Sensitivity of Impulse Responses to Small Low Frequency Co-movements: Reconciling the Evidence on the Effects of Technology Shocks

Nikolay Gospodinov* — Concordia University
Alex Maynard — University of Guelph
Elena Pesavento — Emory University

May 12, 2008

Abstract

We clarify the empirical source of the debate on the effect of technology shocks on hours worked. We argue that the contrasting conclusions from level and difference VAR specifications are not primarily due to the presence or absence of a unit root in hours worked. Rather they are due to a small, but important, low frequency comovement between hours worked and labour productivity growth. This comovement is allowed for in the levels specification, but assumed away in the difference VAR. Our theoretical analysis shows how the empirical results from both side of the debate can be reconciled.

Keywords: Technology shocks, impulse response functions, structural VAR, long-run identification.

JEL Classification: C32, C12, F40.

---

*Corresponding author. Department of Economics. Concordia University, 1455 de. Maisonneuve Blvd. West, Montreal, Quebec, H3G 1M8. nikolay.gospodinov@concordia.ca. Part of this research was done while Elena Pesavento was a Marco Fanno visitor at the Dipartiment di Scienze Economiche at Universita' di Padova. The authors thank Barbara Rossi, Efrem Castelnuovo and seminar partcipants at University of Padova.

1 Introduction and Motivation

An ongoing debate exists regarding the empirical effect of technology shocks on production inputs, such as hours worked. Most standard real business cycle models start with the premise that business cycles result from unexpected changes in production technologies. This has the implication that hours worked and other inputs to production should rise following a positive technology shock. On the other hand, models with frictions, such as sticky prices, often predict an initial fall in hours worked following a productivity shock.

As technology shocks are difficult to measure,[1] following Gali (1999), they are commonly specified as structural shocks in vector autoregressive (VAR) models that are identified via the long-run restriction that only technology shocks have a permanent effect on labour productivity. This identification scheme, an implication of many modern macroeconomic models, has been widely employed in recent years. However, despite its common acceptance, the qualitative results have proven quite sensitive to other aspects of the VAR specification, particularly whether hours worked is specified in levels or differences.

Specifying the VAR in the difference of both hours worked and labour productivity, Gali (1999) and Shea (1999) find that hours worked initially fall following a positive technology shock, a finding which gives support to models with frictions embedded. Other papers have reached similar conclusions (see, for example, Francis and Ramey, 2005; Francis, Owyang and Theodorou, 2003; Basu, Fernald and Kimball, 2006) and this has spurred a line of research aimed at developing general equilibrium models that can account for this empirical finding (Uhlig, 2003; Gali and Rabanal, 2004). However, maintaining the same identification scheme, but allowing hours worked to enter the model in levels, Christiano, Eichenbaum and Vigfusson (2003, 2006) instead provided support for the prediction of standard RBC models, with hours worked rising immediately after a positive productivity shock. Figure 1 plots the estimated impulse response functions (IRFs) based on the

[1] Alexopoulos (2006) and Shea (1999) provide measurements of technological progress based on technology publications and patent data respectively.

levels and differenced specifications with quarterly U.S. data for the period 1948Q2 - 2005Q3.[2] The difference in the impulse response functions are quite striking. Standard unit root and stationarity tests on hours worked, neither of which reject their respective null hypothesis, provide little guidance regarding this specification choice (Christiano, Eichenbaum and Vigfusson, 2006).[3]

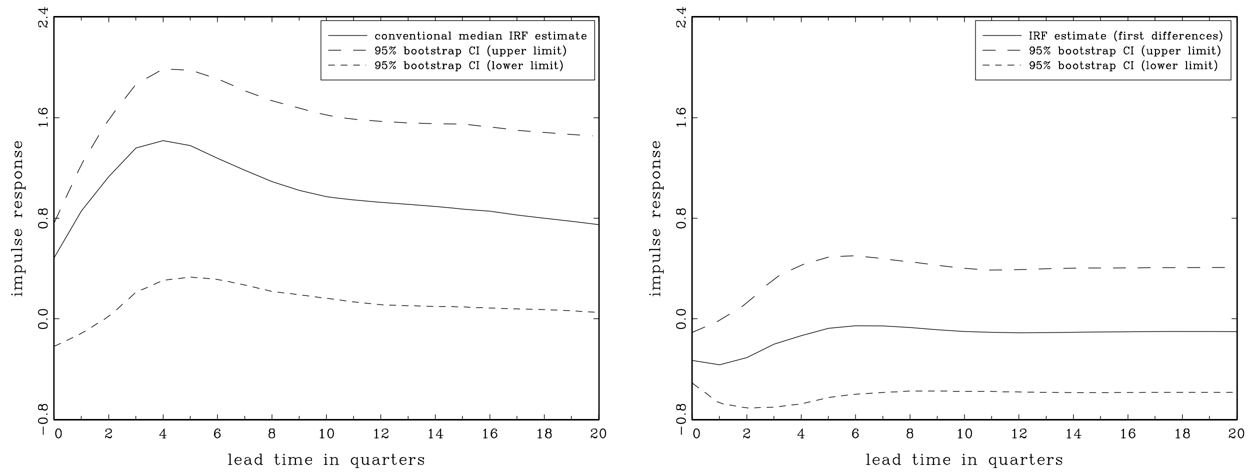

FIGURE 1. Response of hours to a 1% positive technology shock. Left: hours worked in levels; Right: hours worked in first differences

These contradicting results have generated substantial interest in the effect of misspecifying hours worked. Chari, Kehoe and McGrattan (2005) and Christiano, Eichenbaum and Vigfusson (2003, 2006) argue that the differenced VAR is misspecified if hours worked are stationary, which is typically implied by the standard RBC models. Chari, Kehoe and McGrattan (2005) and Christiano, Eichenbaum and Vigfusson (2003) investigate the effect of over-differenced hours on the impulse response analysis and suggest pre-testing and encompassing testing procedures for selecting the stochastic specification for hours worked.

While the levels specification is immune to the aforementioned problems, the finite-lag levels VAR can still be misspecified if the underlying theoretical model implies a dynamic process with an

[2]U.S. data on labour productivity, hours worked in the non-farm business sector and population over the age of 16 from DRI Basic Economics (the mnemonics are LBOUT, LBMN and P16, respectively).

[3]Using a multivariate Bayesian posterior odds procedure, Christiano, Eichenbaum and Vigfusson (2003) find evidence in favor of the levels specification.

infinite lag structure (Chari, Kehoe and McGrattan, 2005; Christiano, Eichenbaum and Vigfusson, 2006). Furthermore, as noticed by Chari, Kehoe and McGrattan (2005) and Christiano, Eichenbaum and Vigfusson (2006), among others, the levels specification tends to produce IRFs with large sampling variability that are nearly uninformative for distinguishing between competing economic theories. Gospodinov (2006) shows that the large sampling uncertainty associated with the IRFs in the levels specification arises from a weak instrument problem when the largest root of hours worked is near the nonstationary boundary. Similarly, Pesavento and Rossi (2006) demonstrate that the impulse response functions from both the levels and difference VAR specifications may be unreliable when the root in hours worked is close to one. In particular, the impulse response functions can be highly biased with confidence intervals that have poor coverage.

Nevertheless, despite the voluminous recent literature on the effects of technology shock on hours worked, there is still little understanding of how such large quantitative and qualitative differences in the impulse responses can be generated. While the literature discussed above points to potential biases in both VAR specifications, it is not clear that such biases are large enough in practice to explain such highly divergent results especially in the short run. In fact, we find that it is nearly impossible to explain these differences based solely on the behavior of hours worked itself and that these differences cannot be justified by small deviations of the largest root of hours worked from unity. As we show later, the seemingly conflicting evidence from the levels and differenced specifications identified with LR restrictions can only be reconciled when these deviations from the exact unit root are accompanied by small low frequency comovements between labour productivity growth and hours worked. We show that this low frequency component, which does not depend on the closeness of the root of hours worked to unity, drives a wedge between the levels and difference specifications with a profound impact on their impulse response functions.

This situation arises when restrictions on the matrix of LR multipliers, which includes low frequency information, are used to identify technology shocks. On a more intuitive level, if hours worked are a highly persistent, but stationary, process, it is possible that labour productivity inherits some small low frequency component from hours worked without inducing any observable changes in its time series properties.

While the levels specification explicitly estimates and incorporates this low-frequency comovement in the computation of the impulse response functions, the difference specification implicitly imposes this element to be zero. It is important to emphasize that this component could be arbitrarily small and yet still produce substantial differences in the impulse responses from the two specifications. Therefore, our results also suggest that a pre-testing procedure for a unit root will be ineffective in selecting a model that approximates well the true IRF when hours worked are close to a unit root process. In this case, the pre-testing procedure would favor the differenced specification, which rules out the above mentioned low frequency correlation, with high probability. This could in turn result in highly misleading IRF estimates.

Another way to look at the problem is to note that the differenced specification ignores possible low frequency comovements between labour productivity and hours worked. Figure 2 plots demeaned hours worked and detrended labour productivity and shows that the two series appear to be inversely related.

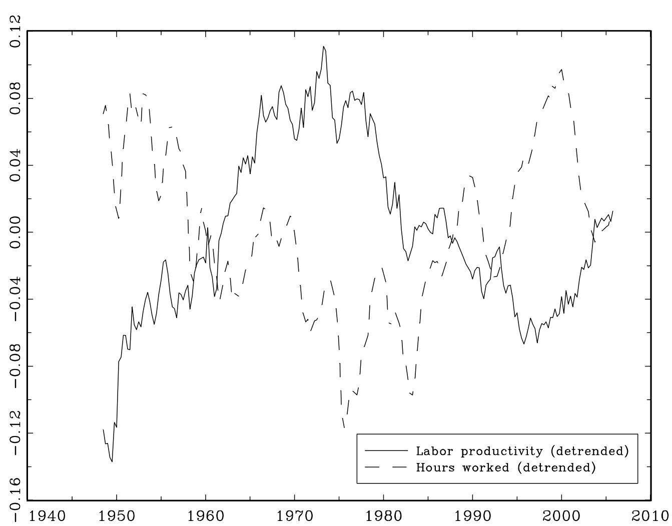

FIGURE 2. Demeaned hours and detrended labour productivity, 1948Q2 - 2005Q3.

We show in our analytical section that this inverse relationship should be translated into low frequency comovement between labour productivity and hours worked, provided that hours worked are stationary. While this comovement is hard to detect by visual inspection of the dynamics of labour productivity growth, Figures 3 reveals that the HP trend of labour productivity growth and hours worked exhibit some similarities and suggest that labour productivity growth may inherit its small low frequency trend component from hours worked.

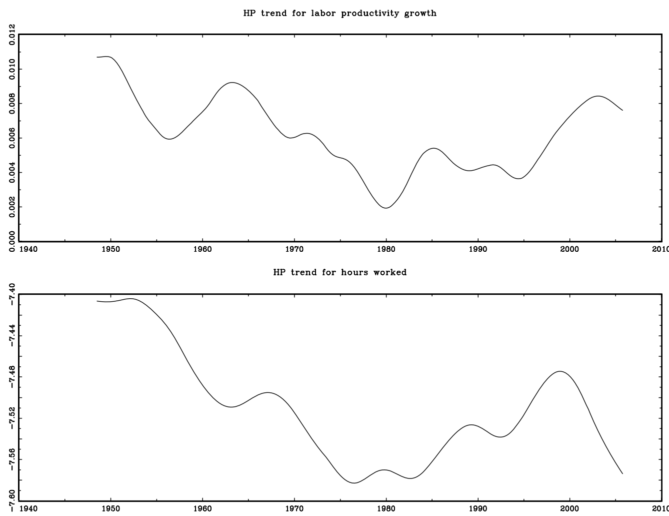

FIGURE 3. HP trends of labour productivity growth (top) and hours worked (bottom).

Fernald (2007) also highlights the sensitivity of the results to low frequency correlation between labour productivity growth and hours worked. Fernald (2007) concentrates on the low frequency correlation due to a similar high-low-high pattern in productivity growth (arising from structural breaks in productivity) and hours per person. In Section 2, we formalize this intuition and present a theoretical model that helps us to identify the possible source of the low frequency correlations and derive the implications for the impulse responses identified with long run restrictions. Section 3 presents the results from a Monte Carlo simulation experiment and Section 4 concludes.

2 Analytical Framework for Understanding the Debate

Our analytical framework and econometric specification is designed to mimic some of the salient features of the data and the implications of the theoretical macro (in particular, RBC) models.

First, we specify labour productivity as an exact unit root process. The RBC model imposes a unit root on technology and the data provide strong empirical support for this assumption. Hours worked exhibit a highly persistent, near-unit root behavior although the standard RBC model implies that they are a stationary process. Since an exact unit root cannot be ruled out as an empirical possibility, we do not take a stand on this issue and consider both the stationary and unit root cases. However, these different specifications (stationary or nonstationary) either allow for or restrict the low frequency comovement between hours worked and labour productivity growth. It turns out that this in turn has profound implications for the impulse response functions.

If hours worked are assumed stationary, the matrix of largest roots of the labour productivity growth and hours worked can contain a non-zero upper off-diagonal element whose magnitude depends on the closeness of the root of hours worked to one. This, typically fairly small, off-diagonal element can produce substantial differences in the shapes and the impact values of the impulse response functions from models that incorporate (levels specification) and ignore (differenced specification) this component.

Alternatively, in the case of an exact unit root for hours worked, the matrix of largest roots specializes to the identity matrix. In this case, there can be no low frequency comovement between hours work and labour productivity growth, ruling this out as an explanation for the difference between the two sets of impulse response functions. It is important to note, however, that this explanation is ruled out only in the case of an *exact* unit root. Our results suggest that this small low-frequency comovement continues to induce large discrepancies between the IRFs of the difference and level VARs, even when the largest root is arbitrarily close to and indistinguishable from unity.[4]

In order to complete the model, we need to adopt an identification scheme that allows us

[4]This argument can also be formalized in a local-unity setting that we consider in Section 2.4. In this setting, the off-diagonal element must itself be vanishing (i.e. local-to-zero), but nonetheless has a critical, non-vanishing, impact on the impulse response functions.

to recover the structural parameters and shocks. We follow the literature and impose the long-run identifying restriction that only shocks to technology can have a permanent effect on labour productivity. In addition, we assume that the structural shocks are orthogonal. In the next subsections, we formalize this analytical framework and work out its implications for the impulse response functions based on levels and differenced specifications.

2.1 Reduced-form model

Consider the reduced form of a bivariate vector autoregressive process $\widetilde{y}_t = (l_t, h_t)'$ of order $p+1$

$$\Psi(L)(I - \Phi L)\widetilde{y}_t = u_t, \tag{1}$$

where $\Psi(L) = I - \sum_{i=1}^{p} \Psi_i L^i = \begin{bmatrix} \psi_{11}(L) & \psi_{12}(L) \\ \psi_{21}(L) & \psi_{22}(L) \end{bmatrix}$, $E(u_t|u_{t-1}, u_{t-2}, ...) = 0$, $E(u_t u_t'|u_{t-1}, u_{t-2}, ...) = \Sigma$ and the matrix Φ is parameterized as $\Phi = \begin{bmatrix} 1 & \delta \\ 0 & \rho \end{bmatrix}$ where δ is small and ρ denotes the largest root of hours worked. This parameterization allows for a small impact of h_t on l_t and directly arises from the eigenvalue decomposition of Φ. More precisely, express Φ as $\Phi = V^{-1}\Lambda V$, where $\Lambda = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}$ contains the largest roots of the system and $V = \begin{bmatrix} 1 & -\gamma \\ 0 & 1 \end{bmatrix}$ is a matrix of corresponding eigenvectors.[5] Simple algebra yields $\Phi = \begin{bmatrix} 1 & -\gamma(1-\rho) \\ 0 & \rho \end{bmatrix}$ so that $\delta = -\gamma(1-\rho)$. When h_t is stationary ($\rho < 1$), we allow for a feedback $\gamma(1-\rho)$ from hours to the difference in productivity. Note that in the exact unit root case, Φ collapses to the identity matrix.[6] The other off diagonal element of V, and therefore of Φ, is set to zero to rule out the possibility of a feedback from the level of productivity to hours worked as this would imply that hours is $I(2)$ when $\rho = 1$

[5] Pesavento and Rossi (2005) use a similar decomposition but they impose diagonality on Φ. In Pesavento and Rossi (2005), the eigenvectors represent possible cointegration relationships.

[6] The persistence in h_t could also be modelled as local to unity $\rho_T = 1 - c/T$ for a fixed constant $c \geq 0$. We will discuss this parametrization in Section 2.4.

and $I(1)$ when $\rho < 1$.

The non-zero off-diagonal element allows a small low frequency component of hours worked to enter labour productivity growth. When the low frequency component is removed from either hours worked (Francis and Ramey, 2005, and Gali and Rabanal, 2004), or from labour productivity growth (Fernald, 2007), or from both (Gambetti, 2007, Pesavento and Rossi, 2005), this coefficient is driven to zero and the estimated IRF resembles the IRF computed from the differenced specification. The above parameterization of Φ can be used to explain this result.

It is convenient to rewrite model (1) in the Blanchard-Quah framework by imposing the exact unit root on productivity so that $\triangle l_t$ is a stationary process. In this case, let $y_t = (\triangle l_t, h_t)'$ and $A(L) = \Psi(L) \begin{bmatrix} 1 & \gamma(1-\rho)L \\ 0 & 1-\rho L \end{bmatrix}$. Then, the reduced form VAR model is given by

$$A(L)y_t = u_t \tag{2}$$

$$y_t = A_1 y_{t-1} + \ldots + A_{p+1} y_{t-p-1} + u_t.$$

2.2 Structural VAR

We denote the structural shocks (technology and non-technology shocks, respectively), by $\varepsilon_t = (\varepsilon_t^z, \varepsilon_t^d)'$, which are assumed to be orthogonal with variances σ_1^2 and σ_2^2, respectively, and relate them to the reduced form shocks by $\varepsilon_t = B_0 u_t$, where $B_0 = \begin{bmatrix} 1 & -b_{12}^{(0)} \\ -b_{21}^{(0)} & 1 \end{bmatrix}$. Premultiplying both sides of (2) by the matrix B_0 yields the structural VAR model

$$B(L)y_t = \varepsilon_t,$$

where $B(L) = B_0 A(L)$.

The matrix of long-run multipliers in the SVAR for y_t is

$$B(I) = \begin{bmatrix} \psi_{11}(1) - b_{12}^{(0)}\psi_{21}(1) & (1-\rho)\left([\gamma\psi_{11}(1) + \psi_{12}(1)] - b_{12}^{(0)}[\gamma\psi_{21}(1) + \psi_{22}(1)]\right) \\ \psi_{21}(1) - b_{21}^{(0)}\psi_{11}(1) & (1-\rho)\left([\gamma\psi_{21}(1) + \psi_{22}(1)] - b_{21}^{(0)}[\gamma\psi_{11}(1) + \psi_{12}(1)]\right) \end{bmatrix}$$

Imposing the restriction that non-technology shocks have no permanent effect on labour productivity renders the matrix $B(I)$ lower triangular. For $\rho < 1$, the LR restriction translates into the restriction $b_{12}^{(0)} = [\gamma\psi_{11}(1) + \psi_{12}(1)]/[\gamma\psi_{21}(1) + \psi_{22}(1)]$.

Suppose now that we run a VAR in first differences. $\triangle\widetilde{y}_t = (\triangle l_t, \triangle h_t)'$. Then, the reduced form becomes

$$\Psi(L) \triangle \widetilde{y}_t = (\rho - 1) \begin{bmatrix} 0 & -\gamma \\ 0 & 1 \end{bmatrix} \widetilde{y}_{t-1} + u_t$$

and the structural form is given by

$$B_0\Psi(L) \triangle \widetilde{y}_t = B_0\,(\rho - 1) \begin{bmatrix} 0 & -\gamma \\ 0 & 1 \end{bmatrix} \widetilde{y}_{t-1} + \varepsilon_t$$

The usual VAR in first differences, ignore the extra term on the right hand side and estimate the misspecified model $B_0\Psi(L) \triangle \widetilde{y}_t = \varepsilon_t$ and a long-run multiplier matrix equal to

$$B(I) = \begin{bmatrix} \psi_{11}(1) - b_{12}^{(0)}\psi_{21}(1) & \psi_{12}(1) - b_{12}^{(0)}\psi_{22}(1) \\ \psi_{21}(1) - b_{21}^{(0)}\psi_{11}(1) & \psi_{22}(1) - b_{21}^{(0)}\psi_{12}(1) \end{bmatrix}.$$

The LR restriction in the misspecified model now implies that $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$ thus ignoring any information contained in the levels and set implicitly this element equal to zero. There are two cases in which the VAR in first difference imposes the correct restriction. The first situation is when $\rho = 1$ so that the extra term drops out and setting $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$ is indeed correct. The second situation is when $\gamma = 0$. In this case hours do not enter the DGP for the difference in productivity and the restriction that a shock to ε_t^d as no LR effect on l_t still implies that the [1,2] element of $B_0\Psi(1)$ is zero so that $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$. So, unless $\rho = 1$ or $\gamma = 0$ the VAR will ignore important information and estimate the wrong indentifying matrix B_0 and consequently the wrong impact effect of a technology shock.

Once the structural parameter $b_{12}^{(0)}$ is obtained (by plugging consistent estimates of the elements of $\Psi(I)$ from the reduced form estimation), the remaining parameters can be recovered from

$B_0 E(u_t u_t') B_0' = E(\varepsilon_t \varepsilon_t')$ or

$$b_{21}^{(0)} = \frac{b_{12}^{(0)} \Sigma_{22} - \Sigma_{12}}{b_{12}^{(0)} \Sigma_{12} - \Sigma_{11}},$$

$$\sigma_1^2 = \Sigma_{11} - 2b_{12}^{(0)} \Sigma_{12} + \left[b_{12}^{(0)} \right]^2 \Sigma_{22}$$

and

$$\sigma_2^2 = \Sigma_{22} - 2b_{21}^{(0)} \Sigma_{12} + \left[b_{21}^{(0)} \right]^2 \Sigma_{11},$$

where Σ_{ij} is the $[ij]$th element of Σ. These parameters can be used consequently for impulse response analysis and variance decomposition.

2.3 Implications for impulse response analysis

The impulse response functions of hours worked to a shock in technology can be computed either from the levels specification (Blanchard and Quah, 1989; Christiano, Eichenbaum and Vigfusson, 2006; among others) or the differenced specification (Gali, 1999; Francis and Ramey, 2005). The levels approach will explicitly take into account and estimate a possible non-zero upper off-diagonal element in Φ but it suffers from some statistical problems when hours worked are highly persistent (Gospodinov, 2006). On the other hand, the differenced approach will produce valid and asymptotically well behaved IRF estimates in the exact unit root case but it ignores any possible information in levels when hours worked is stationary and gives rise to highly misleading IRFs even for very small deviations from the unit root assumption on hours.

Since $b_{12}^{(0)} = [\gamma\psi_{11}(1) + \psi_{12}(1)]/[\gamma\psi_{21}(1) + \psi_{22}(1)]$ and $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$ can produce very different values of $b_{12}^{(0)}$ even for $(\rho - 1)$ very close to zero, the IRFs from these two approaches can be vastly different. For simplicity, take the first-order model where $\Psi(L) = I$. In this case, the two restrictions set the value of $b_{12}^{(0)}$ to γ and 0, respectively, that consequently produce very different values for $b_{21}^{(0)}$ which directly determines the impulse response function since in the first-order model

$$\theta_{hz}^{(l)} = \frac{\partial h_{t+l}}{\partial \varepsilon_t^z} = \left[\Phi^l B_0^{-1} \right]_{21} = \frac{b_{21}^{(0)} \rho^l}{1 - b_{12}^{(0)} b_{21}^{(0)}}. \qquad (3)$$

As it is clear from (3) the impact effect at $l = 0$ does not depend on the value of ρ as $\rho^0 = 1$, but only on the values of $b_{21}^{(0)}$ and $b_{12}^{(0)}$ which themselves depend on $\Psi(1)$ and on γ. Focusing the debate on what is the impact effect of a technology shock on hours worked on the value for ρ is therefore missleading. To visualize the differences in the IRFs from the levels and differenced specifications when Φ is not diagonal, it is instructive to consider the following simplified example. Suppose the true data generating process is a first-order VAR with $\rho = 0.98$, $\gamma = -1$ (which implies an off-diagonal element $\delta = -\gamma(1-\rho) = 0.02$) and $\Sigma = \begin{pmatrix} 1 & -0.2 \\ -0.2 & 0.8 \end{pmatrix}$. From the above formulas, it can be easily inferred that the true values of the parameters that enter the IRF are $b_{12}^{(0)} = 1$, $b_{21}^{(0)} = 0.75$ and $\sigma_1^2 = 1.4$ whereas the differenced approach uses values of $b_{12}^{(0)} = 0$, $b_{21}^{(0)} = -0.2$ and $\sigma_1^2 = 1$. The IRFs based on the levels (true) and differenced specifications are plotted in Figure 4.

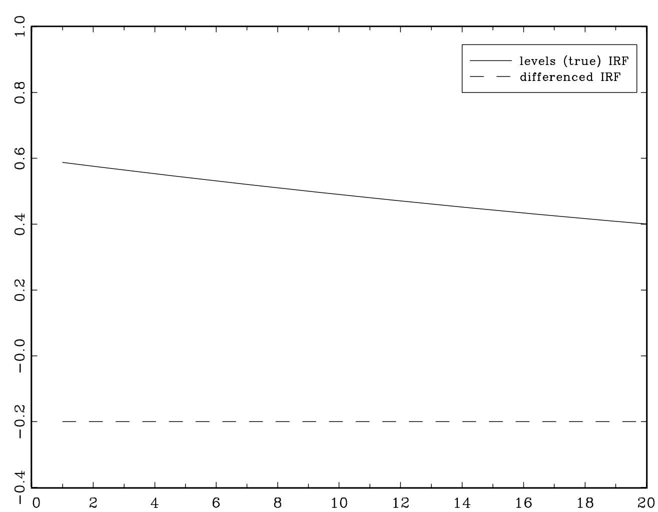

FIGURE 4. Impulse response functions computed form the levels (true) and differenced specifications.

Figure 4 clearly illustrates the large differences in the IRFs from the two specifications that are generated by the presence of a small off-diagonal element δ. Interestingly, the differences between the IRFs do not necessarily disappear as ρ gets closer to one and δ approaches zero. As our analytical framework suggests, these differences arise not only from the deviations of ρ from one but also from the non-zero value of γ that describes the relationship of the variables in levels.

2.4 An alternative parametrization

An alternative parametrization of model (1)is obtained by modeling the largest root in hours as a local to unity process with $\rho = 1 - c/T$ with $c \geq 0$. By letting ρ shrink toward 1 at rate T, we try to capture the idea that in small samples a highly persistent process behaves more like a unit root process than a stationary process. It follows that $\Lambda = \begin{bmatrix} 1 & 0 \\ 0 & 1 - c/T \end{bmatrix}$, $V = \begin{bmatrix} 1 & -\gamma \\ 0 & 1 \end{bmatrix}$ and $\Phi_T = \begin{bmatrix} 1 & -\gamma c/T \\ 0 & 1 - c/T \end{bmatrix}$. In finite samples, as long as $c > 0$, no matter how small, the co-movement between hours and productivity is different than zero, although arbitrarly small.

The reduced form for $y_t = (\triangle l_t, h_t)'$ is now

$$A(L)y_t = u_t$$

with $A(L) = \Psi(L) \begin{bmatrix} 1 & \gamma\,(c/T)\,L \\ 0 & (1 - L) + (c/T)\,L \end{bmatrix}$. When $c = 0$, the unit root case, Φ_T collapses to the identity matrix and there is no feedback from hours to the difference productivity.

This parametrization provides a unifying framework for analyzing both highly persistent ($c > 0$) and unit root ($c = 0$) variables and provides a better approximation to the small sample behavior of hours worked. In addition, in this specifications, the impact of h_{t-1} on $\triangle l_t$ is local-to-zero ($\gamma c/T$) and vanishing at rate T^{-1}, thus capturing the idea that the small frequency co-movement could simply be small sample bias that needs to be taken into account while doing inference in finite samples. Also, writing the model in the local-to-unity form is intuitively appealing since the low frequency correlation between h_{t-1} and $\triangle l_t$ is bound to disappear asymptotically so that hours do not affect productivity growth in the long run.

Under the local-to-unity parameterization, the matrix of long-run multipliers becomes

$$B(I) = \begin{bmatrix} \psi_{11}(1) - b_{12}^{(0)}\psi_{21}(1) & c/T\left([\gamma\psi_{11}(1) + \psi_{12}(1)] - b_{12}^{(0)}[\gamma\psi_{21}(1) + \psi_{22}(1)]\right) \\ \psi_{21}(1) - b_{21}^{(0)}\psi_{11}(1) & c/T\left([\gamma\psi_{21}(1) + \psi_{22}(1)] - b_{21}^{(0)}[\gamma\psi_{11}(1) + \psi_{12}(1)]\right) \end{bmatrix}$$

and the restriction that non-technology shocks have no permanent effect on labour productivity yields $b_{12}^{(0)} = [\gamma\psi_{11}(1)+\psi_{12}(1)]/[\gamma\psi_{21}(1)+\psi_{22}(1)]$ for $c > 0$. Note that when $c = 0$, our assumption of no cointegration between the series implies that $\gamma = 0$ in which case $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$ as in the first differenced case considered above. As a result, the analysis of the shapes of the impulse response functions under the different specifications in Section 2.3 remains unchanged.

3 Monte Carlo Experiment

To demonstrate the differences in the IRF estimators with a non-diagonal Φ, we conduct a small Monte Carlo experiment. 10,000 samples for $y_t = (l_t, h_t)'$ are generated from the VAR(2) model

$$\left[I - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L\right]\left[I - \begin{pmatrix} 1 & \delta \\ 0 & \rho \end{pmatrix} L\right]\begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix},$$

where $\delta = -\gamma(1-\rho)$, $T = 250$ and $(u_{1,t}, u_{2,t})' \sim N(0,\Sigma)$, $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$ with parameter values that are calibrated to match the empirical shape of the IRF of hours worked to a technology shock.[7] The number of lags is assumed known.

In addition to the IRF estimates from the levels and differenced specifications, we consider the IRF estimates from a levels specification with HP detrended productivity growth as in Fernald (2007) and a levels specification with HP detrended hours as in Francis and Ramey (2006).

Figures 5 to 8 show simulation results for the IRFs under four different parameter combinations for ρ and γ, all of which lie in a range of values that is potentially consistent with the actual data. The four panels of each figure correspond to the four different model specifications: a VAR in productivity growth and hours, a VAR in productivity growth and differenced hours, a VAR in HP detrended productivity growth and hours, and a VAR in productivity growth and HP detrended

[7] Note that while the numbers for the short-run dynamics are chosen to match the empirical values estimated from a VAR in levels, in our simulations we will also impose $\rho = 1$ and therefore allowing both specifications (levels and first differences) to be the true DGP.

hours. For each model we show the true IRF (solid line), median Monte Carlo IRF estimate (long dashes), and the 95% Monte Carlo confidence bands (short dashes).

In Figure 5 we consider a stationary but persistent process for hours ($\rho = 0.95$), while allowing a small low frequency component of hours worked to enter labour productivity growth ($\delta = 0.04$). As shown in the figure, the VAR in levels (top-left panel) estimates an IRF that is close to the true IRF, except for a small bias (see Gospodinov, 2006, for an explanation). On the other hand, the VAR with hours in first differences (top-right panel) incorrectly estimates a negative initial impact of the technology shock even though the true impact is positive. These results are in agreement with our discussion in the analytical section.

In Figure 6, we increase the largest root of hours worked from $\rho = 0.95$ to $\rho = 0.97$ and also substantially decrease value of the off-diagonal element from $\delta = 0.04$ to $\delta = 0.015$. Nevertheless, despite these changes, the IRFs shown in the two figures are strikingly similar. This underlines the ability of even a very small low frequency comovement to drive a qualitatively important wedge between the level and difference IRFs. likewise, it illustrates that the largest root need not be far from one for this effect to be important.

The bottom panels of Figures 5 and 6 are also interesting. When the HP filter is used to remove the low frequency component from either labour productivity growth (Fernald, 2007; bottom-left panel) or from hours worked (Francis and Ramey, 2005; bottom-right panel), the estimated IRFs resemble the IRFs computed from the differenced specification. In all three specifications, the removal of the low-frequency component from either series, whether by differencing or by HP filter, eliminates the possibility of any low frequency comovements between the transformed series and this has a profound influence on the IRFs.[8]

[8] After removing the low frequency component, the nature of the IRF changes and it is not completely justifiable to compare the IRFs from the transformed and the original processes but we still report the IRFs on the same graph to illustrate the economically large differences created by a fairly small off-diagonal element. Here, we do not argue if the low frequency components should or should not be removed prior to the IRF analysis but instead provide an analytical framework for explaining and reconciling the conflicting results documented in the empirical literature.

Figure 7 presents the results for the exact unit root case. In this case the matrix of largest roots becomes diagonal, eliminating the low-frequency comovement between hours and productivity growth ($\delta = 0$). As expected, by removing this low frequency comovement, we also eliminate the main qualitative differences between the median IRF response functions from the four models. Despite some small biases, all four median IRFs now correctly sign the impact of the technology shock and come close to tracing out the true IRFs.

Nevertheless, there are still important differences in performance among the four specifications. Not surprisingly, the differenced specification is particularly accurate and produces an unbiased estimator with tight confidence intervals. The estimator from the levels specification exhibits a modest bias that arises from the biased estimation of the largest root of hours and a very large sample uncertainty (Gospodinov, 2006). The estimator from the specification with HP filtered labour productivity growth performs similarly to the differenced estimator although it is slightly biased and more dispersed. As in the previous graphs, the estimator from the model with HP filtered hours resembles initially the shape of the model with filtered labour productivity but its sampling variability quickly shrinks towards zero due to the induced much lower persistence of the transformed process.

In figure 8 we maintain the assumption of a zero off-diagonal element ($\delta = 0$), but return to a persistent but stationary specification for hours worked ($\rho = 0.5$). The median IRFs from all four models are again quite similar both to each other and to the true IRF. In this sense, the basic message from Figures 7 and 8 are similar, despite the fact that hours is nonstationary in Figure 7 but stationary on Figure 8. However, there are still some substantive differences. Most notably, the accuracy of the levels IRF is clearly much improved, with smaller bias and considerably smaller variance.

In summarizing the results from these four tables we note that large qualitative differences in median IRFs for the differenced and levels VARs were observed only in Figures 5 and 6, in which there is a small low-frequency relationship between hours and labor productivity ($\delta \neq 0$). Neither

Figures 7 or 8 show qualitative differences in the median IRFs from the two specifications. Yet in Figure 7 hours has unit root, whereas it is stationary in Figure 8. What the two figures share in common, is the absence of the low frequency comovement of Figures 5 and 6 (i.e. $\delta = 0$). Although the size of the unit root in hours worked has important implications for the bias and sampling distributions of the IRFs, this suggests that it is instead the low frequency comovement that plays the critical role in driving the central qualitative differences between the level and difference specifications.

To better assess the sensitivity of the levels and difference specifications to different values of ρ and δ, we plot in Figures 9 and 10 the true and estimated responses for various degrees of persistence and low frequency comovement. Each line represents values for $\gamma = \{-0.5, -0.2, 0, 0.2, 0.5\}$ which correspond to different off-diagonal elements δ depending on the value of ρ (recall that $\delta = -\gamma(1 - \rho)$). Once again it is clear that while the level specification explicitly estimates and incorporates the different values for δ in the computation of the impulse response functions, the difference specification implicitly imposes this element to be zero which leads to substantial deviations from the true impulse response functions.

Finally, we note that the confidence intervals we report are Monte Carlo confidence intervals which are infeasible since they utilize knowledge of the true data generating process. The bias in the levels VAR and the misspecification in the first difference regressions result in poor coverage of confidence intervals constructed with standard procedures at medium and long horizons (Pesavento and Rossi, 2006) not reflected in our infeasible confidence intervals. At the same time, Figures 5-7 show very well how a wide range of different estimates for the IRF are possible, and that indeed the sampling uncertainty in the levels VAR is larger. At the same time, except for the cases in which either ρ is exactly one, or δ is exactly zero, the true impulse response is never contained in the Monte Carlo confidence bands for the VAR in first differences.

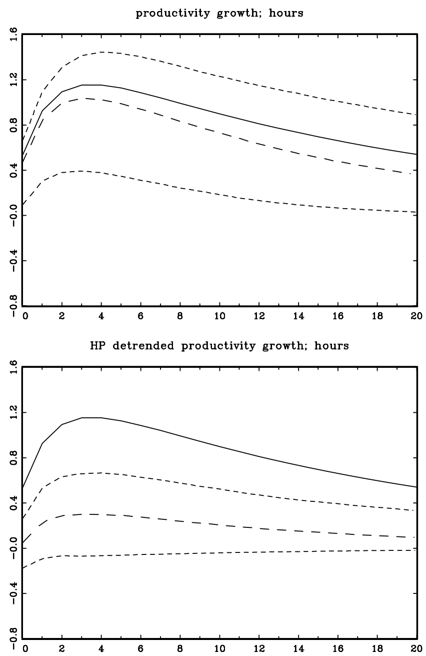

FIGURE 5. $\rho = 0.95, \gamma = -0.8, \delta = 0.04$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands

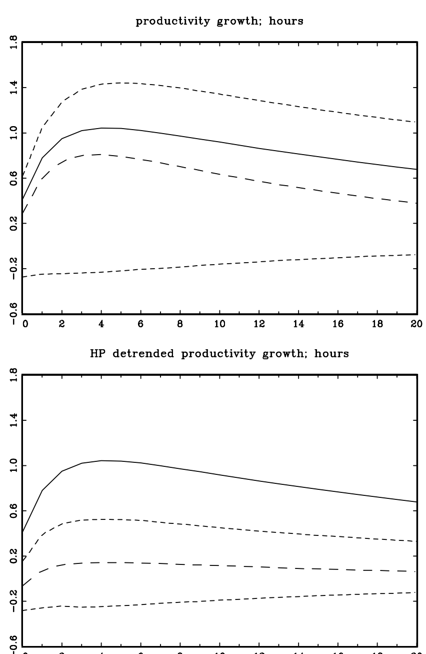

FIGURE 6. $\rho = 0.97, \gamma = -0.5, \delta = 0.015$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

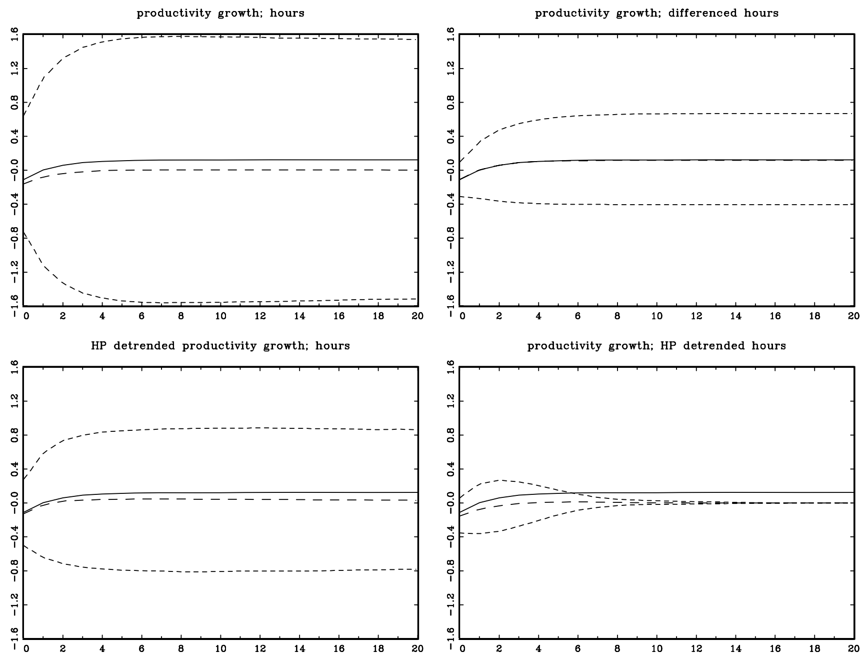

FIGURE 7. $\rho = 1, \delta = 0$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

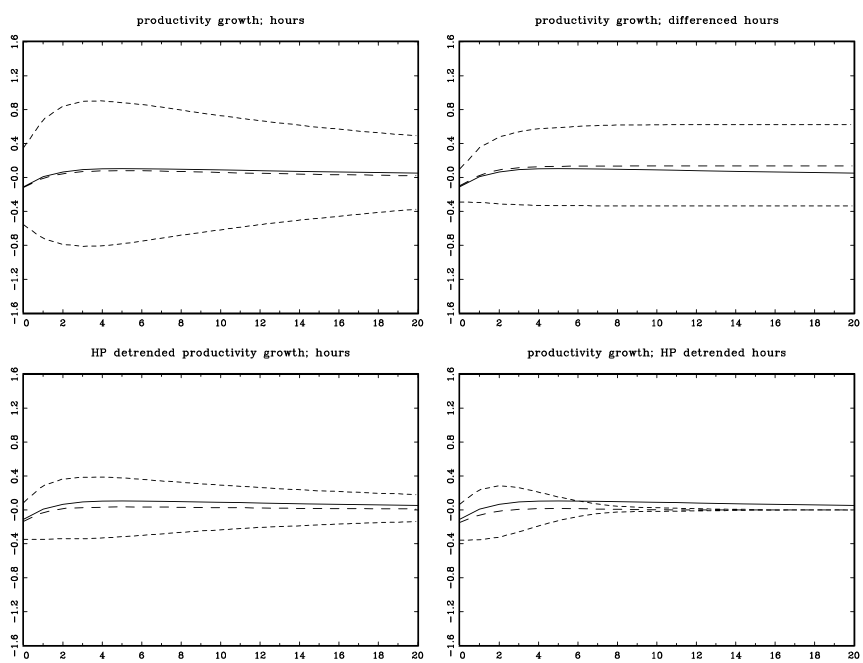

FIGURE 8. $\rho = 0.95, \gamma = 0,\ \delta = 0$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

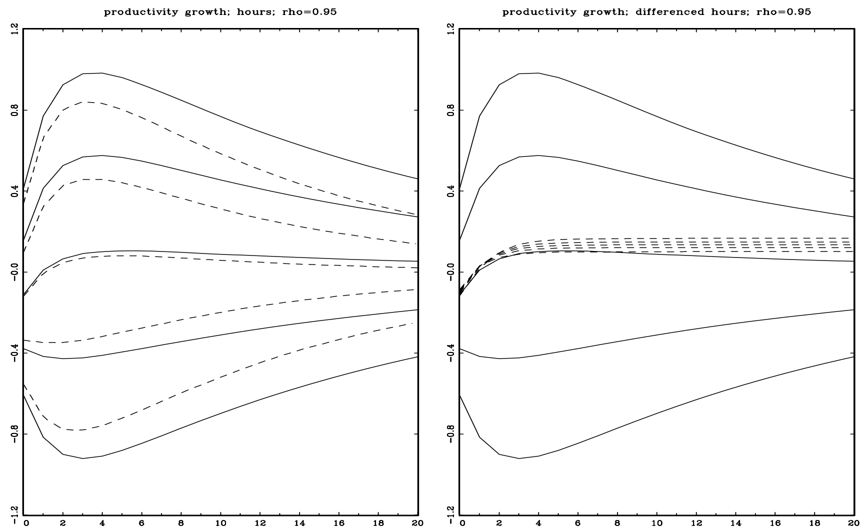

FIGURE 9. $\rho = 0.95, \gamma = \{-0.5, -0.2, 0, 0.2, 0.5\}$. Solid line: true IRF; short dashes: median Monte Carlo IRF estimate.

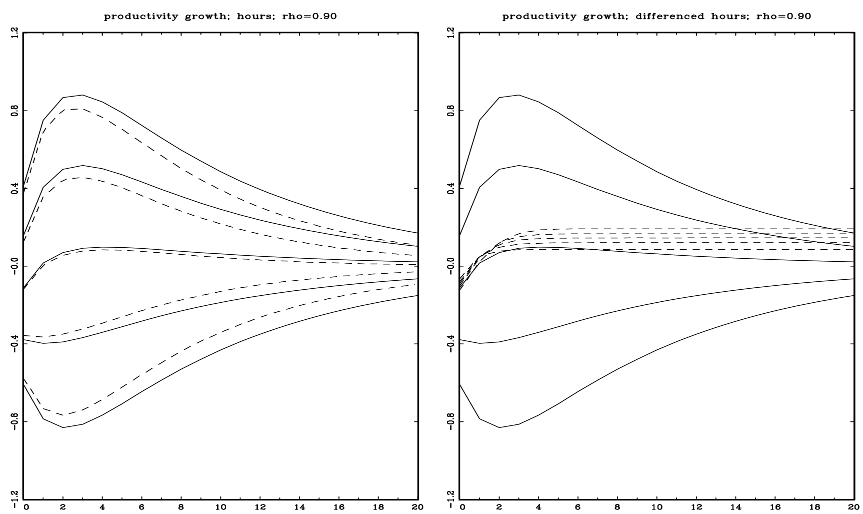

FIGURE 10. $\rho = 0.90, \gamma = \{-0.5, -0.2, 0, 0.2, 0.5\}$. Solid line: true IRF; short dashes: median Monte Carlo IRF estimate.

4 Conclusion

This paper analyzes the source of the conflicting evidence of the effect of technology shock on hours worked reported in several recent empirical studies. While we find that the uncertainty associated with the degree of persistence of hours worked plays an important role in this debate, we argue that the large quantitative and qualitative differences in the IRFs documented in the literature can only arise if it is accompanied by a low frequency comovement between productivity growth and hours worked. This low frequency component depends on the magnitude of the largest root of hours worked and can drive a wedge between the IRF estimates of the levels and difference VAR specifications even when this root is arbitrarily close to one. This implies that pre-testing for a unit root in hours worked may lead to highly misleading IRF inference in the region when the unit root tests have difficulties rejecting the null hypothesis. While the levels VAR appears to provide a more reliable framework for analysis in this setup, it may also produce biased and highly variable IRF estimates especially if hours worked are a near-integrated process. Imposing additional restrictions on the model (see, for example, Gospodinov, 2006) can lead to improved inference for the structural parameters and impulse responses.

References

[1] Alexopoulos, M. (2006), "Read all about it!! What happens following a technology shock?", mimeo, University of Toronto

[2] Basu, S., J.G.Fernald, M.S.Kimball (2006), "Are Technology Improvements Contractionary?", *American Economic Review,* 95(5), 1418-1448.

[3] Blanchard, O., and Quah, D. (1989), "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review,* 79, 655-673.

[4] Chari, V., P. Kehoe, and E. McGrattan (2005), "A Critique of Structural VARs Using Real Business Cycle Theory," *Federal Reserve Bank of Minneapolis, Working Paper* 631.

[5] Christiano, L., M. Eichenbaum and R. Vigfusson (2003), "What Happens After a Technology Shock?", *Board of Governors of the Federal Reserve System, International Finance Discussion Paper* 768.

[6] Christiano, L., Eichenbaum, M., and Vigfusson, R. (2006), "Assessing Structural VARs," in D. Acemoglu, K. Rogoff and M. Woodford, eds., *NBER Macroeconomics Annual* Cambridge, MA: MIT Press.

[7] Fernald, J. (2007), "Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements", *Journal of Monetary Economics,* 54, 2467-2485.

[8] Francis, N., M. Owyang and A. Theodorou (2003), "The Use of Long-Run Restrictions for the Identification of Technology Shocks", *Federal Reserve Bank of St. Louis Review*, 85(6), 53-66.

[9] Francis, N. and V. Ramey (2005), "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited", *Journal of Monetary Economics,* 52, 1379-1399.

[10] Francis, N. and V. Ramey (2006), “A New Measure of Hours Per Capita with Implications for the Technology-Hours Debate,” *mimeo*, University of North Carolina, Chapel Hill.

[11] Gali, J. (1999), “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?”, *American Economic Review*, 89, 249-271.

[12] Gali, J. and P. Rabanal (2004), “Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data?” in B. Bernanke and K. Rogoff, eds., *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press.

[13] Gambetti, L. (2007), “Technology Shocks and the Response of Hours Worked: Time-Varying Dynamics Matter”, *mimeo*, University Pompeu Fabra.

[14] Gospodinov, N. (2006), “Inference in Nearly Nonstationary SVAR Models with Long-Run Identifying Restrictions,” *mimeo*, Concordia University.

[15] Pesavento, E. and B. Rossi (2005), “Do Technology Shocks Drive Hours Up or Down? A Little Evidence From an Agnostic Procedure,” *Macroeconomic Dynamics*, 9, 478-488.

[16] Pesavento, E. and B. Rossi (2006), “Small Sample Confidence Intervals for Multivariate Impulse Response Functions at Long Horizons,”, *Journal of Applied Econometrics*, 21, 1135-1155.

[17] Shea, J. (1999), “What Do Technology Shocks Do?” in B. Bernanke and J. Rotemberg, eds., *NBER Macroeconomics Annual,* Cambridge, MA: MIT Press.

[18] Uhlig, H. (2003), “Do Technology Shocks Lead to a Fall in Total Hours Worked?”, *mimeo.*

[19] Vigfusson, R. (2004), “The Delayed Response to A Technology Shock: A Flexible Price Explanation”, *Board of Governors of the Federal Reserve System, International Finance Discussion Paper* 810.