

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

March 2, 2009

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 can be explained by a small, but important, low frequency co-movement between hours worked and labour productivity growth, which is allowed for in the levels specification but is implicitly set to zero in the differenced VAR. Our theoretical analysis shows that, even when the root of hours is very close to one and the low frequency co-movement is quite small, assuming away or explicitly removing the low frequency component can have large implications for the long-run identifying restrictions, giving 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.

*Corresponding author. Department of Economics. Concordia University, 1455 de. Maisonneuve Blvd. West, Montreal, Quebec, H3G 1M8. nikolay.gospodinov@concordia.ca. The authors thank Efrem Castelnuovo, Yongsung Chang, Leo Michelis, Barbara Rossi and seminar participants at University of Padova and Louisiana State University and conference participants at the 2008 Far Eastern Meetings of the Econometric Society, the 2008 Meetings of the Canadian Econometric Study Group, and ICEEE 2009. 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.

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.¹

As technology shocks are difficult to measure,² they are commonly specified (Gali, 1999, for example) 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. 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 are 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 (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

¹However, Chang, Hornstein, and Sarte (2008) show that once inventories are allowed for the dynamics in both models become more complicated and it is no longer so simple to distinguish between flexible and sticky price models based on the response of hours worked.

²Alexopoulos (2006) and Shea (1999) provide measurements of technological progress based on technology publications and patent data respectively.

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. Chang and Hong (2006) reach similar conclusions using data on industry's total factor productivity despite specifying hours in first differences.

Figure 1 plots the estimated impulse response functions (IRFs) based on the levels and differenced specifications with quarterly U.S. data for the period 1948Q2 - 2005Q3.³ 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).⁴

These contradicting results have generated substantial interest in the effect of misspecifying hours worked. Chari, Kehoe and McGrattan (2008) 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 (2008) 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. Pesavento and Rossi (2005) adopt an agnostic approach that is robust to the degree of persistence in 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 infinite lag structure (Chari, Kehoe and McGrattan, 2008; Christiano, Eichenbaum and Vigfusson, 2006; Ravenna, 2007). Furthermore, as noticed by Chari, Kehoe and McGrattan (2008) 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 com-

³U.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).

⁴Using a multivariate Bayesian posterior odds procedure, Christiano, Eichenbaum and Vigfusson (2003) find evidence in favor of the levels specification.

peting economic theories. Gospodinov (2008) 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 differenced VAR specifications may be unreliable when the root in hours worked is close to one. In particular, the impulse response functions can be severely 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 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 explain these differences based solely on the behavior of hours worked itself and that these differences cannot be justified solely 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 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 imposes this element to be zero. It is important to emphasize that this component could be arbitrarily small and could accompany a largest 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.

Another way to look at the problem is to note that the differenced specification ignores possible low frequency co-movements between labour productivity and hours worked. Figure 2 plots de-meaned hours worked and detrended labour productivity and shows that the two series appear to be inversely related. 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 without inducing any observable changes in its time series properties.

We show in our analytical section that this inverse relationship should be translated into low frequency co-movement between labour productivity growth and hours worked, provided that hours worked are stationary. While this co-movement is hard to detect by visual inspection of the dynamics of labour productivity growth, Figure 3 reveals that the Hodrick-Prescott (HP) trend⁵ 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.

Fernald (2007) also highlights the sensitivity of the results to low frequency correlation between labour productivity growth and hours worked and concentrates on the observed low frequency correlation that is due to a similar high-low-high pattern in productivity growth (arising from structural breaks in productivity) and hours per person. Thus, our findings support those of Fernald (2007) in identifying low frequency correlations as the key to understanding the conflicting results from the levels and differenced specifications. However, our analysis differs from that of Fernald (2007) in a subtle, but fundamentally important, way. In order to generate (with nontrivial probability) differences in the two specifications of a similar magnitude to those found in practice,

⁵Throughout the paper, the value of the smoothing parameter for the HP filter is set to 1,600.

we find it necessary to include this low frequency co-movement in the true generating process. Such a 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 effecting regional transportation or labour search costs, may also have lasting impacts on labour markets. Synthetic simulated series from the RBC model that we briefly consider in Section 3.2 also result in low frequency co-movements similar to those observed in the empirical data.

On the other hand, Fernald (2007) and Francis and Ramey (2006) provide some convincing arguments for why the structural breaks in productivity and hours may have disparate causes. In their analysis, they therefore treat the concurrent nature of the structural breaks in the two series as coincidental. While such a chance occurrence cannot be ruled out in the historical data, in our calibrated simulations we find it nearly impossible to replicate the differences between the two specifications without treating this small low frequency correlation as a real characteristic of the underlying data generating process. Furthermore, when this low frequency component is real, we find that removing or artificially restricting it to zero, as in the difference specification, may lead to substantial biases in the impulse response function, whereas Fernald (2007) comes to the opposite conclusion under the assumption that this correlation is coincidental in nature.

The rest of the paper is organized as follows. 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. Section 4 discusses the main implications of our analysis for empirical work and Section 5 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 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 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 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 difference and level VARs, even when the largest root is arbitrarily close to and indistinguishable

from unity.⁶

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.

2.1 Reduced-form model

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

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

where $\Psi(L) = I - \sum_{i=1}^p \Psi_i L^i = \begin{bmatrix} \psi_{11}(L) & \psi_{12}(L) \\ \psi_{21}(L) & \psi_{22}(L) \end{bmatrix}$, $E(u_t | u_{t-1}, u_{t-2}, \dots) = 0$, $E(u_t u_t' | u_{t-1}, u_{t-2}, \dots) =$

Σ and the matrix Φ is expressed in terms of its eigenvalue decomposition 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 cor-

responding eigenvectors.⁷ Simple algebra yields $\Phi = \begin{bmatrix} 1 & \delta \\ 0 & \rho \end{bmatrix}$, where $\delta = -\gamma(1 - \rho)$, is the

parameter that determines the low frequency co-movement between the variables and ρ denotes

the largest root of hours worked. This parameterization, which arises directly from the eigenvalue

decomposition of Φ , allows for a small (δ) impact of h_t on l_t , provided that ρ is not *exactly* equal

to one. Note that in the exact unit root case, Φ collapses to the identity matrix.⁸ The other off

diagonal element of V , and therefore of Φ , is set to zero to rule out the possibility of feedback from

the level of productivity to hours worked, as this would imply that hours is $I(2)$ when $\rho = 1$ and

⁶This argument can also be formalized in the 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.

⁷Pesavento and Rossi (2006) use a similar decomposition but they impose diagonality on Φ . In Pesavento and Rossi (2006), the eigenvectors represent possible cointegration relationships.

⁸The 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 parameterization in Section 2.4.

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

It is convenient to rewrite model (1) in Blanchard and Quah's (1989) framework by imposing the exact unit root on productivity so that Δl_t is a stationary process. In this case, let $y_t = (\Delta l_t, h_t)'$ 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

$$A(L)y_t = u_t \quad (2)$$

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

The non-zero off-diagonal element $\gamma(1-\rho)L$ 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 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 Φ can be used to explain this result.

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 $B_0 = \begin{bmatrix} 1 & -b_{12}^{(0)} \\ -b_{21}^{(0)} & 1 \end{bmatrix}$. Pre-multiplying 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.⁹ 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 $\Delta\tilde{y}_t = (\Delta l_t, \Delta h_t)'$. Then, the reduced form specializes to

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

and the structural form is given by

$$B_0 \Psi(L) \Delta \tilde{y}_t = \varepsilon_t$$

with a long-run multiplier matrix

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

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

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.

⁹ An alternative formulation of the restriction is that $C(I)$ is lower triangular, where $C(L) = B(L)^{-1}$ describes the moving average representation $y_t = C(L)\varepsilon_t$. Simple matrix algebra shows that the two restrictions are equivalent.

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, 2008). 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 γ does not depend on ρ , these differences can remain large even for $(\rho - 1)$ arbitrarily close, but not equal, to zero. For simplicity, take the first-order model where $\Psi(L) = I$. In this case, the two restrictions set the value of $b_{12}^{(0)}$ to γ and 0, respectively, implying two very different values for $b_{21}^{(0)}$, which, in turn, directly determines the impulse response function, since in the first-order model

$$\theta_{hz}^{(l)} = \frac{\partial h_{t+l}}{\partial \varepsilon_t^z} = \left[\Phi^l B_0^{-1} \right]_{21} = \frac{b_{21}^{(0)} \rho^l}{1 - b_{12}^{(0)} b_{21}^{(0)}}. \quad (4)$$

As it is clear from (4), 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 γ . Focusing the debate on the distance of ρ from one is therefore misleading, provided that ρ is not precisely equal to one.

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 that 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 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, 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 γ , rather than through either δ or ρ . 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 ρ 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.

2.4 An alternative parameterization

The fact that our framework suggests potentially large IRF discrepancies even for values of ρ 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 arises in this case because, in small sample sizes, 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), which shrinks towards unity in larger samples (Phillips, 1987; Chan, 1988). Naturally, this dependence on the sample size is not understood 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 interpreted 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 ρ 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 $\Lambda = \begin{bmatrix} 1 & 0 \\ 0 & 1 - c/T \end{bmatrix}$, $V = \begin{bmatrix} 1 & -\gamma \\ 0 & 1 \end{bmatrix}$ 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 arbitrary small. The reduced form for $y_t = (\Delta l_t, h_t)'$ is now

$$A(L)y_t = u_t$$

with $A(L) = \Psi(L) \begin{bmatrix} 1 & \gamma(c/T)L \\ 0 & (1-L) + (c/T)L \end{bmatrix}$. In the unit root case, $c = 0$ and Φ_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 Δl_t is local-to-zero and vanishing at rate $T^{-1/2}$,¹⁰ capturing the notion that the low frequency co-movement between hours and productivity 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 Δl_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 $c = 0$, the model

¹⁰The level of h affects Δl_t through the term $(\gamma c/T)h_{t-1}$ which is $O_p(T^{-1/2})$ since $T^{-1/2}h_{t-1} = O_p(1)$.

again specializes to the difference 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 2.3 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 is stationary. Thus the finding that hours worked is indistinguishable from a unit root process does not guarantee that the true IRF will be close to IRF from the difference specification.

3 Monte Carlo Experiment

3.1 Model calibrated to US data

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 $\delta = -\gamma(1 - \rho)$, $T = 250$, $(u_{1,t}, u_{2,t})' \sim N(0, \Sigma)$, $\Sigma = \begin{pmatrix} 0.78 & 0 \\ 0 & 0.55 \end{pmatrix}$, and the parameter values are calibrated to match the empirical shape of the IRF of hours worked to a technology shock.¹¹ 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).

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 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),

¹¹Note 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 also impose $\rho = 1$ and therefore allow both specifications (levels and first differences) to be the true DGP.

the 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 (left graph) estimates an IRF that is close, on average, to the true IRF, except for a small bias (see Gospodinov, 2008, 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. 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 co-movement 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 right panels of Figures 5 and 6 are also interesting. 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.¹²

¹²After 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. Nevertheless, we still report the IRFs on the same graph to illustrate the economically large differences created by a fairly small off-diagonal element. We also considered the specification when hours worked are HP-filtered as in Francis and Ramey (2006). The behavior of the IRF estimates in this model is similar to the case of HP-filtered productivity growth. Here, we make no argument as to whether the low frequency components should or should not be removed prior to the IRF analysis. Instead, we provide an analytical framework for explaining and reconciling the conflicting results documented in the empirical literature. We further discuss the implications of low frequency filtering in the next section.

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 co-movement between hours and productivity growth ($\delta = 0$). As expected, by removing this low frequency co-movement, we also eliminate the main qualitative differences between the median IRF response functions from the three models. 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.

Nevertheless, there are still important differences in performance among the three 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 both a modest bias that arises from the biased estimation of the largest root of hours and a very large sample uncertainty (Gospodinov, 2008). 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.

In Figure 8 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, the basic message from Figures 7 and 8 is similar, despite the fact that hours are nonstationary in Figure 7 but stationary in 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 figures, 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 labour productivity ($\delta \neq 0$). Neither Figure 7 nor Figure 8 show qualitative differences in the median IRFs from the levels and differenced specifications. Yet in Figure 7 hours have a unit root, whereas they are stationary in

Figure 8. What the two figures share in common is the absence of the low frequency co-movement of Figures 5 and 6 (i.e. $\delta = 0$). 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 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 , in Figures 9 and 10 we plot the true and estimated responses for various degrees of persistence and 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 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 differenced specification implicitly imposes this element to be zero. This leads to substantial deviations from the true impulse response functions.

We also want to stress that the confidence intervals reported in Figures 5-8 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 5-7 show very well 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 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.

Finally, the differences reported 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 from the different statistical models based on a short-run identification (Cholesky decomposition) scheme, with productivity growth ordered first and hours second. While we recognize that imposing short-run restrictions may be rather ad hoc and may lack a solid theoretical justification, Christiano, Eichenbaum and Vigfusson (2006) demonstrate that the short-run identification scheme produces estimates with appealing statistical properties.¹³ The results from the three models for $\rho = 0.97$ and $\delta = 0.015$ are presented in Figure 11. Unlike the long-run identification scheme (Figure 6), 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.

3.2 RBC model

It is interesting to see if our main conclusions continue to hold if the data are simulated from a dynamic general equilibrium model, in which the persistence and the low frequency co-movements between the variables are implicitly determined. To investigate this, we follow Chari, Kehoe and McGrattan (2008) and Christiano, Eichenbaum and Vigfusson (2006) by generating data from a real business cycle model. The true impulse response functions implied from this structural model are then compared to the estimated impulse responses from a finite-order VAR model. In particular, we use the two-shock CKM specification described in Christiano, Eichenbaum and Vigfusson (2006) as a data generating mechanism (see Christiano, Eichenbaum and Vigfusson, 2006, for details).

Several features of the RBC should be emphasized. First, the RBC model used for simulating the data imposes a unit root on technology while hours worked implied by the model are stationary but highly persistent. As a result, any low frequency co-movements in the model should arise from

¹³Our 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 persistence of the variables and not from structural breaks. Second, the RBC model implies a VARMA (infinite-order VAR) structure for $(\Delta l_t, h_t)'$ and fitting a finite-order VAR model to $(\Delta l_t, h_t)'$ results in biased estimates of the impulse response functions (Chari, Kehoe and McGrattan, 2008; Christiano, Eichenbaum and Vigfusson, 2006; Ravenna, 2007). Although there exist methods for correcting this misspecification bias (for instance, Christiano, Eichenbaum and Vigfusson, 2006), we do not pursue this avenue, since our primary focus in this paper is on the bias that arises in the differenced specification from omitting a possible low frequency co-movement, regardless of whether or not there is an additional source of bias due to lag truncation.

We generate 1,000 samples of 180 observations each and consider both the levels and differenced VAR specifications with four lags. Figure 12 displays time plots of a typical pair of synthetic sequences of demeaned hours and detrended labour productivity generated from the simulated RBC model.¹⁴ The co-movement of the series is similar to that shown in Figure 2 using the actual data. Likewise, Figure 13 displays HP trends of the simulated labor productivity growth and hours worked from the RBC model. The figure again shows a similar low frequency co-movement to that of the real data shown in Figure 3. Thus, the calibrated RBC model appears to produce a low frequency co-movement, similar to the one found in the empirical data. In conjunction with the lag truncation bias, this may help to explain why it produces the large discrepancies in the IRFs of the differenced and level specifications discussed below.

We now turn to the simulated IRFs. Again we consider structural VARs with hours in both differences and levels. Since hours worked is a highly persistent variable, it is tempting to subject this variable to a unit root pre-test and depending on the outcome to model hours either in levels or first differences. Thus, we also report the results from this pre-testing procedure in which the decision of modeling h in levels or first differences is based on an ADF test with 4 lags at 5% significance level.

¹⁴To avoid cherry-picking, we used the last of the 1,000 synthetic series from our simulation. Comparison to other draws indicated that it was not atypical.

The results from the three specifications are presented in Figure 14. As reported elsewhere (Christiano, Eichenbaum and Vigfusson, 2006, for example), the IRF estimates from the levels VAR suffer from an upward bias that is caused by approximating the true VARMA process by a short-order VAR. Using an estimate of the long-run variance matrix as suggested by Christiano, Eichenbaum and Vigfusson (2006) can substantially reduce this bias, although the sampling uncertainty associated with the IRF estimates remains large.¹⁵ As in the previous simulation design, the differenced specification reduces the sampling uncertainty but completely misses the true impulse response due to the omission of important low frequency information. The true impulse response falls entirely outside the 90% Monte Carlo confidence bands obtained from the differenced specification. Due to the relatively high persistence of hours worked, the pre-testing procedure has difficulties rejecting the unit root null and leads to only small improvements over the differenced specification. The estimates are slightly less biased and the confidence bands are wider reflecting the uncertainty regarding the presence of a unit root in hours worked.

In summary, the simulated data from the RBC model show low frequency co-movements similar to those found in the empirical data and produce IRFs in which the levels and difference specifications give widely divergent conclusions. Therefore, although the lag-truncation bias also plays an important role when the data is generated from a dynamic general equilibrium model, we nonetheless re-confirm the central role of the low frequency co-movement in explaining the discrepancy between the IRFs from level and difference specifications.¹⁶

¹⁵Christiano, Eichenbaum and Vigfusson (2006) also find that this bias is substantially smaller after relaxing the assumptions on the variance of the measurement error in the CKM specification. In addition, they consider a specification with wage and price frictions for which the sampling uncertainty is much reduced.

¹⁶Further supporting this conclusion, some results with other specifications of the RBC model (available from the authors upon request) do not seem to produce a negative low frequency co-movement between productivity and hours and also result in smaller discrepancies between the estimated impulse response functions from the levels and differenced specifications.

4 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, when using highly persistent variables in conjunction with long-run identifying restrictions.

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. Since most of the empirical research uses the long-run identifying scheme, our results tend to suggest that the low frequency component affects directly the long-run restrictions, which in turn are used for identifying the business cycle. The low frequency component contains 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 this 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 will be lost and the long-run restriction will be misspecified, which in turn will give 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. Importantly, our results indicate that the cost of falsely removing the low frequency component is larger than the cost of falsely keeping it.

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 difference specification are not robust to small deviations of the largest root from unity, even when those deviations are too small to be empirically detected.

5 Conclusion

This paper analyzes the source of the conflicting evidence on the effect of technology shocks on hours worked reported in several recent empirical studies. Chari, Kehoe and McGrattan (2008) and Ravenna (2007) point out that structural VARs may suffer from a lag truncation bias when the true model has an infinite lag order, as implied by standard real business cycle models. Likewise,

other studies have shown that very different conclusions emerge from structural VARs depending on whether hours worked is treated as a stationary (Christiano, Eichenbaum and Vigfusson 2003, 2006) or nonstationary (Gali, 1999; Francis and Ramey, 2005) variable. While we confirm that both the uncertainty regarding the persistence of hours worked and the bias due to lag truncation play 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 in the presence of a low frequency co-movement between productivity growth and hours worked. This low frequency co-movement drives a wedge between the IRF estimates of the levels and differenced VAR specifications, even when the largest root of hours worked is arbitrarily close to one. This implies that pre-testing for a unit root in hours worked may result in highly misleading IRF inference in the region in which the unit root tests have difficulties rejecting the null hypothesis.

In so much as both studies point to the importance of a low frequency co-movement, our results can also be viewed as supportive of Fernald (2007), who suggests an explanation based on a common high-low-high pattern in both productivity growth and hours per person. On the other hand, while Fernald (2007) interprets this co-movement as resulting coincidentally from a similar, but unrelated, sequence of breaks in both series, we find that it is virtually impossible to generate the observed differences in impulse response functions without the inclusion of a genuine low frequency co-movement in the true data generating process. This leads us to a substantially different interpretation of the conflicting conclusions from the differenced and levels specifications. In Fernald's (2007) framework, the levels specification is misguided, since its long-run identification scheme incorrectly relies on misleading information from a spurious low-frequency co-movement. By contrast, in our framework, it is the difference specification which is biased, since its long-run identification scheme incorrectly ignores the genuine information contained in a true low-frequency co-movement.

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 when the root in hours worked is close or equal to one. Imposing additional restrictions on the model (see, for example, Gospodinov, 2008) can lead to improved inference for the structural parameters and impulse responses. More generally, our results underline and help to explain the potential sensitivity of long-run identifying schemes to uncertainty regarding low frequency dynamics, even when identifying business cycle frequency characteristics.

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., Kimball, M.S., 2006. Are technology improvements contractionary? *American Economic Review* 95, 1418-1448.
- [3] Blanchard, O., 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., Hong, J.H., 2006. Do technological improvements in the manufacturing sector raise or lower employment? *American Economic Review* 96, 352-368.
- [6] Chang, Y., Hornstein, A., Sarte, P.-D., 2008. On the employment effects of productivity shocks: The role of inventories, demand elasticity, and sticky prices. *Mimeo*, University of Rochester.
- [7] Chari, V., Kehoe, P., McGrattan, E., 2008. Are structural VARs with long-run restrictions useful in developing business cycle theory? *Journal of Monetary Economics* 55, 1337-1352.
- [8] Christiano, L., Eichenbaum, M., Vigfusson, R., 2003. What happens after a technology shock? *Board of Governors of the Federal Reserve System, International Finance Discussion Paper* 768.
- [9] Christiano, L., Eichenbaum, M., Vigfusson, R., 2006. Assessing structural VARs. In: D. Acemoglu, K. Rogoff and M. Woodford, eds., *NBER Macroeconomics Annual* Cambridge, MA: MIT Press.
- [10] Fernald, J., 2007. Trend breaks, long-run restrictions, and contractionary technology improvements. *Journal of Monetary Economics* 54, 2467-2485.
- [11] Francis, N., Owyang, M., Roush, J., 2005. A flexible finite-horizon identification of technology shocks. *Federal Reserve Bank of St. Louis, Working Paper* 2005-024E.
- [12] Francis, N., Owyang, M., Theodorou, A., 2003. The use of long-run restrictions for the identification of technology shocks. *Federal Reserve Bank of St. Louis Review* 85(6), 53-66.
- [13] Francis, N., 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., Ramey, V., 2006. A new measure of hours per capita with implications for the technology-hours debate. *Mimeo*, University of North Carolina, Chapel Hill.
- [15] Galí, J., 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review* 89, 249-271.

- [16] Gali, J., Rabanal, P., 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.
- [17] Gospodinov, N., 2008. Inference in nearly nonstationary SVAR models with long-run identifying restrictions. *Journal of Business and Economic Statistics*, forthcoming.
- [18] Pesavento, E., Rossi, B., 2005. Do technology shocks drive hours up or down? A little evidence from an agnostic procedure. *Macroeconomic Dynamics* 9, 478-488.
- [19] Pesavento, E., Rossi, B., 2006. Small sample confidence intervals for multivariate impulse response functions at long horizons. *Journal of Applied Econometrics* 21, 1135-1155.
- [20] Phillips, P.C.B., 1987. Towards a unified asymptotic theory for autoregression. *Biometrika* 74, 535-547.
- [21] Ravenna, F., 2007. Vector autoregressions and reduced form representations of DSGE models. *Journal of Monetary Economics* 54, 2048-2064.
- [22] Shea, J., 1999. What do technology shocks do? In: B. Bernanke and J. Rotemberg, eds., *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press.
- [23] 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. Detrended labour productivity and demeaned hours worked, U.S. data 1948Q2 - 2005Q3.




FIGURE 3. HP trends of labour productivity growth (top graph) and hours worked (bottom graph), U.S. data 1948Q2 - 2005Q3.




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

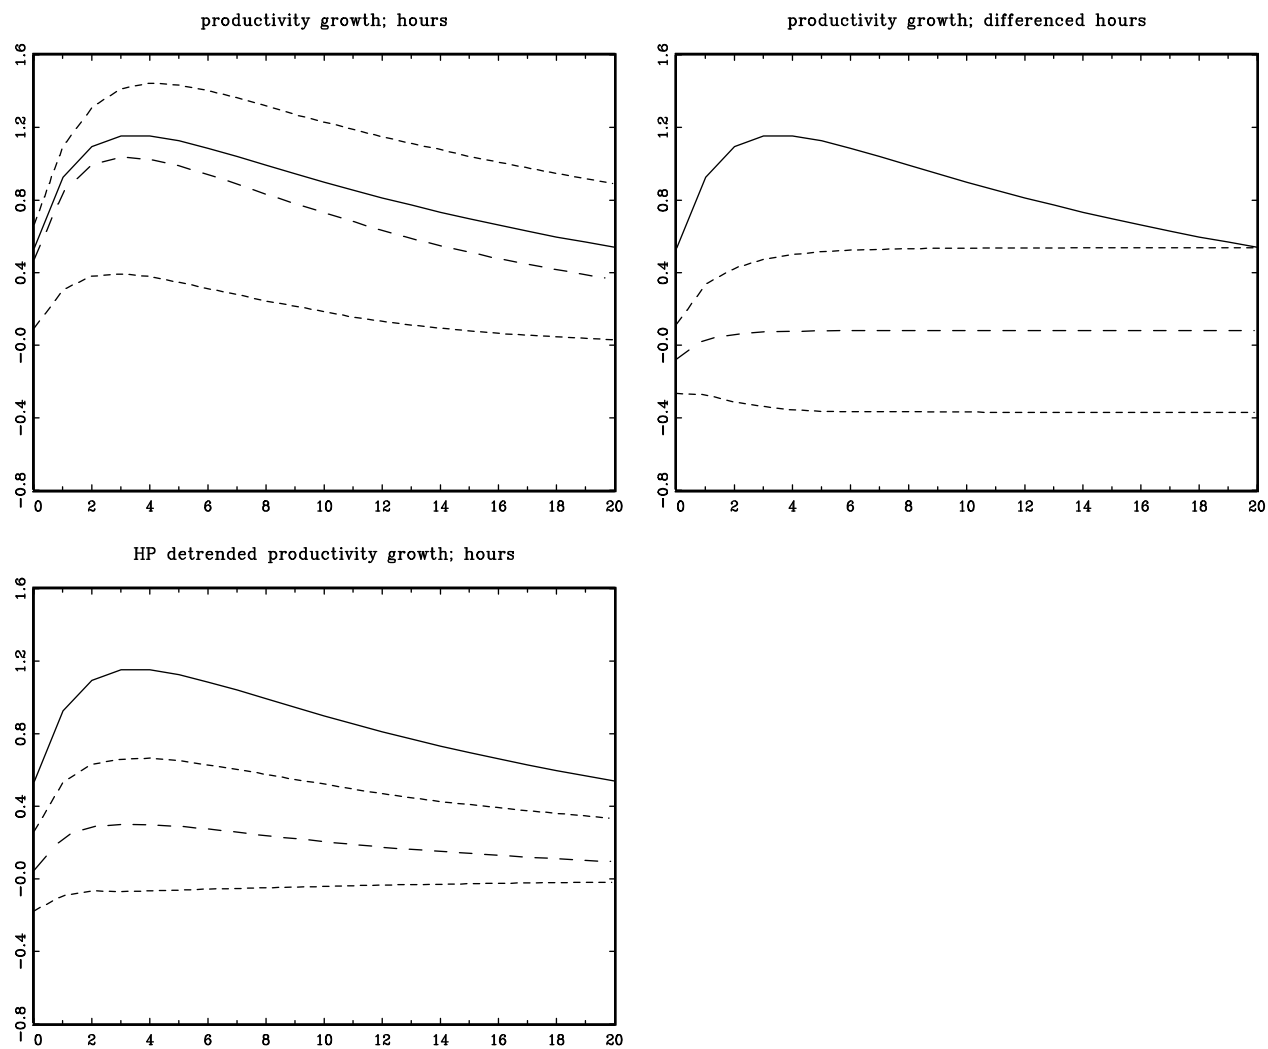


FIGURE 5. $\rho = 0.95, \gamma = -0.8, \delta = 0.04$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

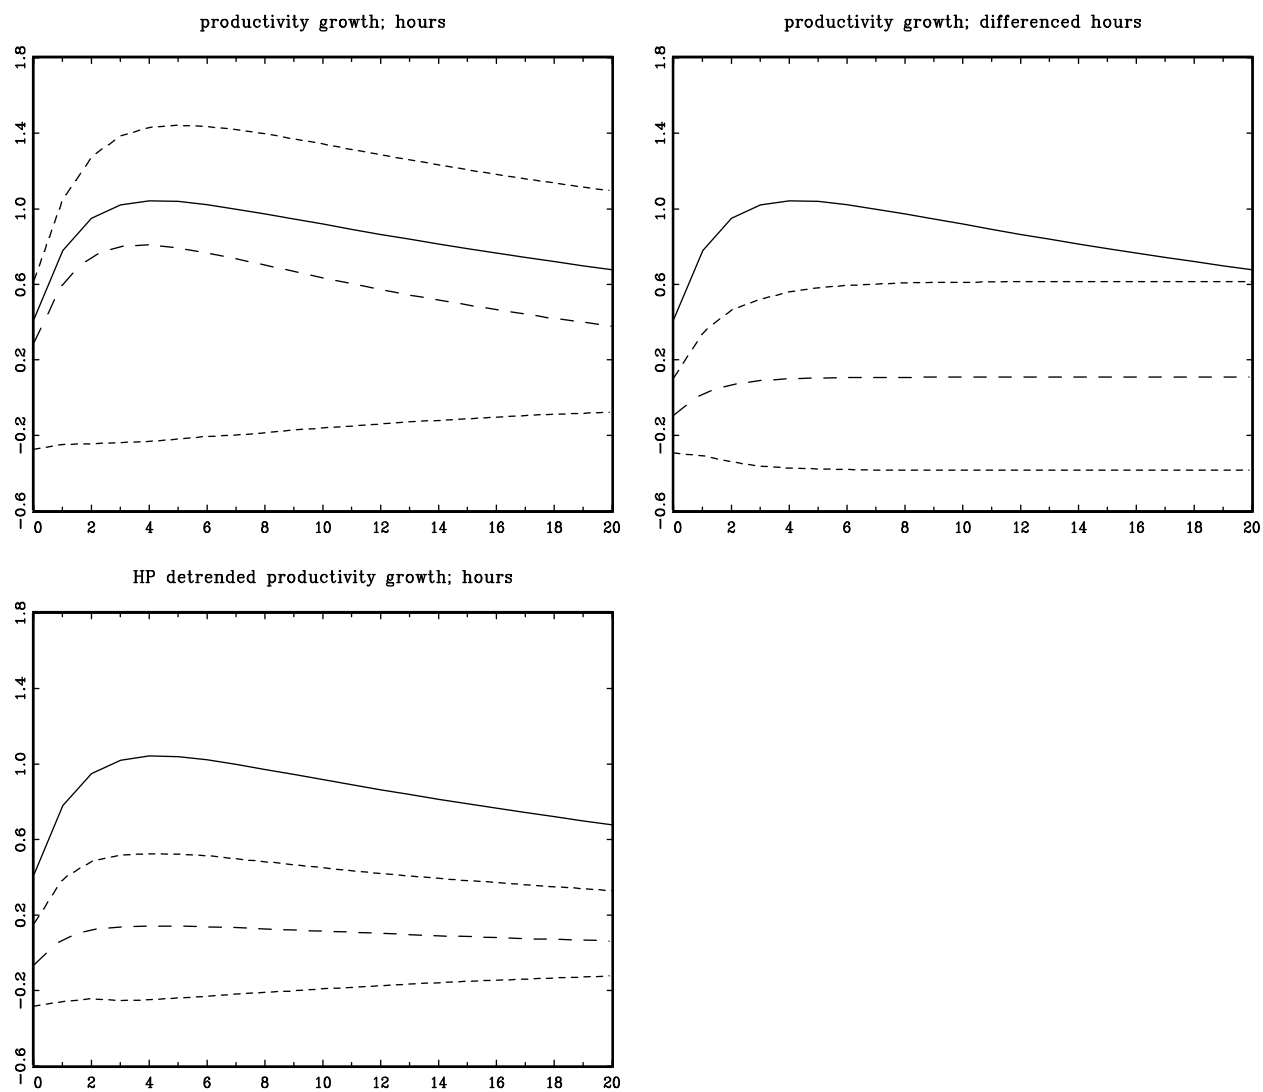


FIGURE 6. $\rho = 0.97, \gamma = -0.5, \delta = 0.015$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

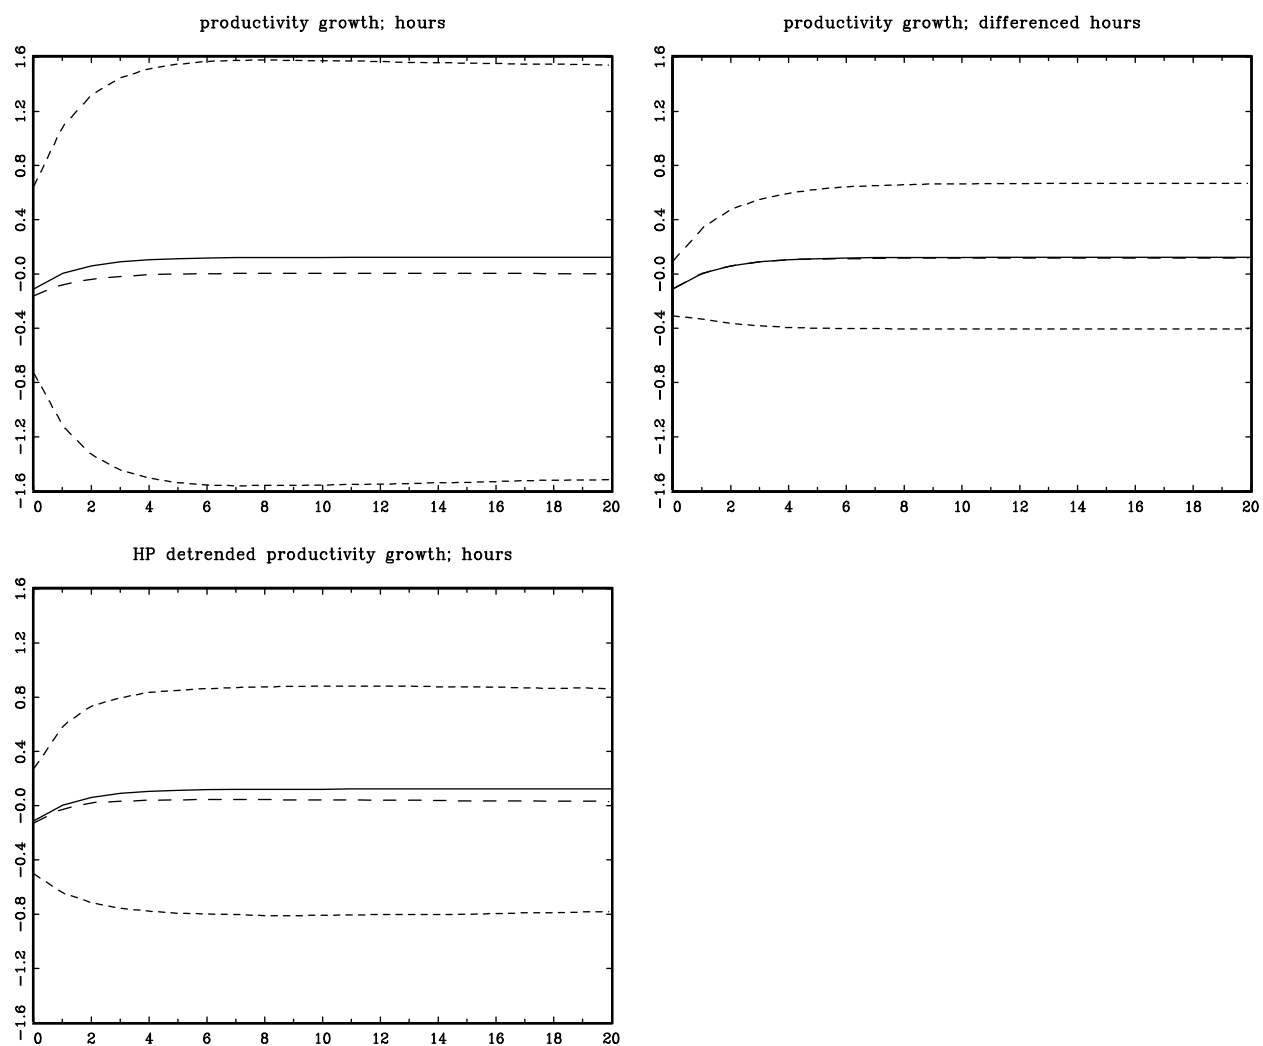


FIGURE 7. $\rho = 1, \delta = 0$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.

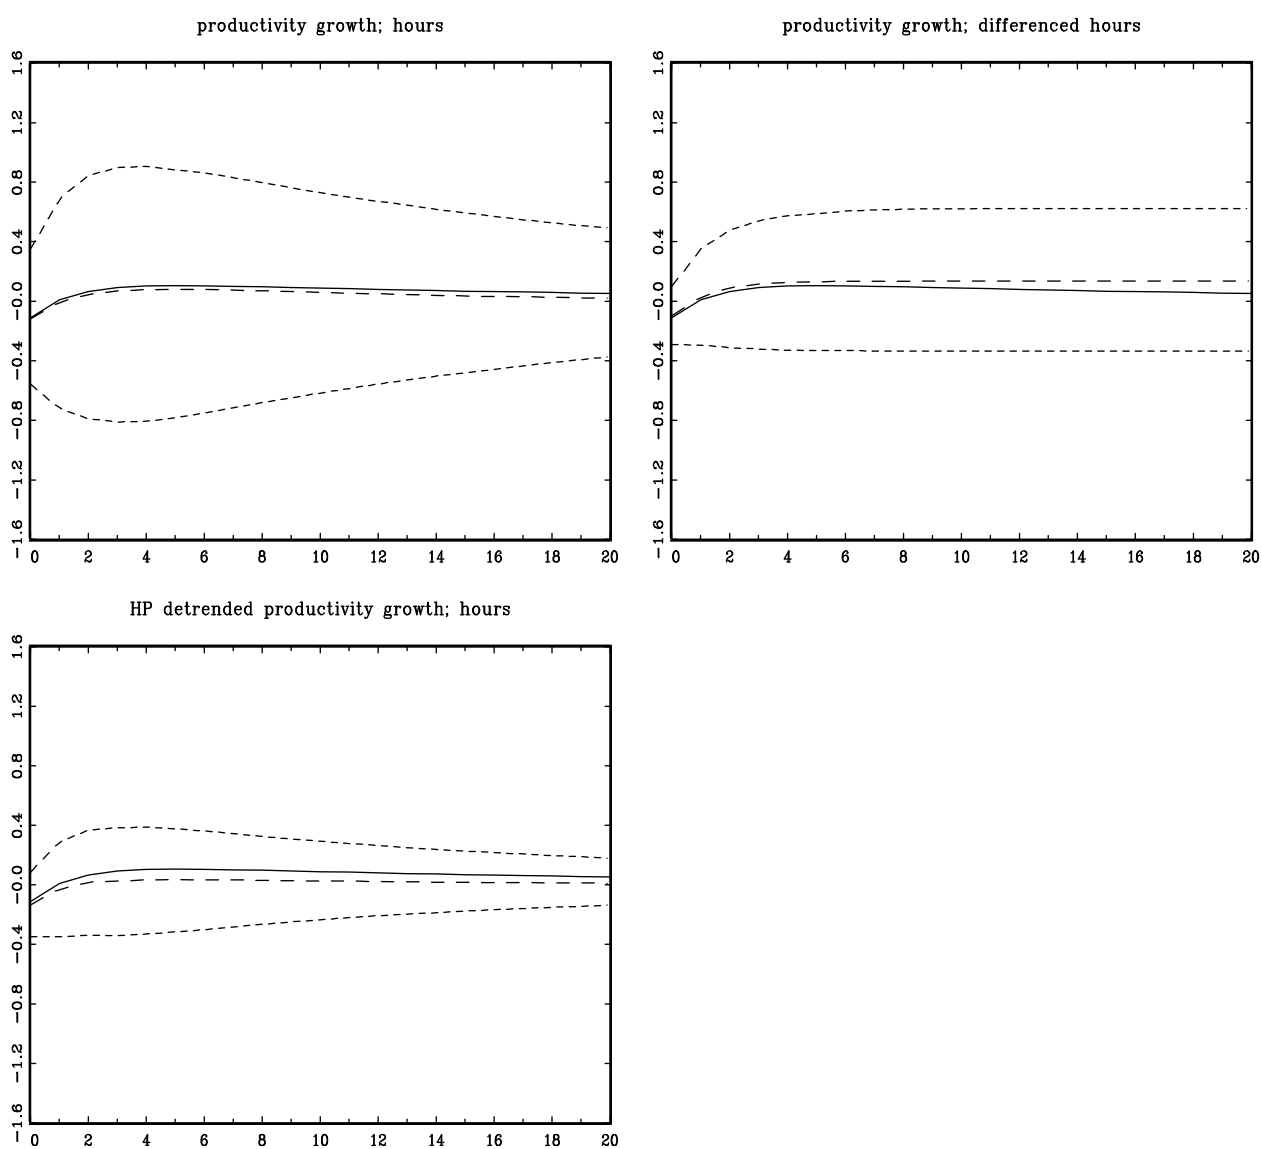


FIGURE 8. $\rho = 0.95, \gamma = 0, \delta = 0$. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.




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.

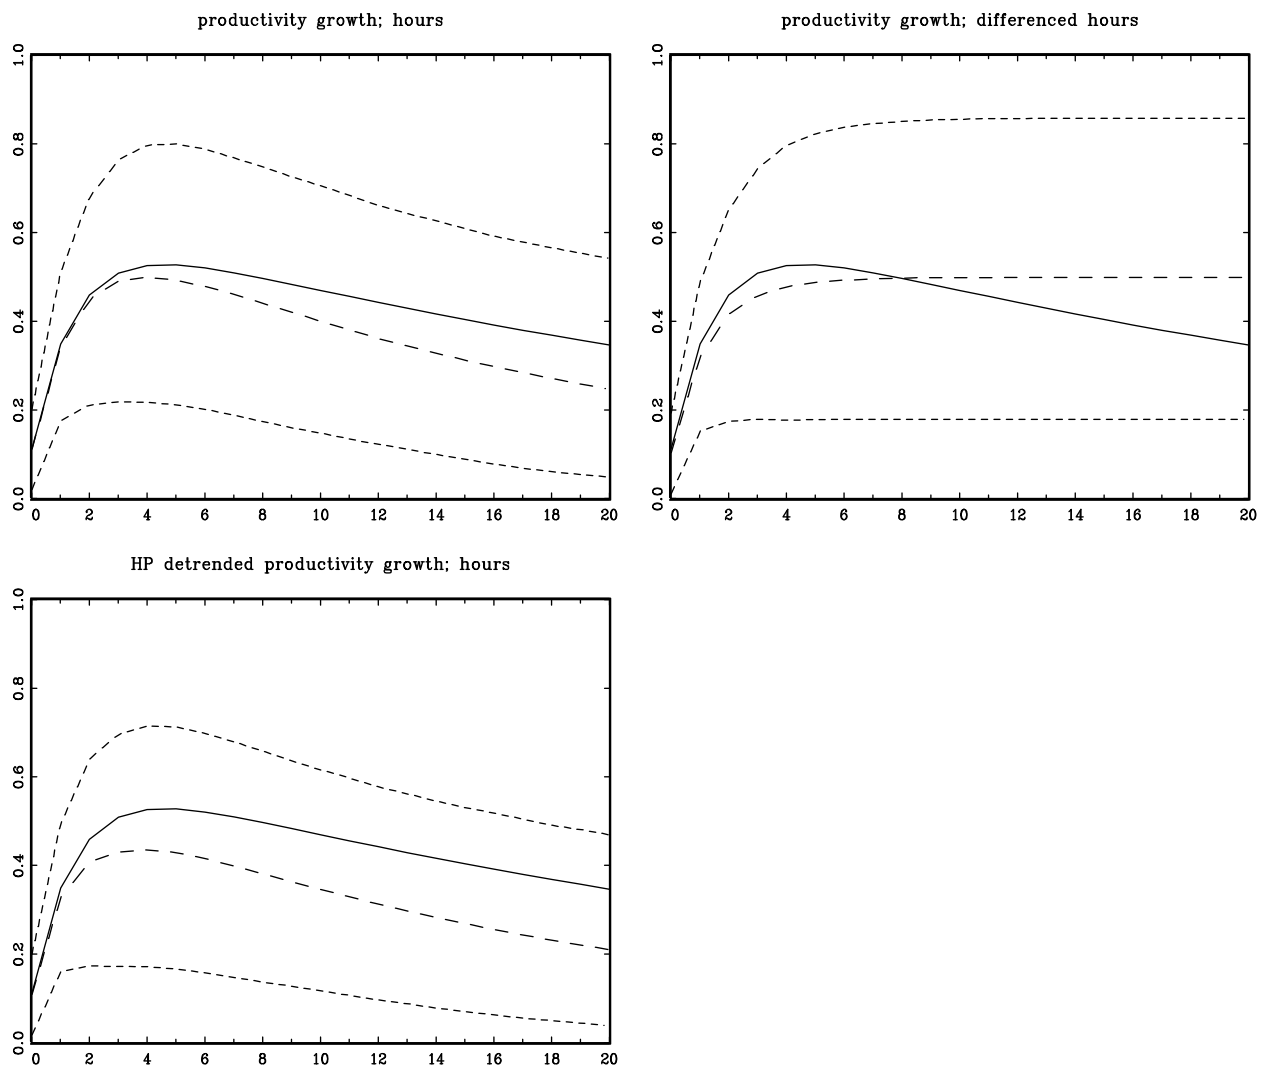


FIGURE 11. Short-run (Cholesky) identification scheme with $\rho = 0.97, \gamma = -0.5, \delta = 0.015$ and covariance between the shocks of 0.1. Solid line: true IRF; long dashes: median Monte Carlo IRF estimate; short dashes: 95% Monte Carlo confidence bands.




FIGURE 12. Detrended labour productivity and demeaned hours worked; simulated data from the CKM specification of the RBC model (Christiano, Eichenbaum and Vigfusson, 2006).

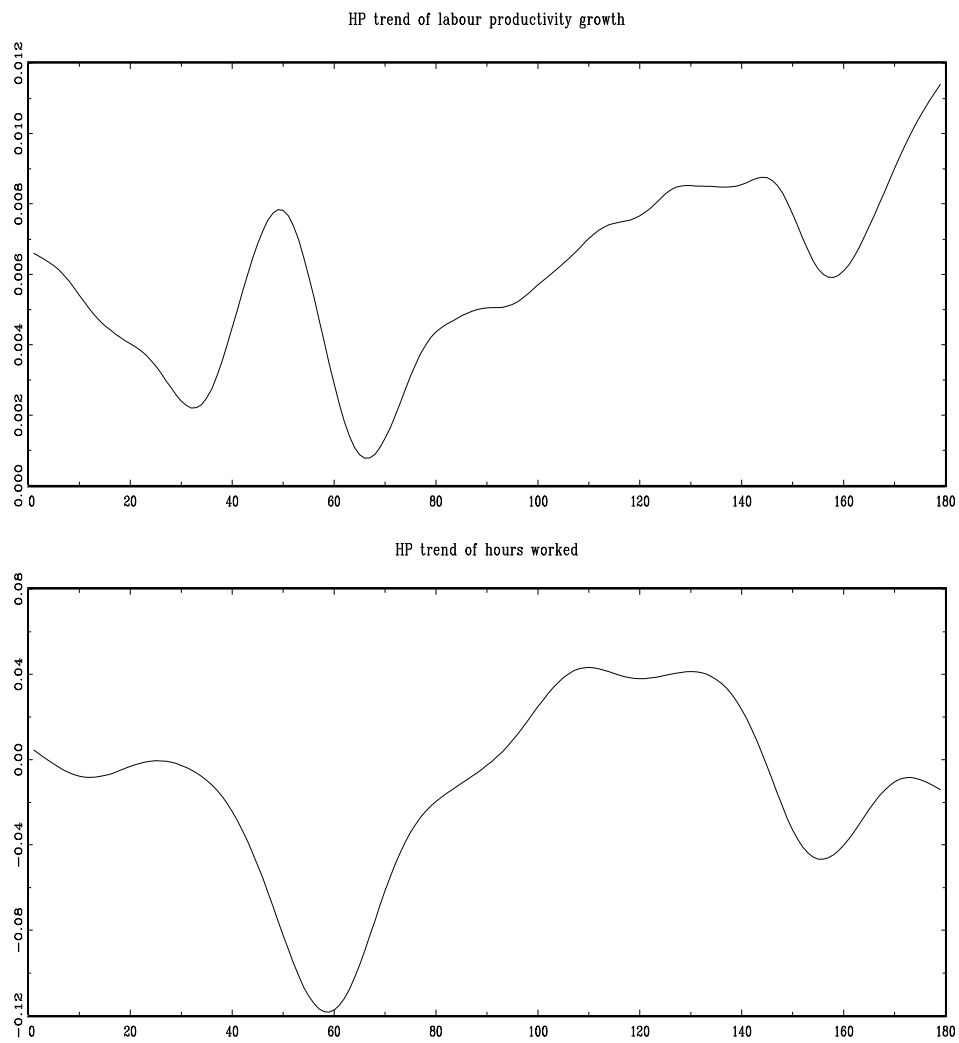


FIGURE 13. HP trends of labour productivity growth (top graph) and hours worked (bottom graph), simulated data from the CKM specification of the RBC model (Christiano, Eichenbaum and Vigfusson, 2006).




FIGURE 14. Monte Carlo IRF estimates and 95% confidence bands from the levels (top graph), differenced (middle graph) and pre-test (bottom graph) VAR specifications on simulated data (1,000 samples of length 180) from the CKM specification of the RBC model (Christiano, Eichenbaum and Vigfusson, 2006).