

On the Sources of the Great Moderation[†]

By JORDI GALÍ AND LUCA GAMBETTI*

The Great Moderation in the US economy has been accompanied by large changes in the 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 an increase in the volatility of hours relative to output, a shrinking contribution of nontechnology shocks to output volatility, and 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: E23, E24, J22, J24)

A large body of empirical research has provided evidence of a substantial decline in the volatility of most US 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) gross domestic product (GDP) and (log) nonfarm business output, before and after 1984, a date which is generally viewed as the starting point of the period of enhanced stability in the US economy. We use quarterly data from the first quarter of 1948 through the fourth quarter of 2005. 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

* Galí: CREI, Universitat Pompeu Fabra, Ramon Trias Fargas 25, 08005 Barcelona, Spain; (e-mail: jgali@crei.cat). Gambetti: Departament d’Economia i Història Econòmica, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain; (e-mail: luca.gambetti@uab.cat). We are grateful to Régis Barnichon, 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 for comments and suggestions. The authors acknowledge 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.

[†] To comment on this article in the online discussion forum visit the articles page at: <http://www.aeaweb.org/articles.php?doi=10.1257/mac.1.1.26>

¹ Early papers on the Great Moderation include those of Chang-Jin Kim and Charles R. Nelson (1999), Margaret M. McConnell and Gabriel Pérez-Quirós (2000), and Olivier J. Blanchard and John A. Simon (2001). A survey of the literature, as well as a discussion of alternative interpretations, can be found in James Stock and Mark W. Watson (2002). Stock and Watson (2005) and Stephen G. Cecchetti, Alfonso Flores-Lagunes, and Stefan Krause (2006) present and discuss some international evidence.

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

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

TABLE 1—THE GREAT MODERATION

	Standard deviation			<i>p</i> -value
	Pre-1984	Post-1984	Post-1984 Pre-1984	
<i>First-difference</i>				
GDP	1.21	0.54	0.44	<0.01
Nonfarm business output	1.57	0.68	0.43	<0.01
<i>BP-filter</i>				
GDP	2.00	0.94	0.47	<0.01
Nonfarm business output	2.59	1.23	0.47	<0.01

Notes: 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 eight-lag window. (See, Priestley 1991, 327.)

deviation for the post-1984 period is less than half that corresponding to the pre-1984 period. Tests of equality of the variance across subperiods 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 20 years is largely the result of smaller shocks impinging on the economy, with structural changes having played at most a secondary role.⁴ A second view, instead, attributes the reduction in aggregate volatility to changes in the economy’s structure and/or in the way policy has been conducted.⁵

In this paper, we provide evidence on some of the changes experienced by the US 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 the observed comovements among output, hours, and productivity; the identification of the sources of those comovements; and 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.

⁴ See, e.g., Alejandro Justiniano and Giorgio Primiceri (2006) and Andres Arias, Gary D. Hansen, and Lee E. Ohanian (2006) for examples of papers making a case for smaller shocks as an explanation for the volatility decline of the past two decades.

⁵ Such explanations include better monetary policy (e.g., Richard Clarida, Galí, and Mark Gertler 2000), improvements in inventory management (e.g., James A. Kahn, McConnell, and Perez-Quirós 2002), financial innovation and better risk sharing (e.g., Karen E. Dynan, Douglas W. Elmendorf, and Daniel E. Sichel 2006), and the optimal response of production and inventories policies to a decline in the persistence of automobile sales (Valerie A. Ramey and Daniel J. Vine 2006).

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

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 and in (log) output, which is given by their sum, as the result of two types of shocks impinging on the economy: technology and nontechnology 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 Timothy Cogley and Thomas J. Sargent (2005), Primiceri (2005), and Luca Benati and Haroon 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 time-varying coefficients overcomes the potential bias caused by the presence of significant low frequency comovements between productivity growth and hours in postwar US data, a problem first diagnosed by John G. Fernald (2007).⁷

In a way consistent with the literature, we uncover a large, (seemingly) permanent, decline in the volatility of output around the mid-1980s. 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 Kevin J. Stiroh (2006), much of that decline appears to be concentrated in the 1980s, 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 nontechnology shocks to the variance of output, both in absolute and relative terms. By contrast, the contribution

⁷ Fernald (2007) 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 Robert Vigfusson (2003).

⁸ Régis Barnichon (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.

of technology shocks to output volatility appears to have 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 nontechnology 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 signs 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 dot-com 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 Jonas D. M. 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 US economy, we believe it can be helpful in 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, because 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 a decline in the size of nontechnology shocks as well as more effective

⁹ Of 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.

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 nontechnology 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 I reports estimates of the standard deviations and correlations of output, hours, and labor productivity and their changes over time. Section II introduces the time-varying VAR approach used to estimate changes over time in conditional second moments and impulse responses, and presents the associated evidence. Section III presents the main empirical findings. Section IV shows the evidence based on the augmented SVAR model. Section V discusses possible interpretations and concludes.

I. The Labor Market and the Great Moderation: Basic Evidence

A. 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-1984 and post-1984 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 subperiods. We use quarterly data covering the sample period 1948:Q1–2005:QIV. All variables refer to the nonfarm business sector.¹⁰ After taking natural logarithms, we report estimates for both first-differenced and BP-filtered data.

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-1984 period. The size of that decline is not proportional, however. 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 relative standard deviations shown in the last three columns of Table 2. 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 of volatility in all variables.

¹⁰ We obtained our raw data from the Haver USECON database. The time series used includes output in the nonfarm business sector (LXNFO) and hours of all persons in nonfarm business (LXNFH). Both variables were normalized by the civilian, noninstitutional population of those 16 years old and older (LNN). Labor productivity was computed as the ratio between the output and hours measured mentioned above. The GDP measure used in Table 1 was drawn from the same database with the mnemonic GDPH.

TABLE 2—CHANGES IN VOLATILITY

	Standard deviation				Relative standard deviation		
	Pre-1984	Post-1984	Post-1984 Pre-1984	<i>p</i> -value	Pre-1984	Post-1984	Post-1984 Pre-1984
<i>First-difference</i>							
Output	1.57	0.68	0.43	0.00	1.00	1.00	1.00
Hours	1.05	0.65	0.61	0.00	0.66	0.95	1.41
Productivity	1.00	0.61	0.62	0.00	0.63	0.89	1.44
<i>BP-filter</i>							
Output	2.59	1.23	0.47	0.00	1.00	1.00	1.00
Hours	2.08	1.39	0.67	0.06	0.79	1.10	1.40
Productivity	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 eight-lag window. See M. B. Priestley (1981, 327).

TABLE 3—CHANGES IN CROSS-CORRELATIONS

	Pre-1984	Post-1984	Change
<i>First-difference</i>			
Output, hours	0.78	0.57	-0.20** (0.08)
Hours, productivity	0.18	-0.41	-0.59** (0.10)
Output, productivity	0.75	0.50	-0.24** (0.11)
<i>BP-filter</i>			
Output, hours	0.89	0.86	-0.02 (0.09)
Hours, productivity	0.18	-0.46	-0.65** (0.15)
Output, productivity	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 eight-lag window. See, e.g., George E. P. Box and Gwilym Jenkins (1976, 376).

**Significant at the 5 percent level.

*Significant at the 10 percent level.

B. 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-1984 and post-1984 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-1984 period. That result is considerably weaker, 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 (2006), 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 instead reflecting changes in either the composition of shocks or in the structure and transmission mechanisms operating in the US 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 a different nature as a source of those changes.

II. A VAR Model with Time-Varying Coefficients and Stochastic Volatility

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

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

$$(1) \quad \mathbf{x}_t = \mathbf{A}_{0,t} + \mathbf{A}_{1,t} \mathbf{x}_{t-1} + \mathbf{A}_{2,t} \mathbf{x}_{t-2} + \dots + \mathbf{A}_{p,t} \mathbf{x}_{t-p} + \mathbf{u}_t,$$

where $\mathbf{A}_{0,t}$ is a vector of time-varying intercepts, and $\mathbf{A}_{i,t}$, $i = 1, \dots, p$, 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

¹¹ As 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 (2007).

sequence of innovations $\{\mathbf{u}_t\}$ follows a Gaussian white noise process with zero mean and time-varying covariance matrix Σ_t , and uncorrelated with all lags of \mathbf{x}_t . Letting $\mathbf{A}_t = [\mathbf{A}_{0,t}, \mathbf{A}_{1,t}, \dots, \mathbf{A}_{p,t}]$, we define $\boldsymbol{\theta}_t = \text{vec}(\mathbf{A}'_t)$, where $\text{vec}(\cdot)$ is the column stacking operator. Conditional on the roots of the associated VAR polynomial being outside the unit circle for all t , we assume $\boldsymbol{\theta}_t$ evolves over time according to the process

$$(2) \quad \boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t,$$

where $\boldsymbol{\omega}_t$ is a Gaussian white noise process with zero mean and constant covariance $\boldsymbol{\Omega}$, and independent of \mathbf{u}_t at all leads and lags.

We model the time variation for Σ_t as follows. Let $\Sigma_t \equiv \mathbf{F}_t \mathbf{D}_t \mathbf{F}'_t$, where \mathbf{F}_t is lower triangular with ones in the main diagonal, and \mathbf{D}_t is a diagonal matrix.¹² Let $\boldsymbol{\gamma}_t$ be a vector containing all the elements of \mathbf{F}_t^{-1} below the diagonal, stacked by rows, and $\boldsymbol{\sigma}_t$ be the vector of diagonal elements of \mathbf{D}_t . We assume

$$(3) \quad \boldsymbol{\gamma}_t = \boldsymbol{\gamma}_{t-1} + \boldsymbol{\zeta}_t$$

$$(4) \quad \log \boldsymbol{\sigma}_t = \log \boldsymbol{\sigma}_{t-1} + \boldsymbol{\xi}_t,$$

where $\boldsymbol{\zeta}_t$ and $\boldsymbol{\xi}_t$ are Gaussian white noise processes with zero mean and (constant) covariance matrices $\boldsymbol{\Psi}$ and $\boldsymbol{\Xi}$, respectively. We assume that $\boldsymbol{\Psi}$ has a block diagonal structure, i.e., all the covariances between coefficients belonging to different equations are zero, and that $\boldsymbol{\Xi}$ is diagonal. Finally, we assume that $\boldsymbol{\xi}_t$, $\boldsymbol{\zeta}_t$, and $\boldsymbol{\omega}_t$ are all mutually independent.

We assume that the vector of VAR innovations \mathbf{u}_t is a (time-varying) linear transformation of the vector of underlying “structural” shocks $\boldsymbol{\varepsilon}_t \equiv [\boldsymbol{\varepsilon}_t^a, \boldsymbol{\varepsilon}_t^d]'$, satisfying $E\{\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t'\} = \mathbf{I}$ for all t , where $\boldsymbol{\varepsilon}_t^a$ represents a technology shock, and $\boldsymbol{\varepsilon}_t^d$ is a non-technology shock (which is occasionally referred to, for convenience, as a “demand” shock). Thus, we assume $\mathbf{u}_t = \mathbf{K}_t \boldsymbol{\varepsilon}_t$ for all t for some nonsingular matrix \mathbf{K}_t satisfying $\mathbf{K}_t \mathbf{K}_t' = \Sigma_t$. 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 \mathbf{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 will see next, that assumption imposes some restrictions that allow us to recover matrix \mathbf{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 $\mathbf{x}_t \equiv [\mathbf{x}'_t, \mathbf{x}'_{t-1}, \dots, \mathbf{x}'_{t-p+1}]'$, $\mathbf{u}_t \equiv [\mathbf{u}'_t, 0, \dots, 0]'$, $\boldsymbol{\mu}_t \equiv [\mathbf{A}'_{0,t}, 0, \dots, 0]'$, and \mathbf{A}_t is the corresponding companion matrix. We use a local approximation of the implied

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

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

$$\frac{\partial \mathbf{x}_{t+k}}{\partial \mathbf{u}'_t} = e_{2,2}(\mathbf{A}_t^k) \equiv \mathbf{B}_{t,k}$$

for $k = 1, 2, \dots$, where $e_{2,2}(\mathbf{M})$ is a function which selects the first two rows and two columns of any matrix \mathbf{M} , and where $\mathbf{B}_{t,0} \equiv \mathbf{I}$. 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

$$\begin{aligned} \frac{\partial \mathbf{x}_{t+k}}{\partial \boldsymbol{\varepsilon}'_t} &= \frac{\partial \mathbf{x}_{t+k}}{\partial \mathbf{u}'_t} \frac{\partial \mathbf{u}_t}{\partial \boldsymbol{\varepsilon}'_t} \\ &= \mathbf{B}_{t,k} \mathbf{K}_t \equiv \mathbf{C}_{t,k} \end{aligned}$$

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

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

$$\tilde{\mathbf{C}}_{t,\infty} \tilde{\mathbf{C}}'_{t,\infty} = \tilde{\mathbf{B}}_{t,\infty} \mathbf{\Sigma}_t \tilde{\mathbf{B}}'_{t,\infty}$$

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

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

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 the Appendix for a detailed description of the method used to estimate that model, which follows Primiceri (2005).

Our analysis 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

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

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

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

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

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

with each of the terms on the right-hand side representing the covariances at time t conditional on technology and nontechnology 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.

III. Changing Labor Market Dynamics and the Great Moderation

A. 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 then stabilizing 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.

¹³ Here, and in subsequent figures, we report statistics starting in 1962:Q1, 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.

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

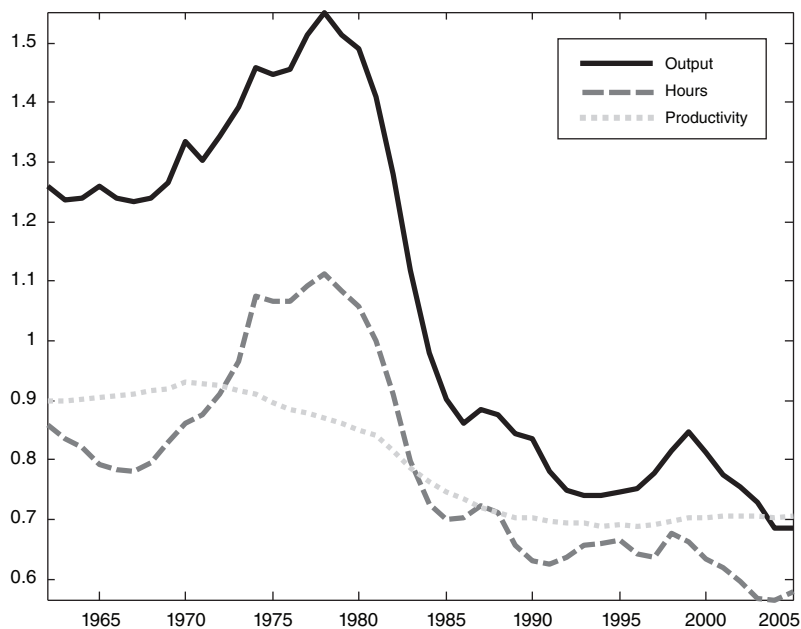


FIGURE 1A. TIME-VARYING STANDARD DEVIATIONS



FIGURE 1B. TIME-VARYING RELATIVE STANDARD DEVIATIONS

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, 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

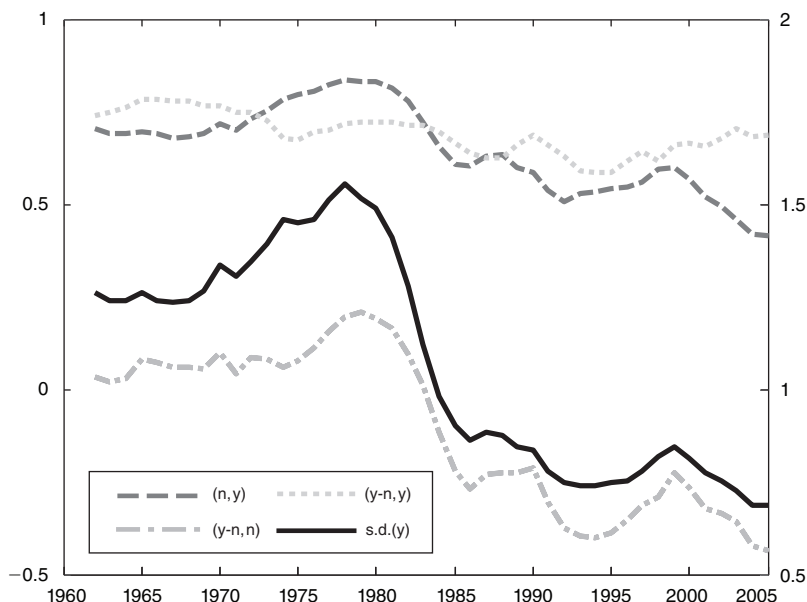


FIGURE 2. TIME-VARYING UNCONDITIONAL CORRELATIONS

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.

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.

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

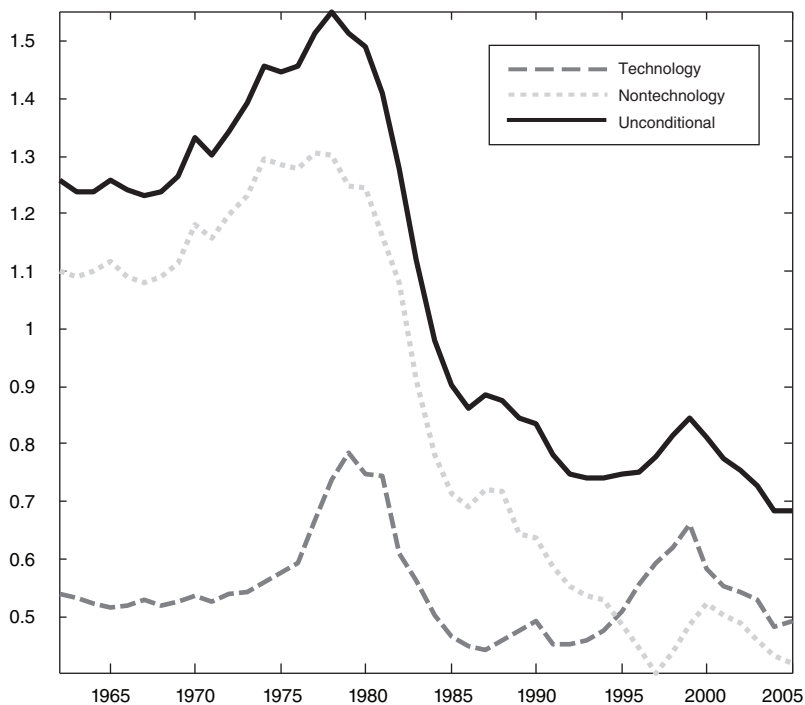


FIGURE 3A. CONDITIONAL STANDARD DEVIATIONS: OUTPUT

B. 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–3C plot the estimates of the (time-varying) standard deviations of each of those variables conditional on technology (dashed line) and nontechnology shocks (dotted line), as implied by our estimated SVAR. In each case, and as reference, we also plot the unconditional standard deviation (solid line).

The pattern that emerges in Figure 3A is unambiguous. The Great Moderation can be largely accounted for by the decline in the contribution of nontechnology 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 nontechnology shocks in the early 1960s, 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 nontechnology shocks explain the bulk of the pattern in the

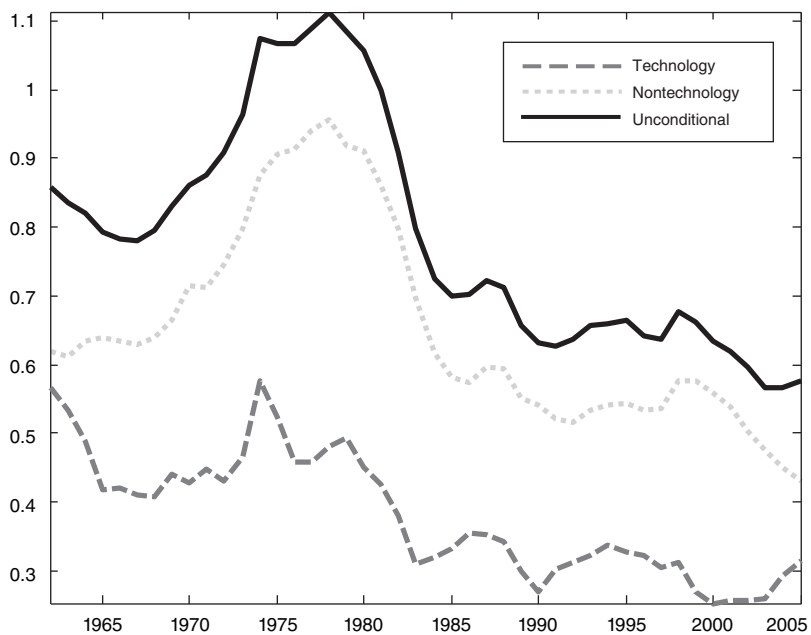


FIGURE 3B. CONDITIONAL STANDARD DEVIATIONS: HOURS

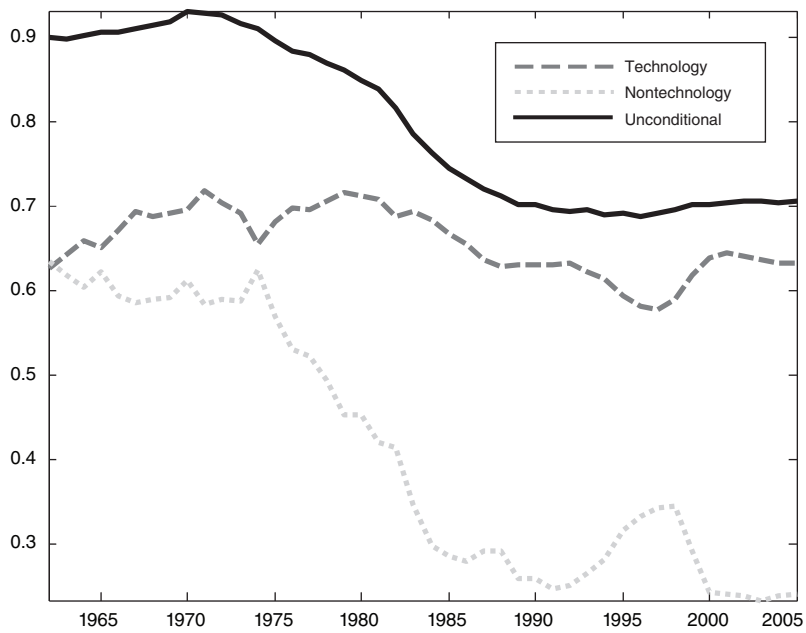


FIGURE 3C. CONDITIONAL STANDARD DEVIATIONS: LABOR PRODUCTIVITY

TABLE 4—CHANGES IN CONDITIONAL VOLATILITY

	Nontechnology shocks				Technology shocks			
	Pre-1984	Post-1984	Post-1984		Pre-1984	Post-1984	Post-1984	
			Pre-1984	<i>p</i> -value			Pre-1984	<i>p</i> -value
<i>First-difference</i>								
Output	1.18	0.56	0.48	0.00	0.46	0.47	1.01	0.91
Hours	0.78	0.59	0.76	0.12	0.39	0.22	0.56	0.00
Productivity	0.53	0.17	0.33	0.00	0.68	0.93	0.88	0.51
<i>BP-filter</i>								
Output	1.97	1.17	0.59	0.07	0.59	0.49	0.83	0.43
Hours	1.60	1.29	0.81	0.40	0.57	0.32	0.56	0.04
Productivity	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 eight-lag window. See Priestley (1991, 327).

standard deviation of hours, including its rise in the 1970s and 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 US 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 nontechnology shocks are responsible for the secular decline over the postwar period in the volatility of labor productivity. Interestingly, the decline in the contribution of nontechnology shocks to that volatility is seen to start in the mid 1970s, well before the onset of the Great Moderation period.

Table 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 nontechnology components of output, hours, and labor productivity for the pre-1984 and post-1984 sample periods. In contrast with the evidence reported in Figures 3A–3C, the statistics reported in Table 4 depend, not only on the estimated moving average coefficients (the $C_{i,k}^{ij}$'s of Section II), 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, nontechnology 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.

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

An important caveat must be raised at this point. So far, our analysis 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 nontechnology 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, henceforth AHO), and which is claimed by the latter authors to fully account for the decline in the cyclical volatility of output. To be more specific, AHO show that the standard deviation of measured total factor productivity (TFP) has declined by a factor of about one-half between the pre-1984 and post-1984 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.

C. Conditional Correlations and Structural Change

In Figures 4A–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 (nontechnology) 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 ,

$$\text{corr}(x_t, z_t) = \lambda_a \text{corr}_a(x_t, z_t) + \lambda_d \text{corr}_d(x_t, z_t),$$

where $\lambda_i \equiv [\sigma_i(x_t)/\sigma(x_t)][\sigma_i(z_t)/\sigma(z_t)]$, and where $\text{corr}_i(x_t, z_t)$ 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.

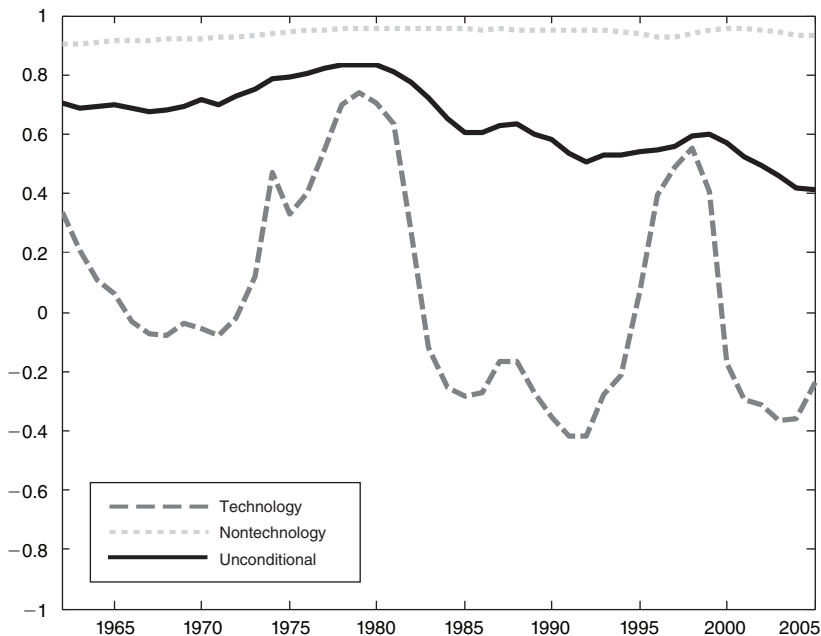


FIGURE 4A. CONDITIONAL CORRELATIONS: HOURS—OUTPUT

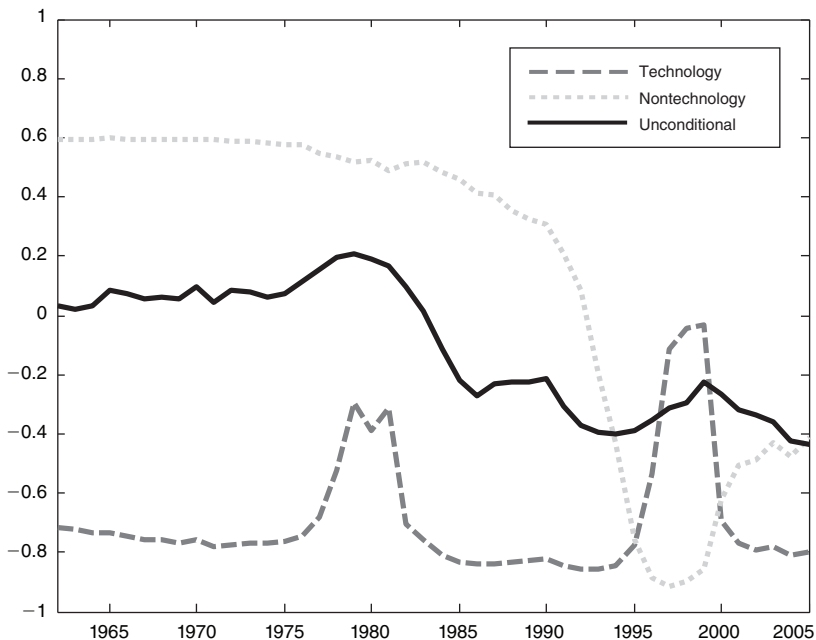


FIGURE 4B. CONDITIONAL CORRELATIONS: LABOR PRODUCTIVITY—HOURS

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 nontechnology shocks (dotted line) with a correlation that fluctuates between positive and (slightly) negative values as a result of technology

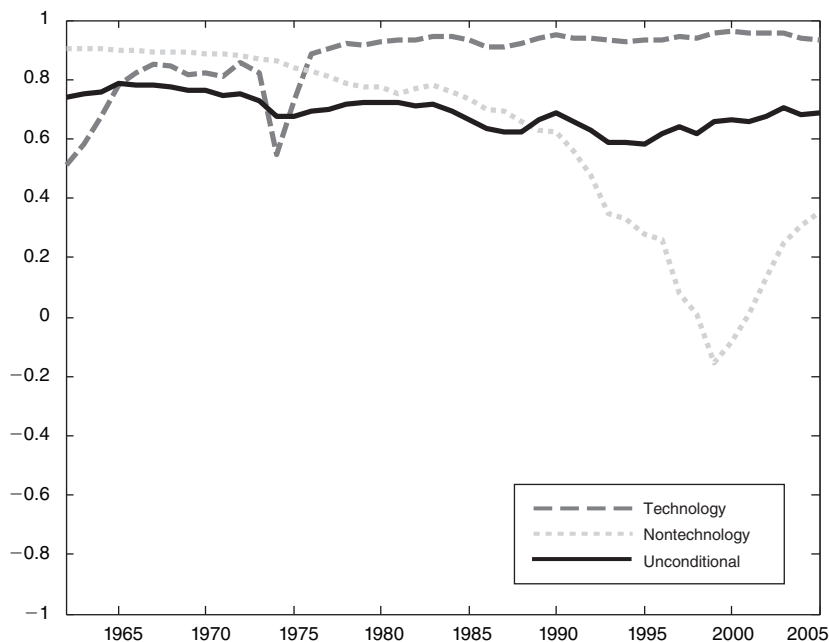


FIGURE 4C. CONDITIONAL CORRELATIONS: LABOR PRODUCTIVITY—OUTPUT

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 nonnegligible 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 nontechnology 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

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

TABLE 5—CHANGES IN CONDITIONAL CORRELATIONS

	Nontechnology shocks			Technology shocks		
	pre-1984	post-1984	change	pre-1984	post-1984	change
<i>First-difference</i>						
Output, hours	0.93	0.95	-0.02 (0.10)	-0.27	-0.65	-0.37** (0.11)
Hours, productivity	0.61	-0.30	-0.93** (0.15)	-0.75	-0.83	-0.07 (0.06)
Output, productivity	0.85	-0.01	-0.87** (0.17)	0.83	0.96	0.12** (0.02)
<i>BP-filter</i>						
Output, hours	0.98	0.98	-0.01 (0.12)	0.06	-0.33	-0.40* (0.23)
Hours, productivity	0.67	-0.57	-1.24** (0.14)	-0.66	-0.73	-0.06 (0.13)
Output, productivity	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 eight-lag window. See, e.g., Box and Jenkins (1976, 376).

**Significant at the 5 percent level.

*Significant at the 10 percent level.

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 induces a negative correlation between hours and labor productivity,

Finally, in Figure 4C, we show 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 nontechnology 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-1984 and post-1984 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 nontechnology 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

¹⁸ Note 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.

The evidence provided above suggests that at least two of the observed changes in unconditional correlations (those involving labor productivity) described in Section I, and earlier in this section, can be attributed to a (large) change in conditional correlations, associated with nontechnology shocks. Furthermore, and as discussed above, the timing of some of those changes matches 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 shocks, hinting instead (though admittedly without proving it) at a potential role for structural change.

D. 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 two of the most significant findings uncovered above, namely, (a) the decline in output volatility resulting from a smaller contribution of nontechnology shocks, and (b) 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 nontechnology shocks. Figure 5A displays the evolution over time of the dynamic response of output to a nontechnology 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 nontechnology 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-1984 and post-1984 periods. Figure 5C shows the difference between those two impulse responses, together with a 68 percent (dashed) and 95 percent (dotted) confidence band 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 is 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 cyclical of labor productivity conditional on nontechnology 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 nontechnology shock has a large and persistent positive effect on labor productivity in the early part of the sample, an observation consistent with the

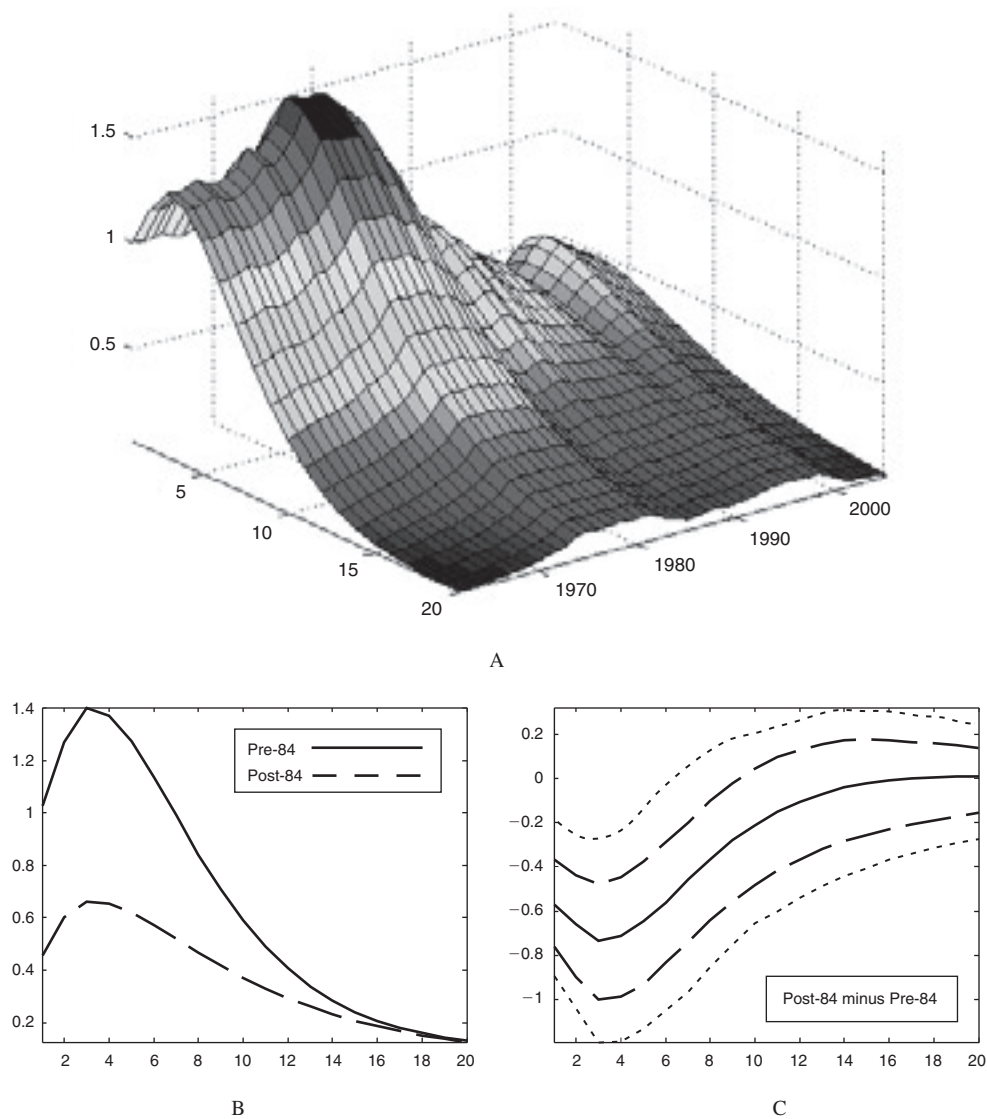


FIGURE 5. Nontechnology Shocks: Output Response

evidence of so-called “short-run increasing returns to labor” (SRIRL) uncovered by a number of economists.¹⁹ Starting in the early 1980s, however, the SRIRL phenomenon vanishes gradually. The response of labor productivity keeps getting smaller over time until eventually it 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-1984 and post-1984 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).

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

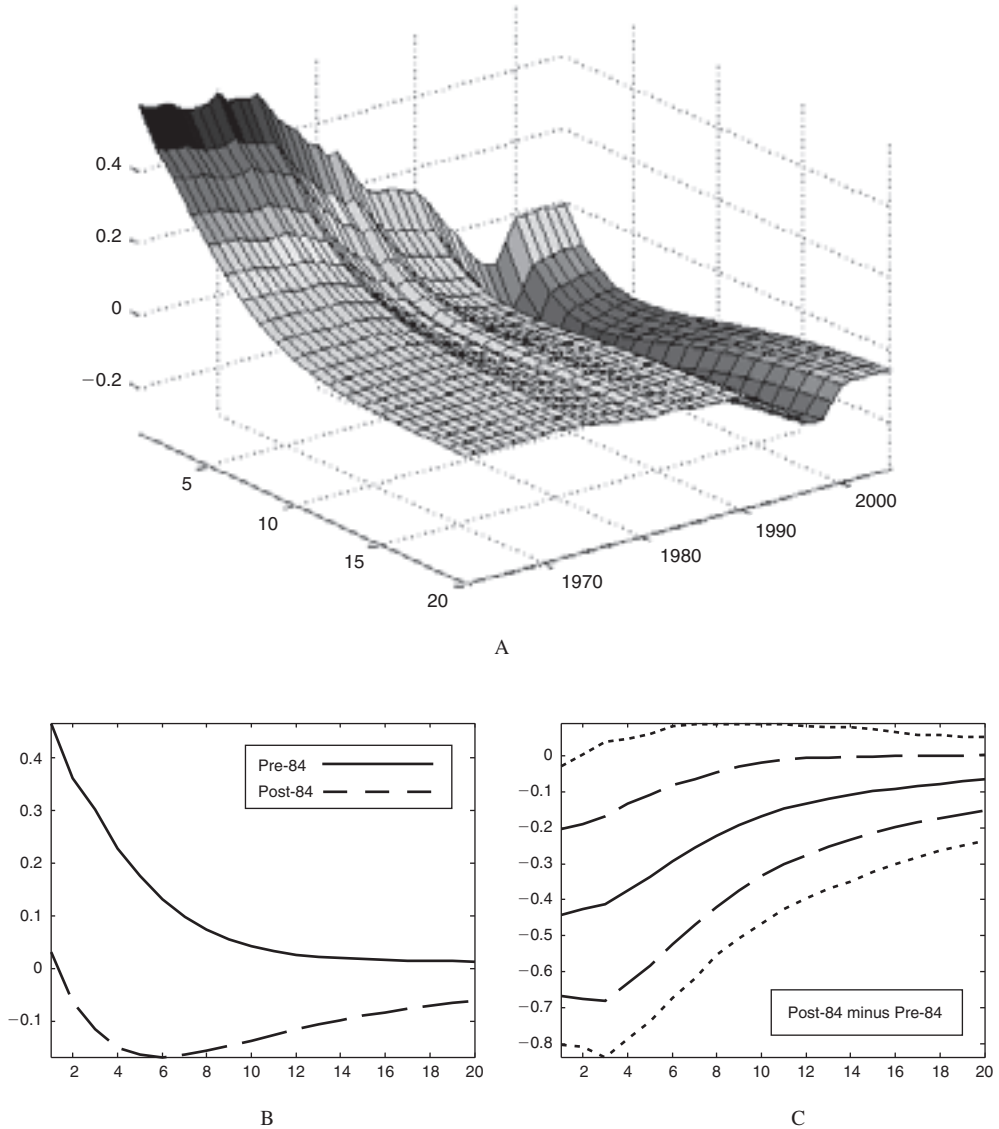


FIGURE 6. Nontechnology Shocks: Labor Productivity Response

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); Susanto Basu, John G. Fernald, and Miles S. Kimball (2005); and Neville 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 the response. In that respect, we note that, some fluctuations notwithstanding, the size of the negative

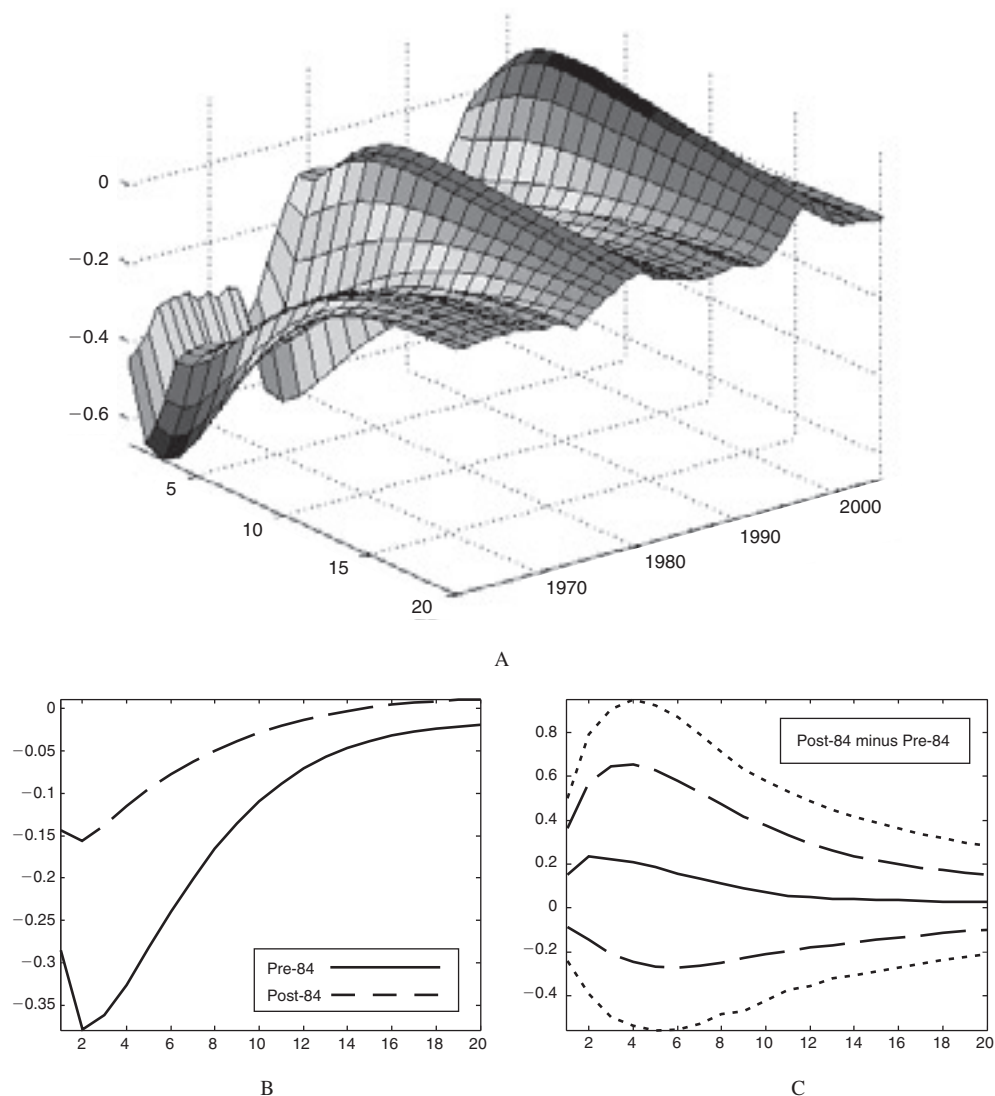


FIGURE 7. TECHNOLOGY SHOCKS: HOURS RESPONSE

response of hours appears to have gone down over time (in absolute value).²⁰ This is reflected in the gap between the “average” impulse responses for the pre- and post-1984 periods shown in Figure 7B, though the gradual change combined with the large confidence bands associated with 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

²⁰ The previous finding accords with the evidence, reported in Galí, David López-Salido, and Javier 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.

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.

E. 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 that 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 nontechnology 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

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. Second, both I-shocks (dashed-dotted line) and nontechnology shocks (dotted line) play an important role in the Great Moderation. Interestingly, however, the patterns of their contribution differ substantially. Roughly speaking, while nontechnology 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.

Thus our augmented model points to an important role of I-shocks as a source of 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 nontechnology shocks to the decline in output volatility appears to be robust to the alternative specification considered here, though the (previously dominant) role

²¹ The 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: $PCDG/PCDGCC96$ (durables), $PCND/PCNDGC96$ (nondurables), $PCESV/PCESVC96$ (services), and $FPI/FPIC1$ (investment).

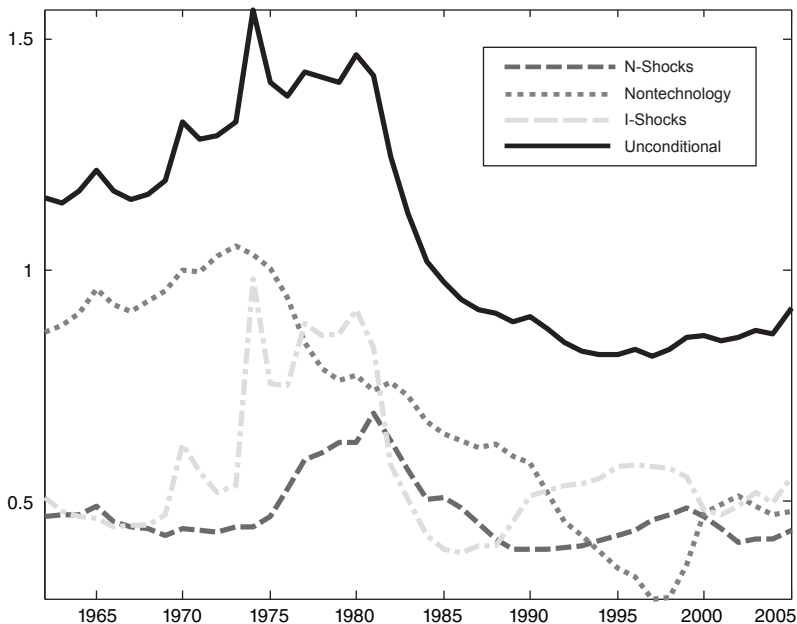


FIGURE 8. AUGMENTED MODEL: CONDITIONAL OUTPUT VOLATILITIES

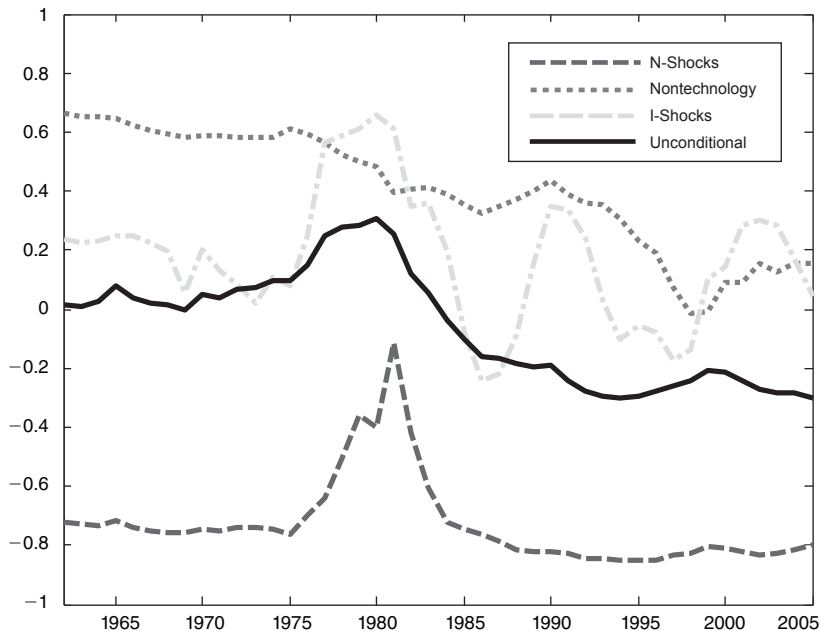


FIGURE 9. AUGMENTED MODEL: CONDITIONAL HOURS-LABOR PRODUCTIVITY CORRELATIONS

of nontechnology shocks in the abrupt volatility decline of the early 1980s 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 nontechnology 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 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.

IV. Tentative Interpretations and Caveats

The remarkable decline in macroeconomic volatility experienced by the US economy since the mid-1980s (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 this paper makes clear, volatility decline has been accompanied by large changes in the patterns of comovements among output, hours, and labor productivity. Those changes are reflected in 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 nontechnology shocks. Second, the Great Moderation period has witnessed a dramatic fall (with sign switch included) in the correlation between hours and labor productivity generated by nontechnology shocks.

The shrinking contribution of nontechnology shocks to output volatility can be due, in principle, to two developments that are not mutually exclusive. First, the average size of the underlying shocks may have become smaller. Second, 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 nontechnology 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-1984 period relative to the earlier period.²³ That evidence is consistent with our finding of a smaller contribution of nontechnology shocks. Yet, and at least in the case of policy shocks, it can hardly be interpreted as being consistent with the

²² Of 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 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, in principle, be caused by nonlinearities 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.

²³ See, in particular, section 5.4 in Stock and Watson (2002) and section 5.D in Frank Smets and Rafael Wouters (2007).

“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 nontechnology 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 nontechnology 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 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 accounts of the role played by the Fed under Alan Greenspan 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 nontechnology 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 independently and identically distributed disturbance uncorrelated with n_t^* . Assuming, for the sake of illustration, a simple production function (in logs) of the form

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

²⁴ See, e.g., John B. Taylor (1999); Clarida, Galí, and Gertler (2000); and Jean Boivin and Marc Giannoni (2006).

²⁵ Evidence of smaller technology shocks in the post-1984 period can be found in Stock and Watson (2002) and Smets and Wouters (2007), among others.

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

where ξ_t represents variations in nonlabor inputs.²⁷ Combining the previous assumptions, we obtain

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

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

In the setup above, a reduction in the degree of labor hoarding γ could potentially account for three of our findings: (a) the increase in the volatility of hours relative to output, (b) the decline in the response of labor productivity to expansionary nontechnology shocks with an eventual switch in the sign of that response (if γ becomes smaller than α), and (c) the shrinking correlation between hours and labor productivity conditional on nontechnology 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 we have also shown that those changes in second moments are largely associated with changes in the economy's response to nontechnology 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 this 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 US 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

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

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

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

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

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

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

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 with 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:QIV. Following Benati and Mumtaz (2007) and Primiceri (2005), we make the following assumptions about prior densities and parameters:

$$p(\theta_0) \propto I(\theta_0) N(\hat{\theta}_{OLS}, \hat{V}(\hat{\theta}_{OLS})),$$

$$p(\log \sigma_0) = N(\log \hat{\sigma}_{OLS}, 10 \times \mathbf{I}),$$

$$p(\gamma_0) = N(\hat{\gamma}_{OLS}, |\hat{\gamma}_{OLS}|),$$

$$p(\Omega) = IW(\bar{\Omega}^{-1}, T_0),$$

$$p(\Psi) = IW(\bar{\Psi}^{-1}, 2), \text{ and}$$

$$p(\Xi_{i,i}) = IG\left(\frac{0.0001}{2}, \frac{1}{2}\right),$$

where $\hat{\theta}_{OLS}$ is the vector of OLS estimates of the VAR coefficients, and $\hat{V}(\hat{\theta}_{OLS})$ is the estimate of their covariance matrix using the initial sample, $\hat{\sigma}_{OLS}$ is a vector containing the elements of the diagonal matrix $\hat{\mathbf{D}}$, and $\hat{\gamma}_{OLS}$ is the element (2, 1) of the lower triangular matrix $\hat{\mathbf{F}}^{-1}$, where $\hat{\mathbf{F}} \hat{\mathbf{D}} \hat{\mathbf{F}}' = \hat{\Sigma}_{OLS}$, and $\bar{\Omega} = 0.005 \times V(\hat{\theta}_{OLS})$, T_0 is the number of observations in the initial sample, and $\bar{\Psi} = 0.001 \times |\hat{\gamma}_{OLS}|$.

B. Estimation

To draw realizations from the posterior density, we use an MCMC, the Gibbs sampler, algorithm which works in an iterative way. Each iteration is done in four steps and consists of 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: $p(\boldsymbol{\theta}^T | \mathbf{x}^T, \boldsymbol{\gamma}^T, \boldsymbol{\sigma}^T, \boldsymbol{\Psi}, \boldsymbol{\Xi}, \boldsymbol{\Omega})$

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

$$(A1) \quad \boldsymbol{\theta}_{t|t+1} = \boldsymbol{\theta}_{t|t} + \mathbf{P}_{t|t} \mathbf{P}_{t|t+1}^{-1} (\boldsymbol{\theta}_{t+1} - \boldsymbol{\theta}_{t|t}),$$

$$(A2) \quad \mathbf{P}_{t|t+1} = \mathbf{P}_{t|t} \mathbf{P}_{t+1|t}^{-1} \mathbf{P}_{t|t},$$

where $p(\boldsymbol{\theta}_t | \boldsymbol{\theta}_{t+1}, \mathbf{x}^T, \boldsymbol{\gamma}^T, \boldsymbol{\sigma}^T, \boldsymbol{\phi}) \sim N(\boldsymbol{\theta}_{t|t+1}, \mathbf{P}_{t|t+1})$.

Step 2: $p(\boldsymbol{\gamma}^T | \mathbf{x}^T, \boldsymbol{\theta}^T, \boldsymbol{\sigma}^T, \boldsymbol{\Psi}, \boldsymbol{\Xi}, \boldsymbol{\Omega})$

This is done following the same procedure described in Primiceri (2005). Conditional on $\boldsymbol{\theta}^T \hat{\mathbf{y}}_t = \mathbf{x}_t - \mathbf{A}_{0,t} - \mathbf{A}_{1,t} \mathbf{x}_{t-1} - \dots - \mathbf{A}_{p,t} \mathbf{x}_{t-p}$ is observable. We can rewrite our system of equations as $\mathbf{F}_t^{-1} \hat{\mathbf{y}}_t = \mathbf{D}_t \boldsymbol{\nu}_t$, where $\boldsymbol{\nu}_t \sim N(0, \mathbf{I})$. Conditional on $\boldsymbol{\sigma}^T$, we use the algorithm of Carter and Kohn to obtain a draw for $\boldsymbol{\gamma}_t$ taking the above system as observational equations and (3) as unobserved states equations. Given that the $\boldsymbol{\gamma}_t$ and the $\boldsymbol{\nu}_t$ are independent across equations, the algorithm can be applied equation by equation. Notice, however, that, in the bivariate case, we have one observable equation and one state.

Step 3: $p(\boldsymbol{\sigma}^T | \mathbf{x}^T, \boldsymbol{\theta}^T, \boldsymbol{\gamma}^T, \boldsymbol{\Psi}, \boldsymbol{\Xi}, \boldsymbol{\Omega})$

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

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

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

We perform 30,000 repetitions. We discard the first 10,000 draws and 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?" National Bureau of Economic Research Working Paper 12079.

- Barnichon, Régis.** 2006. "Productivity, Aggregate Demand and Unemployment Fluctuations." Unpublished.
- ▶ **Basu, Susanto, John G. Fernald, and Miles S. Kimball.** 2006. "Are Technology Improvements Contractionary?" *American Economic Review*, 96(5): 1418–48.
- ▶ **Baxter, Marianne, and Robert G. King.** 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics*, 81(4): 575–93.
- Benati, Luca, and Haroon Mumtaz.** 2007. "U.S. Evolving Macroeconomic Dynamics: A Structural Investigation." European Central Bank Working Paper 746.
- Blanchard, Olivier J., and John A. Simon.** 2001. "The Long and Large Decline in US Output Volatility." *Brookings Papers on Economic Activity*, (1): 135–64.
- ▶ **Boivin, Jean, and Marc Giannoni.** 2006. "Has Monetary Policy Become More Effective?" *Review of Economics and Statistics*, 88(3): 445–62.
- Box, George E. P., and Gwilym M. Jenkins.** 1976. *Time Series Analysis: Forecasting and Control*. Oakland, CA: Holden-Day.
- ▶ **Carter, Chris K., and Robert Kohn.** 1994. "On Gibbs Sampling for State Space Models." *Biometrika*, 81(3): 541–53.
- Cecchetti, Stephen G., Alfonso Flores-Lagunes, and Stefan Krause.** 2006. "Assessing the Sources of Changes in the Volatility of Real Growth." National Bureau of Economic Research Working Paper 11946.
- Christiano, Lawrence J., and Martin Eichenbaum.** 1992. "Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations." *American Economic Review*, 82(3): 430–50.
- Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson.** 2003. "What Happens after a Technology Shock?" National Bureau of Economic Research 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*, 115(1): 147–80.
- ▶ **Cogley, Timothy, and Thomas J. Sargent.** 2005. "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S." *Review of Economic Dynamics*, 8(2): 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*, 53(1): 123–50.
- Fernald, John G.** 2007. "Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements." *Journal of Monetary Economics*, 54(8): 2467–85.
- ▶ **Fisher, Jonas D. M.** 2006. "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy*, 114(3): 413–51.
- ▶ **Francis, Neville, and Valerie A. Ramey.** 2005. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics*, 52(8): 1379–99.
- Galí, Jordi.** 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89(1): 249–71.
- 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(4): 723–43.
- Galí, Jordi, and Pau Rabanal.** 2004. "Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar US Data?" In *NBER Macroeconomics Annual, 2004*, ed. Mark Gertler and Kenneth Rogoff, 225–318. Cambridge, MA: MIT Press.
- Gambetti, Luca.** 2006. "Technology Shocks and the Response of Hours Worked: Time-Varying Dynamics Matter." Unpublished.
- Gelman Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin.** 1995. *Bayesian Data Analysis*. London: Chapman and Hall.
- Gordon, Robert J.** 1990. "Are Procyclical Productivity Fluctuations a Figment of Measurement Error?" Unpublished.
- Hansen, Gary D., and Randall Wright.** 1992. "The Labor Market in Real Business Cycle Theory." *Federal Reserve Bank of Minneapolis Quarterly Review*, 16(2): 2–12.
- Jaquier, Eric, Nicholas G. Polson, and Peter E. Rossi.** 2004. "Bayesian Analysis of Stochastic Volatility Models." *Journal of Business and Economic Statistics*, 12(4): 371–418.
- Justiniano, Alejandro, and Giorgio E. Primiceri.** 2006. "The Time Varying Volatility of Macroeconomic Fluctuations." National Bureau of Economic Research Working Paper 12022.
- Kahn, James A., Margaret M. McConnell, and Gabriel Perez-Quirós.** 2002. "On the Causes of the Increased Stability of the US Economy." *Federal Reserve Bank of New York Economic Policy Review*, 8(1): 183–202.

- ▶ **Kim, Chang-Jin, and Charles R. Nelson.** 1999. "Has the US Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle." *Review of Economics and Statistics*, 81(4): 608–16.
- McConnell, Margaret M., and Gabriel Perez-Quirós.** 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980s?" *American Economic Review*, 90(5): 1464–76.
- Priestley, M. B.** 1981. *Spectral Analysis and Time Series*. San Diego: Academic Press.
- ▶ **Primiceri, Giorgio E.** 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies*, 72(3): 821–52.
- ▶ **Ramey, Valerie A., and Daniel J. Vine.** 2006. "Declining Volatility in the US Automobile Industry." *American Economic Review*, 96(5): 1876–89.
- Sbordone, Argia M.** 1996. "Cyclical Productivity in a Model of Labor Hoarding." *Journal of Monetary Economics*, 38(2): 331–61.
- ▶ **Smets, Frank, and Rafael Wouters.** 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.
- Stiroy, Kevin J.** 2006. "Volatility Accounting: A Production Perspective on Increased Economic Stability." Federal Reserve Bank of New York Staff Report 245.
- Stock, James, and Mark W. Watson.** 2002. "Has the Business Cycle Changed and Why?" In *NBER Macroeconomics Annual, 2002*, ed. Mark Gertler and Kenneth S. Rogoff. Cambridge, MA: MIT Press, 159–218.
- ▶ **Stock, James H., and Mark W. Watson.** 2005. "Understanding Changes in International Business Cycle Dynamics." *Journal of the European Economic Association*, 3(5): 968–1006.
- Taylor, John B.** 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, ed. J. B. Taylor, 319–48. Chicago: University of Chicago Press.