

The effects of technology shocks on hours and output: A robustness analysis*

Fabio Canova
ICREA-UPF and AMeN

David Lopez-Salido
Federal Reserve Board

Claudio Michelacci[†]
CEMFI

This version: February 2008

Abstract

We analyze the effects of neutral and investment-specific technology shocks on hours and output. Long cycles in hours are captured in a variety of ways. Hours robustly fall in response to neutral shocks and robustly 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 for investment specific shocks. Finally, we show that ‘*news shocks*’, i.e. those that generically change expectations about future productivity, are uncorrelated with the estimated technology shocks.

JEL classification: E00, J60, O33.

Key words: Technology disturbances, Structural VARs, Long cycles, News shocks.

*We would like to thank three anonymous referees, the editor of this journal, Thijs Van Rens and Albert Marcet for comments and suggestions; and Neville Francis, Susantu Basu, Sergio Rebelo and Franck Portier for kindly sharing their data. Views expressed are the authors’ and should not be interpreted as those of the the Federal Reserve System, or the Board of Governors.

[†]All authors are also affiliated to CEPR.

1 Introduction

Following the work of Galí (1999, 2005), Christiano et. al. (2003), Uhlig (2004), Francis and Ramey (2005), and more recently Dedola and Neri (2007), there has been a renewed interest in empirically examining the effects of technology shocks on total per-capita hours. This interest is typically motivated by the fact that the dynamics of hours differ if such disturbances occur in a basic RBC model (where hours increase) or in a basic sticky-price model (where hours decrease).

Unfortunately, the available evidence is, at best, mixed since there are a number of reasons that make such a task of applied researchers complicated. First, technology shocks are identified in the data using long run restrictions (see e.g. Galí (1999)). But this form of identification is weak, in the sense of Faust and Leeper (1997), and available samples may be too short or unstable to credibly impose such restrictions (see Erceg, et. al. (2005)). Furthermore, other primitive shocks may have similar long run features as technology shocks (see Uhlig (2004)). Second, different types of technological disturbances may have different effects on hours, making the outcomes of bivariate models, where only one generic technology shock is identified, uninterpretable (see Fisher (2006) and Michelacci and Lopez-Salido (2007)). Third, as recently emphasized by Canova, Lopez-Salido, and Michelacci (2006) and Kehoe and Ruhl (2007), the choice of price deflators may matter for identifying the effects of technological disturbances. Fourth, the response of hours appears depends on a number of auxiliary statistical assumptions, including the treatment of long cycles in hours, the lag length of the empirical model and the horizon at which the identifying restrictions are imposed. Finally, the recent evidence provided by Beaudry and Portier (2005), where generic shocks that change expectations about the future (what they call ‘news shocks’) are the same as the identified technology shocks, makes the interpretation of the latter puzzling.

This paper empirically examines the effects of technology on hours and output addressing these issues in an unified and comprehensive way. In particular, we remove long cycles in hours in a number ways; we separately analyze the dynamics induced by two different technology shocks (neutral and investment specific) and examine their relationship with other shocks identified in the literature. We also

deal with the potential misspecification created by VARs with a limited number of variables and a finite number of lags via a simple prior and we study the robustness of the conclusions to alternative identification approaches and different measures for the price deflators.

We find that once one removes long cycles in hours, all the other pieces of the puzzle become irrelevant; regardless of the lag length, the choice of price deflators, the identification scheme, the presence of omitted variables and other auxiliary statistical assumptions one is forced to make in specifying the VAR, hours robustly fall in response to neutral shocks while they robustly increase in response to investment specific shocks. We find that the contribution of neutral shocks to hours fluctuations is 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 do, at all horizons. The estimated neutral shocks occur at non recessionary periods, while investment specific shocks fail to display significant cyclical features. Our technology shocks are uncorrelated with potentially important omitted variables; they do not stand in for other likely sources of disturbances; they are unrelated to the news shocks identified by Beaudry and Portier (2005), and they differ from the technology shocks one would extract from standard accounting exercises.

Our results complement and qualify several other contributions of the literature. Regarding the issue of long cycles in hours and how they should enter the VAR, two contrasting arguments are typically made. If one conditions the analysis on the models used to interpret the results—Christiano et. al. (2003), Uhlig (2004), Dedola and Neri (2007) —per-capita hours should enter the empirical model in level, since basic RBC and New-Keynesian models produce stationary hours fluctuations, even when technology is non-stationary. If one conditions the analysis on the statistical properties of the data and follows a classical statistical approach —Galí (1999), and Francis and Ramey (2005) —the VAR should include hours in differences. Figure 1, which shows that the standard per-capita hours series displays long but essentially stationary cycles, indicates that both approaches are potentially subject to specification errors. The cycles in hours are of longer duration than those considered in

the business cycle literature reflecting e.g. demographics, long run trends in labor market participation or R&D activities. This paper argues that disregarding them (as one would do by taking hours in levels) or by taking a rough short cut (as one would do by differencing the series) lead to misspecification, inefficiency losses, and potentially uninterpretable results. This view leads to conclusions different from those by Francis and Ramey (2005) who argue that one should find economic reasons leading to a unit root in per-capita hours in order to justify the first differencing of hours.

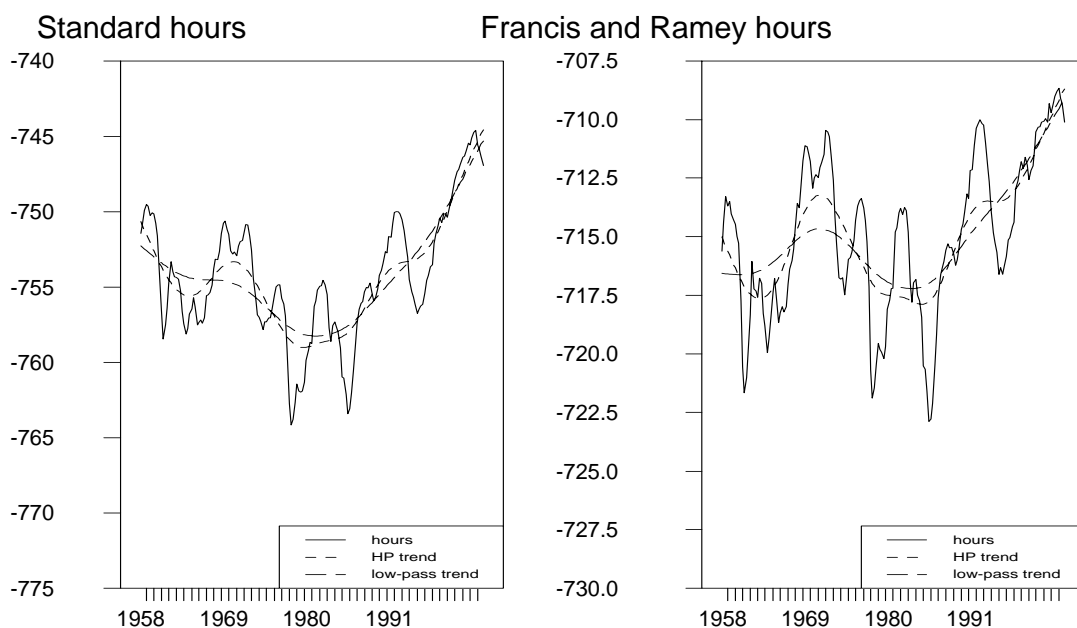


Figure 1: Per-capita hours, HP($\lambda = 6400$) and low pass estimates.

Fernald (2007) has also stressed that, in order to correctly recover the effects of technology shocks on hours, it is important to use time dummies to characterize possible shifts in the variables of the VAR. Here, we provide a robustness analysis—using various measures of hours, different specification choices, medium versus long run restrictions, and several alternative detrending procedure— we analyze the effects on hours by allowing for two different technology shocks, and we examine the

time series features of the estimated technology shocks and their relationship with some interesting economic shocks (including news shocks).

Our evidence is hard to reconcile with simple flexible price and sticky price models.¹ Standard models predict that the percentage of hours and output fluctuations explained by technology shocks should be similar. In Canova et al. (2006) we show that the sign and magnitude of the output and hours responses to technology shocks we report here, and the sign and magnitude of the unemployment and labor market flows responses we report there, are instead consistent with a model of creative destruction, where improvements in the neutral technology trigger adjustments along both the intensive and the extensive margin of the labor market.

The rest of the paper is organized as follows. Section 2 critically summarizes the literature. Section 3 interprets the existing evidence and presents new results when long cycles in hours are accounted for in a variety of ways. Section 4 examines the robustness of the results to a number of changes in the auxiliary assumptions. Section 5 discusses the properties of the estimated technology shocks. Section 6 concludes.

2 The existing evidence

We summarize the current state of the debate in the context of a three equation VAR with labor productivity, the price of investment, both measured in consumption units and per-capita hours. All variables are in logs. A trivariate system is the minimum size system required to recover neutral and investment-specific shocks. Since such a small model is liable to specification errors and we later show how to check for potential omitted variables. The sample runs from 1955:1 to 2000:4—a consistent price of investment series is available only up to that date. 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 in levels (the level system). Investment specific shocks are identified by the requirement that they are the sole source of long run movements in the price of investment while neutral shocks can affect both labor productivity and the price of investment in the long

¹See for example Galí (2005) and Altig et al. (2005).

run. Fisher (2006) has shown how to derive these identifying restrictions in models of both neoclassical and New-Keynesian orientation where the neutral technology and the price of investment are both stochastic and display a unit root. We use 12 lags of each variable in each system and stochastically restrict their decay toward zero assuming that the prior variance of lag j is proportional to j^{-2} . Rather than using a standard lag selection criteria, we allow a generous lag length and a prior to reduce overparametrization, in order to avoid the problems emphasized by Giordani (2004), Chari et. al. (2005), and Fernandez-Villaverde et. al. (2007), who show that a subset of the variables generated by standard models may display decision rules not always representable with a finite order VAR—a problem that may be severe in a three equation system. Unless otherwise stated, all the figures report the point estimate of the dynamic responses and a 90 percent small sample confidence tunnel. When the point estimate lies outside or at the boundary of the tunnel, small sample biases are likely to be important.

The first two boxes of the first row of figures 2 reproduce the response of hours to a neutral shock in the full sample. As documented in the literature, in the level system, per-capita hours positively respond to a neutral shock, the maximum response is delayed by about 5 quarters and the instantaneous impact is insignificant. In the difference system, per-capita hours fall for up to 4 quarters and settle to their long run level from above, but all the responses are generally insignificant.

The sign difference in the point estimates obtained in the level and the difference specification is often attributed to long cycles in hours and to the fact that these movements distort the dynamics in the level system. However, long cycles do not necessarily imply non-stationarity dynamics, as it is commonly assumed (see e.g. Galí (2005)); there could be stationary cycles with long but finite periodicity resulting in standard overdifferencing problems. In addition, since the difference specification emphasizes high frequency hours variability, the importance of measurement error could be magnified. Since the 90 percent tunnel for the difference system is large relative to the one for the level system, this problem is likely to be important. Hence, both systems are misspecified and it is difficult to draw credible conclusions about the conditional dynamics of hours from these two pictures.

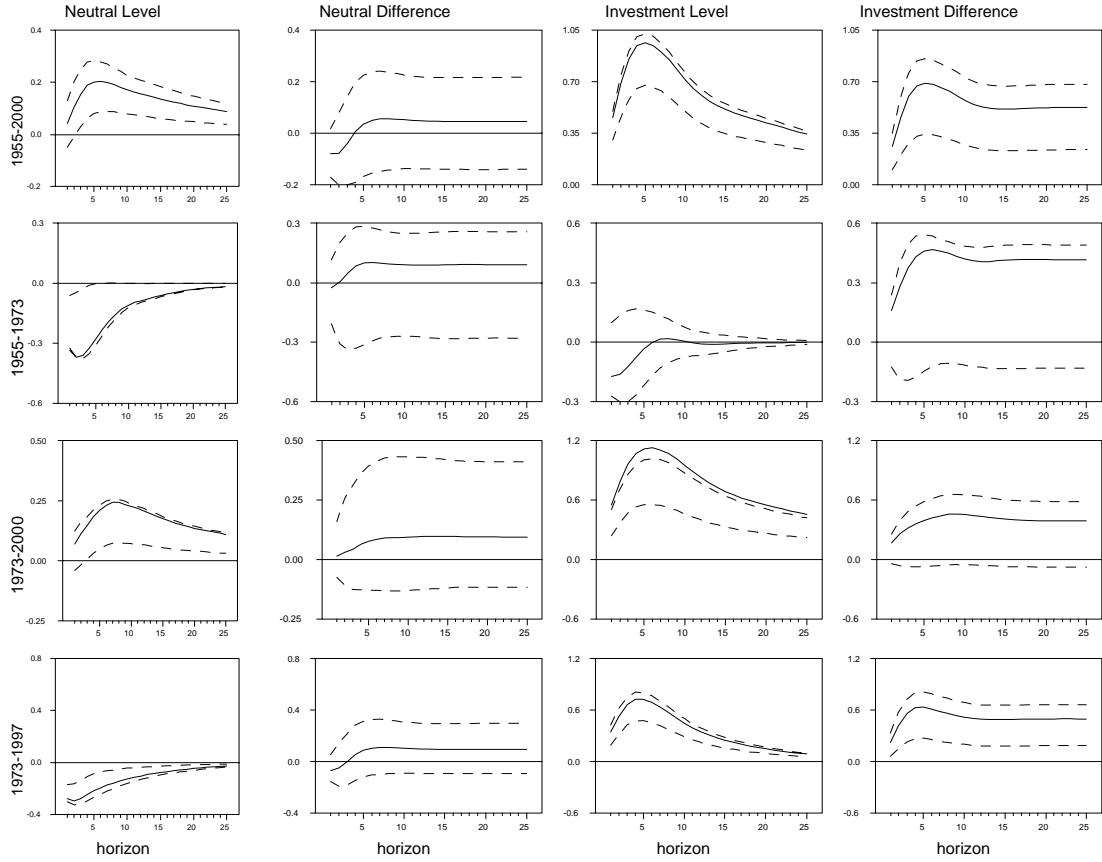


Figure 2: Responses of per-capita hours

One way to take care of long cycles in hours is to split the sample in pieces. We follow Greenwood and Yorokoglu (1997), Fernald (2007) and Canova et. al. (2006), who have suggested that 1973:2 and 1997:1 could be crucial break dates, and we split the sample accordingly (see second and fourth rows of figure 2). In figure 1 these are, approximately, the dates at which the Hodrick Prescott estimate of the long cycles has flex points.

In this period, the dynamics of hours in response to neutral shocks are similar in the two systems. Both in the 1955-1973 and in the 1955-1997 sample, hours instantaneously fall in the level system. The point estimate is persistently negative and significantly so up to four quarters, while for the difference specification, the point estimate turns positive after just two quarters, but remain insignificant at all

horizons. For the 1973-2000 sample, instead hours instantaneously increase both in the level and the difference specification (see third row).

The responses of hours to investment shocks are depicted in the last two columns of figure 2. For the full sample, the level and the difference 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, results differ across specifications. In the level system, hours fall in the 1955-1973 sample, increase in the 1973-2000 and in the 1973-1997 samples, but are significant in the 1955-1973 sample. For the difference system, the instantaneous point estimate is positive in all three subsamples, but responses are significant only in the 1973-1997 sample.

The subsample instability found in the hours responses to both shocks in the level system is consistent with the idea that the relationship between per-capita hours and technology shocks may be changing over time. Subsample evidence, however, is difficult to trust because splitting the sample in pieces introduces large small sample biases. Small sample bias make estimates unreliable for three reasons. First, with small samples, tunnels become larger making point estimates less informative (compare across rows in figure 2). Second, using long run restrictions in a system estimated over a small sample is likely to induce distortions in the structural estimates (see Erceg, et. al. (2005)). Third, small sample bias may interact in an unpredictable way with measurement and aggregation errors, making subsample evidence uninterpretable. Interestingly, no sub-sample instability is present in the difference specification, but here the estimated responses are generally little significant. In sum, neither assuming unit roots nor splitting the sample in pieces seems to be the best way to account for the long cycles in per-capita hours.

3 Long cycles and intercept heterogeneities

The instabilities and sign reversals one finds in columns 1 and 3 of figure 2 (that correspond to the level specification), and the substantial homogeneity in results of columns 2 and 4 (that correspond to the difference specification), are symptomatic of a particular type of data heterogeneity. We will show that sample instabilities are

due to low frequency comovements in the variables of the VAR, without any underlying change in the dynamic response of hours to shocks. For example the difference in the impulse response of hours to neutral shocks in the sub-sample 1973:1997 and in the sub-sample 1973:2000 can be attributed to the highest average productivity growth and hours experienced by the US economy during the productivity revival of the late 90's. If not properly taken into account, these permanent shifts in averages tend to lead to biases in the estimated responses. In the VAR a permanent change in the rate of productivity growth is at least partly identified as a series of neutral technology shocks. Thus, in the sample 1973-2000 when productivity growth is higher than average (as it is in the 1997-2000 period), the 1973-2000 specification identifies a series of positive neutral technology shocks. Since in this period hours are also above average, bias emerges leading to a higher response of hours. To better illustrate the source of bias consider the three clouds of points presented in the left-hand side box of figure 3 which are intended to represent the 1973-1997, the 1955-1973 and the 1997-2000 period, respectively. Across subsamples there is no heterogeneity in the slope of the relationship between the variables on the x and y axis, that are intended to represent the neutral technology component of hours and productivity, respectively. In the first sample, the intercept is positive and large. In the second and third, the intercept is still positive but much smaller. In all three samples, the slope of the relationship is negative and approximately constant.

It is clear that if we pool the first and the third sample together, the slope of the relationship becomes positive (think of this as the subsample 1973-2000) and if we pool together the three samples, the slope of the relationship is again positive. This is exactly the contemporaneous pattern we find in the level system: the 1955-1973 sample give a negative response of hours to neutral shocks, the 1973-2000 and the full sample a positive response. If we difference the variables on the x and y axis (see right hand side box of figure 3), we remove the effects of the mean shifts, independently of the sample we are considering, but the estimate of the magnitude of the slope becomes more noisy due to the magnification of measurement error.




Figure 3: Interpretation of the evidence.

In such a situation, splitting the sample and tracing out the responses separately in each of them is inefficient, as there appears to be little changes in the structural responses, and small sample bias become relevant. To efficiently tackle this issue, we have considered several options. In the first case, the intercept in the per-capita hours equation 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 polynomial). In the third case, we clean the per-capita hours series with a one-sided low-pass filter, which takes away cycles with periodicity higher than 52 quarters (the filter is one minus the low pass filter plotted in figure 1). Finally, in the fourth case the intercept drifts stochastically and potentially continuously over time. In this latter case, we specify an autoregressive mean-reverting law of motion and we use the Kalman filter to recursively estimate it. Note that with all four specifications, small sample bias induced by the use of long run restrictions are considerably reduced, since the full sample of quarterly data is now employed to project estimated VAR coefficients infinitely far into the future.

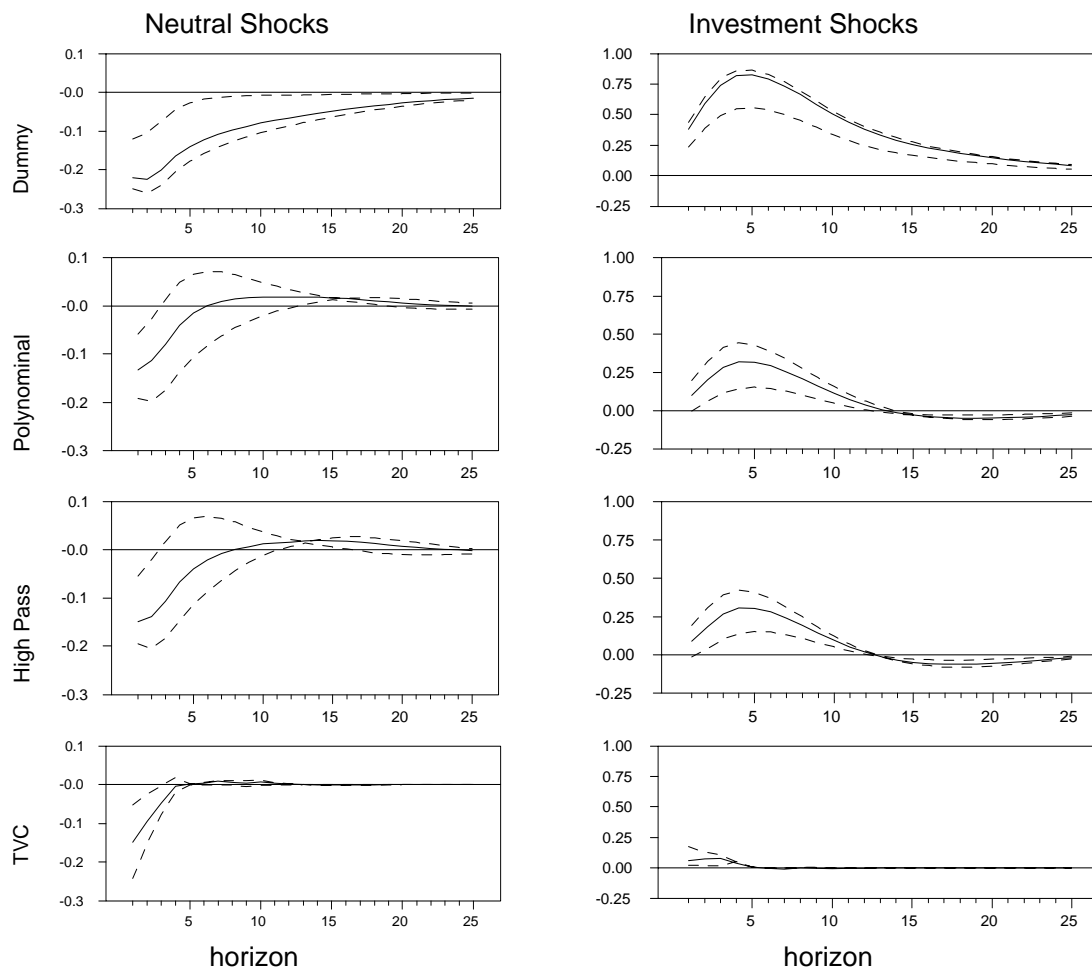


Figure 4: Responses of hours, sample 1955-2000

We plot the responses of per-capita hours in these four specifications in figure 4. The first column refers to neutral shocks and the second to investment shocks. The results are very robust across methods; per-capita hours fall in response to neutral shocks and increase in response to investment specific shocks and the instantaneous response is always significant. Depending on the exact specification, the fall in response to neutral shocks is either persistent (first row) or temporary (next three rows). Note that the hump after 4-5 quarters in the hours response that was present in figure 2—which has been greatly emphasized by Vigfusson (2004)—is no longer

present, at least in response to neutral shocks.

The percentage of the variance of per-capita hours explained by the two shocks is similar in the four specifications; 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 below 10 percent), while investment specific shocks explain between 30-50 percent of the variance of hours at these horizons (20-30 percent with the time varying intercept specification). Interesting, this ordering is reversed for output fluctuations; 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.

Altig, et. al. (2005) have estimated the effects of neutral and investment specific shocks using our same identification approach but a slightly different sample (1959:1-2001:4) and without dealing with the long cycles in hours. 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). However, their numbers are percentages obtained on average at the business cycle frequencies of the spectrum, while here we report percentages obtained on average at business cycle horizons. Moreover, they do not deal with the long cycle in hours, that we show is key to identify the effects of technology shocks.

One question of interest is whether there are instabilities in the responses of hours to technology shocks with any of the specifications in figure 4. The subsample evidence presented in figure 2 was uninterpretable because of the potential interactions between small sample biases and improper treatment of data heterogeneity. Clearly, small sample bias will not disappear. Nevertheless, to the extent that data heterogeneity is now properly taken into account, subsample analysis should be more informative. For the sake of space, figure 5 reports results for the dummy specification. The responses of per-capita hours to neutral shocks are now instantaneously negative in all sub-samples, and they are significant and somewhat persistent in the 1955-1973 and 1973-1997 samples. The responses to investment specific shocks are positive and significant in the 1973-1997 and 1973-2000 samples and similar to those obtained in the full sample. For the 1955-73 sample, the responses are insignificant at all horizons. Hence, once data heterogeneities are taken

into account, the relationship between per-capita hours and technology shocks is relatively stable over time.

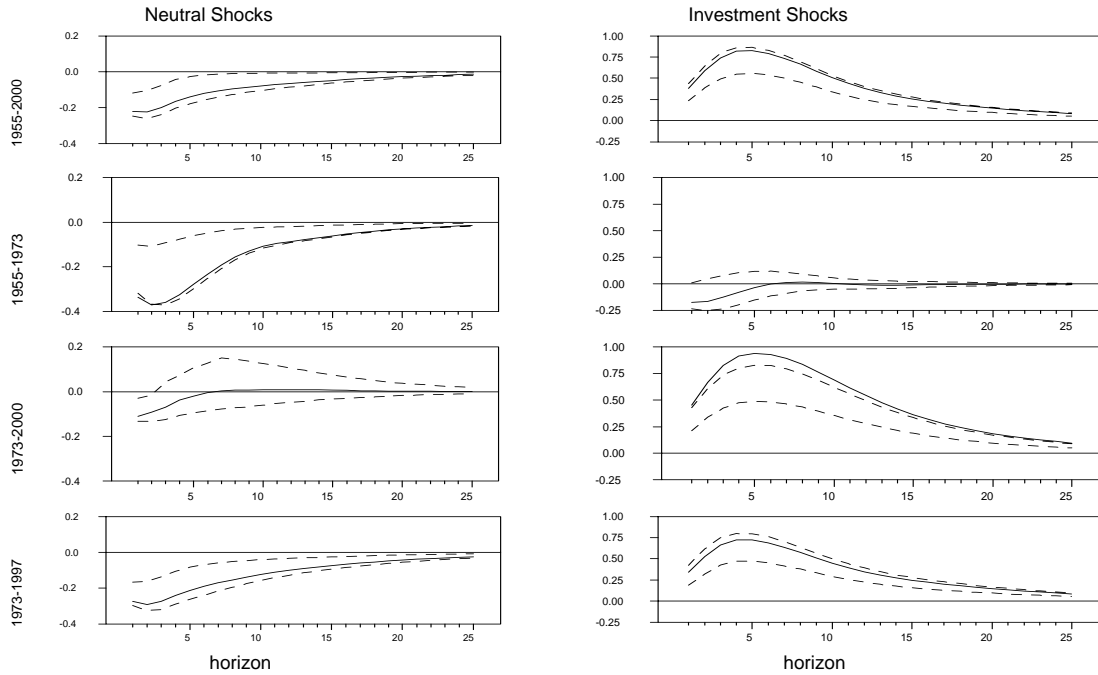


Figure 5: Responses of hours, different samples, dummy specification

4 Robustness

There are many dimensions along which the robustness of the previous conclusions could be examined, making the combination of systems to be estimated quite large. We divide our analysis into three parts. First, we study robustness with regard to the choice of variables, their measurement, and the sample used. Second, we check whether our technology shocks stand-in for omitted variables or other measurable sources of disturbances. Third, we examine whether outcomes are sensitive to the statistical assumptions we have made. Overall, we find that our conclusions are quite robust. For the sake of presentation, we only report results for the dummy specification, as before the exact treatment of long cycles does not seem to matter.

In the systems we have run, labor productivity and the price of investment are measured in consumption units. However, as Canova et. al. (2006) have shown, if foreign goods enter the consumption basket, long run movements in labor productivity can also be driven by external shocks. Since the output basket is less prone to such problems, we have repeated the exercise measuring either labor productivity or both labor productivity and the price of investments with the output deflator. The first row of figure 6 shows that the sign and shape of per-capita hours responses to both shocks remain unchanged.

Next, we have repeated our exercises using the hours series of Francis and Ramey (2006). This series displays less of a trend than the standard one but, it is still not void of long cycles (see right box of figure 1). If we adjust the hours series, the labor productivity series needs to be adjusted as well to make the analysis consistent.² Given the similarities in the paths of the two hours series, the bias in labor productivity introduced by the lack of adjustment is unlikely to be important. The second row of figure 6 shows that, indeed, the qualitative features of per-capita hours responses are unaltered with this new series.

A referee also suggested that the trend in the late 1990s in the hours series could be the result of changes in working age population, and that the trend disappears when the hours series is scaled by labor force participation rather than by working age population. While this alternative series does not match the object of reference in theoretical discussions, which focus on a comprehensive measure of aggregate labor effort (including labor force participation), it is nevertheless worth repeating the estimation with it, since, while free of trends, such a series still displays significant long cycles. As shown in the third row of figure 6, the main features of hours responses are unchanged also with this measure.

Finally, all systems are estimated with data up at 2000:4 since the price of investment series terminates at that date. We have managed to obtain a newly constructed price of investment series, which splices the old series with new data on the price of investment up to 2004:4. The fourth row of figure 6 shows that updating the sample, produces no major changes in the estimated relationship.

²We thank one of the referees for this suggestion.

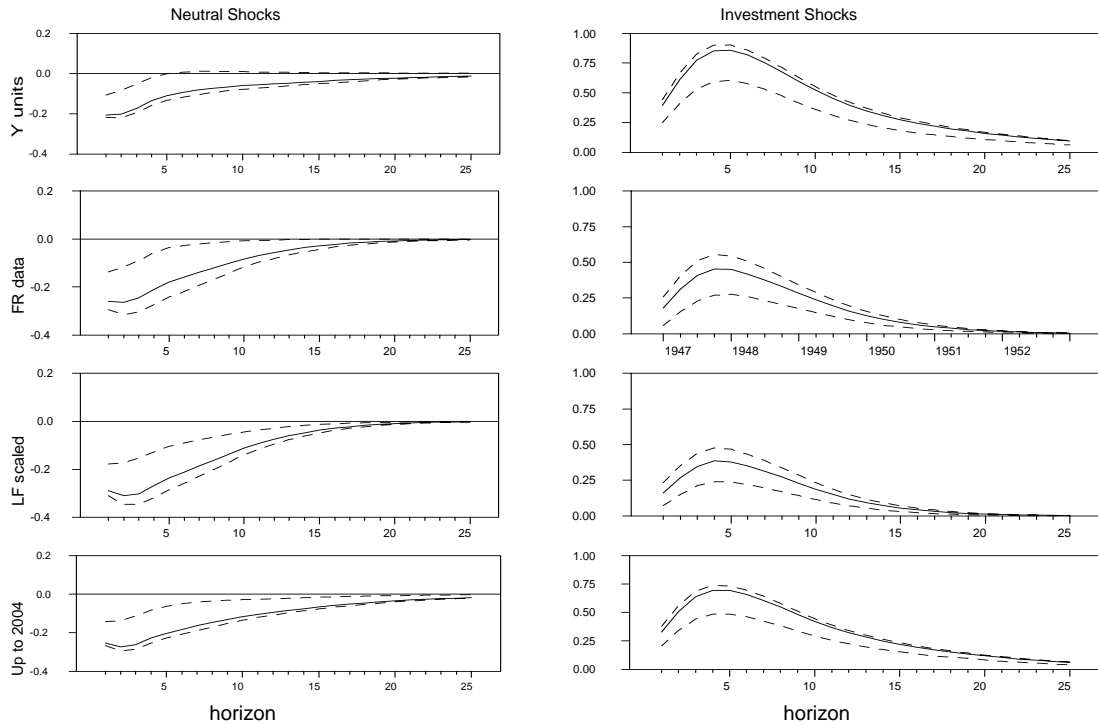


Figure 6: Responses of per-capita hours, Alternative measurements, Dummy specification, 1955-2000

In order to credibly assess the features and the properties of the estimated technology shocks, it is important to make sure that they are not standing-in for other variables or structural shocks. We have checked this in two ways. First, while we have allowed enough lags in each estimated specification to make structural residuals serially uncorrelated, it is always possible that, in a three variable system, omitted variables play a role. For example, Evans (1992) showed that Solow residuals constructed from production functions are correlated with a number of policy variables, therefore making responses to Solow residuals shocks uninterpretable. To check whether omitted variables play a role, we have correlated our two estimated technology shocks with variables which a large class of general equilibrium models driven by neutral and investment specific shocks suggest as being jointly generated with the data we have used. Figure 7 reports the cross-correlations of up to four leads and four lags of the estimated technology shocks with three of these variables

(consumption to output, investment to output, and inflation), and the upper and lower limits of an asymptotic 95 percent confidence tunnel for the null hypothesis of no cross-correlation. Clearly, all three variables fail to be strongly correlated with the estimated shocks. This outcome is confirmed by the results of bivariate Granger causality tests between the potentially omitted variables and the recovered technology shock; lags of consumption to output, investment to output, and inflation do not help in predicting structural residuals. Hence, it is very unlikely that omitted variables play a major role in explaining the results.




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

In addition, we have also correlated our estimated technology shocks with oil price shocks, federal funds futures (FFF) shocks and tax shocks. These disturbances are constructed as the residuals of univariate regressions of each of the three variables on two lags.³ The cross-correlations are all small and never exceed 0.11 in

³The effective tax series is taken from the Congressional Budget Office and is transformed into quarterly frequency using an interpolation routine.

absolute value when we consider up to 4 lags and 4 leads of the disturbances. Hence, our estimated technological shocks do not appear to stand-in for other sources of technological and non-technological disturbances.

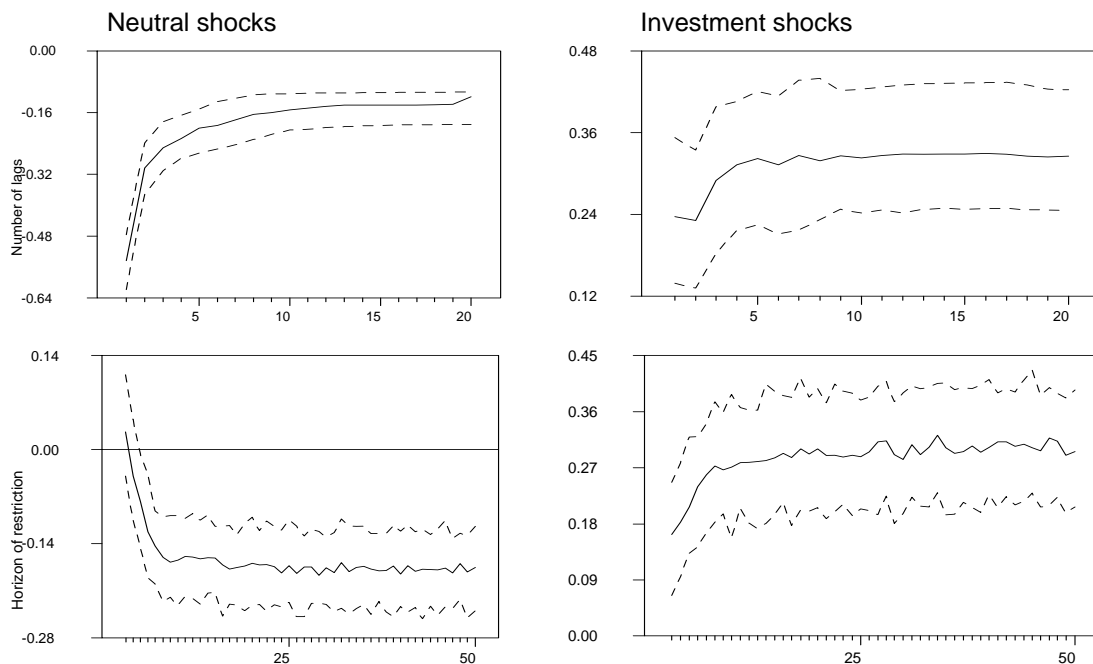


Figure 8: Contemporaneous response of hours, dummy specification, 1955-2000

Another way to examine the potential effect of omitted variables on the relationship between per-capita hours and technology shocks is to check the robustness of the results to changes in the lag length. To the extent that omitted variables result in VAR residuals with MA components, adding lags to the model should help to attenuate the problem. We report median estimates and small sample bands for the contemporaneous response of per-capita hours to the two shocks as the lag length changes in the first row of figure 8. 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 contemporaneous response to neutral shocks becomes greater in magnitude and remains strongly negative.

We have also checked whether the dynamics of hours are robust to the timing and the type of identification restrictions we used. Uhlig (2004) has forcefully argued that disturbances other than technology shocks may have long run effects on labor productivity and that, in theory, there is no horizon at which technology shocks fully account for the variability of labor productivity. This means that the technology shocks we have extracted may have little to do with technological change. Beaudry and Portier (2005) found empirical evidence consistent with this interpretation; technology shocks obtained with long run restrictions on productivity in a bivariate system are strongly correlated with shocks that generically change expectations about the future (what they call news shocks). To study whether this is a problem, we have imposed the restriction that investment specific shocks are the sole source of the fluctuations in the price of investment and that neutral and investment specific shocks are solely responsible for the fluctuations in labor productivity at horizon $k = 1, \dots, 50$. The second row of figure 8 shows the impact response of per-capita hours as the horizon changes: the sign of the response is robust to the horizon at which the restriction is imposed, and it is significant in all specifications, except for the neutral shock identified using horizons shorter than 3 quarters.

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 (2007), for example, assume stationarity of VAR variables and use sign restrictions derived from an RBC model to identify technology shocks. A previous version of the paper has also examined the dynamics of per-capita hours in responses to technology shocks identified via sign restrictions. None of the qualitative conclusions we reach is affected by using this alternative identification strategy.

In sum, if one takes the view that long cycles in per-capita hours can be characterized with time varying VAR intercepts, and that these movements are nearly orthogonal to the (stationary) short run dynamics of the series, all other important specification choices become irrelevant.

5 How do technology shocks look like?

Technology shocks are often hard to interpret, even more so when they are characterized as a unit root process since at each point in time the probability of a technological regress is non-negligible even if a positive drift is allowed. This is a possibility which most likely has never been experienced in the post WWII era in developed countries. 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 stand-in for other sources of structural disturbances. We now study their properties in more details.

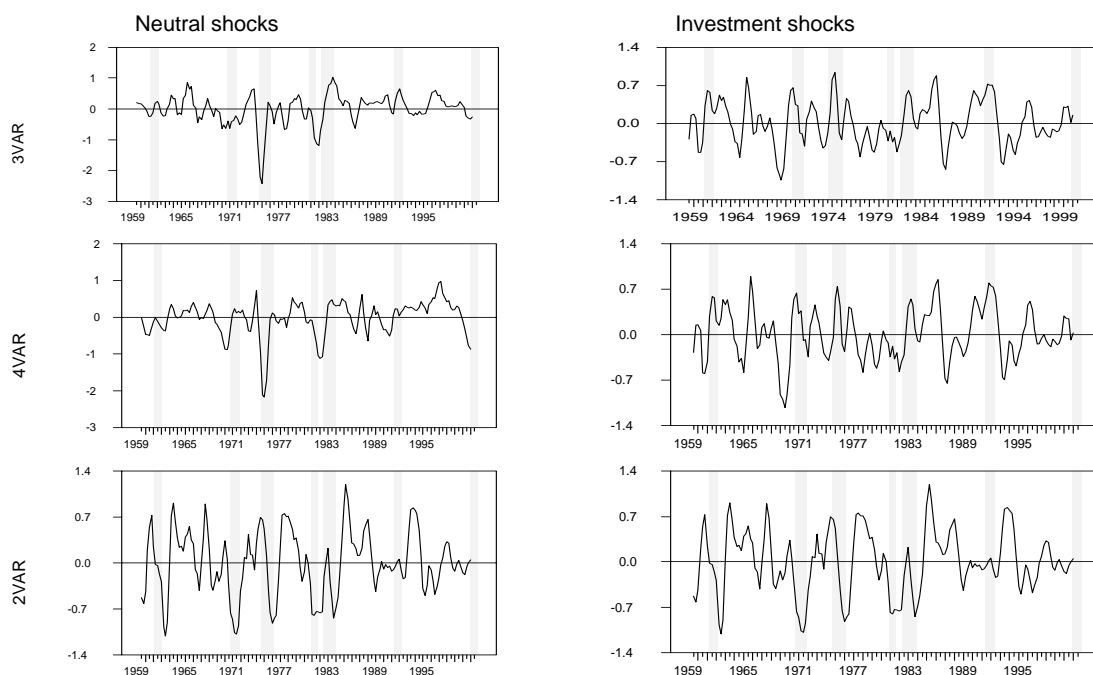


Figure 9: Technology shocks

The first row of figure 9 presents plots of the (smoothed) estimated technology shocks together with NBER recession episodes (shaded areas). Three stark features are evident. First, our neutral shocks display significant cyclicity. In particular,

the series displays troughs which are typically coincident with start of NBER recessions and peaks coincident with start of recoveries. Second, the pattern of ups and downs in our investment shocks only partially coincides with the standard NBER classification. Third, if one excludes the 1975 episode, the volatility of the two shock series is comparable.

Since the neutral disturbance features two deep troughs in 1975 and 1982 and, on average, hours fall in response to neutral shocks, one may be lead to conclude that the two major post WWII recessions were periods where per-capita hours boomed! Such a conclusion is incorrect since, as figure 1 shows, hours fall during these recessions. To assess the role of the two technology shocks in shaping hours fluctuations in specific episodes we present in figure 10 an historical decomposition. Each panel of the figure reports the actual series and the counterfactual series that would have been generated since 1974:1 had only neutral or investment specific shocks being present. Three features are evident. First, the effect of the two shocks on per-capita hours depends on time and on the state of the economy. Second, neutral shocks generate minor fluctuations in per-capita hours, except for the late 1990s. In particular, they generate no fluctuations in the 1975 recession and only a little fall in the 1982 recession. On the contrary, it is at times when per-capita hours peak that neutral shocks induce opposite movements in the counterfactual hours series. Third, investment specific shocks contribute to the fall in hours in the 1975 and 1982 recessions, but a large portion of hours fluctuations in all major historical episodes is left unexplained.

Figure 10 also gives an indication of the nature of the two estimated technology shocks. In fact, the counterfactual price of investment series that investment specific shocks generate is practically identical to the observed price of investment series and the counterfactual labor productivity series that neutral shocks generate is practically identical to the observed labor productivity series. Hence, the two technology shocks are simply univariate shocks to the growth rate of relative price of investment and labor productivity. We next analyze how do these shocks relate to those extracted from standard accounting exercises, which are often used as exogenous forces in calibrated models.

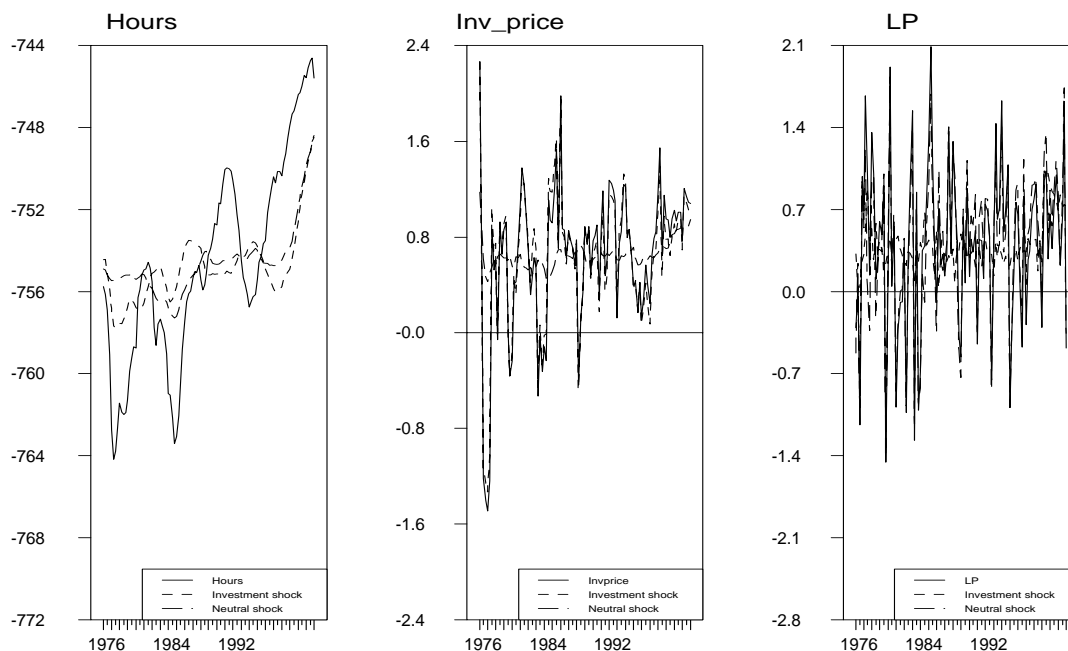


Figure 10: Historical decomposition

Thus, we construct Solow residuals shocks, using a Cobb-Douglas production function, adjusting for capacity utilization, and standard estimates of the labor share. Using the definition of labor productivity and the production function, we have that $\frac{y_t}{n_t} = \left(\frac{u_t k_t}{n_t}\right)^\alpha A_t$, where A_t measures total factor productivity (TFP), N_t is private nonfarm business sector hours, u_t is capacity utilization, and k_t are capital services (both of which are obtained from the US Bureau of Labor Statistics). In the upper left corner of figure 11 we plot neutral 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 innovations in the estimated TFP displays higher volatility (0.79 vs 0.53), that the majority of this volatility is concentrated in the high frequencies of the spectrum, and that TFP shocks are positive at some NBER recession dates. Similar conclusions are obtained if we sample annually our neutral shocks and compare them with the aggregate technology shocks of Basu et al. (2005); the contemporaneous correlation is somewhat lower (0.10) but the ranking of volatilities is unchanged. Hence, if

displaying the right sign in major recessions and being coincident and in phase with the NBER indicator is a plus, our neutral shocks have better features than standard TFP shocks.

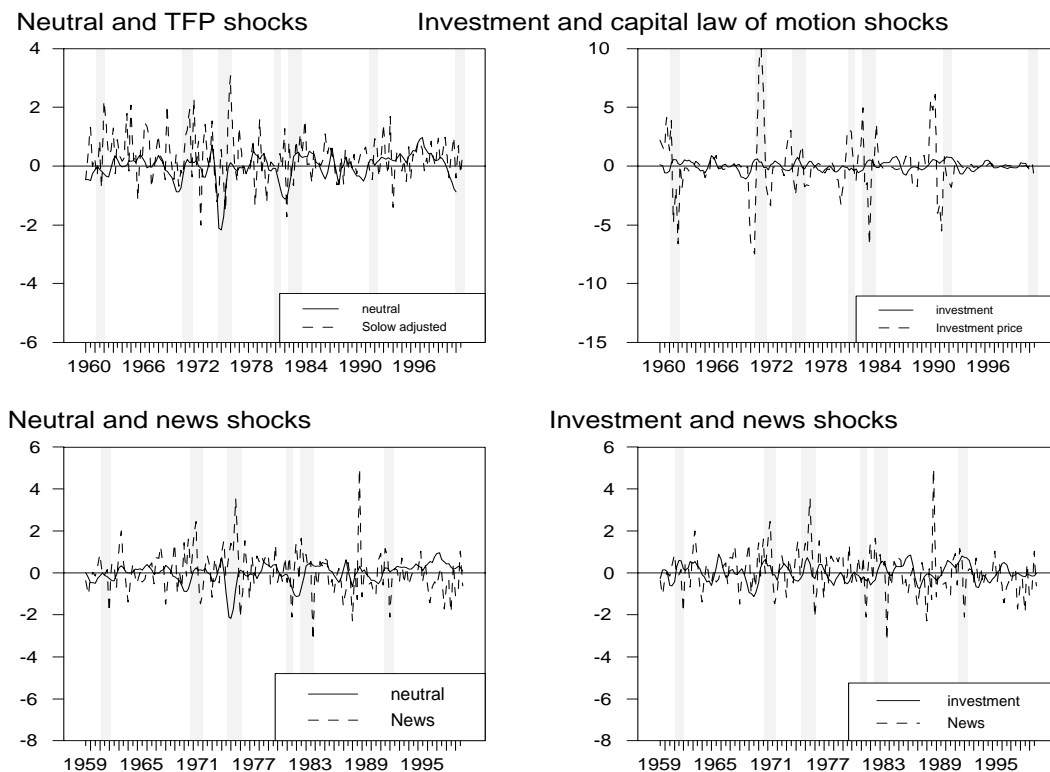


Figure 11: Technology shocks and other shocks

To compare our price investment shock series with an accounting series for disturbances to the price of investment, we use the law of accumulation of capital $k_{t+1} = (1 - \delta)k_t + v_t i_t$, where v_t is the inverse of the price of investment, quarterly measures of the capital stock and of the available investment series. The accounting series for disturbances to the price of investment is obtained by differencing the estimated series for v_t . The upper right corner of figure 11 plots the resulting series and the estimated investment specific shocks. As it was the case with the neutral shocks, the volatility of the accounting series is much higher than the volatility of our estimated price of investment shock series (2.31 vs 0.40); the two series are positively

correlated at leads and lags (maximum effect 0.3 at lag 4), but differ considerably in the last three NBER recessions.

Beaudry and Portier (2005) have shown that there is an almost perfect correlation between technology shocks identified with long run restrictions in a bivariate model with TFP and stock prices and what they call news shocks, i.e. shocks which do not generate any contemporaneous effects on TFP, but instantaneously affect stock prices. Are our two technological disturbances related to news shocks? We graphically show these relationships in the second row of figure 11. It is clear that the correlation between the two series is far from perfect and, if anything, they are moving in opposite direction, especially at NBER recession dates. For example, in the 1975 recession, news shocks are positive, while both our neutral and investment specific shocks are negative. Moreover, even looking forward, news shocks do not capture the dynamics of our two technological shocks. In fact, the regression line between neutral and news shocks has a slope equal to -0.12 (with a *t-statistic* equals to -3.03) and the slope does not change magnitude if we lag the news shock series up to 16 quarters. The slope between investment specific shocks and news shocks is -0.11 (with a *t-statistic* equals to -3.15) and its magnitude falls if we lag the news shock series up to 16 quarters. Hence, the relationship between technology shocks and news shocks identified by Beaudry and Portier (2005) may be spurious. Once proper technology shocks are extracted from the data, news shocks have little to do with them.

To confirm this conclusion we perform two additional exercises. First, we examine whether adding stock prices to our three variable system changes the informational content of our technology shocks. As shown in the second row of figure 9, adding stock prices to the system hardly changes the main features of our technological disturbances. Second, we run two bivariate systems with stock prices and either labor productivity or the price of investment and compare the structural shocks with those obtained in our benchmark VAR. If the results by Beaudry and Portier (2005) are due to the low dimension of the system considered, we should see the time series properties of technology shocks to vary when considering a bivariate system. Indeed, the time series properties of technology shocks do change substantially in bivariate systems (see third row of figure 9). Hence, news shocks correlate with technology

shocks only to the extent that the conditioning set used in the VAR is limited to lagged measures of productivity (or the price of investment) and stock prices.

Finally, a previous version of the paper has also examined how our technology shocks are related to the markup shocks extracted as in Comin and Gertler (2006) and how the resulting capital share, obtained as in Blanchard (1997), relates to the accounting capital share. The contemporaneous correlation of neutral and markup shocks is positive but low (0.12), and the contemporaneous correlation between investment specific and markup shocks is negative but low (-0.19). The two capital share series, while displaying the same cyclical behavior in major historical episodes, differ somewhat in the last few years of the sample.

To sum up, neutral shocks show a marked cyclical pattern but, in general, they have little to do with hours fluctuations except in the late 1990s. Moreover, the negative hours response they induce occur primarily at non-recessionary episodes. Investment specific shocks do not display strong cyclical movements, but they induce movements in per-capita hours in the direction one would expect, especially at recession times. The technology shocks one extracts from a bivariate system, TFP data and standard approaches are substantially different from the shocks we have obtained and there is no evidence that our shocks are related to news shocks.

Since a large portion of the fluctuations in per-capita hours is not due to technology shocks, both on average and in specific historical episodes, what then drives hours fluctuations? Such a question is difficult to answer with our trivariate systems, since the remaining shock is difficult to interpret since it captures all the sources of stationary disturbances in productivity and, potentially, the effects of all omitted stationary variables. Nevertheless, an analysis of its features may shed light on its nature and its relationship with hours fluctuations. It turns out that such a shock is negatively and significantly correlated with the news shock contemporaneously (estimate is -0.03 and t-stat -2.47) and that the maximum correlation is lagged 7 quarters (estimate is 0.20); it is also somewhat correlated with oil shocks (maximum correlation is 0.17 with 7 leads of oil shocks) and with FFF innovations (maximum correlation is 0.16 with 2 lags of FFF shocks), but it is unrelated to inflation, to the consumption to output or the investment to output ratios which may proxy for goods demand driven fluctuations. We also find that this shock is quite cyclical;

i.e. it displays troughs at major NBER recessions and induces positive per-capita hours and output responses.

6 Discussion and conclusions

The previous evidence substantially qualifies what is available in the literature. We show that the presence of long cycles in hours has little effects on the sign of the dynamic relationship between investment specific shocks and hours, while we confirm it could drive the sign of the dynamic relationship between neutral shocks and hours. We also show that neither differencing the hours series, nor taking the series in levels is the right approach to follow. The first choice induces overdifferencing and enhances the importance of measurement error, the second introduces important bias in the estimation. Allowing the intercept of the hours series in the VAR to vary over time is instead a much more efficient approach. In addition, given that the observed long cycles in hours are relatively stable over time, low frequency movements in hours can not be generated by permanent changes in taxes, in the relative importance of wealth and substitution effects or in the size of the government employment sector. This also suggests that models should not necessarily provide economic reasons leading to a unit root in per-capita hours, as part of the literature seems to argue; see for example Francis and Ramey (2005).

In the literature, the empirical evidence on the relationship between per-capita hours and technology shocks seems to depend on a number of specification choices. Our analysis shows that such an outcome is due to an inappropriate treatment of the long cycles in hours. Once these cycles are taken care in any reasonable way, the data robustly suggest a number of interesting facts which we believe are useful for guiding business cycle analysis. In particular:

- Per-capita hours fall in response to neutral shocks and increase in response to investment shocks.
- Neutral shocks explain a small portion of per-capita hours fluctuations and a much larger portion of output fluctuations; the opposite is true for investment specific shocks.

- The negative response of per-capita hours to neutral shocks primarily occurs at non-recessionary times.
- Neutral shocks are more cyclical than investment specific shocks and display upturns and downturns which match the NBER classification.
- The technology shocks we recover are uncorrelated with news shocks or technology shocks extracted with accounting exercises.
- Shocks other than technological disturbances are crucial in explaining the dynamics of per-capita hours since the mid-1950s.

Future research should try to directly relate these findings to theories of the business cycle by simulating models and running trivariate VARs on similar sample sizes. Some of our results may be hard to reconcile with standard models, both of flexible price and sticky price orientations. In particular, models with a standard production function and preferences may have hard time to generate different signs and different magnitudes in hours in response to neutral and investment specific shocks. They may also have a hard time to reproduce the relative size of hours (and output) fluctuations explained by the two technology shocks we found in the data. Canova, et. al. (2006) instead discuss a model where improvements in the neutral technology cause Schumpeterian creative destruction and trigger adjustments along both the intensive and the extensive margins of the labor market. They show that once the technology shocks extracted from the VAR are fed into such a model, the qualitative and quantitative features of the responses of hours, unemployment, finding and separation rates are accurately reproduced.

There are many other dimensions along which our work could be extended. First, it would be interesting to try to explain what drives long cycles in hours. While Comin and Gertler (2006) have made a step in that direction, much work still needs to be done. Second, one could try to relate the dynamics of hours to the “Great Moderation” literature (see e.g. Canova, et. al. (2007)). Gambetti (2005) has done some steps in this direction by estimating an empirical model where the dynamic relationships and the variances of the shocks are allowed to change over time. He argues that the relationship between technology shocks and hours displays

a significant change since the late 1990's, which could be consistent with the evidence presented in figure 2, and which deserves further investigation. Finally, while most analyses concentrate on the US, one would like to know if our evidence also holds across countries. Since, series for hours and the price of investment are hard to obtain for many countries this may require considerable work. All in all, these extensions help to provide a more complete picture of the dynamics induced by various technology shocks and suggest ways to account for them with fully specified dynamic models.

References

- [1] Altig D, Christiano J, Eichenbaum M, Linde J. 2005. Firm-Specific Capital, Nominal Rigidities and the Business Cycle. NBER Working Paper 11034.
- [2] Basu S, Fernald J, Kimball M. 2006. Are technology improvements contractionary? *American Economic Review* **96**: 1418-1448.
- [3] Beaudry P, and Portier F. 2005. Stock prices, news and economic fluctuations. *American Economic Review* **96**: 1293-1307.
- [4] Blanchard O. 1997. The Medium Run. *Brookings Papers on Economic Activity* **2**: 89-158.
- [5] Canova F, Lopez-Salido D, Michelacci C. 2006. Schumpeterian technology shocks. Available at www.crei.cat/people/canova/
- [6] Canova F, Gambetti L, Pappa E. 2007. The Structural Dynamics of Output Growth and Inflation: Some international Evidence. *Economic Journal* **117**: C167-C191.
- [7] Chari V, Kehoe P, McGrattan E. 2005. A critique of SVAR using Real Business Cycle theory. Fed of Minneapolis, working paper 631.
- [8] Christiano L, Eichenbaum M, Vigfusson R. 2003. What happens after a technology shock. Available at <http://www.faculty.econ.northwestern.edu/faculty/christiano/research.htm>.
- [9] Comin D, Gertler M. 2006. Medium Term Business Cycles. *American Economic Review* **96**: 523-551.
- [10] Dedola L, Neri S. 2007. What does a technology shock do? A VAR Analysis with model-based sign restrictions. *Journal of Monetary Economics* **54**: 512-549.
- [11] Erceg C, Guerrieri L, Gust C. 2005. Can Long Run Restrictions Identify technology shocks. *Journal of the European Economic Association* **3**: 1237-1278.

- [12] Evans, C. 1992. Productivity Shocks and Real Business Cycles. *Journal of Monetary Economics* **29**: 191-208.
- [13] Faust J, Leeper E. 1997. Do Long Run Restrictions give reliable Results? *Journal of Business and Economic Statistics* **15**: 343-353.
- [14] Fernald J. 2007. Trend Breaks, Long Run Restrictions, and the Contractionary Effects of Technology Shocks. *Journal of Monetary Economics* **54**: 2467-2485.
- [15] Fernandez-Villaverde J, Rubio-Ramirez J, Sargent T, Watson M. 2007. The ABC (and D) for understanding VARs, *American Economic Review* **97**: 1021-1026.
- [16] Fisher J. 2006. The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. *Journal of Political Economy* **114**: 413-451.
- [17] Francis N, Ramey V. 2005. Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited. *Journal of Monetary Economics* **52**: 1379-1399.
- [18] Francis N, Ramey V. 2006. Measures of per-capita hours and their implications for the technology-hours debate. University of North Carolina, manuscript.
- [19] Galí J. 1999. Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?, *American Economic Review* **89**: 249-271.
- [20] Galí J. 2005. Trends in hours, Balanced Growth, and the Role of Technology in the Business Cycle. *Federal Reserve Bank of St. Louis Review*, July-August, 459-486.
- [21] Galí J, Rabanal P. 2004. Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data? *NBER Macroeconomics Annual* **20**: 225-288.
- [22] Gambetti L. 2005. Technology shocks and the response of hours worked: time varying dynamics matter. Available at <http://www.econ.upf.es/gambetti>.

- [23] Giordani P. 2004. An alternative explanation of the price puzzle. *Journal of Monetary Economics* **51**: 1271-1296.
- [24] Greenwood J, Yorokoglu M. 1997. 1974. *Carnegie-Rochester Conference Series on Public Policy* **46**: 49-95.
- [25] Kehoe T, Ruhl K. 2007. Are Shocks to the Terms of Trade Shocks to Productivity?. NBER Working Paper 13111.
- [26] Michelacci C, Lopez-Salido D. 2007. Technology Shocks and Job Flows. *Review of Economic Studies* **74**: 1195-1227.
- [27] Uhlig H. 2004. Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association* **2**: 361-371.
- [28] Vigfusson R. 2004. The delayed response to technology shock: a flexible price explanation. Federal Reserve Board, International Finance discussion paper 810.