On the robust effects of technology shocks on hours worked and output

Fabio Canova
ICREA, Universitat Pompeu Fabra, AMeN, CEPR

David Lopez-Salido
Board of Governors and Bank of Spain

Claudio Michelacci
CEMFI

February 2006

Abstract

We analyze the effects of neutral and investment-specific technology shocks on hours and output in VARs where hours are considered both in level and first difference. Subsample analysis shows instabilities in the relationships with the level specification. Low frequency movements in hours are captured via switching intercepts or polinominals. Hours fall in response to neutral shocks and increase in response to investment specific shocks. The percentage of the variance of hours (output) explained by neutral shocks is small(large); the opposite is true when considering investment specific shocks. News shocks are uncorrelated with neutral and investment specific shocks.

JEL classification: E00, J60, O33.

Key words: Technology disturbances, structural VARs, low frequency movements, news shocks.

1 Introduction

There has been a renewed interest in the literature in examining the effects of technology shocks on either total or per-capita hours worked following the work of Gali (1999), Christiano et. al. (2003), Uhlig (2004), Dedola and Neri (2004), Francis and Ramey (2005), among others. This interest is typically motivated by the fact that the dynamics of hours differ if a technological disturbance occurs in a basic RBC model (hours increase) or in a basic sticky-price model (hours decrease). Such a manichean view is probably unhealthy, as there are version of RBC models in which hours fall (see e.g. the models analyzed by Francis and Ramey (2005)) and parameterizations of New-Keynesian models with capital in which hours increase in response to technological disturbances (see e.g. Pappa (2004)). Nevertheless, since it is hard to find sharp and contrasting predictions for these two types models, the temptation to use the empirical response of hours to assert the relative plausibility of the two theories for business cycle fluctuations is strong.

Unfortunately, the empirical evidence on the relationship between technology shocks and hours worked is mixed. Four main reasons contribute to make the task of applied investigators complicated: first, technology shocks are difficult to identify and only their long run implications are considered robust to measurement and specification errors; second, several types of technological disturbances may drive fluctuations in hours worked; third, the choice of price deflators may be crucial in assessing the effects of technological disturbances; fourth, the response of hours worked appears to depend on a number of auxiliary assumptions, including the treatment of the low frequency movements in hours, the length of the VAR - theoretical models used to guide the interpretation of the results may not be approximable with finite length VARs - and the horizon at which the identifying restrictions are imposed.

Regarding the low frequency movements in hours, two contrasting arguments are typically made. If we condition the analysis on the models one uses to interpret the results as, for example, Christiano et. al. (2003), Uhlig (2004), Dedola and Neri (2004) do, per-capita hours should be enter the VAR in level since both RBC and New-Keynesian models produce stationary fluctuations in this variable, even when the technology is drifting in a non-stationary fashion. However, if one runs unit roots tests, she will have hard time to reject the hypothesis that per-capita hours are non-stationary. Therefore, if we condition the analysis on the statistical properties of the data and follow a classical statistical approach, as e.g. Gali (1999), and Francis and Ramey (2005) do, the VAR should include hours in differences. The picture is complicated by the fact that unit root tests are unlikely to be conclusive with the available sample size and, as one can easily see in figure 1, per-capita hours display a non-linear U-shaped trend which makes the outcome of unit root tests dubious. Hence, both approaches are potentially subject to specification errors and it is not clear which strategy one should follow. The alternative approach used by Fernald (2004) - break the sample in pieces - is economically more appealing since hours and labor productivity share low frequency movements, but it is equally problematic since large short sample biases will result.

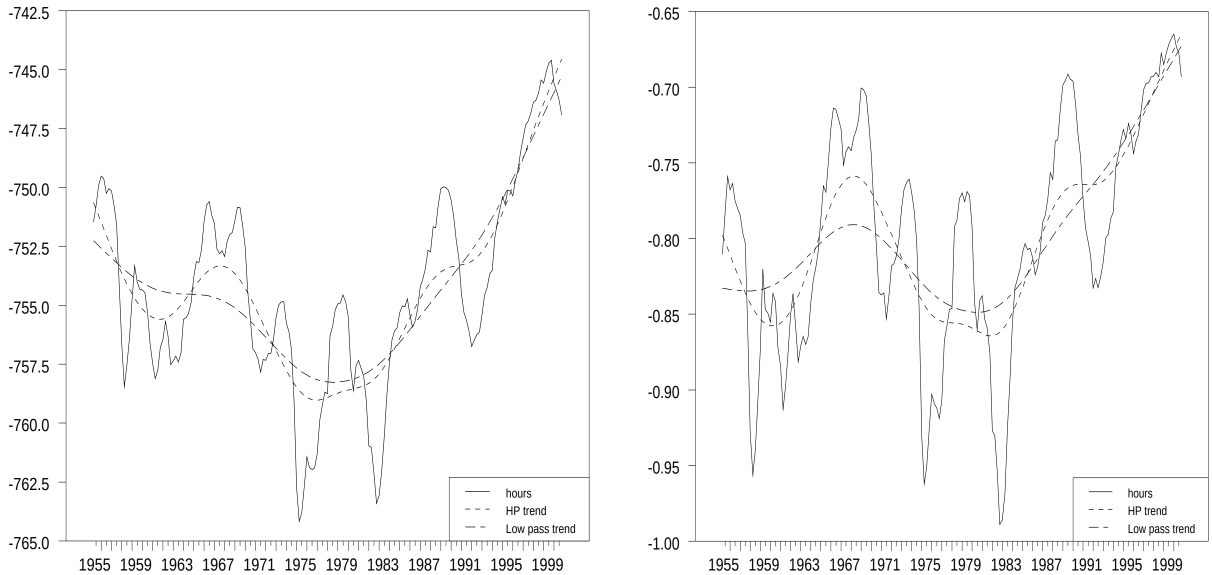

Standard hours series

Francis and Ramey hours series

Figure 1: Per-capita hours and various trends

The issue of the length of VAR has been brought back to the attention of applied researchers by the recent contributions of Giordani (2003) and Chari et. al. (2005), who show that a subset of the variables generated by a standard models may have a solution which is not always representable with a finite order VAR, and Rubio et. al. (2005), who provide conditions to construct valid VAR approximations. For the debate on the dynamics of hours in response to technology shocks, this issue is important because estimated VARs typically include only a small subset of the potentially interesting variables and because the lag length may be too short to capture the implied dynamics when the true DGP is an VARMA with large MA roots (see e.g. the simple model used by Gali (1999)).

Basu et al. (2005) have argued that, in measuring the effects of technological disturbances on hours worked, it matters whether the CPI or the GDP deflator is used to compute the price of investments. Theoretically, the choice depends on whether the economy is treated as closed or open, since in the latter case CPI and GDP deflators may considerably differ in the long run. Empirically, it is often unclear which choice is made. Hence, it is difficult to gauge whether the results hinge on this assumption or not.

Finally, the fact that different primitive shocks may have similar long run implications as technology shocks, that long run restrictions are weak in the sense of Faust and Leeper (1997), and that the available samples are too short to credibly impose such restrictions (as recently shown by Erceg, et al. (2005)) has prompted researchers to search for alternative identification strategies (see e.g. Uhlig (2004), Dedola and Neri (2004)). Furthermore, the evidence in Beaudry and Portier (2005) is, at the minimum, puzzling for this line of research.

This paper examines the effects of technology on hours worked directly addressing all the problems which make the empirical evidence difficult to interpret. In particular, we deal both with the presence of low frequency movements in hours and the potential misspecifications created by the use of VARs with a limited number of variables and a finite number of lags; we examine the robustness of the conclusions to alternative identification approaches and different measures for the price of investment goods; and we separately analyze the dynamics induced by neutral and investment specific technology shocks. We show that once one takes care of low frequency movements in hours worked in any reasonable matter, all the other pieces of the puzzle become irrelevant: regardless of the selected lag length, the presence (or absence) of omitted variables, the choice of price deflators, the identification scheme and other auxiliary statistical assumptions one is forced to made in specifying the VAR, hours fall in response to neutral technology shocks and increase in response to investment specific shocks and both responses tend to be persistent. We also demonstrate that the contribution of neutral technology shocks to the fluctuations in hours worked is typically small, while the contribution of investment specific shocks is substantial. Interestingly, the relative importance of the two shocks for output fluctuations is reversed: neutral shocks explain about twice as much of the forecast error variance of output than investment specific shocks at all horizons. We show that the time path of our estimated neutral shocks has peaks and troughs in correspondence of NBER peaks and throughs, while this is not always the case for investment specific shocks. Also, we show that neither of our two shocks is highly correlated with news shocks at any horizon ranging from contemporaneous up to 4 years and, if anything, displays a mild negative correlation with them.

What conclusions one draws from these results? Although the combined effect of the two technology shocks on per-capita hours worked and output is significant, it is far from clear which theoretical paradigm is more appropriate to explain the data. The dynamics of hours in response to neutral technology shocks are inconsistent with the implications of a basic RBC models, but in contrast with a simple New-Keynesian argument, neutral shocks generate important output fluctuations and, perhaps more importantly, they look like business cycle shocks. The dynamics of hours in response to investment specific shocks are in line with the predictions of a basic RBC model, but such shocks explain little of output fluctuations and fail to display a strong cyclical behavior, in particular, in important recessions episodes.

The rest of the paper is organized as follows. Section 2 critically summarizes the existing evidence and explains why outcomes are uninterpretable. Section 3 provides new evidence when low frequency movements in hours worked are accounted for in a variety of ways. Section 4 shows that our results are robust to a number of changes in the auxiliary assumptions. Section 5 provides some discussion and comparisons with the literature and section 6 looks into the properties of the estimated shocks. Section 7 concludes.

2 The existing evidence

We start by summarizing existing evidence and the current state of the debate. We do so in the context of a three equations VAR model composed of labor productivity, the price of investment, both measured in consumption units, and per-capita hours, which is the minimum sized system required to distinguish between neutral and investment specific shocks. Since such a small model is liable to specification errors, we show later how to check for potential omitted variables in a simple and visually appealing way. The sample we use runs from 1955:1 to 2000:4. We present results when all three variables enter the VAR in first difference (the difference system) and when the first two enter in first difference and per-capita hours are in logs (the level system). Neutral shocks are identified by the requirement that they are the sole source of labor productivity movements in the long run while investment specific shocks are identified by the requirement that they are the sole source of movements in the price of investment in the long run. Fisher (2003) has shown how to derive these restrictions from two basic equations embedded in models of both neoclassical and New-Keynesian orientations: a production function with unit root disturbances and a capital accumulation equation where the price of investment in terms of consumption goods is stochastic and displays a unit root. We use 12 lags of each variable in each of the estimated systems, stochastically restrict their decay toward zero as lag length increases and construct standard errors bands using standard bootstrap methods.

The first row of figures 2 reproduces what is more or less known in the literature: in the level system, per-capita hours positively respond to a neutral shock, the maximum response appears to be delayed by about 5 quarters and the instantaneous impact is not significant. In the difference system, per-capita hours fall for up to 4 quarters and then settle to their long run (positive) level, but responses are insignificant at all horizons. Interestingly, and despite these differences, neutral disturbances account for less than 2 percent of the variability of hours worked at horizons ranging from 8 to 24 quarters in both specifications.

The dynamics of output in response to the two shocks are similar across specifications, quite standard and therefore omitted: output instantaneously increases, displays a small hump after about 4 quarters and then settles to its long run value.

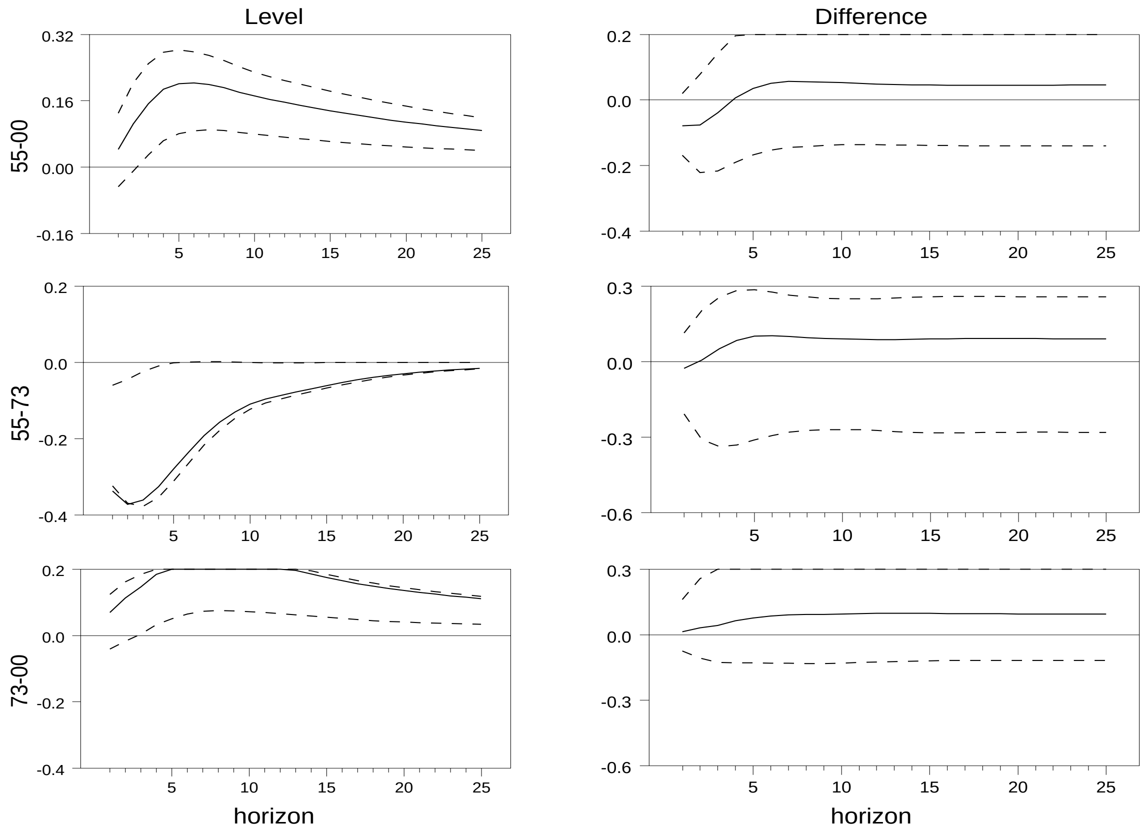

Figure 2: Responses of hours, neutral shocks.

The sign difference in the point estimates across specifications is typically attributed to low frequency movements in hours worked and the fact that, when these movements are not appropriately captured, they distort the dynamics in the level specification.

The right panel of figure 1, which shows a predominant U-shaped pattern in the log-level of the series, confirms the results obtained with the level system are probably dubious. However, such a pattern does not necessarily imply the existence of permanent shocks nor of non-stationarity dynamics, as it is commonly assumed (see e.g. Gali (2005)) - there could be stationary cycles with long but finite periodicity, undetectable with 45 years of quarterly data - resulting in standard overdifferencing problems. In addition, a difference specification, emphasizing high frequency hours variability, may to magnify the importance of measurement errors and produce in standard error bands which are large relative to those obtained with the level specification (see figure 2). Hence, both the difference and the level specifications are probably misspecified and it is difficult to draw conclusions on the relative credibility of the results they produce.

One alternative way to take care of low frequency movements in per-capita hours is split the sample in pieces. We follow the literature and choose 1973:2 as a break point, a date that many take as critical to understand the dynamics of the productivity series (see Greenwood and Yorokoglu (1997)) and that others (see e.g. Fernald (2004)) consider as well important to understand the dynamics of per-capita hours.

The second and the third rows of figure 2 report the responses of per-capita hours in two subsamples for the two specifications. The results appear to have considerably changed. For the sample 1955-1973, hours instantaneously fall in both specifications: in the level system significantly so; in the differences system responses continue to be insignificant. The point estimate of the response of hours in the level specification is persistently negative and significantly so up to 4 quarters, while for the difference specification, the point estimate turns positive after just one quarter, but remain insignificant at all horizons. For the 1973-2000 sample, hours instantaneously increase in both specifications. In the level system, responses are significant only after 4 quarters and display a clear hump, while in the difference specification responses are positive all horizons, albeit insignificant. Simple inspection of figure 1 indicates that the trend in the hours series may have changed once again somewhere in 1990s. We have repeated estimation using data up to 1992:1 and 1997:1 for the second subsample: the responses of hours to neutral shocks remain positive in both specifications and significant only in the level system after about 4 quarters.

How does one interpret this evidence? Typically, the considerable subsample instability found in the level system is considered not genuine and thought to reflect improper treatments of low frequency movements in per-capita hours. But should one really trust subsample evidence? Probably not because to take care of one potential misspecification problem (low frequency movements), an approach which splits the sample in pieces introduces potentially large small sample biases. Small sample biases make estimates unreliable for many reasons. First, point estimates less informative than otherwise: this is evident comparing standard error bands across rows in figure 2. Second, identifying long run restrictions imposed in a system estimated over a small sample are likely to induce serious distortions in the structural estimates (see Erceg, et al (2005)). Third, small sample biases may interact in an unpredictable way with measurement and aggregation errors, making the subsample evidence uninterpretable.

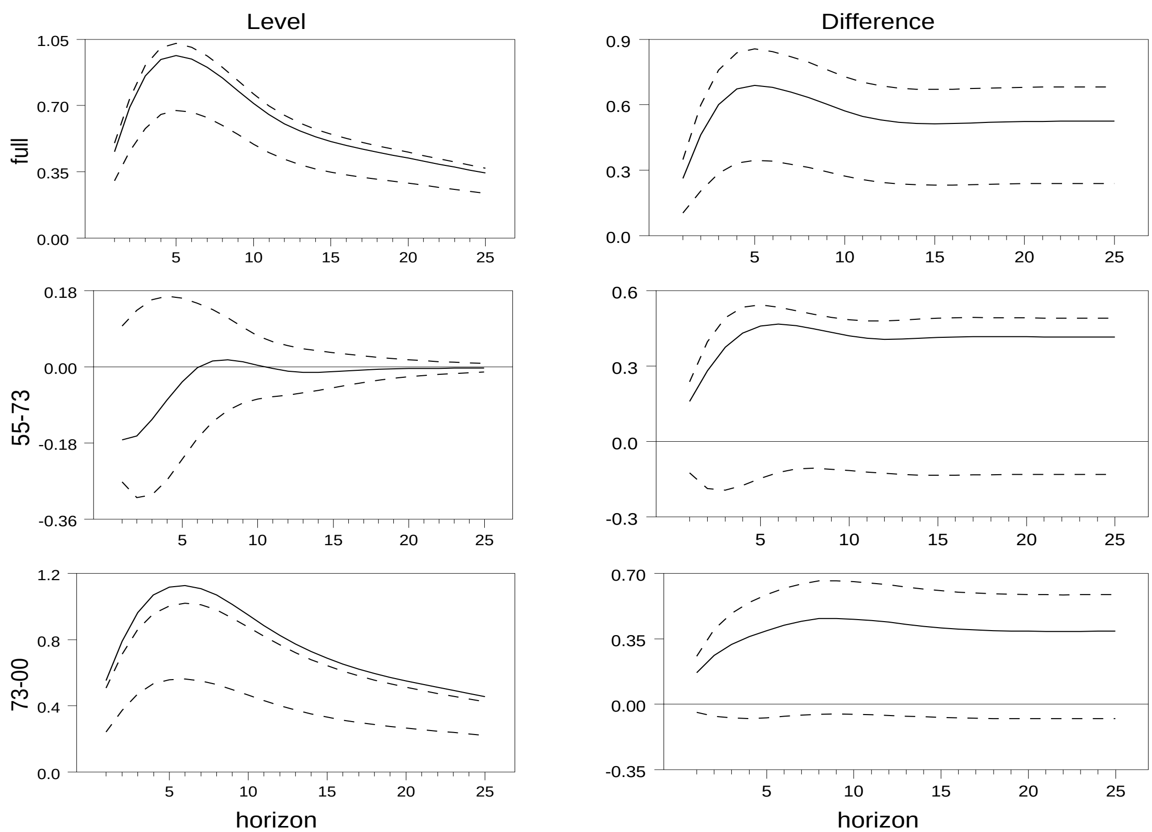

Figure 3: Responses of hours, investment specific shocks.

Neutral technology shocks are not the only possible source of long run fluctuations in real variables. For example, Uhlig (2004) indicates that tax shocks may have roughly the same effects as neutral shocks and Fisher (2003), following the lead of Greenwood et al. (2001), suggests that investment specific shocks could be an important source of technological disturbances. What are the effects of investment specific shocks on per-capita hours?

Figure 3 presents per-capita hours responses for the full sample and the same two subsamples analyzed in figure 2 for the level and the difference system. For the full sample, the two specifications agree: hours responses display a hump shaped pattern; the increase is instantaneously significant; and the magnitude of the effect is roughly similar. When we split the sample the results differ across specifications: hours fall in the 1955-1973 sample and increase in the 1973-2000 sample in the level system but in the first sample estimates are never statistically significant. For the difference specification, the sign of the instantaneous point estimate is the same across subsamples, but none of the responses is significant. Output responses to investment specific shocks almost entirely replicate the responses of hours and the significance of the responses is also similar. The only exception is the sample 1973-2000 in the difference system: here output responses are positive and significant all horizons.

Investment specific shocks appear to be an important source of cyclical variations in per-capita hours. In fact, they explain 16-21 percent of the variability of hours in the first difference specification and 49-65 of the variability of hours in the level specification at horizons ranging from 8 to 24 quarters.

While it is difficult to find a simple and convincing story which can account for the complexity of patterns found in figures 2 and 3, it is apparent that small sample biases could play an important role. Hence, subsample evidence can not be used to support or contradict the claim that low frequency movements in hours drive the results in the level specification.

In conclusions, neither assuming unit roots nor splitting the sample in piece seems the best way to account for the particular U-shaped low frequency movements that per-capita hours worked display. In the next section we describe approaches more appropriately deal with this issue.

3 An alternative treatment of trends

The cyclical components that result from taking the difference between the actual series and, e.g., the two trends we have presented in figure 1, seem to be reasonably stable over time. Therefore, splitting the sample and tracing out the dynamics separately in each of them is inefficient, as there appears to be little changes in the cyclical patterns of hours. One way to take care of low frequency movements in per-capita hours consistent with this observation, is to allow the intercept of the hours equation to vary with time but restrict the slopes to be time invariant. We have considered several options. In the first case, the intercept is deterministically broken at 1973:2 and 1997:1 (the dummy specification); in the second case the intercept is allowed to be a deterministic function of time (up to a third order polynominal); in the third case we prefilter the hour series with a high-pass filter, which takes away cycles with periodicities higher than 52 quarters and in the fourth one the intercept is allow to drift stochastically and potentially continuously over time. In this latter case, we specify an autoregressive mean-reverting law of motion and use the Kalman filter to recursively estimate it over the full sample. Note that in all cases, biases induced by the use of long run identification restrictions are also considerably reduced, since the full sample of quarterly data is now employed to project estimated VAR coefficients infinitely far into the future.

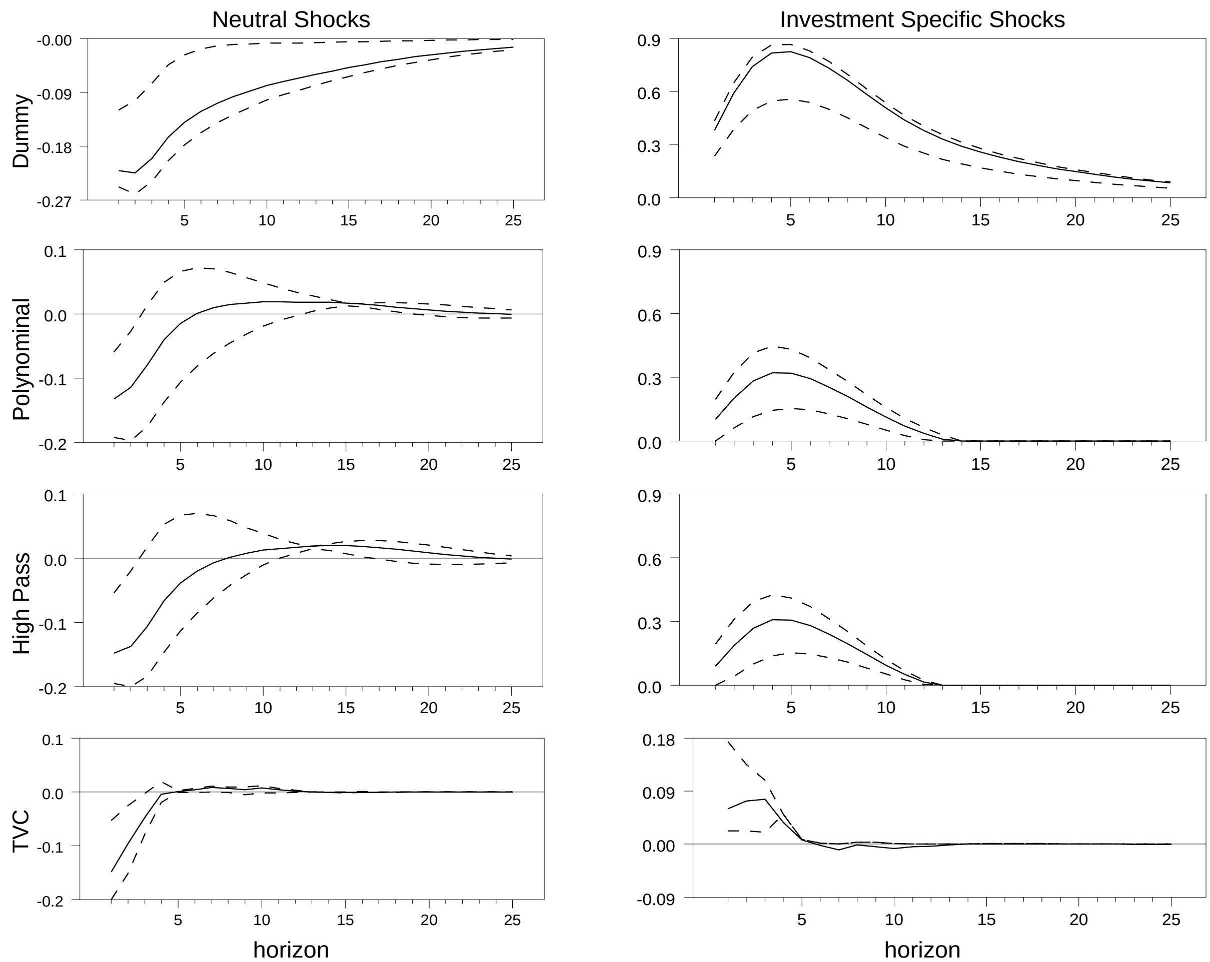

Figure 4: Responses of hours, various detrending, 1955-2000

We present the response of hours in these specifications for the 1955-2000 sample in figure 4: the first column refers to neutral shocks and the second to investment specific shocks. The results are sharp and consistent across trend specifications: per-capita hours fall in response to neutral shocks and increase in response to investment specific shocks. Depending on the exact specification, the fall in response to neutral shocks is either persistent (first row) or temporary (next three rows). Nevertheless, the instantaneous response is always significant; in no case a hump results after 4-5 quarters; and when the responses turn positive they are hardly ever significant.

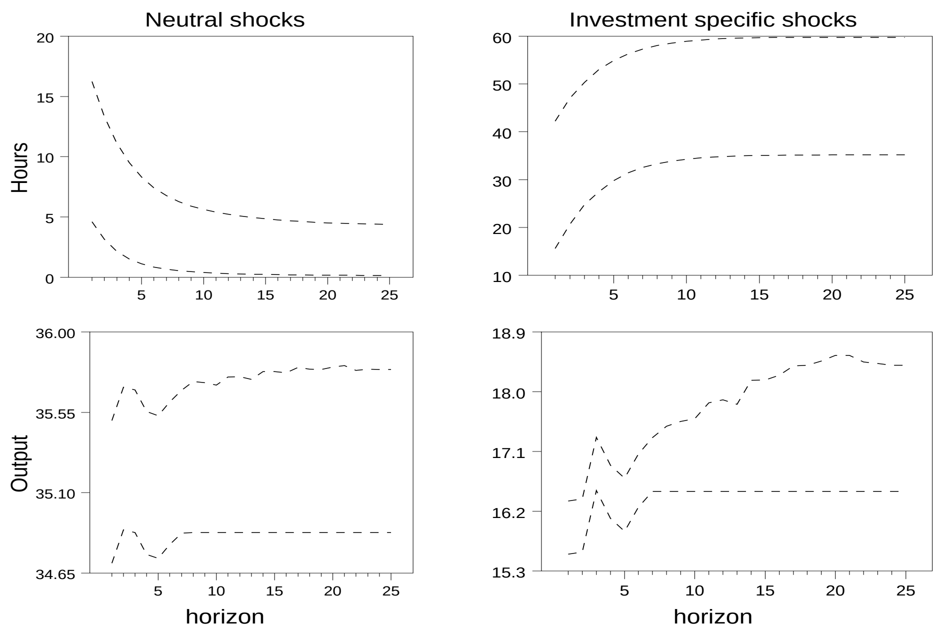

Figure 5: Variance decomposition, dummy specification, 1955-2000, 68 percent bands

The percentage of the variance of per-capita hours explained by the shocks is similar in the four cases: neutral disturbances have negligible effects on per-capita hours at horizons varying between 8 and 24 quarters (the upper 95th percentile of the distribution is always less than 10 percent) while investment specific shocks explain between 30-50 percent of the variance of hours at these horizons (20-30 percent with the continuously varying intercept specification). For output fluctuations the opposite appears to be true: neutral shocks explain on average about 35 percent of output fluctuations and investment specific shocks only about 18 percent of output fluctuations, regardless of the horizon we consider (see figure 5 for the dummy specification).

Are these results robust? There are many dimensions along which robustness could be checked, making the combination of system to be estimated quite large. We divide our analysis into two parts. First, we study the robustness of the results to the choice of variables included in the system and their measurement. Second, we check whether outcomes are sensitive to the statistical assumptions we have made. As it will be clear, our conclusions on the role of technology shocks are quite robust.

4 Robustness

To start with we analyze the robustness of the results to the measurement of real variables. In the systems we have run, labor productivity and the price of investment are all measured in consumption units. However, it is easy to show that if foreign goods enter the consumption basket, labor productivity can be moved by shocks other than neutral shocks in the long run and, e.g., external shocks can play an important role. Since the output basket is less prone to such problems, we have repeated the exercise measuring either labor productivity or both the labor productivity and the price of investments with the output deflator. Conclusions similar are in both cases. We report the results obtained when both variables are measured in output units in the first row of figure 6: the sign and the shape of both responses is unchanged.

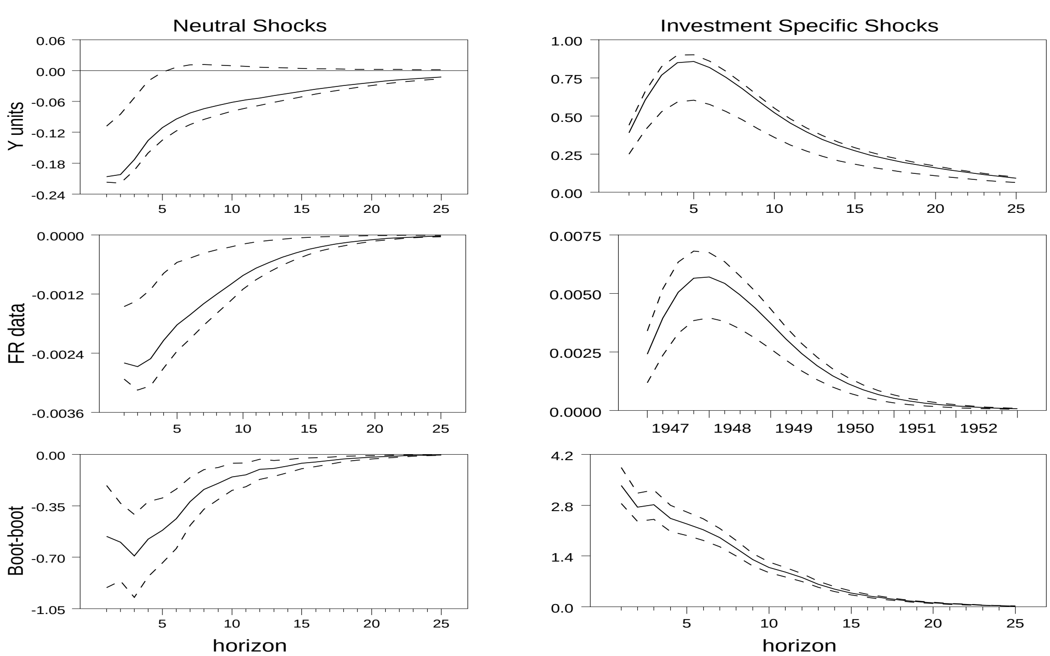

Figure 6: Responses of per-capita hours, alternative measurements, dummy specification, 1955-2000

Second, we have repeated our exercises using the recently constructed hours series of Francis and Ramey (2005). One of the problems of the original hour series is that low frequency movements, correlated with low frequency movements in labor productivity, are very evident. The series of Francis and Ramey displays less low frequency movements but, contrary to what the authors claim, it is still not void of them (see right panel of figure 1). Therefore, it can be used to check if our treatment of low frequency movements is appropriate. The results obtained running the system with the new series are in the second row of figure 6: clearly none of the conclusions is altered if the new per-capita series is employed.

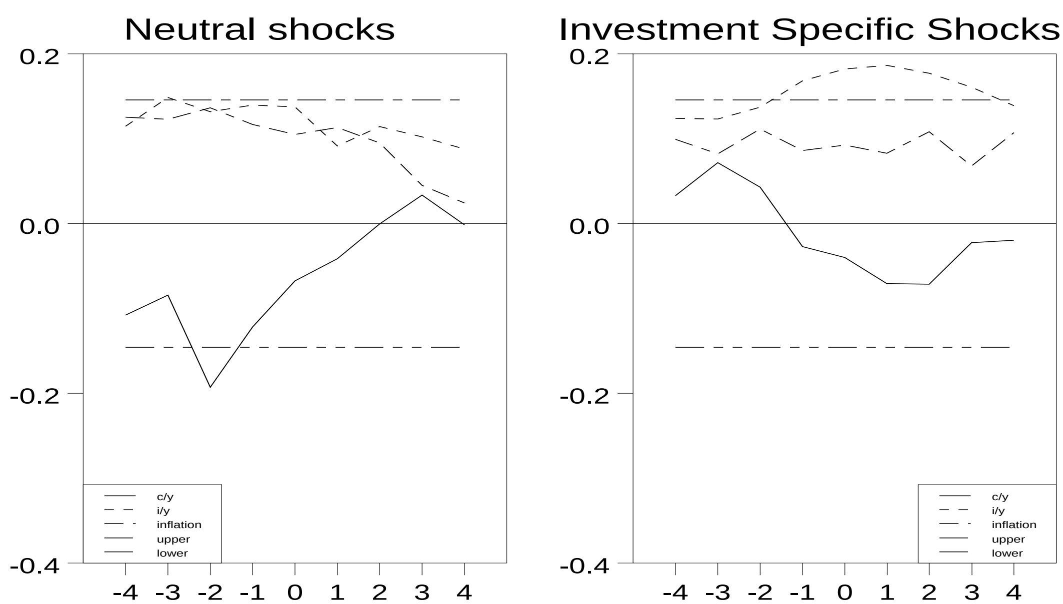

Figure 7: Cross correlation structural shocks-omitted variables, dummy specification, 1955-2000

Third, while we have allowed long enough lags in each estimated specification, it always possible that omitted variables play an important role in the results. For example, Evans (1992) showed that Solow residuals are correlated with a number of policy variables, therefore making responses to Solow residuals shocks uninterpretable. To check for this possibility we have correlated our two estimated technology shocks with variables which a large class of general equilibrium model driven by neutral and investment specific shocks suggests as being jointly generated with the data we have used. Figure 7, which reports the correlations of consumption to output, investment to output, and inflation with the estimated structural shocks up to 4 leads and 4 lags and the upper and lower limits of an asymptotic 95 percent confidence band for the dummy specification, suggests that this is not the case. In fact, in none of 4 estimated system we find that omitted variables are significantly correlated with the structural shocks we extract: the maximum point estimate of these cross-correlation is -0.20 and occurs with the dummy specification.

We have also correlated our estimated technology shocks with oil price, federal funds future and tax shocks. These are constructed as the residuals of univariate regressions of each the three variables on two lags each. The effective tax series is taken from the Congressional Budget Office and is transformed into quarterly frequency using an interpolation routine. The correlations are all small and never exceed 0.11 in absolute value for cross correlations up to 4 lags and 4 leads. Therefore, our technological disturbances do not appear to stand-in for other sources of technological and non-technological disturbances either.

Another way to examine the potential effects of omitted variables is to check the robustness of the results to changes in the lag length. To the extent that omitted variables results in VAR residuals with MA components, adding lags to the model should help to eliminate this problem (as shown in Chari, et. al. (2005)). We report median estimates and standard error bands for the contemporaneous response of per-capita hours worked to the two shocks in figure 8 for the dummy specification. It is clear that the sign of the responses is very robust to the choice of lag length. Interestingly, when a short lag length is used the response to neutral shocks become more strongly negative.

Fourth, we modified the way standard error bands are computed. OLS estimates of VAR coefficients are typically biased downward when small sample sizes are used. If these are used to construct bootstrap samples, they produce standard error for impulse responses which are wrongly centered and possibly biased. While our sample is not necessarily short, systems with a large number of lags are subject to the same problems. We have therefore repeated the computation of standard error bands using the bootstrap-after-the-bootstrap approach of Kilian (1998) for each of the systems we have presented in figure 4. None of our conclusions change. For illustration purposes, we present the responses in the dummy specification in the third row of figure 6. Apart from the fact that bands are tighter and that the hump in response to investment specific shock has disappeared, no other qualitative difference is visible.

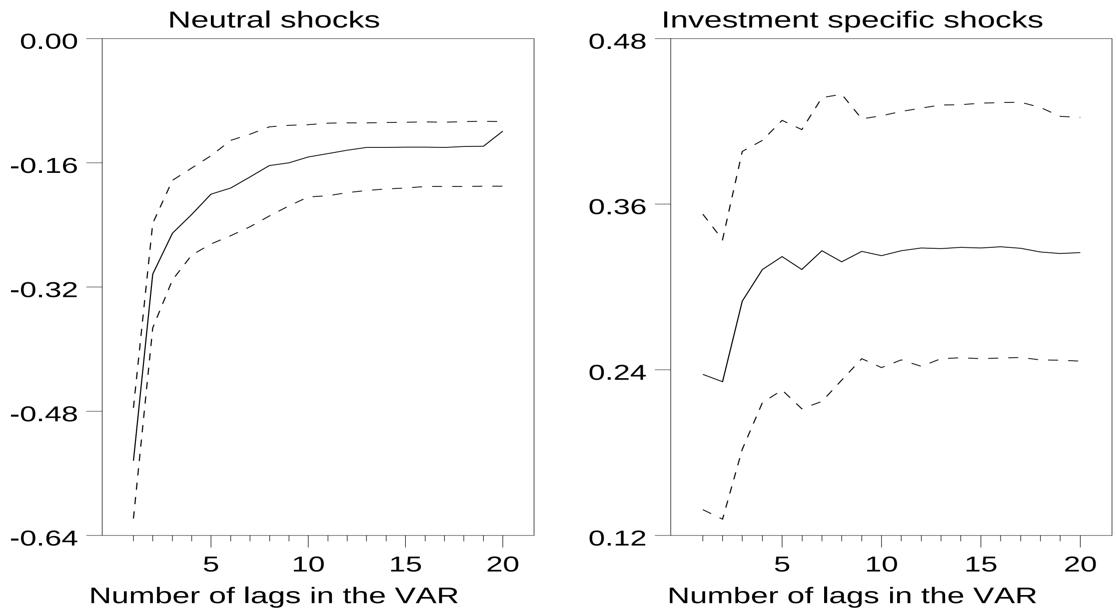

Figure 8: Contemporaneous response of hours, various lags, dummy specification, 1955-2000

Fifth, we have altered the identification restrictions. Uhlig (2004) has forcefully argued that disturbances other than neutral technology shocks may have long run effects on labor productivity and that, in theory, there is no horizon at which neutral shocks fully account for the variability of labor productivity. Therefore, the neutral shocks we have extracted may not be structural and the fact that they generate minor fluctuations for per-capita hours may be the result of mixing up, for example, technology and other shocks. Beaudry and Portier (2005) have found empirical evidence consistent with this interpretation: neutral technology shocks obtained in this way are highly correlated with "information" shocks extracted using shock run restrictions and a bivariate system with total factor productivity and stock prices. To study the extent of these potential problems, we have imposed the restriction that the two shocks are the sole source of fluctuations in labor productivity and the price of investment at horizons varying $k = 1, \ldots 50$. We report in figure 9 the impact response of hours worked in these systems for the dummy specification. It is clear that the sign of the response is robust to the horizon at which the restriction is imposed, it is significant in all the specifications we consider, except for neutral shock at horizons shorter than 3 quarters.

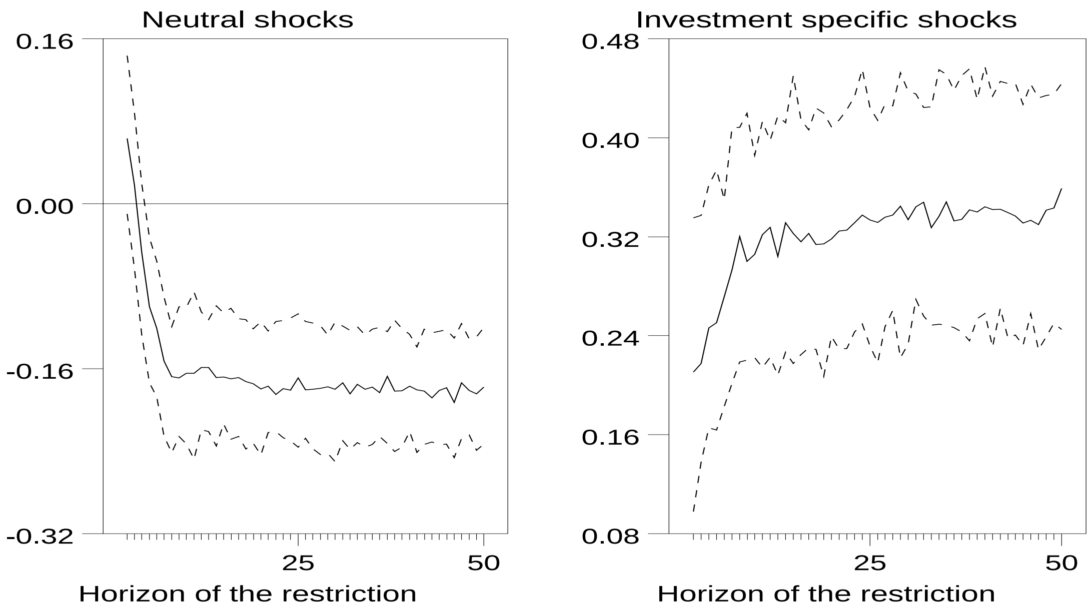

Figure 9: Contemporaneous response of hours, identification restrictions at various horizons, dummy specification, 1955-2000

Long run restrictions are vacuous if the series they constrain are stationary around some deterministic trend or simply nearly integrated. When this is the case one needs to devise alternative restrictions to identify the two shocks of interest. Dedola and Neri (2004), for example, take the point of view that, given that unit root tests still have low power, assuming stationary time series in the VAR may not be such a bad assumption and use sign restrictions derived from an RBC model to identify the two shocks. Are results robust to these alternative specification choices? Are the responses of hours to the two technology shocks independent of whether we condition on the price of investment and labor productivity having or not having a unit root and, in the latter case, identifying shocks using sign restrictions?

The sign restrictions we choose are simple and are likely to be shared by a number of models with different micro-fundations: we require that in response to a neutral shocks, productivity responses are contemporaneously non-negative while the response of the (relative) price of investment is contemporaneously positive; and that in response to an investment specific shock, productivity responses are contemporaneously positive and price of investment contemporaneously negative. Note that we are extremely parsimonious in the number of restrictions we employ: only the contemporaneous effects are constrained; lagged effects are totally unrestricted.

Figure 10 presents standard error bands for the responses (in this case computed via Monte Carlo integration) together with the median response of per-capita hours to the two shocks. Clearly, the qualitative pattern of the responses is unchanged: hours fall for 4 quarters in response to a neutral shock (although less persistently so); increase in response to an investment specific shock and display a hump shaped pattern. Perhaps unsurprisingly, it takes much longer for responses to settle to their long run equilibrium with this specification.

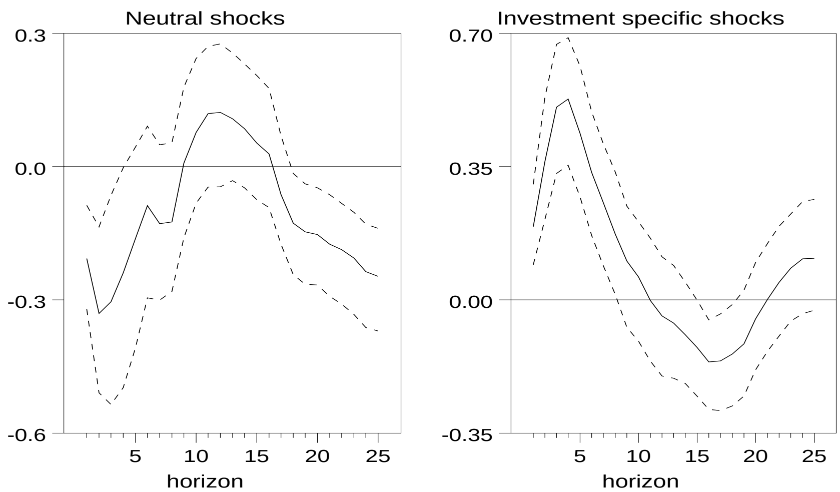

Figure 10: Hours responses, sign restrictions, dummy specification

In sum, if one takes the view that low frequency movements in per-capita hours can be characterized with time varying intercepts, and that these movements are nearly orthogonal to the (stationary) short run dynamics of the series, all other important specification choices become irrelevant. In particular, the exact measurement of the variables of the VAR, presence of omitted variables, the specification of insufficient lags, the use of inappropriate standard error bands and identification schemes, and debatable choices about the presence or absence of unit roots in the variables of the system are unimportant in characterizing the responses of hours worked, to both neutral and investment specific shocks.

5 Some discussions

The evidence we have collect strongly suggests that positive neutral shocks induce a fall in per-capita hours and, for most of the cases, we find that the effect is persistent. These shocks have an important effect on output volatility but they appear to be a minor source of hours variations at business cycle horizons. On the other hand, positive investment specific shocks tend to increase per-capita hours worked and generate a hump after about 4 quarters. Furthermore, they are important sources of hours but not of output variations at business cycle horizons. Altig, et. al. (2005) have estimated the effects of neutral and investment specific shocks on a number of US variables using a slightly different sample (1959:1-2001:4) but the same long run methodology to identify the two shocks. They report that the contribution of both shocks to per-capita hours and output volatility at business cycle frequencies is roughly the same (about 15 percent). The two sets of results are not necessarily in contrast which each other: in fact, they report percentages obtained on average at the business cycle frequencies of the spectrum, while here we report percentages obtained on average at business cycle horizons. Hence, the larger explanatory power that neutral shocks have for output and that investment specific shocks have for per-capita hours must come from either the high portion or the low portion of the spectrum of these variables.

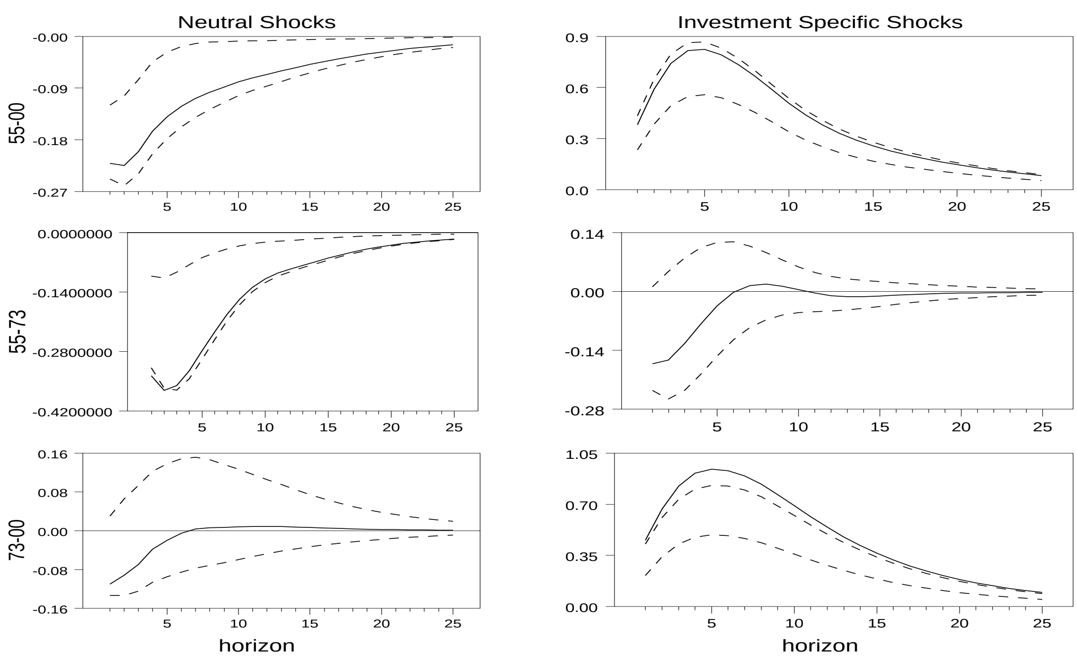

Figure 11: Structural responses, different samples, dummy specification

One question of interest is whether there are time instabilities in the response of hours to technology shocks. The subsample evidence presented in figures 2 and 3 was unintepretable because of the potential interactions between small sample biases and improper treatment of low frequency movements in the per-capita hour series. Clearly, small sample biases will always be present in short subsamples. Nevertheless, to the extent that low frequency movements are now properly taken into account, subsample analysis could be more informative. Figure 11 reports some subsample evidence. The responses of per-capita hours to neutral shocks are negative in both subsamples and positive; they are significant and somewhat persistent in the 1955-1973 sample but always insignificant in the 1973-2000 sample. The responses to investment specific shocks are positive and significant in the 1973-2000 sample and are similar to those obtained in the full sample. For the 1955-73 sample, the responses are insignificant at all horizons. Interestingly, while subsample estimates of the hours responses to neutral shocks are different from those reported in figure 2, estimates of the response to investment specific shocks are very similar to those of figure 3. Hence, while low frequency movements in hours matter for the dynamics in response to neutral shocks, they are apparently unimportant in deciding the sign and the significance of the responses to investment specific shocks.

Overall, the evidence suggest that there are potentially interesting changes in the business cycle relationship between hours and technology shocks, which demand further investigation.

6 How do technology shocks look like?

Technology shocks are often hard to interpret and they are even more so when they are characterized as unit root processes - at each point in time the probability of a technological regress is non-negligible even if a drift is allowed. We have shown that the shocks we have extracted are less than the usual black-box disturbances as they do not correlate with variables potentially omitted from the specification and they do not necessarily stand-in for other sources of structural disturbances. Nevertheless, to be able to understand the transmission pattern they induce, it is necessary to study in more details their properties. In particular, one would like to know what is their time series profile, if they display ups and downs which roughly match the ups and downs of US business cycles, etc.

To start with we present in figure 12 plots of the (smoothed) estimated technology shocks together with NBER recession episodes. Three stark features are evident. First, the time series properties of our neutral technology shocks match pretty well NBER phases. In particular, the series displays deep throughs which are typically coincident with start of NBER recessions. Second, the patterns of ups and downs in the price of investment shocks only partially coincide with standard NBER classifications. In general, they show somewhat higher frequency movements and fail to display any special pattern in the 1990-1991 and the 2000 recessions. Third, if one excludes the 1974 episode, the volatility of the two shock series is comparable. Hence, while the neutral shocks look like business cycle disturbances, this is not necessarily the case for investment specific shocks.

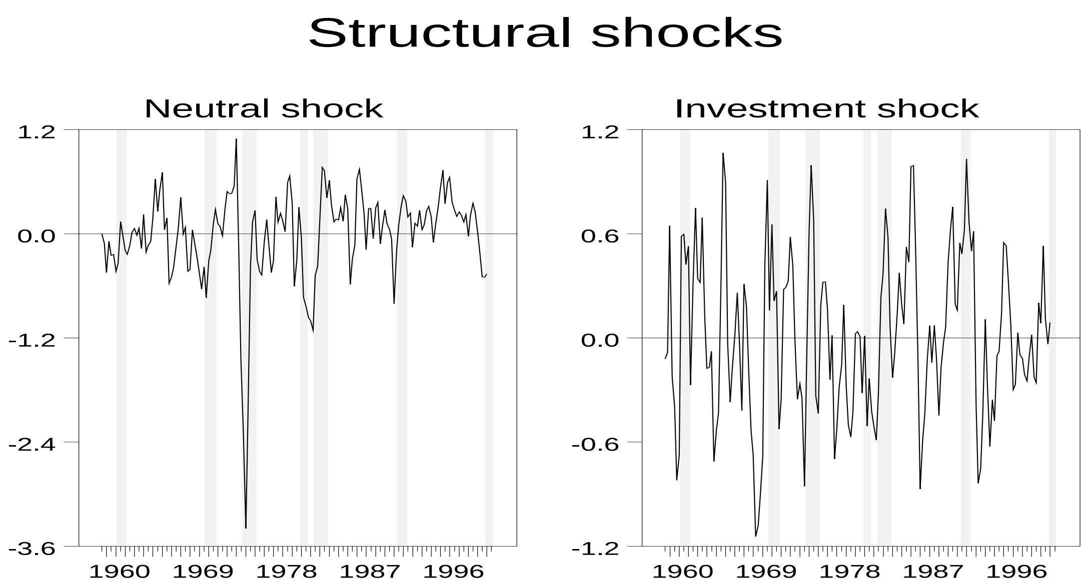

Figure 12: Structural shocks

Beaudry and Portier (2005) have shown that there is a surprising and almost perfect correlation between neutral technology shocks identified with long run restrictions in a bivariate model with total factor productivity (TFP) and stock prices and what they call ”news” shocks, i.e. shocks which do not generate any contemporaneous effects on TFP productivity, but instantaneously affect stock prices. How are our technological disturbances related to news shocks? Is there a sense in which technology shocks have a delayed effect on productivity so that stock prices shocks rationally capture this effect? We graphically show our two technology shocks and the news shock of Beaudry and Portier in figure 13. Even without fancy statistics, it is clear at that the correlation between the two series is far from perfect and, if anything, it looks as if that they are moving in opposite direction, especially at NBER recession dates. For example, in the 1974 recession, news shocks are positive, while both our neutral and investment specific shocks are strongly negative and, even looking forward, news shocks do not seem to capture any of the dynamics of our two technological shocks. In fact, the regression line between neutral and news shocks has a slope equal to -0.12 (t-stat=-3.03) and the slope does not change magnitude if we lag the news shock series up to 16 quarters while the slope between investment specific shocks and news shocks is -0.11 (t-stat= -3.15) and its magnitude falls if we lag the news shock series up to 16 quarters.

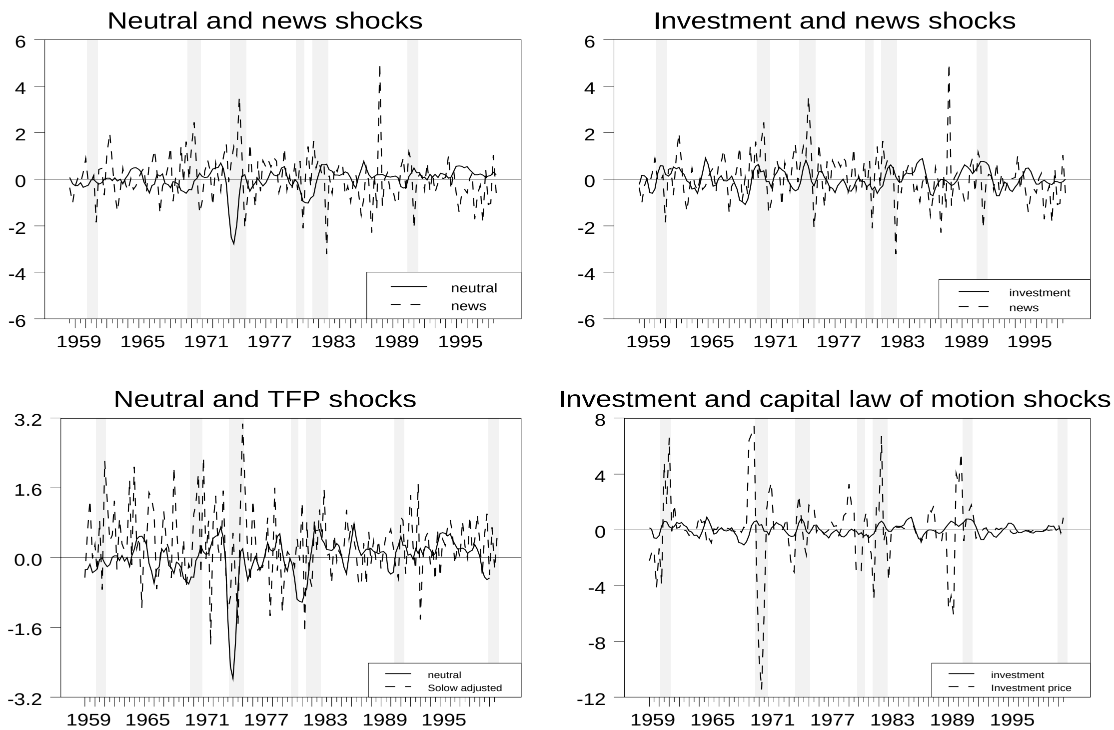

Figure 13: Structural and other shocks

It is also useful to compare our estimated neutral technological disturbances with Solow residuals shocks constructed using a standard Cobb-Douglas production function, adjusting for capacity utilization, and standard estimates of the capital and labor share. In fact, using the definition of labor productivity and a production function, we have that $\frac{y_t}{n_t} = (\frac{ku_t k_t}{n_t})^{\alpha} z_t$, where z_t is a measure of total factor productivity (TFP) , n_t is private nonfarm business sector hours, ku_t is capacity utilization, and k_t are capital services (both of which are obtained from the US Bureau of Labor Statistics). In the bottom left corned of figure 13 we plot our neutral shocks and TFP shocks obtained differencing the estimated TFP series. It is clear that the correlation of the two series is low (the maximum value occurs contemporaneously and it is only 0.19, that the innovations in the estimated TFP display higher volatility

(0.79 vs 0.53), that the majority of this volatility seems to be concentrated in the high frequencies of the spectrum, and that TFP shocks are positive at some NBER recession dates. Hence, our neutral shocks seem to have much better features than standard TFP shocks.

Finally, our price investment shock series can compared to a series for disturbances to the price of investment reconstructed using (interpolated) quarterly measures of the capital stock, a quarterly depreciation rate and the available investment series. In fact, $k_{t+1} = (1-\delta)k_t + v_t i_t$ where v_t is the inverse of the price of investment. We plot our price of investments shock series and the difference version of the reconstructed series for v_t in the lower right panel of figure 13. As it was the case with the neutral shocks, the volatility of the reconstructed series is much higher than the volatility of our estimated price of investment shock series (2.31 vs 0.40). The two series are negatively correlated at leads and lags (maximum effect -0.34 at lead 4) and both seem not to be in phase with the last two NBER recessions.

In sum, the evidence of this section suggests three important facts. First, the technology shocks one extracts from bivariate systems, TFP data and standard approaches are substantially different from the shocks we have obtained. Second, our shocks imply TFP and capital services series which are also considerably different from those one would obtained using standard accounting exercises. Third, while our neutral shocks look like business cycle shocks, and the investment shocks have a somewhat cyclical behavior, the technology shocks obtained in standard decompositions have uninterpretable time series patterns. As a consequence, the responses of per-capita hours to these shocks are also economically uninterpretable.

7 Conclusions

In this paper we have examined the effects of technology on per-capita hours carefully examining a number of issues which make the empirical evidence typically difficult to interpret. In particular, we dealt both with the presence of low frequency movements in hours and the potential misspecifications created by the use of VARs with a limited number of variables and a finite number of lags and by an incorrect use of bootstrap methods to construct confidence intervals for impulse responses; we have examined the robustness of the conclusions to alternative identification approaches and to the measurement of the variables of the system; and we have separately analyzed the dynamics induced by neutral and investment specific shocks. We show that if one takes reasonable care of the low frequency movements in hours worked, all the other pieces of the puzzle become irrelevant: regardless of the lag length of the VAR, the presence (or absence) of omitted variables, the identification scheme, the exact measurement of the variables used and other auxiliary statistical assumptions one is forced to make in specifying the VAR, hours fall in response to neutral technology shocks and increase in response to investment specific shocks and that the effects are persistent. We also demonstrate that the contribution of neutral technology shocks to the fluctuations in hours worked is typically small, while the contribution of investment specific shocks is substantial. Interestingly, the relative importance of the two shocks for output fluctuations is reversed: neutral shocks explain about twice as much as the forecast error variance of output than investment specific shocks at all horizons. We show that the time path of our estimated neutral shocks has peaks and troughs in correspondence of NBER peaks and throughs, while this is not necessarily the case for investment specific shocks. Also, we show that neither of our two shocks is correlated with news shocks and it does not stand in for any other likely sources of technological and non-technological disturbances.

While we have argued that it is unwise to use the evidence on the dynamic response of hours to technology shocks to disentangle sources of cyclical fluctuations, it is also true that our investigation offers some important insight for this theoretical debate. First of all, it appears that shocks that drive fluctuations in hours per-capita and in output are distinct: both simple RBC and New-Keynesian models typically fail to display this feature and therefore are inappropriate to interpret the complexities of the data. Second, technological disturbances may robustly drive up or down hours worked depending on their nature. Once again, simplistic models which do not distinguish between neutral or investment specific shocks can not be used to interpret the data. Third, while neutral shocks appear to have a marked cyclical pattern, investment specific shocks tend to display more high frequency fluctuations. All in all, while one can not deny that the combined effect of technology shocks on hours worked and output is significant, neither the simplest RBC nor the simplest New-Keynesian models can fully explain the evidence we have presented.

There are many dimensions along which our work could be extended or qualified. For example, although visual evidence suggests that the cyclical component of per-capita hours is stable, one may want to study the responses to technology shocks in systems where both the intercept and the slope are allowed to vary over time (as in Gambetti (2005)). This could be important in understanding the dynamics of macro variables in the late 1990s, where both labor productivity and hours displayed unusual features. One could also extend our analysis to consider both the intensive and the extensive margins of labor market (hours vs employment responses) and the flows in and out of employment (finding and separation rates). Evidence on these issues can help us to determine which theoretical story of transmission is more appropriate to explain the facts. Finally, the fact that the cyclical response of per-capita hours and output to the two shocks have different magnitude suggests the presence of an important cyclical wedge between the two. Providing a convincing theoretical reason for this wedge is a challenging but probably rewarding exercise.

References

[1] Altig, D. Christiano, J. Eichenbaum, M. and Linde, J. (2005), “Firm-Specific Capital, Nominal Rigidities and the Business Cycle”, NBER Working Paper 11034.

[2] Beaudry, P. and Portier, F. (2005), “ Stock prices, news and economic fluctuations, forthcoming, *American Economic Review*

[3] Basu, S, Fernald, J, Fisher, J. and Kimball, M. (2005) ”Sector Specific Technological Change ”, Boston College, manuscript.

[4] Chari, VV, Kehoe, P. and McGrattan, E. (2005) A critique of SVAR using Real Business Cycle theory, Fed of Minneapolis, working paper 631.

[5] Christiano, L, Eichenbaum, M. and Vigfusson, R.(2003) ”What happens after a technology shock”, Northwestern university, http://www.faculty.econ.northwestern.edu/faculty/christiano/research.htm.

[6] Dedola, L. an Neri, S. (2004) What does a technology shock do? A VAR Analysis with model-based sign restrictions, EABCN discussion paper.

[7] Erceg, C., Gorrieri, L. and Gust, C. (2005) ”Can Long Run Restrictions Identify technology shocks”, forthcoming,*Journal of the European Economic Association*.

[8] Evans, C. (1992) ”Productivity Shocks and Real Business Cycles”, *Journal of Monetary Economics*, 29, 191-208.

[9] Fernald, J. (2004), “Trend Breaks, Long Run Restrictions, and the Contractionary Effects of Technology Shocks”, Federal Reserve Bank of San Francisco, working paper 2005-21..

[10] Fisher, J. (2003), “Technology Shocks Matter”, Federal Reserve Bank of Chicago, working paper 2-14

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

[12] Francis, N. and Ramey, V. (2006), “Measures of per-capita hours and their implications for the technology-hours debate”, University of North Carolina, manuscript.

[13] Faust, J. and Leeper, E. (1997) "Do Long Run Restrictions give reliable Results",*Journal of Business and Economic Statistics*, 15, 343-353.

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

[15] Galí, J. (2005), "Trends in hours, Balanced Growth, and the Role of Technology in the Business Cycle ", *Federal Reserve of St. Louis Review*, July-August, 459-486.

[16] Galí, J. and Rabanal, P., (2004), "Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data?", *NBER Macroeconomics Annual*, 225-288.

[17] Gambetti, L. (2005) Technology shocks and the response of hours worked: time varying dynamics matter, http:// www.econ.upf.es/gambetti.

[18] Giordani, P. (2004) An alternative explanation of the price puzzle, *Journal of Monetary Economics*, 51, 1271-1296.

[19] Greenwood, J. Hercowitz, Z. and Krusell, P. (1997) "Long-Run Implications of Investment-Specific Technological Change", *American Economic Review,* 87, 342-362.

[20] Greenwood, J. and M. Yorokoglu (1997), "1974", Carnegie-Rochester Conference Series on Public Policy 46, 49-95.

[21] Kilian, L. (1999) "Finite sample properties of bootstrap and percentile t-bootstrap confidence interval for Impulse responses", *Review of Economics and Statistics*, 81, 652-660.

[22] Pappa, E. (2004) The labor market effects of fiscal shocks: New Keynesian or RBC transmission?, IGIER working paper 293.

[23] "Uhlig, H. (2004) Do technology shocks lead to a fall in total hours worked" *Journal of the European Economic Association,* 2, 361-371.

[24] Rubio, J, Sargent, T and Villaverde, J (2005) The ABC (and D) for understanding VARs, Federal Reserve of Atlanta, working paper 2005-9.

[25] Vigfusson, R. (2004) The delayed response to technology shock: a flexible price explanation, Federal Reserve Board, International Finance discussion paper 810.