (0.23)
<i>Hours, Productivity</i>	0.67	-0.57	-1.24** (0.14)	-0.66	-0.73	-0.06 (0.13)
<i>Output, Productivity</i>	0.80	-0.41	-1.21** (0.16)	0.69	0.89	0.19** (0.09)

Note: Test of equality of correlations across the two subsamples based on the asymptotic standard errors of estimated correlations computed using an 8-lag window.(see, e.g., Box and Jenkins (1976), p. 376). One asterisk denotes significance at the 10 percent level. Two asterisks indicate significance at the 5 percent level.

Figure 1a
Time-Varying Standard Deviations




Figure 1b
Time-Varying Relative Standard Deviations




Figure 2
Time-Varying Unconditional Correlations




Figure 3a
Conditional Standard Deviations: Output




Figure 3b
Conditional Standard Deviations: Hours




Figure 3C
Conditional Standard Deviations: Labor Productivity




Figure 4a
Conditional Correlations: Hours - Output




Figure 4b

Conditional Correlations: Labor Productivity - Hours




Figure 4c

Conditional Correlations: Labor Productivity - Output

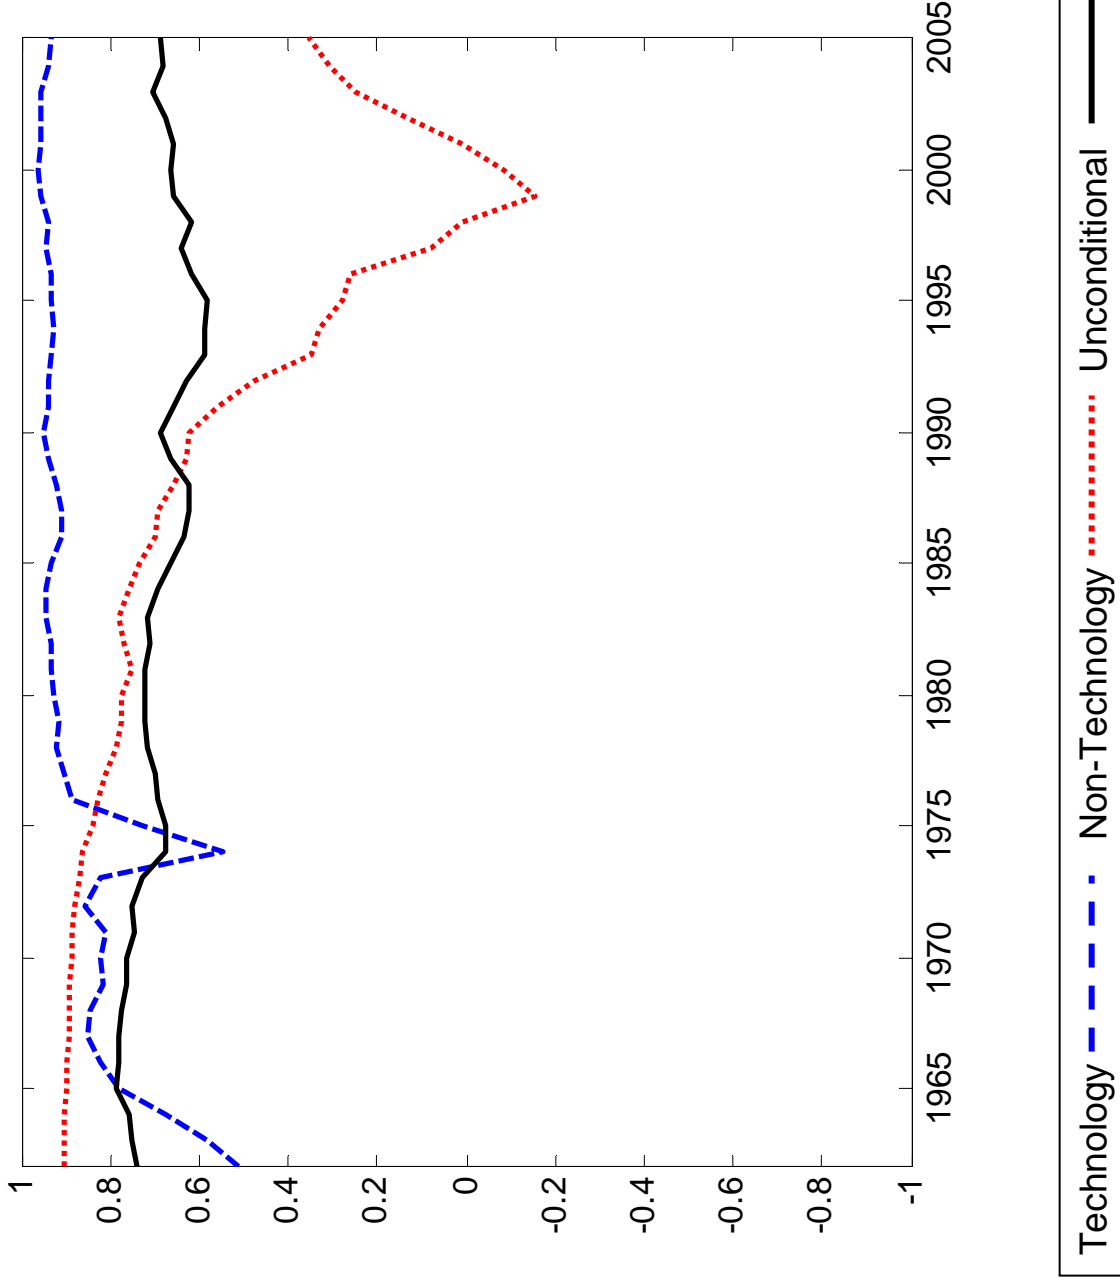


Figure 5a
Non-Technology Shocks: Output Response






Figure 5b



Pre-84
Post-84

Figure 5c



Post-84 minus Pre-84

Figure 6a
Non-Technology Shocks: Labor Productivity Response

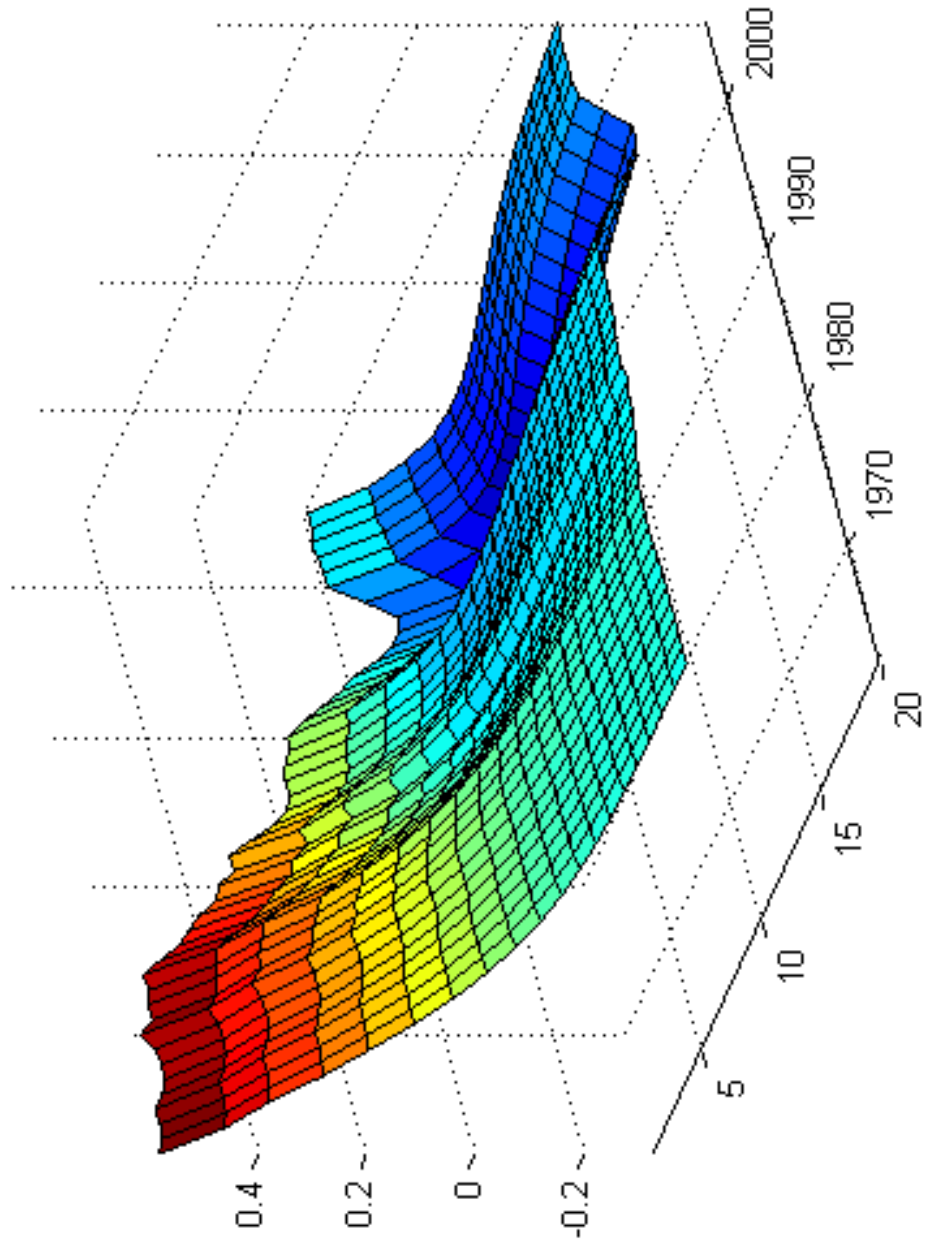
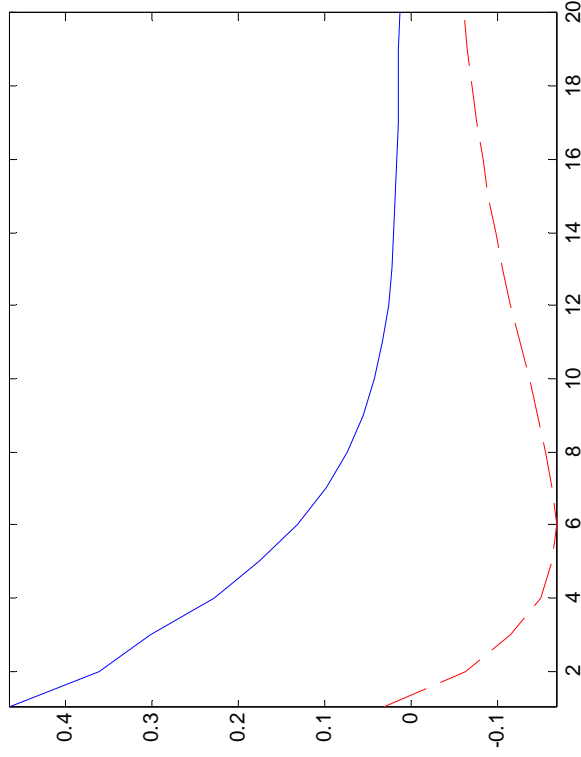



Figure 6b



Pre-84
Post-84

Figure 6c



Post-84 minus Pre-84

Figure 7a
Technology Shocks: Hours Response




Figure 7b





Figure 7c




Figure 8
Augmented Model: Conditional Output Volatilities




Figure 9

Augmented Model: Conditional Hours-Labor Productivity Correlations

