Broken trend stationarity of hours worked

Marcos Sanso-Navarro[*]
Universidad de Zaragoza

January 2010

Abstract

The estimated impact of a technology shock on hours worked using structural vector autoregressions depends to a great extent on whether or not hours worked is considered to be integrated of first order. It is shown in this paper that the widely analyzed time series of hours worked per capita in the U.S. business sector evolves around a broken linear trend. When this fact is taken into account, the unit root null is rejected by recently proposed tests. Therefore, it can be stated that empirical specifications with hours in first differences are not recommended. It seems more appropriate to control for the presence of this shift in the deterministic component. We also draw this conclusion from a bivariate model for both productivity growth and hours worked. Our results suggest that technology improvements have a negative but non-significant effect on hours only in the very short run. This impact later becomes positive and statistically significant after five periods.

JEL codes: C12, C22, E00.

Keywords: Structural Change, Unit Roots, Technology shocks, SVAR.

1 Introduction

A burgeoning literature in recent years has tried to disentangle the effects that technology improvements have on hours worked. It was motivated by Galí (1999) who, using Structural Vector Autoregressions (SVARs) with long-run restrictions, found that a positive technology shock has a negative impact on hours. This result was interpreted as contrary to standard Real Business Cycle (RBC) models, which assign technological shocks a central role in aggregate fluctuations. More specifically, a key prediction of these models is a positive co-movement of output, labour input and productivity in response to technology improvements. The negative response of hours worked was considered to be more in accordance with the sticky prices New Keynesian (NK) models. Using a very similar approach, this result was challenged by Christiano et al. (2003) whose estimated effect of technology improvements on hours worked was positive.

[*]Address: Departamento de Análisis Económico. Facultad de Ciencias Económicas y Empresariales. Gran Vía 2. 50005 Zaragoza (Spain). E-mail: marcossn@unizar.es. Tel: (34) 976 761 000 Ext. 4629. Fax: (34) 976 761 996. Homepage: http://dae.unizar.es/marcossn/index.html.

The main difference between these two studies is how hours worked was specified in the empirical model. While Galí (1999) considered that hours is a non-stationary variable and introduced it in first differences, Christiano et al. (2003) argued that it is more reasonable to think that it is stationary and worked with it in levels.

The cause behind this uncertainty about the most appropriate specification for hours worked is the tests applied are not able to reject the null hypothesis of a unit root. In addition, although there is no *a priori* reason to think that the hours worked per capita time series is non-stationary, most studies work with this variable in first differences. This may be a consequence of the little effort made to empirically corroborate this property. Whelan (2009) has recently advocated for the non-stationarity of hours worked per capita in the light of results obtained from the application of univariate unit root and stationarity tests. One reasonable explanation for concluding that hours is an I(1) variable might be the low power of conventional univariate unit root tests. Kappler (2009) has tried to overcome this limitation by the use of panel methods. However, his results do not find evidence against the unit root null. As an alternative, Gil-Alana and Moreno (2009) have used a multivariate fractional integration approach obtaining more evidence in favour of the stationarity of hours worked.

The argument put forward in this paper is that the lack of evidence against the non-stationarity of hours worked derives from the fact that the unit root tests applied to date have neglected the presence of a shift in the deterministic component of this variable, leading to a reduction in their power (Perron, 1989). The presence of low frequency movements in the time series of hours worked per capita in the U.S. has already been discussed because there was an interest in explaining the different evolution of this series after the mid-70s in comparison with those in the European countries. A summary of this literature can be found in Shields and Shields (2008). Of all the determinants these authors looked at, they conclude that the decrease experienced by U.S. hours worked per capita until the beginning of the 80s was mainly a consequence of an increase in the marginal tax rates that raised the price of leisure for families. The explanation given for the later upward trend focuses on the rising health care costs that were passed from the employers onto their workers. In a study more in line with the present one, Francis and Ramey (2009) argue that the U-shape followed by U.S. hours worked per capita is determined by the movement of the baby boom generation through the labour market and changes related to the public and non-profit sectors.

This paper does not try to throw further light on the determinants of the evolution followed by hours worked per capita in the U.S. We only intend to establish the statistical validity of a changing trend for this variable and analyze its influence on unit root testing and subsequent empirical modelling. As noted by Perron (2006), structural breaks complicate more than just the testing for unit roots. In addition, testing for shifts in the deterministic component of a given time series depends on its order of integration. On the one hand, if the variable is considered to be difference-stationary when it is, in fact, stationary in levels, structural break tests will suffer from low power. On the other, structural break test statistics applied to the time series in levels will have different limiting distributions depending on its order of integration. This circular problem between tests regarding the parameters of the trend function and unit root tests has motivated the appearance of procedures to test for the stability of the trend function of a univariate time series that are robust to whether the errors are I(0) or I(1).

The suitability of a broken linear trend for the widely analyzed time series of hours worked per capita in the U.S. business sector is going to be tested for using the statistic proposed in Perron and Yabu (2009). Once the presence of a broken linear trend is established, the unit root null hypothesis will be assessed under this specification of the deterministic component using the tests developed in Carrión-i-Silvestre et al. (2009). Compared to conventional univariate unit root tests, the latter have the virtue of considering the presence of shifts at unknown dates both under the null and the alternative hypotheses. This gives the tests the correct size and higher power. Our results give evidence that hours per capita evolves around a broken linear trend in a stationary manner. Therefore, it is concluded that introducing this variable in first differences into SVARs does not make sense and can lead to biases (Christiano et al., 2003; Erceg et al., 2005).

In line with the suggestions in Ng and Vogelsang (2002) and Lütkepohl and Krätzig (2004), a further contribution of this paper is to analyze the effects that a technology shock has on hours worked by simultaneously controlling for the presence of shifts in the deterministic components of the two variables involved in the bivariate SVAR commonly used in this context. We try to complement the attempts made by previous studies that only deal with shifts in the deterministic component of one of the two time series (Fernald, 2007; and Canova et al., 2009). Before doing so, the presence of shifts in the deterministic components has been statistically checked in a two-equation framework using the test proposed by Qu and Perron (2007). The impulse response function analysis shows that a technology improvement has an immediate negative impact on hours. After that, hours worked per capita increase. These results are similar to those found in Basu et al. (2006) both using an augmented growth-accounting approach and introducing a 'purified' technology measure into the SVAR framework. The main virtue of their approach is that it does not depend on the way hours worked are specified and their findings are interpreted as consistent with Dynamic General Equilibrium (DGE) models with sticky prices. In addition, this pattern for the impulse response function of hours to a technology shock is common in recent related studies (Francis and Ramey, 2009; Fève and Guay, 2009).

The rest of the paper is structured as follows. After describing the data source and the variables analyzed in Section 2, the presence of a broken linear trend in the U.S. hours worked per capita time series is established in Section 3. In addition, this specification for the deterministic component is compared with the alternatives already used in the related literature and the unit root null hypothesis is tested. Section 4 estimates the effect that a positive technology shock has on hours worked controlling for the presence of deterministic shifts in both productivity growth and hours worked. Finally, Section 5 concludes.

2 Data source

The variables analyzed in this paper are an updated version of those in Christiano et al. (2003) that have later been used in a substantial number of studies. They have been constructed with data extracted from the Federal Reserve Bank of St. Louis (FRED®) which, in turn, are compiled from the Bureau of Labor Statistics (U.S. Department of Labor). The data has a quarterly frequency and covers the period from 1948:Q1 to 2007:Q4.

Hours worked refers to the business sector of the economy (HOABS). It is expressed as an index (1992=100) and has previously been seasonally adjusted. This variable has been converted into per capita terms by dividing it by the civilian non-institutional population over 16 (CNP16OV). The latter has been changed to a quarterly frequency by averaging the monthly observations and then transformed into an index with the same base year as the variable hours worked. The Household Survey used to construct the population time series is subject to revisions related to population controls which affect data back to 2000[1] causing a break at this date. Because we are interested in changes in the deterministic component of the hours per capita time series constructed from this population measure, it is convenient to avoid this shift induced by a methodological change. Instead of to the smoothed version of the population time series used by Fernald (2007), we have opted for correcting the estimated magnitude of the level shift automatically detected by the TRAMO/SEATS software in the first quarter of the year 2000.

As has already been noted, we will also study to what extent the properties of the deterministic component of hours worked per capita affect the results obtained from the analysis of the effects that a technology shock has on this variable. For this reason, the output per hour worked in the U.S. business sector (OPHPBS) will also be used below. As in the case of hours worked, this variable is reported in terms of the year 1992 and its seasonal component has been removed. Finally, note that the analysis will be carried out with the natural logarithm of hours worked per capita and the growth rate of output per hour (productivity, hereafter) in percentage terms.

3 Unit root and shift in trend testing

The presence of a unit root in hours worked per capita will first be analyzed by means of standard univariate unit root tests. The alternatives implemented are those discussed in Ng and Perron (2001) because they have good size and power. One way to achieve these is by the application of a local-to-unity Generalized Least Squares (GLS) detrending method (Elliot et al., 1996) before testing the unit root null hypothesis.

Let us consider that the data generating process of the variable under scrutiny (y_t) can be expressed as the sum of a deterministic (d_t) and an autoregressive (u_t) component:

$$\begin{aligned} y_t &= d_t + u_t \\ u_t &= \alpha u_{t-1} + v_t \end{aligned} \tag{1}$$

for $t = 1, ..., T$ where it is assumed that $E(u_0^2) < \infty$. The disturbance term is defined as $v_t = \theta(L)e_t = \sum_{i=0}^{\infty} \theta_i e_{t-i}$, with $\sum_{i=0}^{\infty} i\,|\theta_i| < \infty$ and $\{e_t\} \sim iid(0, \sigma_e^2)$.

The deterministic component in (1) is given by $d_t = \psi z_t$. As is common practice for the hours per capita time series, it is considered to be made up of both a constant and a trend. That is, $z_t = (1, t)'$. Therefore, $\psi = (\mu, \beta)$ is a vector of unknown parameters. Considering that $-1 < \alpha \leq 1$, the null hypothesis of interest is $\alpha = 1$ against the alternative $\alpha < 1$.

[1] This information can be found in http://research.stlouisfed.org//fred2/series/CNP16OV.

For a given time series $\{x_t\}_{t=0}^{T}$, let $x^{\bar{\alpha}} = (x_0^{\bar{\alpha}}, x_t^{\bar{\alpha}}) = (x_0, (1 - \bar{\alpha}L)x_t)$, $t = 1, ..., T$ for $\bar{\alpha} = 1 + \frac{\bar{c}}{T}$. $\bar{c}$ is known as the non-centrality parameter which, given the specification chosen for the deterministic component, has been set to -13.5. The proposal of Ng and Perron (2001) is to test for a unit root in the GLS detrended time series for y_t ($\tilde{y}_t$), which is given by:

$$\tilde{y}_t = y_t - \hat{\psi} z_t \tag{2}$$

where $\hat{\psi}$ is obtained from the OLS regression of $y^{\bar{\alpha}}$ on $z^{\bar{\alpha}}$.

Ng and Perron (2001) also proposed a modification to the information criteria of reference to choose the truncation lag used to augment the auxiliary regression or to estimate the spectral density estimator at frequency zero. The latter allows us to control for the presence of autocorrelation that induces size distortions in unit root tests. This modification consists of imposing the null hypothesis and including a stochastic term in the penalty factor. Following the suggestion in Perron and Qu (2007), OLS detrended data have been considered when calculating the modified Akaike criterion. The maximum number of lags allowed has been set to $int(12 \cdot (\frac{T}{100})^{\frac{1}{4}})$.

[Insert Table 1 here]

Our results from the application of the standard univariate GLS detrending-based unit root tests to the hours worked per capita time series are presented in the second column of Table 1. Consistent with the findings in previous studies, the null hypothesis of a unit root cannot be rejected at conventional significance levels by any of the applied tests. As has been the custom in the related literature, this would lead us to think of introducing this variable in first differences in multivariate systems that do not allow for the presence of cointegration relationships.

It is important to adequately model the deterministic components when dealing with time series in order to obtain consistent estimates and good forecasts. As it is widely known that the existence of structural breaks in the trend function is a problem of long-horizon data, the evolution of hours worked per capita during the sample period has been plotted in Figure 1. It can be observed that this time series evolves around a more complicated trend than a simple linear one. As has previously been reported in the literature, it follows a decreasing trend until the mid-70s and an increasing path after the early 80s. This evolution can be considered to be a reasonable explanation for the inability of the standard unit root tests to reject the unit root null because they have low power in the presence of changes in the deterministic components (Perron, 1989).

[Insert Figure 1 here]

As noted by Perron (2006), structural breaks complicate more than just the testing for unit roots. In addition, testing for shifts in the deterministic component depends on the order of integration of the variable analyzed. On the one hand, if the variable is considered to be difference-stationary when it is, in fact, stationary in levels, structural break tests will suffer from low power. On the other, structural break test statistics applied to the time series in levels will have different limiting distributions depending on its order of integration. This circular problem between tests regarding the parameters of the trend function and unit root tests has motivated the appearance of procedures to test for the stability of the trend function of a univariate time series that are robust to whether the errors are I(0) or I(1). Given the virtues of these procedures, they are intended to be used before studying the integration order of a given time series. This is the case of the methods proposed by Vogelsang (2001) and Harvey et al. (2009). They suggest constructing weighted averages of test statistics using scaled unit root and stationarity tests, respectively, in such a way that the asymptotic critical values are the same for the stationary and non-stationary cases. The test that has been applied in this paper is that proposed by Perron and Yabu (2009). Contrary to those mentioned before, it does not involve any random scaling and this test is the best choice in both the I(0) and I(1) cases, although it is not necessarily true for the local-to-unity case. Furthermore, the finite sample properties of the unit root tests in the presence of a broken trend that will be described and implemented below are known provided that this procedure has been applied as a pre-test.

The specification for the shift experienced by the deterministic component of the hours worked per capita time series is that corresponding to the most general case. The latter allows for both a structural change in the intercept and the slope parameters. According to (1), it is going to be specified that:

$$\begin{aligned} z_t^1 &= (1, DU_t, t, DT_t)' \\ \psi^1 &= (\mu_0, \mu_1, \beta_0, \beta_1) \end{aligned} \tag{3}$$

where

$$\begin{aligned} DU_t &= 1(t > T_b) \\ DT_t &= 1(t > T_b)(t - T_b) \end{aligned} \tag{4}$$

T_b is the break date and $1(\cdot)$ is the indicator function. Thus, DU_t and DT_t imply a change in the intercept and the slope parameters, respectively, after the break date. That is, they are a level and a trend shift dummy. The null hypothesis of interest is that of no trend shift ($H_0 : \mu_1 = \beta_1 = 0$).

The test proposed by Perron and Yabu (2009) (Exp-W_{FS}, hereafter) is based on a Feasible GLS procedure that uses a super-efficient estimate of the autoregressive parameter α that relies on a truncation when it is equal to one. In addition, the finite sample properties of this method are improved using a bias correction. For an unknown break, the Exp functional of Wald tests for all possible break dates (Andrews and Ploberger, 1994) has very similar limiting distributions in the I(0) and I(1) cases. For this reason, it is suggested that the largest critical value should be taken in order to work with a robust statistic. Nevertheless, it should be noted that this makes the Exp-W_{FS} test conservative, especially when the sum of the autoregressive parameters is close to but not equal to one.

The resulting value of the Exp-W_{FS} test statistic from its application to the hours worked per capita in the U.S. business sector time series is reported in the lower panel of Table 1. As is commonly done, a trimming of 15% both at the beginning and the end of the sample have been considered. The obtained test statistic is 2.59 so, according to the asymptotic critical values reported in Perron and Yabu (2009) for the trimming parameter and trend shift model specified, the null hypothesis of no broken trend can be rejected at the 10% significance level. More specifically, the highest p-value that should be considered corresponds to the I(1) case and is equal to 0.09. If hours worked per capita were a stationary time series, the relevant p-value would be equal to 0.06. Given this evidence, it seems appropriate to take this trend shift into account when testing for the unit root null.

Kejriwal and Perron (2009) have proposed that the Exp-W_{FS} test can be implemented in a sequential manner in order to look for the presence of further trend shifts. Using the break date estimated by the minimization of the sum of squared residuals (SSR) from a regression of the hours worked time series on a constant, a time trend, a level shift dummy and a slope shift dummy (see (3) and (4)), located in 1983:Q4, it is not possible to reject the null hypothesis of a single break in favour of the presence of two breaks.

Most of the existing literature dealing with the issue of unit root testing in the presence of breaks in the deterministic component only considered them under the alternative hypothesis of stationarity. Carrión-i-Silvestre et al. (2009) have proposed an extension to the GLS detrending-based tests described at the beginning of this section to allow for multiple breaks at unknown dates both under the null and the alternative hypotheses. In line with our previous results, only the case of a single trend shift is going to be analyzed here. The non-centrality parameter and the critical values are obtained by simulations because they depend on the number and location of the break. This parameter has been set to -18.02.

Resulting unit root test statistics allowing for a trend shift both under the null and the alternative are reported in the third column of the upper panel of Table 1. It can be observed that the null hypothesis of hours per capita evolving around a broken linear trend in a non-stationary manner can be rejected at the 5% significance level by all the tests applied. Moreover, the p-value is lower than 0.01 in three of the seven alternatives. Consequently, it can be stated that the widely used time series of hours worked per capita in the U.S. business sector can be considered to be broken trend stationary. Therefore, our evidence does not justify the first-differenced specification of this variable that is commonly found in the literature.

[**Insert Table 2 here**]

Galí (1999) was the first to suggest controlling for the presence of a deterministic component in the hours worked time series in terms of a simple linear trend. Taking this idea as their starting point, later studies have used a quadratic trend (Christiano et al., 2003; Galí and Rabanal, 2005) or two level shifts (Canova et al., 2009). These alternatives have been compared with the broken linear trend proposed here using a simple adjustment and diagnostic statistics. The latter are reported in Table 2. Reinforcing the previous analysis, the best fit is obtained for the broken linear trend specification, whose coefficient of determination and log-likelihood are the highest. This specification also leads to the lowest standard errors, SRR and Akaike information criterion.

The second best option is that of a quadratic trend while the worst fit corresponds to the simple linear trend case.

Galí and Rabanal (2005) were not able to obtain evidence against the unit root null with a quadratic trend specification and an OLS detrending approach. Ayat and Burridge (2000) and Harvey et al. (2008) have proposed applying the Augmented Dickey-Fuller (ADF) test statistic to GLS detrended data and this specification of the deterministic component[2]. The test statistic obtained following this approach is -3.52 which, compared with the finite sample critical values reported in the first of these studies[3], allows the rejection of the unit root null at the 5% significance level. Therefore, it can be concluded that we have been able to find additional evidence against the non-stationarity of hours worked with a quadratic specification of the trend, although it is slightly weaker to when the deterministic component is a broken linear one.

4 Effects of productivity shocks on hours

In the light of the evidence presented in the previous section, it can be stated that working with hours per capita as if they were difference stationary might not be correct and can lead to biased estimations when trying to determine the effects of technological improvements on hours worked (Christiano et al., 2003; Erceg et al., 2005). The alternative proposed here in order to carry out this analysis consists of explicitly specifying the level and trend shift in hours in its corresponding equation in the bivariate system rather than trying to achieve stationarity by first-differencing this variable.

As noted in Ng and Vogelsang (2002), although there is evidence of mean and/or trend shifts in many macroeconomic time series, vector autoregression (VAR) estimations usually ignore them. However, it seems more appropriate to introduce breaks into multivariate analyses when they are observed at a univariate level. These authors also demonstrated that inference based on the estimated VAR is invalid when mean shifts are omitted. The importance of correctly specifying the deterministic components in VARs has also been emphasized by Lütkepohl and Krätzig (2004).

As has already been mentioned, controlling for the presence of the deterministic component of hours worked when analyzing the effect of a technological shock on this variable is not new. Furthermore, Fernald (2007) took into account two level shifts in U.S. business sector productivity growth in 1973:Q2 and 1997:Q2 in the same context[4]. Canova et al. (2009) later checked the robustness of the estimated effects to several detrending alternatives for both productivity

[2] The non-centrality parameter $\bar{c}$ has been set to -18.5.

[3] Appendix B, page 95.

[4] This author suggested that these changes implied a high-low-high pattern in the productivity growth time series. This evolution is supposed to be related to unusual historical events with persistent and non-permanent effects like steam power, electricity, the interstate highway system or information technologies. Another interpretation, given by Francis and Ramey (2009), is that they reflect the entry of the baby boom generation into the labour market.

growth and hours worked. Nevertheless, to date, no study has exhaustively modelled the presence of these changing deterministic components in both productivity growth and hours worked at the same time.

Before estimating the VAR model to analyze the effects that a technology shock has on hours, the empirical importance of the deterministic shifts in the bivariate system is going to be established. This will be done by the application of the procedure of Qu and Perron (2007) that allows us to test for the presence of structural changes with restrictions on the parameters in linear multivariate regression models. Following the findings in Fernald (2007), two level shifts in productivity growth and, in line with previous results, a trend shift in hours worked have been specified. The coefficients related to the autoregressive lags are considered to be the same throughout the sample period while the possibility of heteroskedasticity across subsamples has been allowed. The estimation is carried out using a quasi-maximum likelihood approach assuming serially uncorrelated normal errors. The null hypothesis is that of no structural change and the test statistic is the maximal value of the likelihood ratio over all possible sample partitions (sup LR).

[**Insert Table 3 here**]

Table 3 reports the value of the sup LR test statistic for four different VAR orders, the estimated break dates using a dynamic algorithm and their corresponding 90% confidence intervals. It can be observed in the second column of this table that the null hypothesis of no structural change can be rejected at the 1% significance level regardless of the number of autoregressive lags considered. Once the presence of the deterministic shifts is taken into account, the Akaike information criterion leads us to set the autoregressive order to 3. In this case, the estimated break dates for the two level shifts are located in 1973:Q1 and 1995:Q3, very similar to those found in Fernald (2007). More interestingly, the estimated break date for the trend shift in hours is the same as to that obtained with the univariate approach of the previous section. Given the statistical significance of these deterministic shifts, they have been included in their corresponding equation of the bivariate VAR.

VARs are known to be 'reduced form' models because they basically summarize the dynamic properties of the data. However, our interest lies in disentangling the effects that a shock in one of the variables has on the rest of the system. These can be determined by the use of structural VARs where identification focuses on the errors. These errors are interpreted as a linear combination of exogenous shocks. Under the assumption of orthogonality, it is possible to analyze the dynamic impact of isolated impulses. It should be emphasized once again that, although the deterministic terms are not affected by these shocks, it is necessary to adjust for their presence before analyzing the dynamic interactions between the variables involved.

Because our two variables of interest are stationary, they can be expressed as a distributed lag of two types of shocks:

$$\begin{bmatrix} \Delta p_t \\ n_t \end{bmatrix} = \begin{bmatrix} B_{11}(L) & B_{12}(L) \\ B_{21}(L) & B_{22}(L) \end{bmatrix} \begin{bmatrix} \varepsilon_t^p \\ \varepsilon_t^n \end{bmatrix} \tag{5}$$

Δp_t is productivity growth (in percentage terms) and n_t is hours worked per capita (in natural logarithms) ε_t^p and ε_t^n are the technology and non-technology shocks, respectively.

The main issue when dealing with SVARs is the identification of the shocks. Following Galí (1999), this has been achieved by imposing long-run restrictions *à la* Blanchard and Quah (1989)[5]. Specifically, it is assumed that non-technology shocks do not affect productivity growth. This identification restriction implies in (5) that $B_{12}(L) = 0$. That is to say, the matrix of long-horizon multipliers is lower-triangular.

The dynamic effects of structural shocks have been investigated by impulse response functions. They are plotted in Figure 2 for all possible combinations of shocks and affected variables. Hall (1992)'s equal-tailed percentile 90% bootstrap confidence bands, calculated using 2000 replications, are also reported.

[Insert Figure 2 here]

The estimated effects of non-technology shocks on productivity and hours are similar to those in Galí (1999). This is true for both the first-differenced and the detrended specifications of hours worked considered by this author. A positive non-technological shock increases productivity at first but this effect becomes negative after three periods and disappears with time (Figure 2b). The influence of this kind of shock on hours is always positive and the estimated impulse response function is hump-shaped (Figure 2d). The impact of a positive technology shock on productivity is also similar to those already established in the literature because technology improvements always increase productivity in the subsequent periods (Figure 2a).

However, our findings with respect to the effect of a positive technology shock on hours worked[6] (Figure 2c) suggest that, after an immediate negative impact, the impulse response function becomes positive after two periods. It can be observed that these estimated effects are not significantly different from zero until the fifth period after the shock takes place. These results resemble those in Basu et al. (2006) who, using an augmented growth-accounting framework found a very similar pattern for the response of hours to technology improvements. It should be emphasized that their conclusions did not depend on the way hours were specified. Moreover, the same effects were obtained when using a 'purified' technology series in a SVAR with long-horizon restrictions. These authors interpreted their findings as consistent with DGE models with sticky prices.

Also using a SVAR framework, Francis and Ramey (2009) found a similar response of hours worked to a positive technological shock using new measures of hours and productivity that try

[5] Although this approach has been criticized by Chari et al. (2008), its usefulness in identifying technological shocks and relating the empirical evidence to the RBC and NK paradigms has been corroborated by Erceg et al. (2005) and Francis and Ramey (2005).

[6] The same analysis has been carried out for the first-differenced and level specifications of hours and neglecting the presence of deterministic shifts with our updated sample period. As could be expected, the first-differenced specification leads us to conclude that hours decrease after a positive technology shock while the level specification suggests that the response is positive.

to correct for demographics and sectoral labour shifts. The main difference observed is that our results are of a higher statistical significance in comparison with those obtained from their non-differenced specification. In addition, our results closely resemble those recently obtained by Fève and Guay (2009) using a two-stage approach in a bivariate SVAR. Contrary to the evidence reported in the present paper, these two studies still consider empirical specifications with hours in first differences.

Finally, it is worth mentioning that this paper has focused on the most simple (but also the most commonly applied) specification when trying to clarify the evidence regarding the RBC and NK type of models. Nonetheless, the use of bivariate SVARs allows us to highlight the basic issues in a simple setting and compare the results with a greater number of related studies. The suitability of changing trends in SVARs involving more variables (Erceg et al., 2005) or different technological shocks (Fisher, 2006) remains an open question for further research.

5 Conclusions

This paper has empirically established that the deterministic component of the widely analyzed time series of hours worked per capita in the U.S. business sector is better specified as a broken linear trend. Taking this into account, recent tests give evidence against the unit root null hypothesis for this variable. Therefore, it is concluded that there is no reason to introduce hours worked in first differences into SVARs when estimating the effects of a technology shock. As an alternative, and trying to complement previous attempts in the literature, we propose to control for the presence of shifts in the deterministic components of both productivity growth and hours worked, simultaneously. Our results contrast with the predictions of standard RBC models, but only in the very short run and with little statistical significance.

References

[1] Andrews, D.W.K. and W. Ploberger, 1994. Optimal Tests When a Nuisance Parameter is Present Only Under the Alternative. Econometrica 62, 1383-1414.

[2] Ayat, L. and P. Burridge, 2000. Unit root tests in the presence of uncertainty about the non-stochastic trend. Journal of Econometrics 95, 71-96.

[3] Blanchard, O. and D. Quah, 1989. The dynamic effects of aggregate demand and supply disturbances. American Economic Review 79, 655–673.

[4] Basu, S., J.G. Fernald and M.S. Kimball, 2006. Are Technology Improvements Contractionary? American Economic Review 96, 1418-1448.

[5] Canova, F., D. López-Salido and C. Michelacci, 2009. The effects of technology shocks on hours and output: A robustness analysis. Forthcoming in Journal of Applied Econometrics.

[6] Carrión-i-Silvestre, J.L., D. Kim and P. Perron, 2009. GLS-based unit root tests with multiple structural breaks both under the null and the alternative hypotheses. Econometric Theory 25(6), 1754-1792.

[7] Chari, V.V., P.J. Kehoe and E.R. McGrattan, 2008. Are structural VARs with long-run restrictions useful in developing business cycle theory? Journal of Monetary Economics 55, 1337–1352.

[8] Christiano, L.J., M. Eichenbaum and R. Vigfusson, 2003. What Happens After a Technology Shock? NBER Working Paper No. 9819.

[9] Elliott, G., T.J. Rothenberg and J.H. Stock, 1996. Efficient Tests for an Autoregressive Unit Root. Econometrica 64, 813–836.

[10] Erceg, C., L. Guerrieri and C. Gust, 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] Fève, P. and A. Guay, 2009. The Response of Hours to a Technology Shock: A Two-Step Structural VAR Approach. Journal of Money, Credit and Banking 41(5), 987-1013.

[13] Francis, N. and V.A. Ramey, 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 V.A. Ramey, 2009. Measures of Per Capita Hours and their Implications for the Technology-Hours Debate. Journal of Money, Credit and Banking 41(6), 1071-1097.

[15] Galí, J., 1999. Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations? American Economic Review 89, 249-271.

[16] Galí, J. and P. Rabanal, 2005. Technology shocks and aggregate fluctuations: how well does the real business cycle model fit postwar U.S. data? NBER Macroeconomics Annual, 225–288.

[17] Gil-Alana, L. A. and A. Moreno, 2009. Technology Shocks and Hours Worked: A Fractional Integration Perspective. Macroeconomic Dynamics 13(5), 580-604.

[18] Hall, P., 1992. The Bootstrap and Edgeworth Expansion. Springer-Verlag, New York, USA.

[19] Harvey, D.I., S.J. Leybourne and A.M.R. Taylor, 2008. Testing for the unit roots and the impact of quadratic trends, with an application to relative primary commodity prices. Granger Centre Discussion Paper 08/04.

[20] Harvey, D.I., S.J. Leybourne and A.M.R. Taylor, 2009. Simple, robust, and powerful tests of the breaking trend hypothesis. Econometric Theory 25, 995-1029.

[21] Kappler, M., 2009. Do hours worked contain a unit root? Evidence from panel data. Empirical Economics 36(3), 531-555.

[22] Kejriwal, M. and P. Perron, 2009. A Sequential Procedure to Determine the Number of Breaks in Trend with an Integrated or Stationary Noise Component. Purdue University Economics Working Paper No. 1217.

[23] Lütkepohl, H. and M. Krätzig, 2004. Applied Time Series Econometrics. Cambridge University Press. New York, USA.

[24] Ng, S. and P. Perron, 2001. Lag length selection and the construction of unit root tests with good size and power. Econometrica 69, 1519–1554

[25] Ng, S. and T.J. Vogelsang, 2002. Analysis of Vector Autorregressions in the presence of shifts in mean. Econometric Reviews 21, 353-381.

[26] Perron, P., 1989. The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis. Econometrica 57, 1361-1401.

[27] Perron, P., 2006. Dealing with Structural Breaks. In Patterson, K. and T.C. Mills (eds.) Palgrave Handbook of Econometrics, Vol. 1: Econometric Theory. Palgrave Macmillan, 278-352.

[28] Perron, P. and Z. Qu, 2007. A simple modification to improve the finite sample properties of Ng and Perron's unit root tests. Economics Letters 94, 12-19.

[29] Perron, P. and T. Yabu, 2009. Testing for Shifts in Trend with an Integrated or Stationary Noise Component. Journal of Business and Economics Statistics 27(3), 369-396.

[30] Qu, Z. and P. Perron, 2007. Estimating and testing structural changes in multivariate regressions. Econometrica 75, 459-502.

[31] Shields, G. and M.P. Shields, 2008. Four themes to explain the rise in hours worked in the United States. Applied Economics Letters 15, 1163-1166.

[32] Vogelsang, T., 2001. Tests for a shift in trend when serial correlation is of unknown form. Mimeo, Department of Economics, Cornell University.

[33] Whelan, K.T., 2009. Technology shocks and hours worked: checking for robust conclusions. Journal of Macroeconomics 31(2), 231-239.

Table 1: Unit root and trend shift testing. Hours worked per capita, 1948:Q1-2007:Q4.

	No structural break	Single structural break
Augmented Dickey-Fuller (ADF)	-1.37	-4.00***
Elliot-Rothenberg-Stock (P_T)	20.38	5.26**
Modified P_T (MP_T)	19.34	5.22**
Phillips-Perron (Z_α)	-4.72	-33.40***
Modified Z_α (MZ_α)	-4.68	-33.23**
Modified Sargan-Bhargava (MSB)	0.32	0.12**
Modified Z_t (MZ_t)	-1.51	-4.03***
Exp-W_{FS}	—	2.59*
p-value I(0)	—	0.06
p-value I(1)	—	0.09
Break date	—	1983:Q4

Note: Hours worked per capita are expressed in natural logarithms. They are equal to the hours worked in the U.S. business sector divided by the civilian non-institutional population over 16 (1992=100). Standard univariate unit root tests are those based on the quasi-GLS detrending method discussed in Ng and Perron (2001). The deterministic component is made up of a constant and a trend. The number of augmentation lags have been selected using the MAIC criterion calculated following the suggestion in Perron and Qu (2007). The Exp-W_{FS} test is that of Perron and Yabu (2009) for model III with a trimming of 15%. The break date has been estimated by minimizing the sum of squared residuals from a regression of hours worked on a constant, a time trend, a level shift and a slope shift dummy. Unit root tests allowing for a break are those proposed by Carrión-i-Silvestre et al. (2009). ***, ** and * denote rejection of the null at the 1, 5 and 10% significance level, respectively.

Table 2: Comparison of raw deterministic component adjustments. Hours worked per capita, 1948:Q1-2007:Q4.

	Linear trend	Quadratic trend	Two level shifts	Broken linear trend
Constant	4.71^{***}	4.78^{***}	4.70^{***}	4.77^{***}
Trend	$-4.22\cdot10^{-4}$ ***	$-2.31\cdot10^{-3}$ ***	—	$-1.48\cdot10^{-3}$ ***
$(\text{Trend})^2$	—	$7.90\cdot10^{-6}$ ***	—	—
Level shift dummy (I)	—	—	-0.08^{***}	$7.94\cdot10^{-2}$ ***
Trend shift dummy	—	—	—	$1.63\cdot10^{-3}$ ***
Level shift dummy (II)	—	—	0.04^{***}	—
Adjusted R^2	0.26	0.62	0.40	0.73
Standard error	0.05	0.03	0.04	0.03
SSR	0.57	0.29	0.46	0.21
Log-likelihood	384.89	464.61	409.67	505.45
AIC	-3.19	-3.85	-3.39	-4.18

Note: Hours worked per capita are expressed in natural logarithms. They are equal to the hours worked in the U.S. business sector divided by the civilian non-institutional population over 16 (1992=100). The two level shifts take place in 1973:Q2 and 1997:Q2 and are based on the estimation by Fernald (2007) for the corresponding productivity growth time series. The break date for the broken linear trend specification is based on the results reported in Table 1. Inferences drawn from Newey-West standard errors. *** denotes significant at the 1% significance level.

Table 3: Deterministic shifts testing in a bivariate VAR model of productivity growth and hours worked per capita, 1948:Q1-2007:Q4.

VAR order	sup LR test	1st level shift in productivity growth	Trend shift in hours worked	2nd level shift in productivity growth
1	95.95***	1973:Q1 [1969:Q1 , 1979:Q2]	1982:Q4 [1981:Q1 , 1983:Q2]	1991:Q4 [1988:Q4 , 1995:Q1]
2	89.59***	1973:Q1 [1969:Q1 , 1978:Q3]	1982:Q4 [1981:Q3 , 1983:Q2]	1995:Q3 [1990:Q3 , 1999:Q2]
3	85.21***	1973:Q1 [1969:Q3 , 1978:Q1]	1983:Q2 [1981:Q4 , 1983:Q3]	1995:Q3 [1990:Q4 , 1998:Q4]
4	78.92***	1977:Q3 [1974:Q4 , 1978:Q4]	1986:Q2 [1985:Q1 , 1986:Q4]	1995:Q3 [1989:Q4 , 1999:Q3]

Note: The sup LR test statistic is that proposed by Qu and Perron (2007) to test for structural changes in multivariate frameworks. A restricted version where only the deterministic components are allowed to change has been considered. The trimming parameter has been set to 0.15 and the test is robust to the presence of heteroskedasticity. Break dates estimated using a dynamic programming algorithm based on a quasi-maximum likelihood estimation assuming uncorrelated normal errors. 90% confidence intervals reported. *** denotes rejection of the null at the 1% significance level.

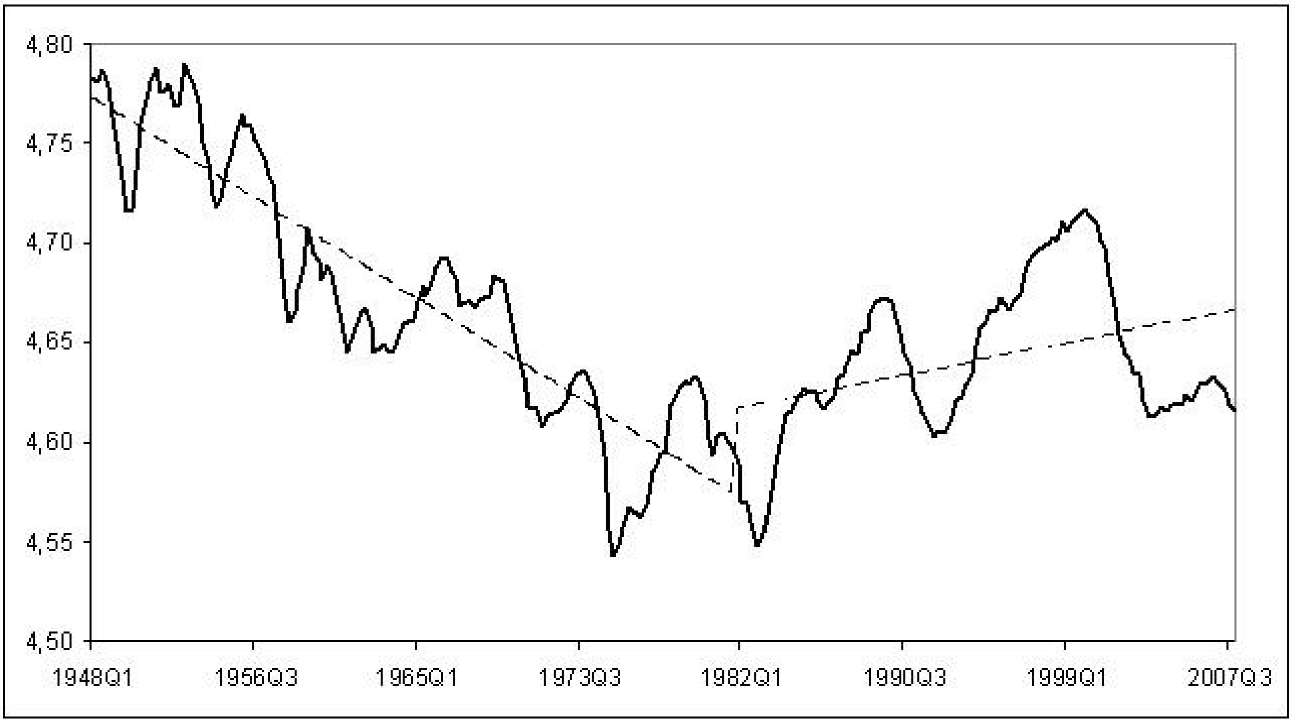

Figure 1: Hours worked per capita (1992=100, in natural logarithms) and adjusted broken linear trend, 1948:Q1-2007:Q4.

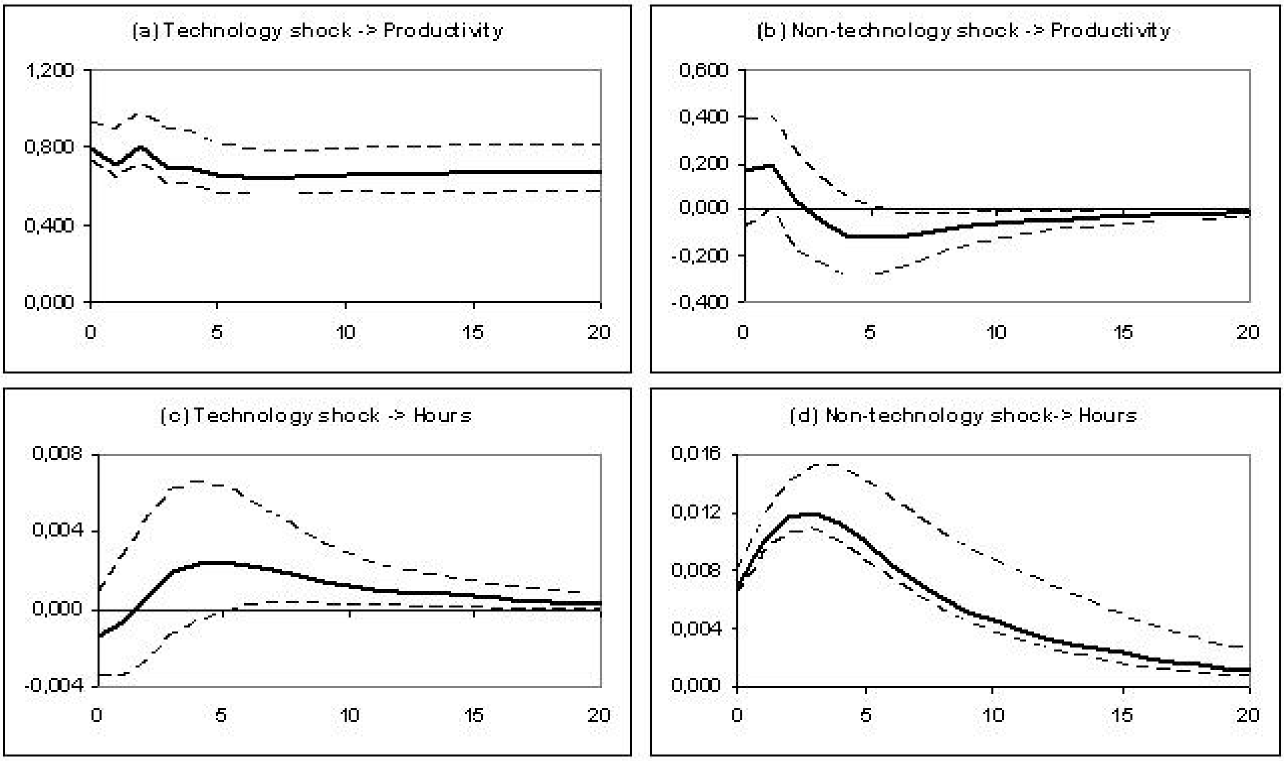

Figure 2: Impulse-response functions from a bivariate SVAR with productivity growth and (log) hours per capita controlling for deterministic shifts in both time series. Hall (1992)'s equal-tailed percentile 90% bootstrap confidence bands reported (2000 replications, dotted lines).