

# SENSITIVITY OF IMPULSE RESPONSES TO SMALL LOW FREQUENCY CO-MOVEMENTS: RECONCILING THE EVIDENCE ON THE EFFECTS OF TECHNOLOGY SHOCKS

Nikolay Gospodinov\*  
Concordia University and CIREQ

Alex Maynard†  
University of Guelph

Elena Pesavento‡  
Emory University

## Abstract

This paper clarifies the empirical source of the debate on the effect of technology shocks on hours worked. We find that the contrasting conclusions from levels and differenced VAR specifications, documented in the literature, can be explained by a small low frequency co-movement between hours worked and productivity growth which gives rise to a discontinuity in the solution for the structural coefficients identified by long-run restrictions. While the low frequency co-movement is allowed for in the levels specification, it is implicitly set to zero in the differenced VAR. Consequently, even when the root of hours is very close to one and the low frequency co-movement is quite small, removing it can give rise to biases large enough to account for the empirical difference between the two specifications.

*Keywords:* Technology shocks, impulse response functions, structural VAR, long-run identification, low frequency co-movement.

*JEL Classification:* C32, C51, E32, E37.

---

\*Department of Economics. Concordia University, 1455 de. Maisonneuve Blvd. West, Montreal, Quebec, H3G 1M8 Canada (nikolay.gospodinov@concordia.ca)

†Department of Economics, University of Guelph, Guelph, Ontario, N1G 2W1 Canada (maynarda@uoguelph.ca)

‡Department of Economics, Emory University, Atlanta, Georgia, 30322-2240 USA (epesave@emory.edu)

# 1 Introduction

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. The dynamic responses in both models becomes more complicated once inventories are included (Chang, Hornstein, and Sarte, 2009). As technology shocks are difficult to measure, they are commonly specified as structural shocks in vector autoregressive (VAR) models that are identified via the long-run (LR) restriction that only technology shocks have a permanent effect on labour productivity (Gali, 1999, for example; see Alexopoulos (2006), Shea (1999) and Basu, Fernald and Kimball (2006) for alternative measures of technology shocks). The LR 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 some aspects of the VAR specification, particularly whether hours worked are modeled in levels or first 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 embedded frictions. Other papers have reached similar conclusions (see, for example, Francis and Ramey, 2005; Basu, Fernald and Kimball, 2006; among others) and this has spurred a line of research aimed at developing general equilibrium models (Gali and Rabanal, 2004) or alternative finite-horizon identification schemes (Uhlig, 2004; Francis, Owyang and Roush; 2005) that can account for this empirical finding.

However, maintaining the long-run identification restriction, but allowing hours worked to enter the model in levels, Christiano, Eichenbaum and Vigfusson (2003, 2006) provide support for the

prediction of standard RBC models, with hours worked rising immediately after a positive productivity shock. Christiano, Eichenbaum and Vigfusson (2003) argue strongly in favor of the levels specification and report that the levels specification encompasses the estimated impulse response function of the differenced specification, but not vice-versa. More recently, Fernald (2007) provides both intuition and simulation results in support of the claim that the results from the levels specification are mechanically driven by a common high-low-high pattern which arises from two level-shift breaks in both productivity growth and hours worked, occurring in the early 1970s and mid-1990s. After removing sub-sample means, Fernald (2007) finds that the levels and differenced specifications produce similar impulse response functions, both qualitatively matching the original results of Gali (1999). As we discuss below, an important assumption made in this framework is that the similar timing and direction of the breaks in productivity and hours is treated as coincidental.

Our paper takes up the challenge of reconciling the conflicting empirical findings reported in the literature and contributes to the understanding of this debate in several respects. We demonstrate analytically that the extreme sensitivity to different model specifications, first noted by Christiano, Eichenbaum and Vigfusson (2003), appears to be due to a discontinuity in the solution for the structural coefficients implied by the long-run restriction. Interestingly, we find that this discontinuity arises only in the presence of a low frequency correlation between hours worked and productivity growth. This draws a tight link between the apparently conflicting results of Christiano, Eichenbaum and Vigfusson (2003), who argue that the differenced SVAR is misspecified, and Fernald (2007) who argue that the levels specification is misleading without accounting for structural breaks. Although they point in opposite directions, both sets of results, at least implicitly, rely on similar low frequency correlations.

We argue that the difference in conclusions cannot be determined solely on the basis of empirical methods, such as unit-root pre-tests or Hodrick-Prescott (HP) pre-filtering. Instead, the appropriate conclusions that one draws from any of these approaches rests critically on the eco-

conomic assumptions made about the source of the low frequency correlation. If these low frequency correlations are treated as a true feature of the data generating process, as is implicitly the case in Christiano, Eichenbaum and Vigfusson (2003), then over-differencing improperly removes this low frequency correlation, thereby corrupting the long-run identification of the difference specification. Such true low frequency co-movement may be plausible if technological changes have long-lasting effects on the underlying structure of the labour market. For example, technological improvements give rise to greater efficiency in household production, leading to increased female labour market participation. Likewise, technological innovations affecting regional transportation or labour search costs, may also have lasting impacts on labour markets. On the other hand, Fernald (2007) provides some convincing arguments for why the similar timing of the structural breaks in productivity and hours may be coincidental, arising from disparate causes. Although Francis and Ramey (2009) instead demonstrate that this low frequency behavior may be driven by common demographic and sectoral employment changes, they argue that it violates the long-run identifying assumption and should thus be treated as low frequency noise. In either case, it is the presence of this low frequency correlation that corrupts the long-run identification and renders the unmodified levels specification misleading.

The popularity of the LR identification scheme derives in large part from its robustness to model specification, in the sense that it often remains valid for a wide variety of macroeconomic models. However, the implementation of the LR restriction also relies on the low frequency properties of the data. Our results, which illustrate the possibility of discontinuity in this dependence, reinforce the conclusions from the empirical literature suggesting that the long-run identifying scheme can be far less robust to assumptions on the low frequency properties of the variables. Of course, there may still be many cases in which robust empirical results can be obtained; for example, if the variables in question are clearly stationary or if low frequency correlations are not present. Nevertheless, we echo the recommendation made by Fernald (2007) that empirical researchers should check carefully

the robustness of their results to alternate assumptions on the low frequency properties of the data.

In our analysis, we assume that the data can be well approximated by a finite order invertible VAR. Therefore the issues discussed here are distinct from those addressed by Erceg, Guerrieri and Gust (2005), Chari, Kehoe and McGrattan (2008), and Christiano, Eichenbaum and Vigfusson (2006), who debate the success of finite structural VARs in approximating the infinite order models that are implied by economic theory. Nonetheless, our main conclusions continue to hold if the data are generated from a calibrated dynamic general equilibrium model. Simulated draws from the calibrated RBC model exhibit a low frequency co-movement similar to that found in the actual data, and the results for a fixed VAR lag are qualitatively similar to those presented Section 4 (simulation results are available in an online appendix).

The rest of the paper is organized as follows. Section 2 briefly reviews some empirical evidence and provides the intuition behind our findings. In Section 3, we formalize this intuition and present a theoretical model that helps us to identify the possible source of low frequency correlations and derive the implications for the impulse responses identified with long run restrictions. Section 4 presents the results from a Monte Carlo simulation experiment. Section 5 discusses the main implications of our analysis for empirical work and Section 6 concludes.

## **2 Illustrative Example and Intuitive Arguments**

To put the subsequent discussion in the proper empirical context, we present in Figure 1 the estimated impulse response functions (IRFs) based on the levels and differenced specifications with quarterly U.S. data for the period 1948Q2 - 2005Q3. 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 LBOU, LBMN and P16, respectively). The difference in the impulse response functions is quite striking. 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 quali-

tative differences in the impulse responses can be generated. While the literature attributed these discrepancies 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 justify these differences solely by the behavior of hours worked itself and, in particular, by small deviations of the largest root of hours worked from unity.

It is well known, for example, that over-differencing, and misspecification in general, can lead to biased results. However, what is indeed surprising is that a seemingly very minor, even undetectable, misspecification in the difference specification, may lead to a very substantial bias in the resulting impulse response function. 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). Pesavento and Rossi (2005) provide confidence intervals on the largest autoregressive root in hours worked using inversions of four different unit root tests. All four confidence intervals include unity and, in two cases, the lower bound on the largest root exceeds 0.980 (in the other two cases, it exceeds 0.925). On the face of it, this hardly appears to be a case in which over-differencing would lead to large misspecification errors. In fact, in a reduced-form near unit root model, the specification error that arises from over-differencing is second-order. Nevertheless, Christiano, Eichenbaum and Vigfusson (2003) report quite a large specification error in their calibrated simulations. This provocative result has yet to be satisfactorily explained in an econometric sense.

Another way to look at the problem is to note that the differenced specification ignores possible low frequency co-movements between labour productivity growth and hours worked. Figure 2 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. On a more intuitive level, if hours worked are a highly persistent, but stationary, process, it is possible that labour productivity growth inherits some small low frequency component

from hours without inducing any observable changes in its time series properties.

In fact, 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 co-movements between labour productivity growth and hours worked. We show that this low frequency co-movement drives a wedge between the levels and differenced 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. While the levels specification explicitly estimates and incorporates this low frequency co-movement in the computation of the impulse response functions, the differenced specification restricts this element to be zero. It is important to emphasize that this component could be arbitrarily small and could accompany an autoregressive (AR) root arbitrarily close to one, 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. In the next section, we provide more formal arguments for explaining and reconciling the conflicting empirical evidence from the levels and differenced specifications.

### **3 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 macroeconomic (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 co-movement between hours worked and labour productivity growth. It turns out that this has crucial implications for the impulse response functions.

If hours worked are 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 IRFs from models that incorporate (levels specification) or ignore (differenced specification) this component.

If hours worked have an exact unit root, the matrix of largest roots specializes to the identity matrix. In this case, there can be no low frequency co-movement 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 co-movement can continue to induce large discrepancies between the IRFs of the differenced and levels VARs, even when the largest root is arbitrarily close to and indistinguishable from unity.

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



### 3.1 Reduced-form model

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

$$\mathbf{\Psi}(L)(\mathbf{I} - \mathbf{\Phi}L)\tilde{\mathbf{y}}_t = \mathbf{u}_t, \quad (1)$$

where  $E(\mathbf{u}_t | \mathbf{u}_{t-1}, \mathbf{u}_{t-2}, \dots) = \mathbf{0}$ ,  $E(\mathbf{u}_t \mathbf{u}_t' | \mathbf{u}_{t-1}, \mathbf{u}_{t-2}, \dots) = \mathbf{\Sigma}$ ,  $\sup_t E[\|\mathbf{u}_t\|^{2+\xi}] < \infty$  for  $\xi > 0$ , and  $\mathbf{\Psi}(L) = \mathbf{I} - \sum_{i=1}^p \mathbf{\Psi}_i L^i = \begin{bmatrix} \psi_{11}(L) & \psi_{12}(L) \\ \psi_{21}(L) & \psi_{22}(L) \end{bmatrix}$  is a finite-order lag polynomial whose roots lie strictly outside the unit circle. The matrix  $\mathbf{\Phi}$  can be expressed in terms of its eigenvalue decomposition as  $\mathbf{\Phi} = \mathbf{V}^{-1} \mathbf{\Lambda} \mathbf{V}$ , where  $\mathbf{\Lambda} = \begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}$  contains the largest roots of the system and  $\mathbf{V} = \begin{bmatrix} 1 & -\gamma \\ 0 & 1 \end{bmatrix}$  is a matrix of corresponding eigenvectors (see, for example, Pesavento and Rossi, 2006). Simple algebra yields  $\mathbf{\Phi} = \begin{bmatrix} 1 & \delta \\ 0 & \rho \end{bmatrix}$ , where  $\delta = -\gamma(1 - \rho)$ , is the parameter that determines the low frequency co-movement between the variables and  $\rho$  denotes the largest root of hours worked.

This parameterization, which arises directly from the eigenvalue decomposition of  $\mathbf{\Phi}$ , allows for a small ( $\delta$ ) impact of  $h_t$  on  $l_t$ , provided that  $\rho$  is not exactly equal to one. Note that in the exact unit root case,  $\mathbf{\Phi}$  collapses to the identity matrix. The other off-diagonal element of  $\mathbf{V}$ , and therefore of  $\mathbf{\Phi}$ , is set to zero as it would otherwise imply that hours is  $I(2)$  when  $\rho = 1$  and  $I(1)$  when  $\rho < 1$ . In principle, the model can also be generalized to include a non-zero (but asymptotically vanishing) feedback from the level of productivity to hours worked. Simple algebra (available from the authors upon request) shows that this parameterization does not materially affect our subsequent analysis. For this reason, we set the lower off-diagonal element of  $\mathbf{\Phi}$  to zero without any loss of generality. It is important to note, however, that the zero lower off-diagonal restriction on matrix  $\mathbf{\Phi}$  does not rule out a feedback from productivity growth to hours worked in higher-order ( $p > 1$ ) VAR models. Thus, it has no implication for the direction of causality implied by the low frequency correlation between hours worked and productivity growth. For example, in a VAR(2) model, the lagged productivity growth is allowed to affect hours worked through the possibly non-zero coefficient  $\psi_{21}$ .

It is convenient to rewrite model (1) in Blanchard and Quah's (1989) framework by imposing the exact unit root on productivity so that  $\Delta l_t$  is a stationary process. In this case, let  $\mathbf{y}_t = (\Delta l_t, h_t)'$  and  $\mathbf{A}(L) = \mathbf{\Psi}(L) \begin{bmatrix} 1 & \gamma(1-\rho)L \\ 0 & 1-\rho L \end{bmatrix}$ . Then, the reduced form VAR model is given by

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{u}_t \quad (2)$$

or  $\mathbf{y}_t = \mathbf{A}_1\mathbf{y}_{t-1} + \dots + \mathbf{A}_{p+1}\mathbf{y}_{t-p-1} + \mathbf{u}_t$ . The non-zero off-diagonal element  $\gamma(1-\rho)L$  allows for the possibility that a small low frequency component of hours worked affects labour productivity growth. When the low frequency component is removed from either hours worked (Francis and Ramey, 2009, and Gali and Rabanal, 2004) or labour productivity growth (Fernald, 2007), this coefficient is driven to zero and the estimated IRF resembles the IRF computed from the differenced specification. The above parameterization of  $\mathbf{\Phi}$  can be used to explain this result.

### 3.2 Structural VAR

We denote the structural shocks (technology and non-technology shocks, respectively), by  $\boldsymbol{\varepsilon}_t = (\varepsilon_t^z, \varepsilon_t^d)'$ , which are assumed to be orthogonal with variances  $\sigma_1^2$  and  $\sigma_2^2$ , respectively, and relate them to the reduced form shocks by  $\boldsymbol{\varepsilon}_t = \mathbf{B}_0\mathbf{u}_t$ , where  $\mathbf{B}_0 = \begin{bmatrix} 1 & -b_{12}^{(0)} \\ -b_{21}^{(0)} & 1 \end{bmatrix}$ . Pre-multiplying both sides of (2) by the matrix  $\mathbf{B}_0$  and defining  $\mathbf{B}(L) = \mathbf{B}_0\mathbf{A}(L)$  yields the structural VAR model

$$\mathbf{B}(L)\mathbf{y}_t = \boldsymbol{\varepsilon}_t.$$

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

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

Imposing the restriction that non-technology shocks have no permanent effect on labour productivity renders the matrix  $\mathbf{B}(1)$  lower triangular. For  $\rho < 1$ , this 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 one assumes  $\rho = 1$  and let  $\Delta\tilde{\mathbf{y}}_t = (\Delta l_t, \Delta h_t)'$ . Then the reduced form specializes to

$$\mathbf{\Psi}(L) \Delta\tilde{\mathbf{y}}_t = \mathbf{u}_t \quad (3)$$

and the structural form is given by  $\mathbf{B}_0\boldsymbol{\Psi}(L)\Delta\tilde{\mathbf{y}}_t = \boldsymbol{\varepsilon}_t$  with a long-run multiplier matrix

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

Note that the LR restriction implies  $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$  and even if the upper right element of  $\boldsymbol{\Phi}$  is non-zero, the differenced VAR would ignore any information contained in the levels by implicitly setting this element to zero.

Once the structural parameter  $b_{12}^{(0)}$  is obtained (by plugging consistent estimates of the elements of  $\boldsymbol{\Psi}(1)$  from the reduced form estimation), the remaining parameters can be recovered from  $\mathbf{B}_0E(\mathbf{u}_t\mathbf{u}_t')\mathbf{B}_0' = E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t')$  or

$$b_{21}^{(0)} = \frac{b_{12}^{(0)}\Sigma_{22} - \Sigma_{12}}{b_{12}^{(0)}\Sigma_{12} - \Sigma_{11}}, \quad \sigma_1^2 = \Sigma_{11} - 2b_{12}^{(0)}\Sigma_{12} + [b_{12}^{(0)}]^2\Sigma_{22} \quad \text{and} \quad \sigma_2^2 = \Sigma_{22} - 2b_{21}^{(0)}\Sigma_{12} + [b_{21}^{(0)}]^2\Sigma_{11},$$

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

### 3.3 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 or the differenced specification. The levels approach will explicitly take into account and estimate a possible non-zero upper off-diagonal element in  $\boldsymbol{\Phi}$ , but it suffers from some statistical problems when hours worked are highly persistent. Christiano, Eichenbaum and Vigfusson (2003) note that the levels specification tends to produce IRFs with large sampling variability that are nearly uninformative for distinguishing between competing economic theories. Gospodinov (2010) shows that this large sampling uncertainty arises from a weak instrument problem when the largest root of hours worked is near the nonstationary boundary. 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 low frequency co-movement between hours and labour productivity growth when hours worked is stationary. It can therefore give 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  $b_{12}^{(0)}$ , the IRFs from these two approaches can be vastly different. In fact, because the value of  $\gamma$  does not depend on  $\rho$ , these differences can remain large even for  $(\rho - 1)$  arbitrarily close, but not equal, to zero. For simplicity, take the first-order model where  $\Psi(L) = \mathbf{I}$ . In this case, the two restrictions set the value of  $b_{12}^{(0)}$  to  $\gamma$  and 0, respectively, implying two very different values for  $b_{21}^{(0)}$ , which, in turn, directly determines the IRF. In particular, in the first-order model, the impulse response of hours at time  $t + k$  to a unit shock in technology at time  $t$  is given by

$$\left[\Phi^k \mathbf{B}_0^{-1}\right]_{21} = \frac{b_{21}^{(0)} \rho^k}{1 - b_{12}^{(0)} b_{21}^{(0)}}. \quad (4)$$

As it is clear from (4), the impact effect at  $k = 0$  does not depend on the value of  $\rho$  as  $\rho^0 = 1$ , but only on the values of  $b_{21}^{(0)}$  and  $b_{12}^{(0)}$ , which themselves depend on  $\Psi(1)$  and  $\gamma$ . Focusing the debate on the distance of  $\rho$  from one is therefore misleading, provided that  $\rho$  is not precisely equal to one.

Interestingly, the differences between the IRFs do not necessarily disappear as  $\rho$  gets closer to one and  $\delta$  approaches zero. As our analytical framework suggests, they can remain substantial even for values of  $\rho - 1$  and  $\delta = -\gamma(1 - \rho)$  arbitrarily close, but not equal, to zero. This is because, provided that  $\rho < 1$ , the size of this discrepancy depends on the co-movement through the parameter  $\gamma$ , rather than through either  $\delta$  or  $\rho$ . At a more intuitive level, the reason that the short-horizon IRFs can be highly sensitive to even small low frequency co-movements accompanying small deviations of  $\rho$  from one, is that they are identified off of long-run identification restrictions, which depend entirely on the zero frequency properties of the data. As reported below, a similar sensitivity does not arise when short-run identification restrictions are employed.

### 3.4 An alternative parameterization

The fact that our framework suggests potentially large IRF discrepancies even for values of  $\rho$  quite close to one is practically relevant, precisely because this is the case in which unit root tests have the greatest difficulty detecting stationarity. The low power of the unit root test in this case arises

because, in small samples, the resulting process for hours may behave more like a unit root process than like a stationary series. This concept has been formalized in the econometrics literature by the near unit root or local-to-unity model, in which  $\rho = 1 - c/T$  for  $c \geq 0$  is modelled as a function of the sample size  $T$  and shrinks towards unity as  $T$  increases (Phillips, 1987; Chan, 1988). Naturally, this dependence on the sample size should not be interpreted as a literal description of the data, but rather as a device to approximate the behavior of highly persistent processes in small samples. What makes this modelling device particularly relevant, is that, for small values of the local-to-unity parameter  $c$ , it describes a class of alternatives to  $\rho = 1$  against which unit root tests have no consistent power. Intuitively,  $c = T(1 - \rho)$  can be viewed as measuring the distance of the root from one relative to the sample size. Small values of  $c$  correspond to cases in which  $T$  is relatively small and  $\rho$  is relatively close to one, so that unit root tests have low power and the difference specification is likely to be employed when computing IRFs.

An alternative parameterization of the model in (1) is therefore obtained by modeling the largest root in hours as a local-to-unity process with  $\rho = 1 - c/T$  with  $c \geq 0$ . Then, it follows that  $\mathbf{\Lambda} = \begin{bmatrix} 1 & 0 \\ 0 & 1 - c/T \end{bmatrix}$ ,  $\mathbf{V} = \begin{bmatrix} 1 & -\gamma \\ 0 & 1 \end{bmatrix}$  and  $\mathbf{\Phi}_T = \begin{bmatrix} 1 & -\gamma c/T \\ 0 & 1 - c/T \end{bmatrix}$ . In finite samples, as long as  $c > 0$ , the feedback from hours to productivity is non-zero, yet arbitrarily small. The reduced form for  $\mathbf{y}_t = (\Delta l_t, h_t)'$  is now

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{u}_t$$

with  $\mathbf{A}(L) = \mathbf{\Psi}(L) \begin{bmatrix} 1 & \gamma(c/T)L \\ 0 & (1-L) + (c/T)L \end{bmatrix}$ . In the unit root case ( $c = 0$ ),  $\mathbf{\Phi}_T$  collapses to the identity matrix, the variables are not cointegrated and there is no feedback from hours to productivity growth. Thus, the impact of  $h_{t-1}$  on  $\Delta l_t$  is local-to-zero and vanishing at rate  $T^{-1/2}$ , since the level of  $h$  affects  $\Delta l_t$  through the term  $(\gamma c/T)h_{t-1}$  which is  $O_p(T^{-1/2})$  since  $T^{-1/2}h_{t-1} = O_p(1)$  in the local-to-unity setup. This captures the notion that the low frequency co-movement between hours and productivity growth must be small if the root of hours is close to unity. Writing the model in the local-to-unity form is also intuitively appealing since the low frequency correlation

between  $h_{t-1}$  and  $\Delta l_t$  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

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

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  $c = 0$ , the model again specializes to the differenced VAR specification in (3), for which the LR specification implies  $b_{12}^{(0)} = \psi_{12}(1)/\psi_{22}(1)$ . As a result, the analysis of the shapes of the impulse response functions under the different specifications in Section 3.3 remains unchanged. This confirms the finding that substantial differences in IRFs can arise even within this class of models for which unit root tests are not powerful enough to detect that hours worked are stationary. Thus, the stylized fact that hours worked are indistinguishable from a unit root process does not guarantee that the true IRF will be close to the IRF from the differenced specification.

## 4 Monte Carlo Experiment

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

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

where  $\delta = -\gamma(1 - \rho)$ ,  $T = 250$ ,  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$ , and the parameter values are calibrated to match the empirical shape of the IRF of hours worked to a technology shock.

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 consider both  $\rho < 1$  and  $\rho = 1$  and therefore allow both specifications (levels and first differences) to be the true DGP. The lag order of the VAR 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). The smoothing parameter for the HP filter is set to 1,600.

Figures 3 to 5 show simulation results for the IRFs under four different parameter configurations for  $\rho$  and  $\gamma$ , all of which lie in a range of values that is potentially consistent with the actual data. The three panels of each figure correspond to the different model specifications: a VAR in productivity growth and hours, a VAR in productivity growth and differenced hours and a VAR in HP detrended productivity growth and hours. For each model we show the true IRF (solid line), the median Monte Carlo IRF estimate (long dashes), and the 95% Monte Carlo confidence bands (short dashes).

In Figure 3 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 (left graph) estimates an IRF that is close, on average, to the true IRF, except for a small bias (see Gospodinov, 2010, for an explanation). On the other hand, the VAR with hours in first differences (middle graph) incorrectly estimates a negative initial impact of the technology shock even though the true impact is positive. This underlines the ability of even a small low frequency co-movement to drive a qualitatively important wedge between the IRFs based on the levels and differenced models.

The lower panel of Figure 3 also provides interesting information. When the HP filter is used to remove the low frequency component from labour productivity growth (Fernald, 2007), the estimated IRF resembles the IRF computed from the differenced specification. The graphs clearly demonstrate that the removal of the low frequency component, by either differencing or HP filtering, eliminates the possibility of any low frequency co-movements between the transformed series and this has a profound influence on the IRFs. We also considered the specification when hours worked are HP-filtered as in Francis and Ramey (2009). The behavior of the IRF estimates in this model is similar to the case of HP-filtered productivity growth.

Figure 4 presents the results for the exact unit root case. In this case, the matrix of largest roots becomes diagonal, eliminating the low frequency co-movement between hours and productivity growth ( $\delta = 0$ ). Despite some small biases, all median IRF estimates now correctly sign the impact of the technology shock and come close to tracing out the true IRFs. Not surprisingly, the differenced specification is particularly accurate and produces an unbiased estimator with tight confidence intervals. The estimator from the levels specification exhibits both a modest bias that arises from the biased estimation of the largest root of hours and a very large sample uncertainty (Gospodinov, 2010). The estimator using HP filtered labour productivity growth performs similarly to the differenced estimator, although it is slightly biased and more dispersed.

In Figure 5, we maintain the assumption of a zero off-diagonal element ( $\delta = 0$ ) and return to a persistent but stationary specification for hours worked ( $\rho = 0.95$ ). The median IRFs from all models are again quite similar, both to each other and to the true IRF. In this sense, beside having smaller bias and variance, the basic message from Figures 4 and 5 is similar, despite the fact that hours are nonstationary in Figure 4 but stationary in Figure 5.

In summarizing the results from these four figures, we note that large qualitative differences in median IRFs for the differenced and levels VARs were observed only in Figure 3, in which there is a small low frequency relationship between hours and labour productivity ( $\delta \neq 0$ ). Neither Figure 4 nor Figure 5 show qualitative differences in the median IRFs from the levels and differenced specifications. Yet, in Figure 4, hours have a unit root, whereas they are stationary in Figure 5. While small sample bias is present and affects the precision of the estimation, our simulations show that the small sample bias and persistence of the non-technology shocks alone are not enough to generate the substantial differences in the impulse responses that we find in practice. What Figures 4 and 5 share in common is the absence of the low frequency co-movement of Figure 3. Although the size of the unit root in hours worked has important implications for the sampling distributions of the IRFs, these results suggest that it is the low frequency co-movement that plays the critical



role in driving the qualitative differences between the level and differenced specifications.

We also want to stress that the confidence intervals reported in Figures 3-5 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). This is not reflected in our infeasible confidence intervals. At the same time, Figures 3 and 4 show clearly how a wide range of different estimates for the IRF are possible, and that the sampling uncertainty in the levels VAR is indeed larger. At the same time, except for the cases in which either  $\rho$  is exactly one or  $\delta$  is exactly zero, the true impulse response is never contained in the Monte Carlo confidence bands for the VAR in first differences.

To better assess the sensitivity of the levels and differenced specifications to different values of  $\delta$ , we plot in Figure 6 the true and estimated responses for  $\rho = 0.95$  and various degrees of low frequency co-movement. Each line represents values for  $\gamma = \{-0.5, -0.2, 0, 0.2, 0.5\}$ , which correspond to different off-diagonal elements  $\delta$  depending on the value of  $\rho$  (recall that  $\delta = -\gamma(1 - \rho)$ ). Once again, it is clear that, while the level specification explicitly estimates and incorporates the different values for  $\delta$  in the computation of the impulse response functions, the differenced specification implicitly imposes this element to be zero. This leads to substantial deviations from the true impulse response functions.

Finally, the differences in the IRFs for the various model specifications are expected to arise only in the case of long-run identification restrictions that are directly affected by the inclusion or the omission of the low frequency component. In order to verify this conjecture, we estimate the IRFs based on a short-run identification (Choleski decomposition) scheme, with productivity growth ordered first and hours second. We emphasize that our short-run identifying scheme is used only to illustrate the relative insensitivity of the IRFs to the low frequency co-movement when they are identified by short-run restrictions. We do not advocate its use in practice since it has

no clear theoretical justification. (See Christiano, Eichenbaum and Vigfusson (2006) for a more sophisticated, model-based, short-run identification scheme.)

The results from the three models for  $\rho = 0.95$  and  $\delta = 0.04$  are presented in Figure 7. Unlike the long-run identification scheme (Figure 3), the IRF estimates for all specifications are very close to the true IRF and fall inside the 95% Monte Carlo confidence bands. This suggests that the short-run identification scheme is robust to the presence or absence of low frequency co-movements, which is not the case with identifying restrictions that are based on long-run information.

The simulation results so far are obtained under the maintained hypothesis of an underlying low frequency co-movement in the data for  $\delta \neq 0$ , which is implicitly assumed to be structural in nature. An online appendix (available on the authors' webpages) presents additional simulation results from two observationally equivalent structural break models that give rise to the common high-low-high pattern observed by Fernald (2007). These simulations are intended to provide some insight into the reasons underlying the differences between our results and the results in Fernald (2007). The simulation results clearly indicate that imposing the assumption that the timing of the breaks is coincidental (Fernald, 2007) or exogenously driven (Francis and Ramey, 2009) is crucial for producing evidence supporting the break removal or HP filtering prior to the IRF analysis advocated by Fernald (2007) and Francis and Ramey (2009). By contrast, the results from a co-break model, where the similar magnitude and timing of the breaks is driven by a common underlying component, are qualitatively similar to those presented above, in which removing the low frequency component leads to substantial deviations of the IRF estimates from their true values.

## 5 Discussion of Results

The analytical and numerical results presented above clearly suggest that some seemingly innocuous transformations of the data can lead to vastly (qualitatively and quantitatively) different policy recommendations. The main objective of this paper is to illustrate and identify the source of these

differences. At the same time, several interesting observations and remarks emerge from our analysis that highlight some potential pitfalls in empirical work with structural dynamic models that use highly persistent variables in conjunction with long-run identifying restrictions. Following Blanchard and Quah (1989), long-run identifying restrictions have been popular tools in applied macroeconomics (see for example Rogers (1999), Lastrapes (1992) and Clarida and Gali (1994) among others), while ever since Nelson and Plosser (1982), the low frequency properties of macroeconomic data series has been widely debated. Therefore, these more general insights are likely also to be useful in other contexts.

First, it is common practice in macroeconomics to remove low frequency components by applying the HP filter when focusing on business cycle frequencies. For example, Fernald (2007) argues that the low frequency component is not important for business cycle analysis. The effect of technology shocks on hours worked is typically evaluated at business cycle frequency and it is reasonable to assume that the removal of low frequency components will not affect the conclusions. We agree with this position, provided that the structural shocks are identified using short- or medium-run restrictions. We argue that the low frequency component contains important long-run information that, while not directly relevant at business cycle frequencies, affects in a fundamental way the long-run restrictions. Therefore, omitting or explicitly removing the low frequency correlation can result in misspecification of the long-run restriction and hence the business cycle component that is of primary interest to the analysis. In contrast, the low frequency component does not seem to matter for the short-run restrictions and the transformations applied to the data do not affect the impulse responses that they identify, as illustrated in our simulation section.

Although the analogy is not exact, the removal of low frequency components bears some similarities to ignoring the long-run information contained in the error-correction term in cointegrated models. The cointegration information does not directly affect the business cycle analysis, but is essential to the long-run equilibrium. If we use short-run restrictions, the cointegration information

can be left out without serious consequences. If the data are subjected to differencing (filtering) prior to the analysis, the long-run information contained in the cointegrating relationship is lost and the long-run restriction is misspecified, which in turn gives rise to misleading results.

Second, it is well known that a highly persistent linear process often exhibits dynamics that are observationally equivalent to dynamics generated by a long memory, structural break or regime-switching process. Therefore, it is difficult to statistically distinguish between these processes in finite samples and commit to a particular specification. In our context, it is hard to determine if the low frequency component (for example, the U shape in hours worked) and co-movement are spurious or not. Fernald (2007) convincingly illustrates the cost of falsely keeping the low frequency component if this co-movement is spurious. Our results indicate that there is an equally large cost of falsely removing it when the co-movement is a true feature of the correctly identified model. Ultimately, the researcher has to take a stand on whether the long-restriction applies to the original or filtered data. Our analysis in the previous sections provides important information on the sensitivity (robustness) of the different statistical transformations of the data to misspecification of the long-run restriction.

Finally, pre-testing procedures that are used to determine which specification is more appropriate perform poorly, especially when the data are highly persistent. Our analysis suggest that large differences in the IRFs arise even when the largest root is arbitrarily close to one, in which case the pre-testing procedure selects the differenced specification with probability approaching one. Put another way, we find that, when identified by LR restrictions, the IRFs from the differenced specification are not robust to small deviations of the largest root from unity, even when those deviations are too small to be empirically detected.

## 6 Conclusion

This paper analyzes the source of the conflicting evidence from structural VARs identified by long-run restrictions on the effect of technology shocks on hours worked reported in several recent empirical studies. We show analytically that the extreme sensitivity of the results to different model specifications can be explained by a discontinuity in the solution for the structural coefficients implied by the long run restrictions, which arises only in the presence of a low frequency correlation between hours worked and productivity growth. The critical mechanism underlying the difference between the levels and differenced specifications, is that the differenced specification restricts this correlation to zero when solving for structural model, whereas the levels specification allows it to enter in unrestricted manner. Consequently, it may not be surprising that alternative filtering approaches reported in the literature, such as HP filtering and trend-break removal, which eliminate this low frequency correlation, provide evidence supportive of the differenced VAR. We also demonstrate that low frequency correlations capable of causing strong discrepancies between the two specifications are compatible with autoregressive roots in hours worked that are indistinguishable from one. This sharp discontinuity implies that one cannot rely on univariate unit root tests to resolve this debate, since they are not designed to discriminate between exact and near-unit root models.

Fernald (2007) also highlights the role of an observed low frequency correlation in the data, modeled as a common U-shaped pattern driven by structural breaks, and provides a number of convincing empirical exercises to illustrate its importance. While this insightful analysis clearly demonstrates the empirical link between the low frequency correlation and the conflicting results, to date there has not been a full theoretical understanding of why this low frequency correlation plays such an important role. We fill this gap by showing, in a more general analytic framework, that the key role played by this low frequency correlation is to create a discontinuity between the structural solutions of the differenced and levels specifications.

We argue that the importance of the low frequency correlation cannot by itself resolve the debate, because, depending on the way it is modeled, it may lead to biases in either the levels or differenced specification. However, in conjunction with Fernald (2007) and Francis and Ramey (2009), our results help to clarify the terms of the debate. We demonstrate that if there is a true low frequency correlation in the population model that is correctly identified by the long-run identification restriction, then any procedure which removes this low frequency correlation, whether by differencing, HP filtering, or trend-break removal would result in a substantial bias. In fact, we find that one cannot reproduce the discrepancy in the results with any reasonable probability in a correctly identified model, without introducing such a true population correlation.

The reason that this finding might seem to be at odds with those of Fernald (2007) and Francis and Ramey (2009), who both argue that the levels specification is biased, is that neither of these papers perceive the observed correlation as an inherent feature of the correctly identified model. Fernald (2007) argues that the observed low frequency correlation is purely coincidental as a similar pair of breaks in both series occur due to historical happenstance. Francis and Ramey (2009) treat the observed correlation as a population characteristic explained by common low frequency trends in demographic and public employment, but argue that the long-run restriction does not hold until these low frequency trends are purged from the data. In our view, the debate therefore hinges on the interpretation given to this low frequency correlation. If one has reasons to believe that there is a genuine low frequency co-movement in a correctly identified model, this would support the findings of Christiano, Eichenbaum and Vigfusson (2003). On the other hand, if one is convinced either that the observed correlation is coincidental (Fernald, 2007) or that it is due to factors that violate the identification restriction (Francis and Ramey, 2009), then the results may be interpreted as supporting the earlier findings of Gali (1999). More generally, our results help to explain the empirically observed sensitivity of long-run identifying schemes to uncertainty regarding low frequency dynamics, even when identifying characteristics at business cycle frequencies.

## Acknowledgements

The authors thank the Editor (Jonathan Wright), an Associate Editor, two anonymous referees, Michelle Alexopolous, Efrem Castelnuovo, Yongsung Chang, Francesco Furlanetto, Leo Michelis, Barbara Rossi, Shinichi Sakata, seminar participants at University of Padova, University of Cyprus, University of British Columbia, Carleton University, Queens University and Louisiana State University and conference participants at the 2008 Far Eastern Meetings of the Econometric Society, the 2008 Meeting of the Canadian Econometric Study Group, ICEEE 2009, CIREQ Time Series Conference and the 2009 Joint Statistical Meetings for helpful comments and suggestions. Part of this research was done while Elena Pesavento was a Marco Fanno visitor at the Dipartimento di Scienze Economiche at Universita' di Padova. Nikolay Gospodinov gratefully acknowledges financial support from FQRSC and SSHRC.

## References

- [1] Alexopoulos, M. (2006), “Read All About It!! What Happens Following a Technology Shock?” Mimeo, University of Toronto.
- [2] Basu, S., Fernald, J. G., and Kimball, M. S. (2006), “Are Technology Improvements Contractionary?” *American Economic Review*, 95, 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] Chan, H. C. (1988), “The Parameter Inference for Nearly Nonstationary Time Series,” *Journal of the American Statistical Association*, 83, 857-862.
- [5] Chang, Y., Hornstein, A., and Sarte, P.-D. (2009), “On the Employment Effects of Productivity Shocks: The Role of Inventories, Demand Elasticity, and Sticky Prices,” *Journal of Monetary Economics*, 56, 328-343.
- [6] Chari, V., Kehoe, P., and McGrattan, E. (2008), “Are Structural VARs with Long-Run Restrictions Useful in Developing Business Cycle Theory?” *Journal of Monetary Economics*, 55, 1337-1352.
- [7] Christiano, L., Eichenbaum, M., and Vigfusson, R. (2003), “What Happens After a Technology Shock?” Board of Governors of the Federal Reserve System, International Finance Discussion Paper 768.
- [8] Christiano, L., Eichenbaum, M., and Vigfusson, R. (2006), “Assessing Structural VARs,” in *NBER Macroeconomics Annual*, eds. D. Acemoglu, K. Rogoff and M. Woodford, Cambridge: MIT Press.
- [9] Clarida, R., and Gali, J. (1994), “Sources of Real Exchange Rates Fluctuations: How Important are Nominal Shocks?” *Carnegie-Rochester Conference Series in Public Policy*, 41, 1-56.
- [10] Erceg, C., Guerrieri, L., and Gust, C. (2005), “Can Long-Run Restrictions Identify Technology Shocks?” *Journal of the European Economic Association*, 3, 1237-1278.
- [11] Fernald, J. (2007), “Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements,” *Journal of Monetary Economics*, 54, 2467-2485.
- [12] Francis, N., Owyang, M., and Roush, J. (2005), “A Flexible Finite-Horizon Identification of Technology Shocks,” Federal Reserve Bank of St. Louis, Working Paper 2005-024E.
- [13] Francis, N., and Ramey, V. (2005), “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *Journal of Monetary Economics*, 52, 1379-1399.
- [14] Francis, N., and Ramey, V. (2009), “Measures of Hours Per Capita with Implications for the Technology-Hours Debate,” *Journal of Money, Credit, and Banking*, 41, 1071-1097.



- [15] Gali, J. (1999), “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, 89, 249-271.
- [16] Gali, J., and Rabanal, P. (2004), “Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data?” in *NBER Macroeconomics Annual*, eds. B. Bernanke and K. Rogoff, Cambridge: MIT Press.
- [17] Gospodinov, N. (2010), “Inference in Nearly Nonstationary SVAR Models with Long-Run Identifying Restrictions,” *Journal of Business and Economic Statistics*, 28, 1-12.
- [18] Lastrapes, W. D. (1992), “Sources of Fluctuations in Real and Nominal Exchange Rates,” *Review of Economic and Statistics*, 74, 530-539.
- [19] Nelson, C., and Plosser, C. (1982), “Trends and Random Walks in Macroeconomic Time-Series: Some Evidence and Implications,” *Journal of Monetary Economics*, 10, 139-162.
- [20] Pesavento, E., and Rossi, B. (2005), “Do Technology Shocks Drive Hours Up or Down? A Little Evidence from an Agnostic Procedure,” *Macroeconomic Dynamics*, 9, 478-488.
- [21] Pesavento, E., and Rossi, B. (2006), “Small Sample Confidence Intervals for Multivariate Impulse Response Functions at Long Horizons,” *Journal of Applied Econometrics*, 21, 1135-1155.
- [22] Phillips, P. C. B. (1987), “Towards a Unified Asymptotic Theory for Autoregression,” *Biometrika*, 74, 535-547.
- [23] Ravenna, F. (2007), “Vector Autoregressions and Reduced Form Representations of DSGE Models,” *Journal of Monetary Economics*, 54, 2048-2064.
- [24] Rogers, J. H (1999), “Monetary Shocks and Real Exchange Rates,” *Journal of International Economics*, 49, 269-288.
- [25] Shea, J. (1999), “What Do Technology Shocks Do?” in *NBER Macroeconomics Annual*, eds. B. Bernanke and J. Rotemberg, Cambridge: MIT Press.
- [26] Uhlig, H. (2004), “Do Technology Shocks Lead to a Fall in Total Hours Worked?” *Journal of the European Economic Association*, 2, 361–371.




Figure 1. Response of hours worked to a 1% positive technology shock, U.S. data 1948Q2 - 2005Q3.

Top graph: hours worked in levels; Bottom graph: hours worked in first differences.




Figure 2. HP trends of labour productivity growth (top graph) and hours worked (bottom graph), U.S. data 1948Q2 - 2005Q3.

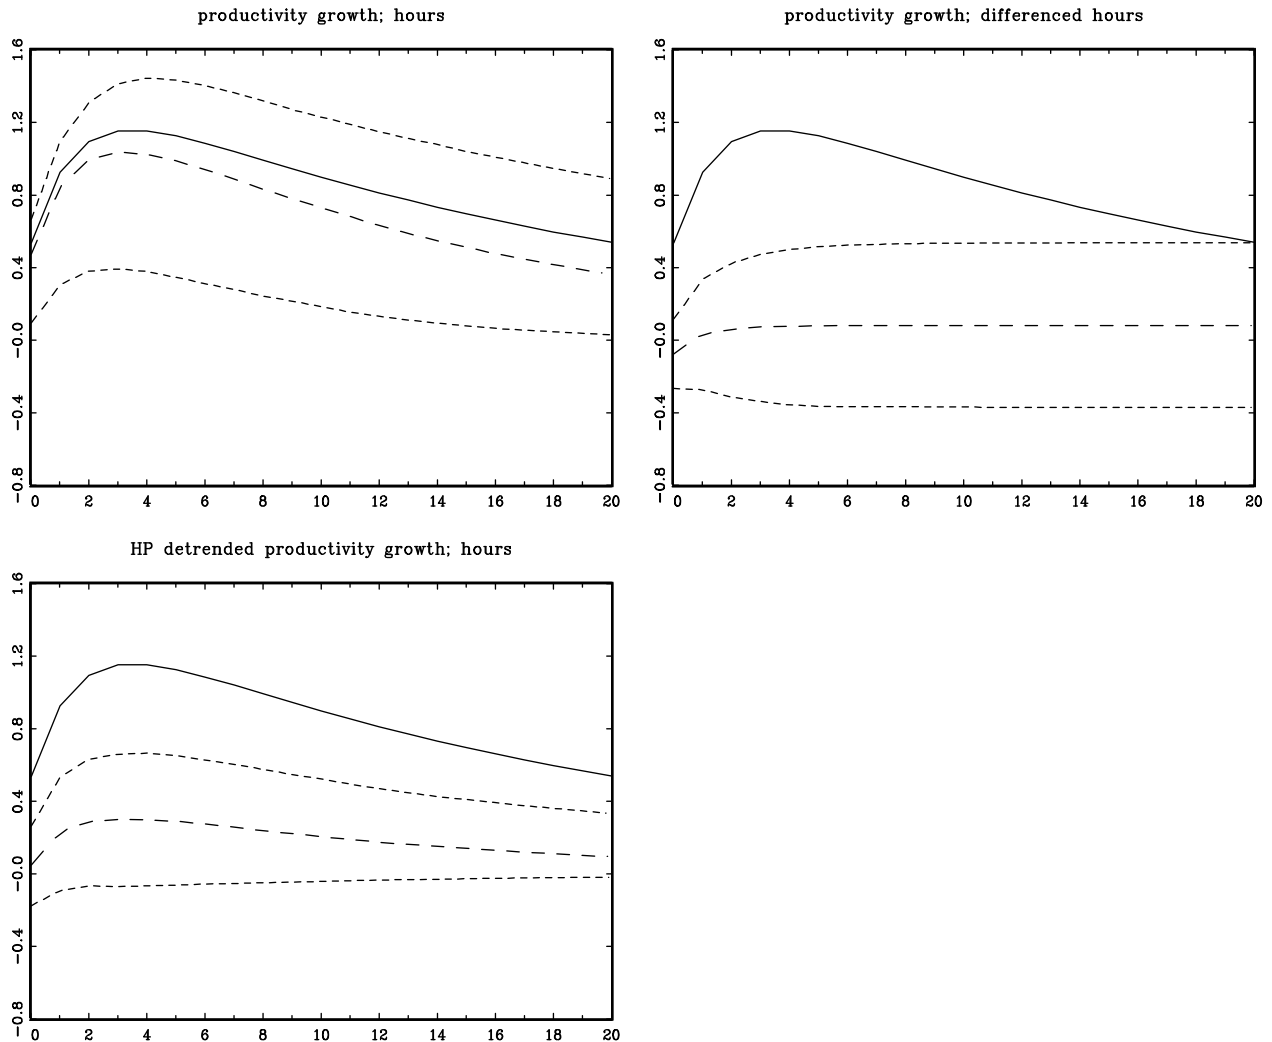


Figure 3. Response of hours to a positive technology shock (long-run identification) with data simulated from the model  $\left[ \mathbf{I} - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L \right] \left[ \mathbf{I} - \begin{pmatrix} 1 & \delta \\ 0 & \rho \end{pmatrix} L \right] \begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}$ , where  $\rho = 0.95$ ,  $\delta = 0.04$  ( $\gamma = -0.8$ ),  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$  and  $T = 250$ . Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

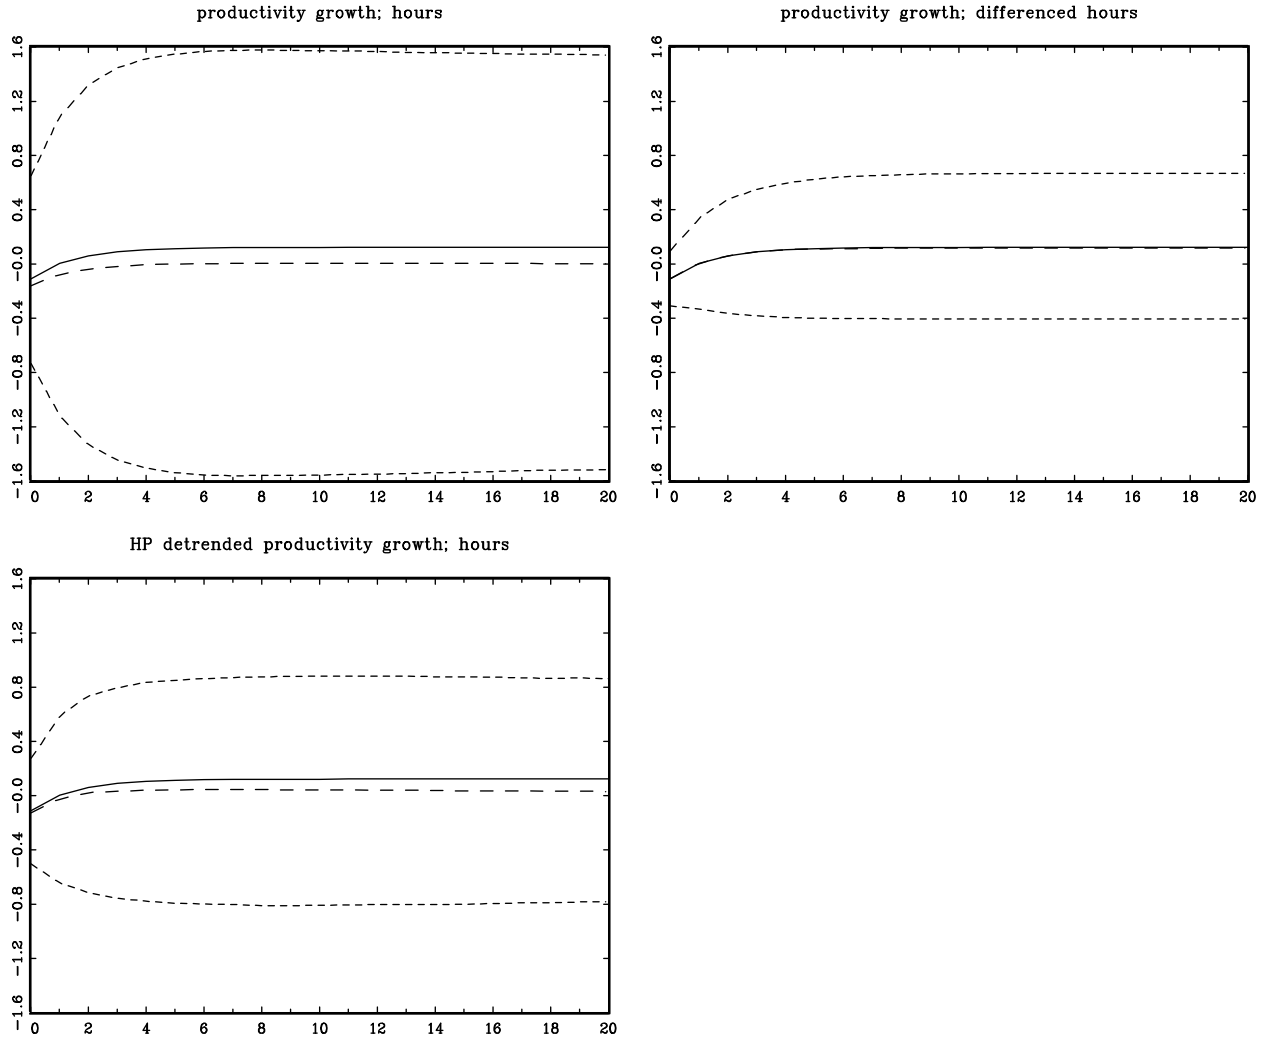


Figure 4. Response of hours to a positive technology shock (long-run identification) with data simulated from the model  $\left[ \mathbf{I} - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L \right] \left[ \mathbf{I} - \begin{pmatrix} 1 & \delta \\ 0 & \rho \end{pmatrix} L \right] \begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}$ , where  $\rho = 1$ ,  $\delta = 0$  ( $\gamma = 0$ ),  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$  and  $T = 250$ . Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

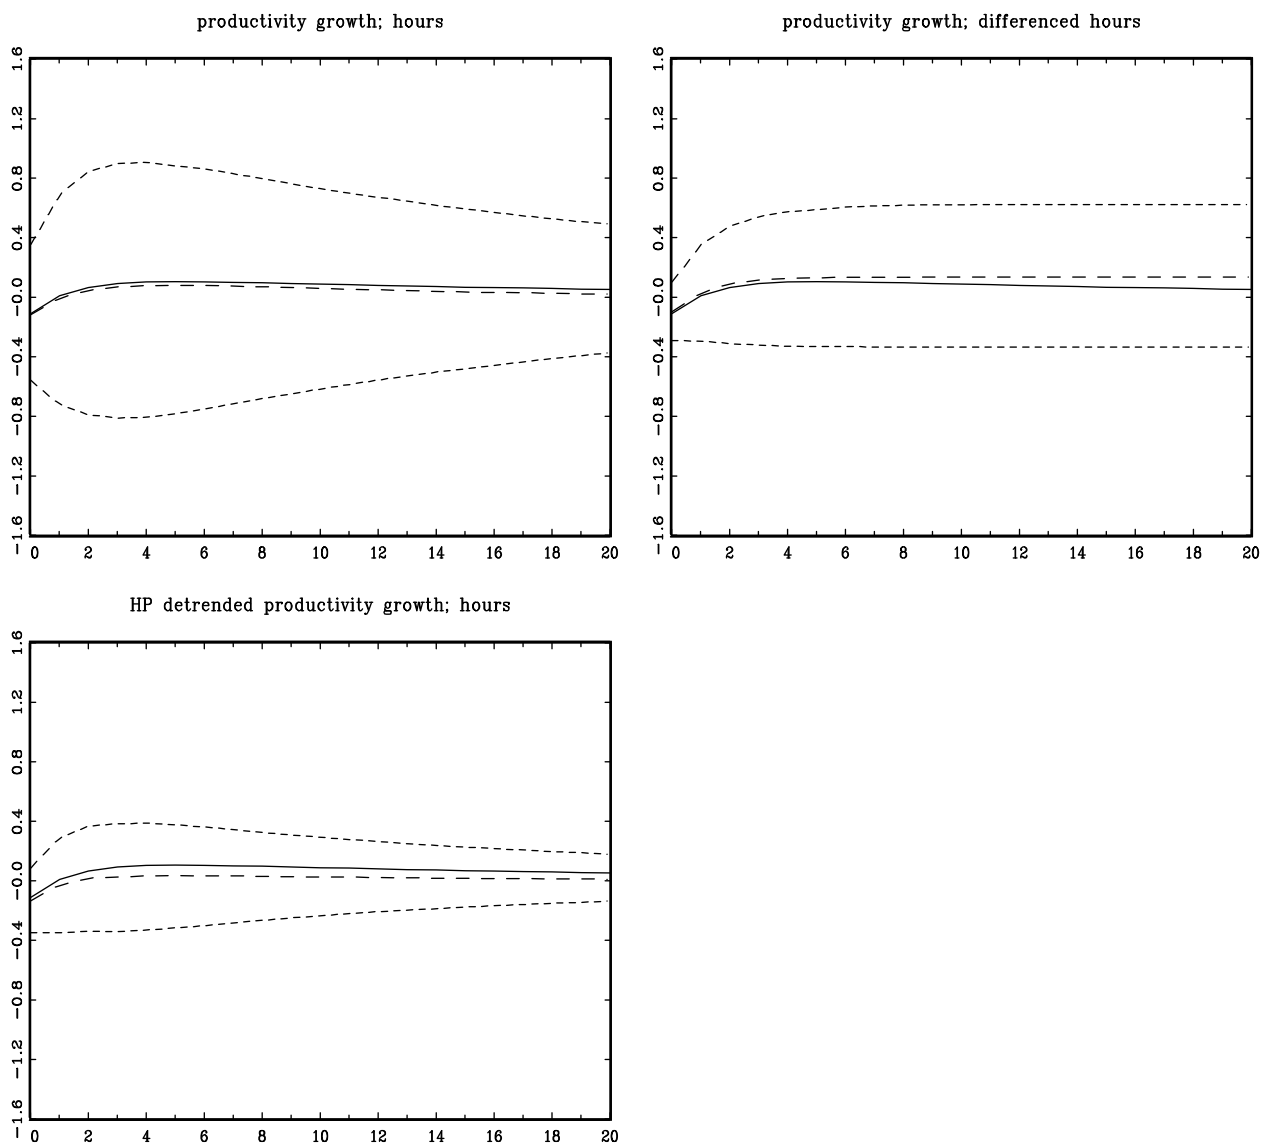


Figure 5. Response of hours to a positive technology shock (long-run identification) with data simulated from the model  $\left[ \mathbf{I} - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L \right] \left[ \mathbf{I} - \begin{pmatrix} 1 & \delta \\ 0 & \rho \end{pmatrix} L \right] \begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}$ , where  $\rho = 0.95$ ,  $\delta = 0$  ( $\gamma = 0$ ),  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$  and  $T = 250$ . Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

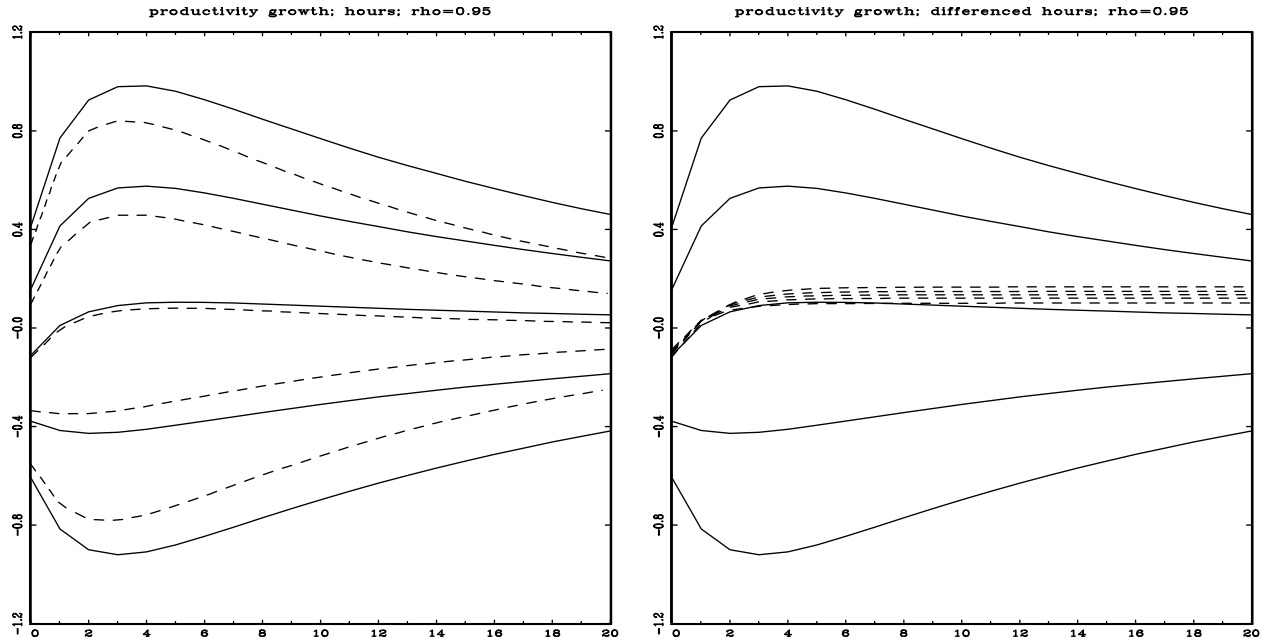


Figure 6. Response of hours to a positive technology shock (long-run identification) with data simulated from the model  $\left[ \mathbf{I} - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L \right] \left[ \mathbf{I} - \begin{pmatrix} 1 & -\gamma(1-\rho) \\ 0 & \rho \end{pmatrix} L \right] \begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}$ , where  $\rho = 0.95$ ,  $\gamma = \{-0.5, -0.2, 0, 0.2, 0.5\}$ ,  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$  and  $T = 250$ . Solid line: true IRF; short dashes: median Monte Carlo IRF estimate.

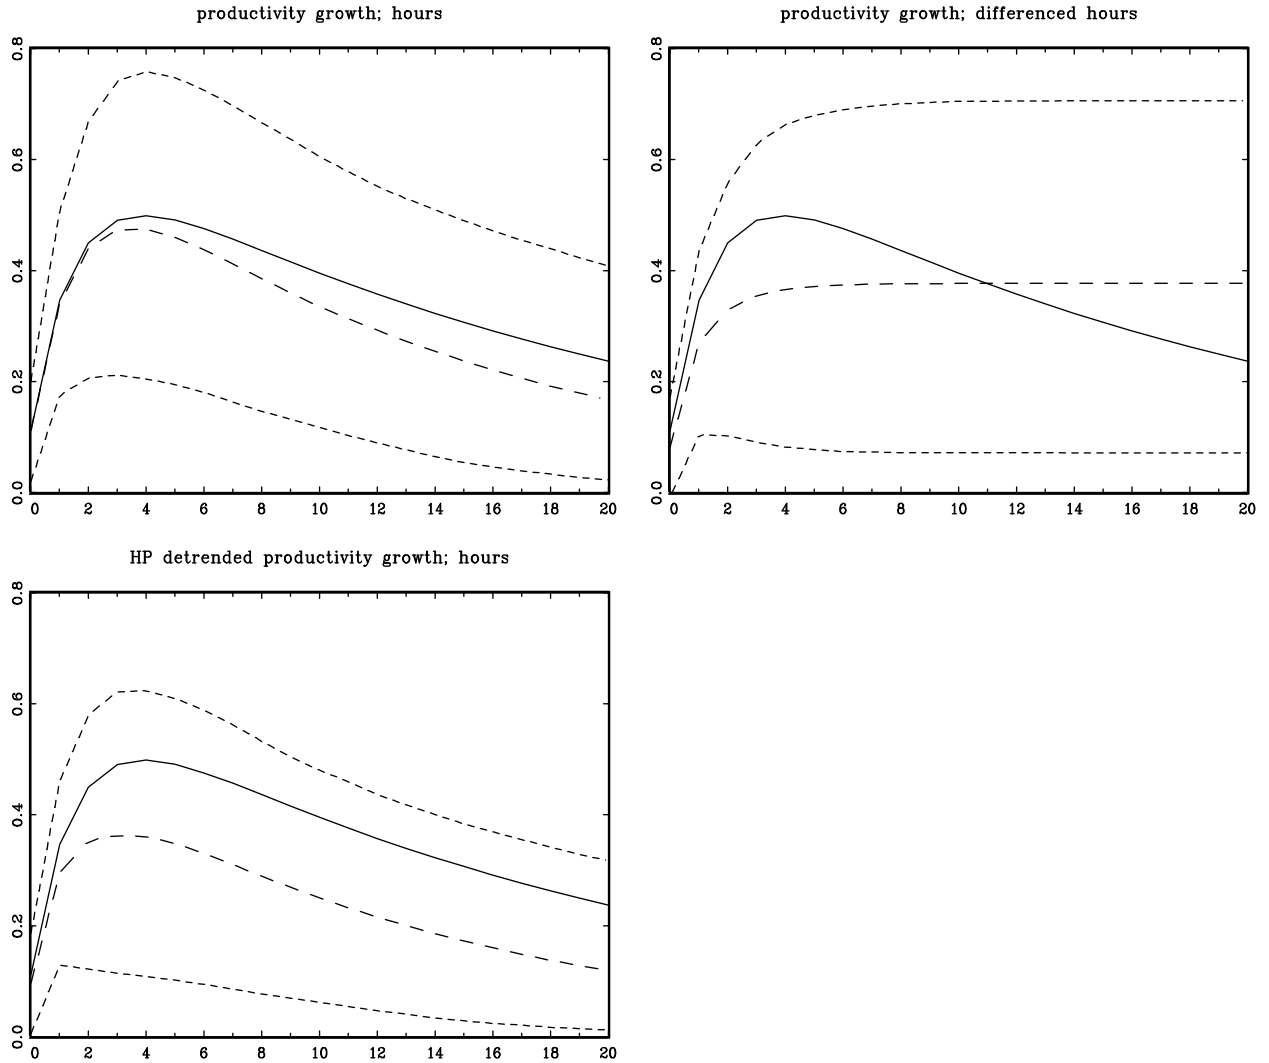


Figure 7. Response of hours to a positive technology shock (short-run (Choleski) identification) with data simulated from the model  $\left[ \mathbf{I} - \begin{pmatrix} -0.05 & -0.08 \\ 0.2 & 0.55 \end{pmatrix} L \right] \left[ \mathbf{I} - \begin{pmatrix} 1 & \delta \\ 0 & \rho \end{pmatrix} L \right] \begin{pmatrix} l_t \\ h_t \end{pmatrix} = \begin{pmatrix} u_{1,t} \\ u_{2,t} \end{pmatrix}$ , where  $\rho = 0.95$ ,  $\delta = 0.04$  ( $\gamma = -0.8$ ),  $(u_{1,t}, u_{2,t})' \sim iidN(\mathbf{0}, \Sigma)$ ,  $\Sigma = \begin{pmatrix} 0.78 & 0.1 \\ 0.1 & 0.55 \end{pmatrix}$  and  $T = 250$ . Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.