JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ. (2009)
Published online in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/jae.1090

THE EFFECTS OF TECHNOLOGY SHOCKS ON HOURS AND OUTPUT: A ROBUSTNESS ANALYSIS

FABIO CANOVA,[a,b,c]* DAVID LOPEZ-SALIDO[c,d] AND CLAUDIO MICHELACCI[c,e]

[a] *ICREA-UPF, Barcelona, Spain*
[b] *AMeN, Barcelona, Spain*
[c] *CEPR, London, UK*
[d] *Federal Reserve Board, Washington, DC*
[e] *Cemfi, Madrid, Spain*

SUMMARY

We analyze the effects of neutral and investment-specific technology shocks on hours and output. Long cycles in hours are removed 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. '*News shocks*' are uncorrelated with the estimated technology shocks. Copyright © 2009 John Wiley & Sons, Ltd.

1. INTRODUCTION

There has been a renewed interest in empirically examining the effects of technology shocks on total per-capita hours following the work of Galí (1999, 2005), Christiano *et al.* (2003), Uhlig (2004), Francis and Ramey (2005) and, more recently, Dedola and Neri (2007). This interest is typically motivated by the fact that the dynamics of hours differ if such disturbances occur in a basic real business cycle (RBC) model (where hours increase) or in a basic sticky-price model (where hours decrease). Unfortunately, the available evidence is, at best, mixed and its interpretation controversial.

This state of affairs is due, in part, to the complicated task that applied researchers face. First, technology shocks are identified using long-run restrictions (see, for example, Galí, 1999). However, it is known that identification via long-run restrictions 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 to technology shocks (see Uhlig, 2004). Second, different types of technological disturbances may have different effects on hours, making the outcomes of standard bivariate models, where only one generic technology shock is identified, uninterpretable (see Fisher, 2006; Michelacci and Lopez-Salido, 2007). Third, as recently emphasized by Canova *et al.* (2006) and Kehoe and Ruhl (2007), the choice of price deflators may matter for correctly recovering the effects of technological disturbances. Fourth, the response of hours may depend 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,

* Correspondence to: Fabio Canova, Department of Economics, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain. E-mail: fabio.canova@upf.edu

the recent evidence provided by Beaudry and Portier (2005), where generic shocks that change expectations about the future (which they call 'news shocks') are the same as the identified technology shocks, makes the interpretation of the latter problematic.

This paper empirically examines the effects of technology on hours and output, addressing these issues in a unified and comprehensive manner. 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 and study the robustness of results to alternative identification schemes and different measures for the price deflators.

We find that once we remove long cycles in hours, all the other pieces of the puzzle become irrelevant; regardless of the lag length, choice of price deflators, identification scheme, presence of omitted variables and other auxiliary statistical assumptions, hours robustly fall in response to neutral shocks and robustly increase in response to investment-specific shocks. 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 as investment-specific shocks of the forecast error variance of output at all horizons. Estimated neutral shocks have peaks and troughs which occur in correspondence of National Bureau of Economic Research (NBER) peaks and troughs, 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; are unrelated to the news shocks of Beaudry and Portier (2005); and differ from the technology shocks one would extract from accounting exercises.

Our results complement and qualify several other contributions in the literature. Regarding the issue of long cycles in hours and how they should be dealt with in 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; Francis and Ramey, 2005), the VAR should include hours in differences. Figure 1 shows that the standard per-capita hours series displays long but essentially stationary cycles. These cycles are of longer duration than those considered in the business cycle literature and may reflect, for example, demographics, 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) leads to misspecification, efficiency losses, and potentially uninterpretable results. Hence our conclusions differ from those of Francis and Ramey (2005) who, in order to justify a specification in first differences, argue that one should find economic reasons that lead to a unit root in per-capita hours.

Fernald (2007) has also stressed the importance of removing long cycles in the variables of the VAR, in order to correctly recover the effects of technology shocks on hours. Relative to that paper, we emphasize that shifts do not necessarily appear in the autocovariance function of hours; provide an extensive robustness analysis; analyze the effects on hours of two different types technology shocks; and examine the time series features of the estimated technology shocks and their relationship with other interesting economic shocks.

DOI: 10.1002/jae

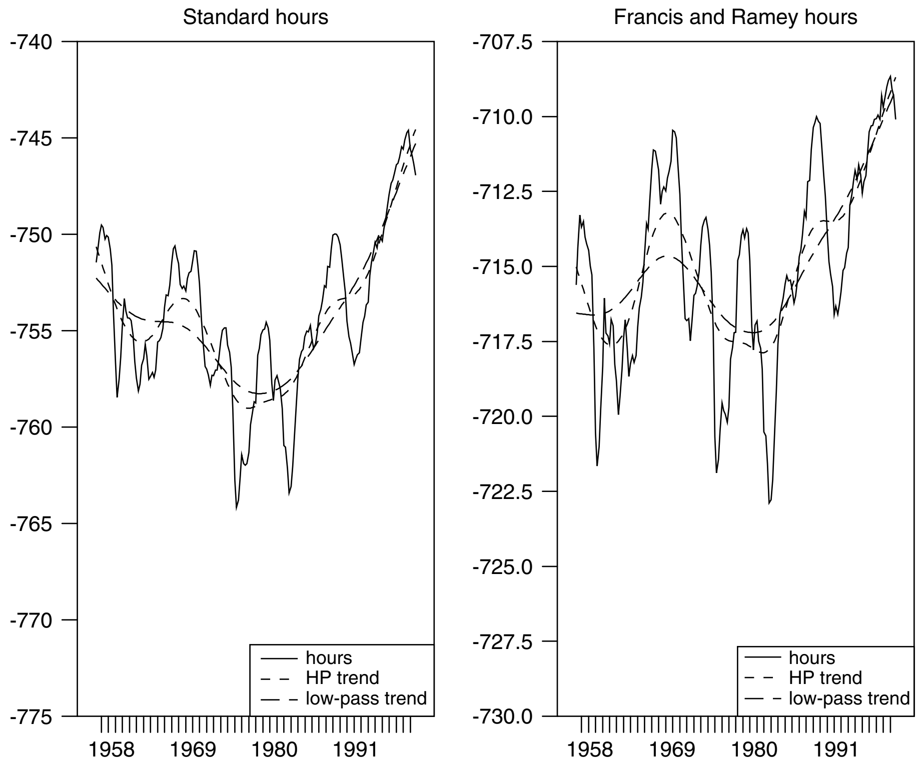

Figure 1. Per-capita hours, HP ($\lambda = 6400$) and low-pass estimates

Contrary to Fisher (2006), Galí (2005), and Altig *et al.* (2005), who advocated simple neoclassical flexible and sticky price models to interpret the facts, we interpret the evidence as pointing to an alternative direction. The neoclassical framework in fact predicts 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 discuss in that paper, are instead consistent with a Schumpeterian model of creative destruction, where improvements in the neutral technology trigger adjustments along 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. Section 4 examines the robustness of our findings. Section 5 discusses the properties of estimated technology shocks. Section 6 concludes.

2. THE EXISTING EVIDENCE

We summarize the current state of the debate using a VAR with labor productivity and the price of investment, both measured in consumption units and per-capita hours. All variables are in logs.

DOI: 10.1002/jae

A trivariate specification is the minimum size system required to recover 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. 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 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 assuming 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 run. Fisher (2006) has shown how to derive these 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}. We use a generous lag length and reduce overparametrization with a prior, rather than using a standard lag selection criteria, 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 that are not always representable with a finite order VAR—a problem that may be severe if the theory has implications for many variables and the VAR has three equations. Unless otherwise stated, the figures report the point estimate of the dynamic responses and a 90% 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 Figure 2 report the response of hours to a neutral shock in the full sample. As documented in the literature, per-capita hours positively respond to a neutral shock in the level system, 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 responses are generally insignificant.

The sign difference in the point estimates obtained in the two specifications is often attributed to long cycles in hours and to the fact that these movements distort the conditional dynamics of the level system. However, long cycles do not necessarily imply non-stationarity dynamics as commonly assumed (see, for example, 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% tunnel for the difference system is large relative to that for the level system, this problem is likely to be important. Hence both systems are likely to be misspecified and it is difficult to draw credible conclusions about the conditional dynamics of hours from these two pictures.

One way to take care of long cycles in hours is to split the sample into 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. From Figure 1 one can see that these are, approximately, the dates at which the Hodrick Prescott estimate of the long cycles has flex points.

The dynamics of hours in response to neutral shocks are similar in the two systems (see second and third rows of Figure 2) and in both the 1955–1973 and 1973–1997 samples hours instantaneously fall. In the level system the point estimate is persistently negative, and significantly so for a number of quarters; in the difference specification, the point estimate turns positive after two quarters but remain insignificant at all horizons. If we had considered the 1973–2000 sample rather than the 1973–1997 sample, we would have found that hours instantaneously increase in

DOI: 10.1002/jae

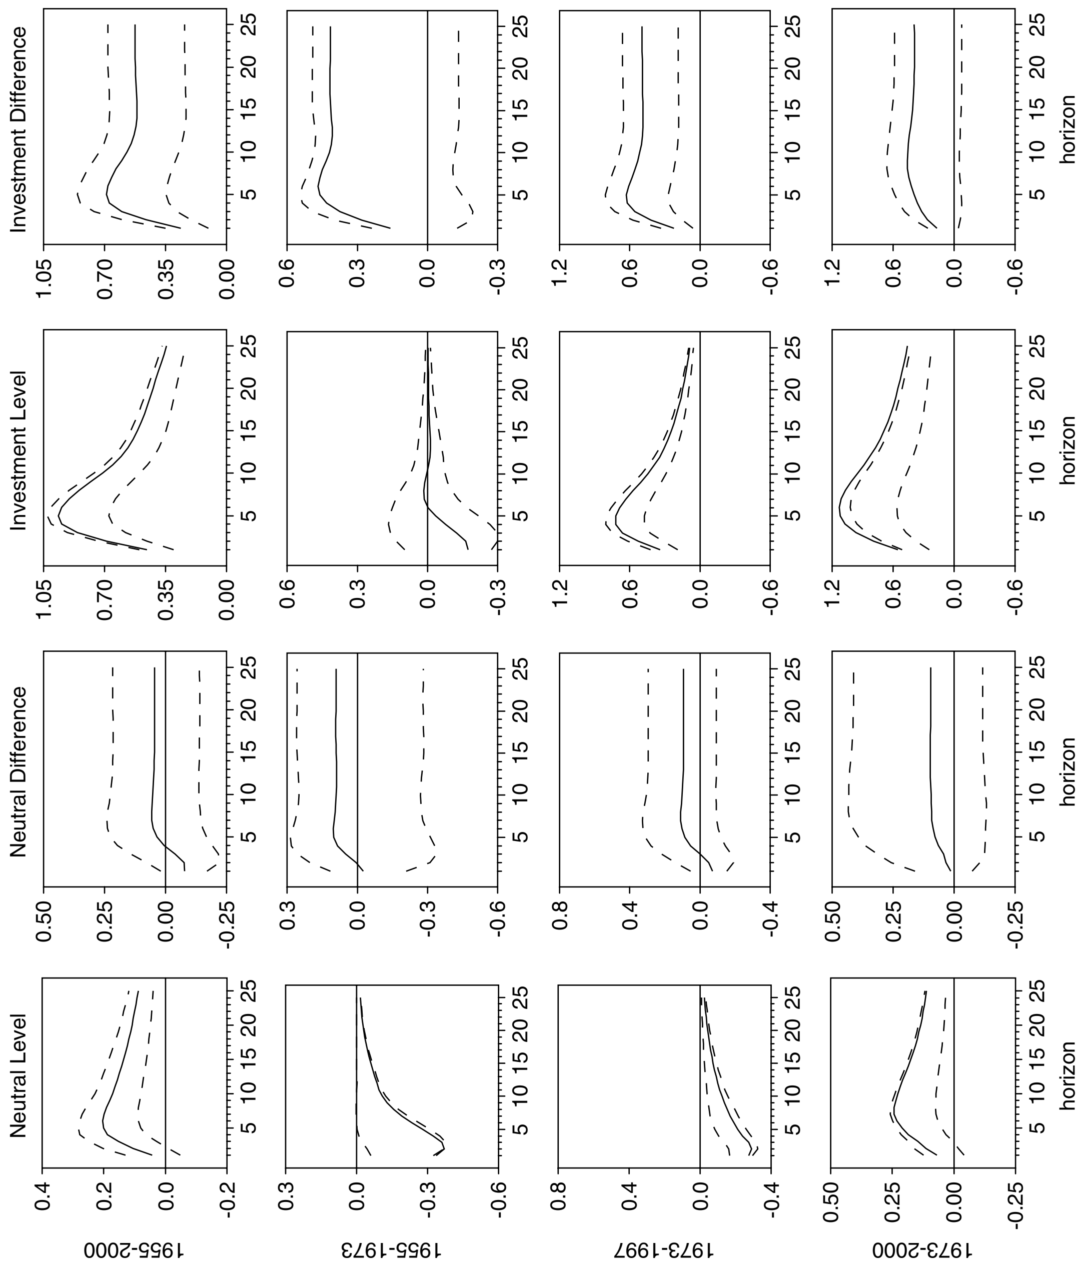

Figure 2. Responses of per-capita hours. —— : Median responses - - - - : 90 percent tunnel

DOI: 10.1002/jae

both specifications even though the increase is significant only in the level system (see third row of Figure 2).

The responses of hours to investment shocks are depicted in the last two columns of Figure 2. For the full sample, the level and 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 insignificantly fall in the 1955–1973 sample, and significantly increase in the 1973–1997 and 1973–2000 samples. For the difference system, the point estimate is positive in all three subsamples, but responses are significant only in the 1973–1997 sample.

The subsample instability found in the level system for both shocks may lead researchers to believe that the relationship between per-capita hours and technology shocks has changed over time. What is puzzling for this interpretation is that no subsample instability is present in the difference specification. Subsample evidence, however, is difficult to trust because small-sample biases may be important. Small-sample biases 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 biases may interact in an unpredictable way with measurement and aggregation errors, making subsample evidence uninterpretable. In sum, neither assuming unit roots nor splitting the sample into pieces seems to be the best way to account for long cycles in per-capita hours.

3. LONG CYCLES AND INTERCEPT HETEROGENEITIES

The instabilities and sign reversals one finds in the level specification (columns 1 and 3 of Figure 2) and the substantial homogeneity but insignificance of the results with the difference specification (columns 2 and 4) are symptomatic of a level shift in the data. To see why this may be the case, consider the three clouds of points presented on the left-hand side box of Figure 3, which are intended to represent the 1973–1997, 1955–1973 and 1997–2000 impact relationship between neutral shocks (on the x-axis) and per-capita hours (y-axis). Note that in the first sample the intercept is positive and large. In the second and third samples 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 (think of this as the subsample 1973–2000) or the three samples together without taking into account intercept heterogeneity, the slope of the relationship becomes positive. This is exactly the pattern of contemporaneous responses we found in the level system: in the 1955–1973 and 1973–1997 samples the response of hours to neutral shocks is negative; in the 1973–2000 and in the full sample the response is positive. If we difference the variables we remove the effects of the mean shifts, independently of the sample. Hence, the slope of the relationship is estimated to be negative, but uncertainty about its magnitude is large, possibly due to the magnification of measurement errors (see right-hand side box of Figure 3).

While mechanically it is clear how the pattern we observe in Figure 2 can happen, it is worth describing why such an outcome is to be expected when analyzing the relationship between technology shocks and hours. Take, for example, the subsamples 1973–1997 and 1973–2000. Here the differences in the sign of the hours responses to technology shocks can be attributed to the higher level that average productivity growth and hours experienced during the productivity

DOI: 10.1002/jae

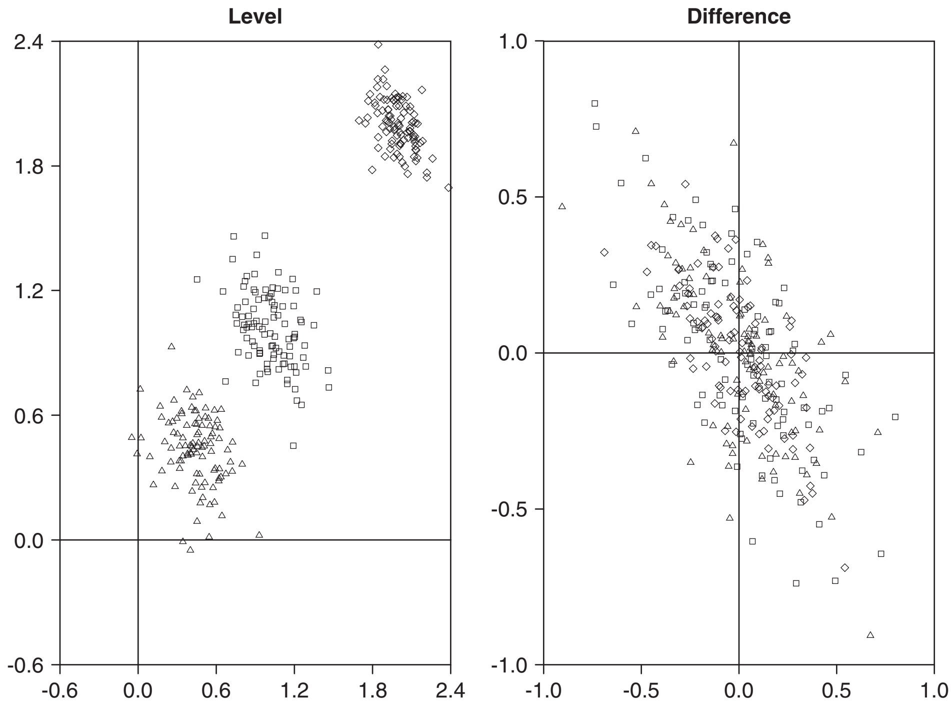

Figure 3. Interpretation of the evidence

revival of the late 1990s. If not properly accounted for, these shifts lead to biases in the estimated responses because changes in the average rate of productivity growth are at least partly identified as a series of neutral technology shocks. Thus, in the sample 1973–2000, when productivity growth is higher on average than in the 1973–1997 sample, the specification identifies a series of positive neutral technology shocks. Since in this period hours are also above average, a bias emerges.

If the situation described in Figure 3 is correct, splitting the sample and tracing out the responses separately in each of them is inefficient, since there are small changes in the structural relationship, and may induce small-sample biases. Hence, to efficiently take care of the long cycles in hours, 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 (which are those displayed as low-pass trend 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 in all specifications the small-sample biases induced by the use of long-run restrictions are considerably reduced, since the full sample of quarterly data is now employed.

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; per-capita hours fall in response to neutral shocks and increase in response to investment-specific

DOI: 10.1002/jae

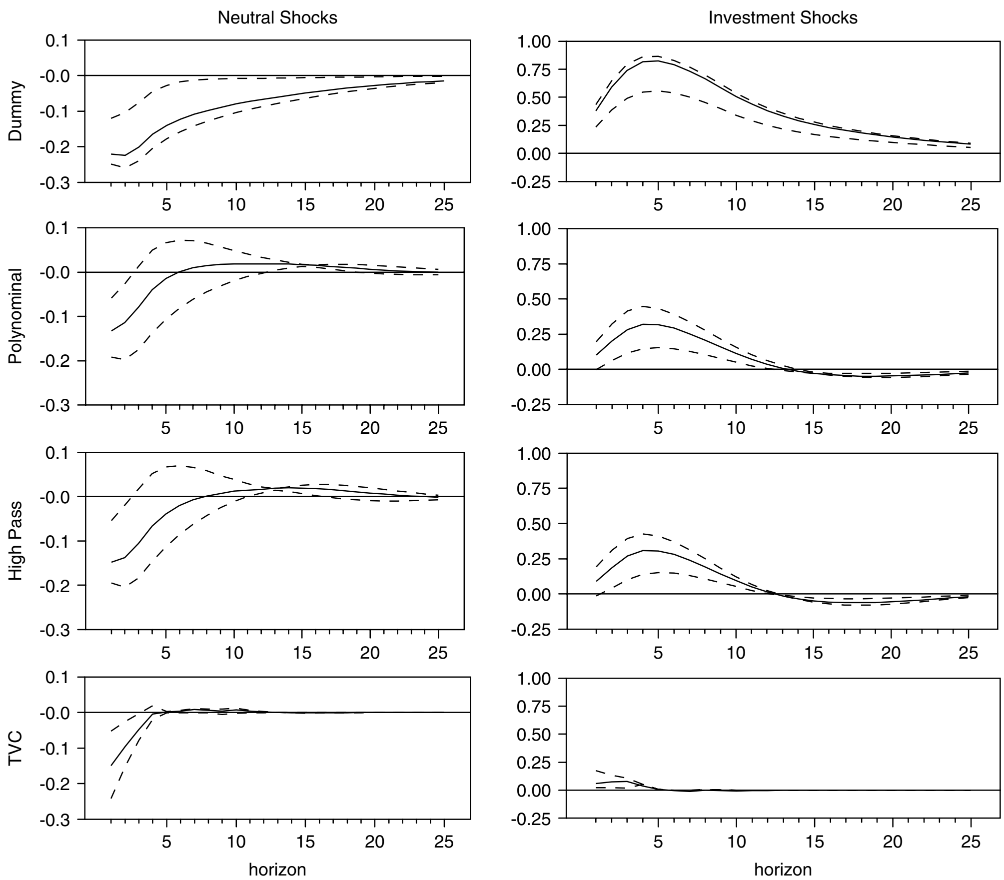

Figure 4. Responses of per-capita hours, sample 1955–2000. ——— : Median responses - - - - : 90 percent tunnel

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 also that the hump present in Figure 2—which has been greatly emphasized by Vigfusson (2004)—is absent here.

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%), while investment-specific shocks explain between 30% and 50% of the variance of hours at these horizons (20–30% with the time-varying intercept specification). Interestingly, this ordering is reversed for output fluctuations; regardless of the horizon, neutral shocks explain on average about 35% and investment specific shocks only about 18% of output fluctuations (see Figure 5 for the dummy specification).

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). They report

DOI: 10.1002/jae

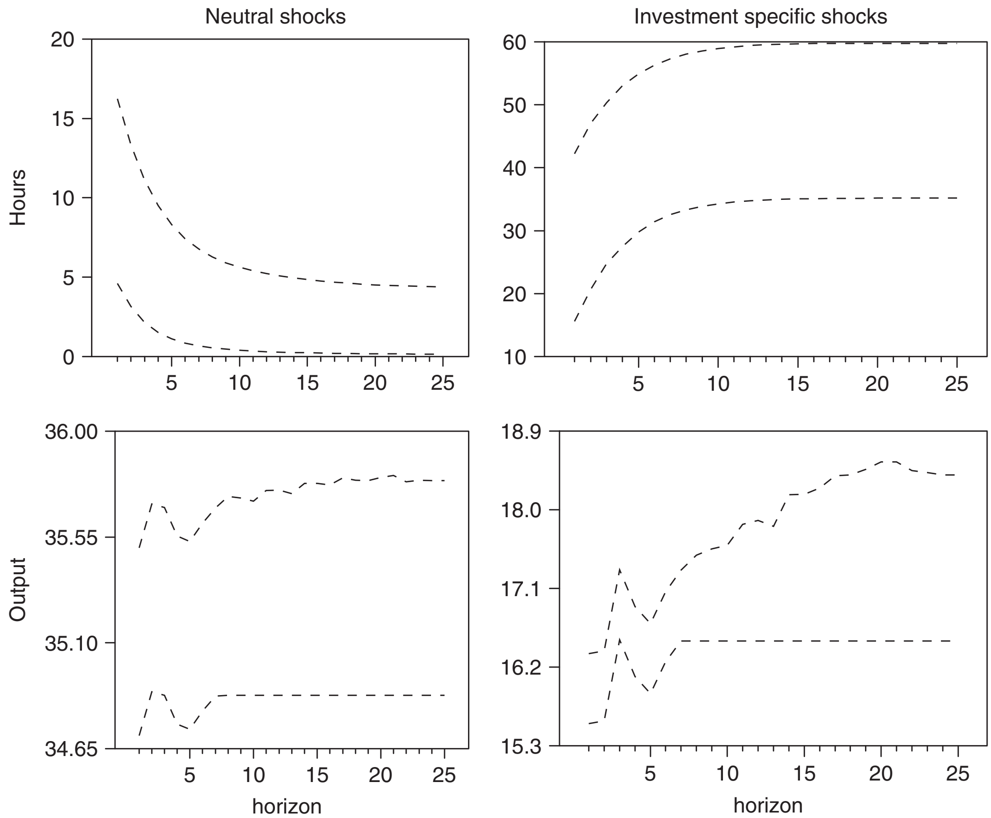

Figure 5. Variance decomposition, dummy specification, 90% bands

that the contribution of both shocks to per-capita hours and output volatility at business cycle frequencies is roughly the same (about 15%). 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 neglect long cycles in hours, which we show are key to understanding 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 original subsample evidence was uninterpretable because of the potential interactions between small-sample biases and improper treatment of data heterogeneities. Clearly, small-sample biases will not disappear. Nevertheless, to the extent that data heterogeneity is properly accounted for, subsample analysis should be more informative. Figure 6 reports results for the dummy specification and shows that the relationship between per-capita hours and technology shocks is stable over time. The responses of per-capita hours to neutral shocks are instantaneously negative and significant in all subsamples 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 1955–1973 sample, the responses are insignificant at all horizons.

DOI: 10.1002/jae

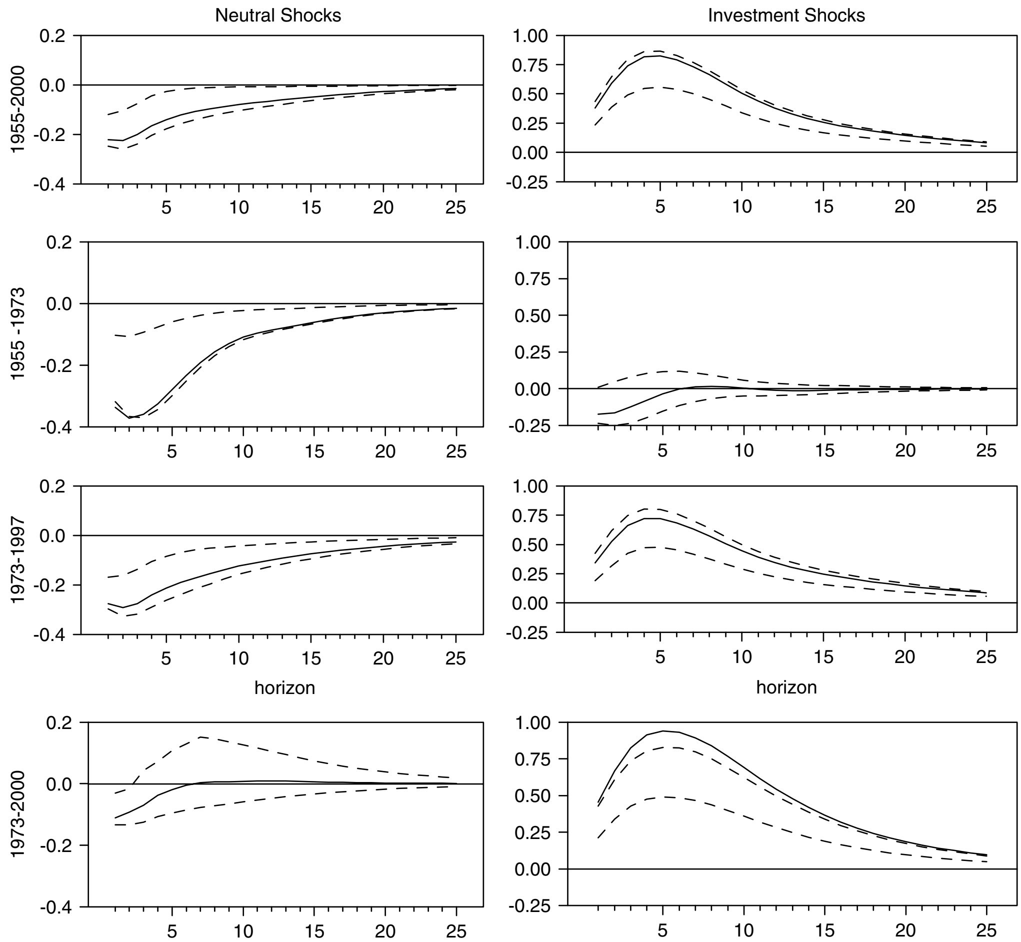

Figure 6. Responses of per-capita hours, different samples, dummy specification. ——— : Median responses - - - - : 90 percent tunnel

4. ROBUSTNESS

There are many dimensions along which the robustness of our 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 respect 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, our conclusions are quite robust. For the sake of presentation, we only report results for the dummy specification. As before, the exact way long cycles in hours are dealt with does not 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, external shocks could move labor productivity in the long run. Since the

DOI: 10.1002/jae

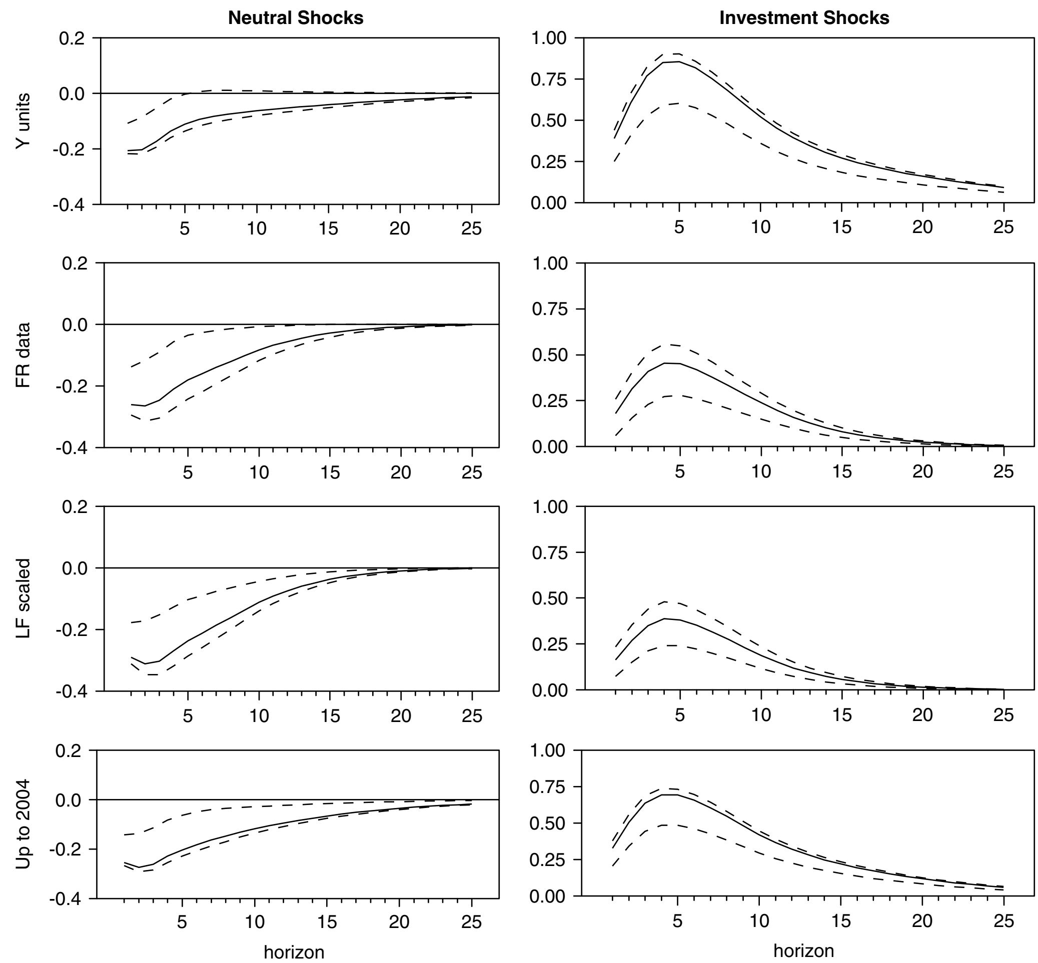

Figure 7. Responses of per-capita hours, alternative measurements, dummy specification, 1955–2000. ——— : Median responses - - - - : 90 percent tunnel

output basket is less prone to such problems, we have repeated estimation measuring either labor productivity or both labor productivity and the price of investments in output units. The first row of Figure 7 shows that the sign and shape of per-capita hours responses to both shocks are 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-hand 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. However, 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 7 shows that, indeed, the qualitative features of per-capita hours responses are unaltered when this new series is used in the VAR.

A referee suggested that the trend in the hours series in the late 1990s could be the result of changes in working age population, and that the trend disappears when the hours series is scaled

DOI: 10.1002/jae

by labor force participation rather than by working age population. This alternative series does not match the object of reference in the theoretical discussions, which focus on a comprehensive measure of aggregate labor effort (including labor force participation). Nevertheless, it is worth repeating the estimation since, while free of trends, such a series still displays long cycles. The third row of Figure 7 shows that 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 7 shows that updating the sample produces no major changes in the estimated relationship.

In order to credibly claim that our analysis is informative about the relationship between technology shocks and per-capita hours, it is important to make sure that the estimated technology shocks are not standing in for missing variables or other structural shocks. While we have allowed sufficient lags in each estimated specification, and this 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. To check whether omitted variables play a role, we have correlated our two estimated technology shocks with a set of 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 8 reports the cross-correlations of up to four leads and four lags of the estimated technology shocks with consumption to output, investment to output, and inflation, and the upper and lower limits of an asymptotic 95% 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; lags of consumption to output, investment to output, and inflation do not help in predicting structural residuals. Hence it is unlikely that omitted variables play a major role in explaining the results.

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

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 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 9. The sign of the responses is very robust to the choice of lag length.

We have also checked whether the dynamics of hours are robust to the timing and the type of identification restrictions we use. 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. Beaudry and Portier (2005) found empirical evidence consistent with this

DOI: 10.1002/jae

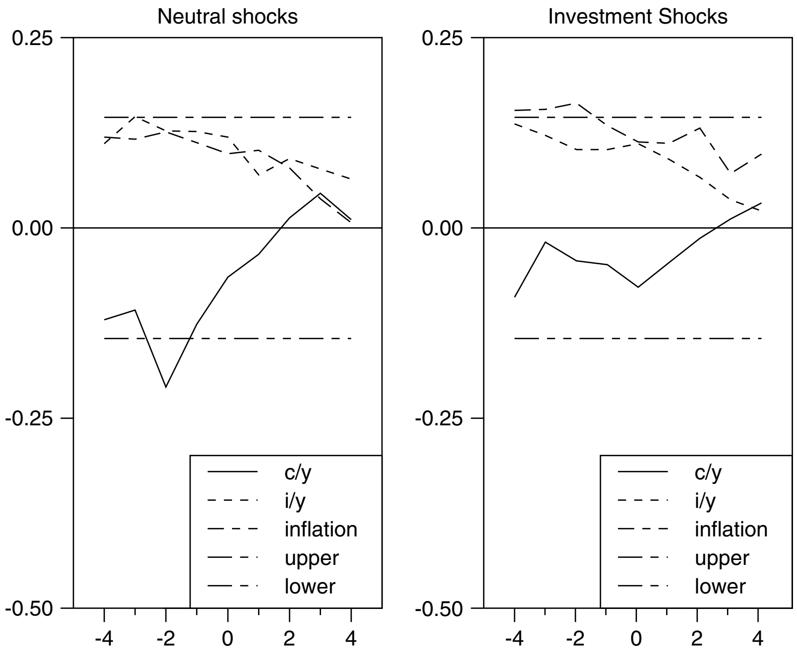

Figure 8. Cross-correlation structural shocks-omitted variables, dummy specification, 1955–2000. c/y : consumption to output ratio i/y : investment to output ratio

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 (which they call news shocks). To study whether this is a problem, we have imposed our identification restrictions at varying horizons. The second row of Figure 9 shows the impact response of per-capita hours as the horizon changes: the sign of the response is robust 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 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 are affected by 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. WHAT 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

DOI: 10.1002/jae

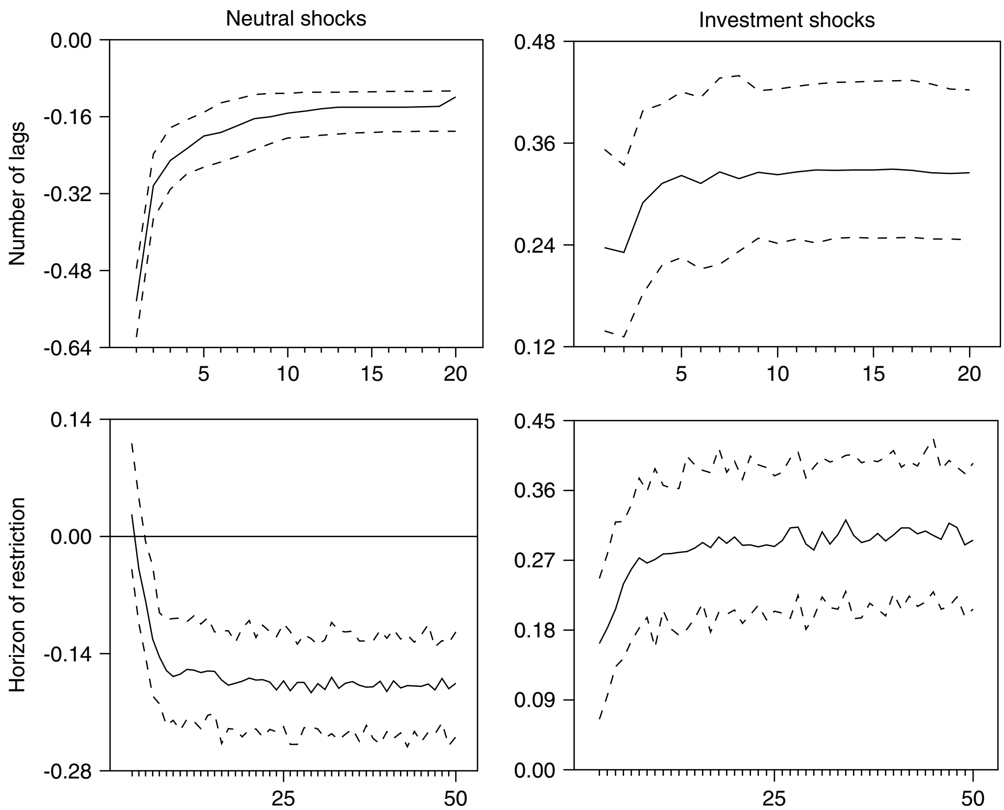

Figure 9. Contemporaneous response of per-capita hours, dummy specification, 1955–2000

even if a positive drift is allowed for. 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 detail.

The first row of Figure 10 plots the (smoothed) estimated technology shocks together with NBER recession episodes (shaded areas). Three stark features are evident. First, our neutral shocks display significant cyclicality. Furthermore, peaks are typically coincident with start of NBER recessions and troughs 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 shocks 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 led to conclude that the two major post-World War II 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 these episodes we have computed an historical decomposition (see Figure 11 in Canova *et al*., 2008). Four features emerge. First, the effect of the two shocks on per-capita hours depends on time and on the state of the economy. Second, neutral

DOI: 10.1002/jae

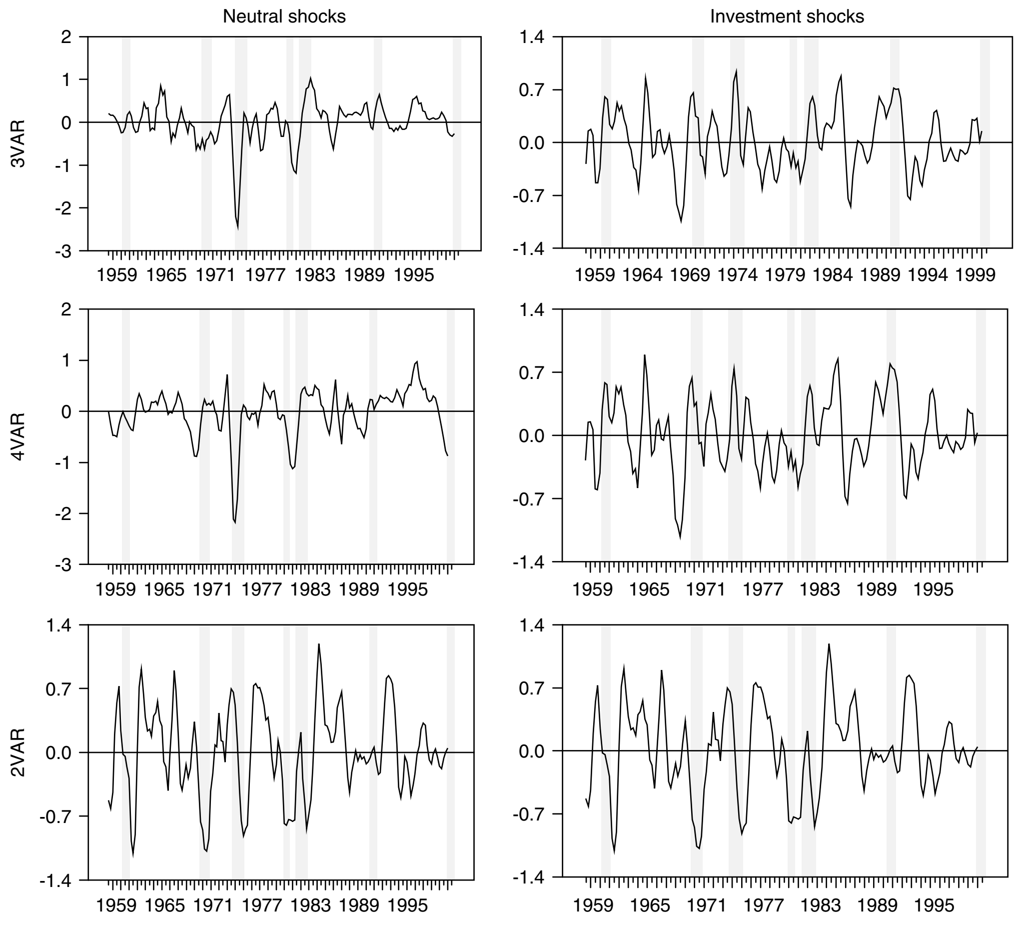

Figure 10. Time path of technology shocks

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 small fall in the 1982 recession. On the other hand, it is 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 the 1975 and 1982 recessions, but a large portion of hours fluctuations in both episodes is left unexplained.

Fourth, the counterfactual price of investment series generated by investment-specific shocks is nearly identical to the observed price of investment series and the counterfactual labor productivity series produced by neutral shocks is virtually identical to the observed labor productivity series.

We next analyze how these shocks relate to those extracted from standard accounting exercises, which are often used as exogenous forces in calibrated models. 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

DOI: 10.1002/jae

are obtained from the US Bureau of Labor Statistics). The correlation of the two series is low (the maximum value occurs contemporaneously and it is only 0.19). We also find 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 in some NBER recessions. Similar conclusions are obtained if we sample our neutral shocks annually and compare them with the aggregate technology shocks of Basu and Kimball *et al.* (2006); 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.

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. As with neutral shocks, the volatility of the accounting series is much higher than the volatility of our estimated investment shocks (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 estimated technological disturbances related to their news shocks? The correlation between news and technology shocks is far from perfect and, if anything, they move in opposite direction, especially at NBER recession dates. For example, in the 1975 recession, news shocks are positive, while both our neutral and investment shocks are negative. Even looking forward, news shocks do not capture the dynamics of our two technology shocks. In fact, the regression line between neutral and news shocks has a slope equal to -0.12 (with a t-statistic of -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 of -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.

To confirm this conclusion we perform two additional exercises. First, we have examined 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 10, 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, we should see in the time series properties of technology shocks obtained bivariate systems to change. Indeed, the time series properties of technology shocks do change substantially in bivariate systems (compare the first and the third row of Figure 10). Hence news shocks correlate with technology shocks only to the extent that the conditioning set is limited to lagged measures of productivity (or the price of investment) and stock prices.

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 during non-recessionary episodes. Investment-specific shocks do not

DOI: 10.1002/jae

display strong cyclical features, 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.

6. DISCUSSION AND CONCLUSIONS

The evidence this paper presents substantially qualifies what is available in the literature. We show that the presence of long cycles in hours has little consequence 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 splitting the sample in pieces is the right approach to follow. Allowing the intercept of the hours series in the VAR to vary over time is instead a much more natural choice. In addition, given that the observed long cycles in hours are relatively stable over time, they cannot 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 as, for example, Francis and Ramey (2005) seem to argue.

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 of in any reasonable way, the data robustly suggest a number of interesting facts. 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 relate these finding to theories of the business cycle by simulating models and running trivariate VARs on similar sample sizes. Some of our results are hard to reconcile with standard models, both of flexible price and sticky price orientations. In particular, models with standard preferences and production function may have a hard time generating different signs and different magnitudes in hours responses to neutral and investment-specific shocks and may also have a hard time reproducing the relative size of hours (and output) fluctuations explained by the two technology shocks we found in the data. Canova *et al*. (2006) 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 a VAR are fed into such

DOI: 10.1002/jae

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, for example, Canova *et al*., 2007). Gambetti (2005) has attempted to do this by estimating an empirical model where the dynamic relationships and variances of the shocks are allowed to change over time. He argues, consistent with the evidence presented in Figure 2, that the relationship between technology shocks and hours displays a significant change since the late 1990s and such a result deserves further investigation. Finally, while most studies concentrate on the USA, 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 preparatory work.

ACKNOWLEDGEMENTS

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. The views expressed in the paper are of the authors and not those of the Federal Reserve System, or the Board of Governors. Canova acknowledges the financial support of the Spanish Ministry of Education, through grant SEJ-2004-21682-E, and of the Barcelona Economic Program (CREA).

REFERENCES

Altig D, Christiano J, Eichenbaum M, Linde J. 2005. Firm-specific capital, nominal rigidities and the business cycle. *NBER Working Paper 11034*.

Basu S, Fernald J, Kimball M. 2006. Are technology improvements contractionary? *American Economic Review* **96**: 1418–1448.

Beaudry P, Portier F. 2005. Stock prices, news and economic fluctuations. *American Economic Review* **95**: 1293–1307.

Canova F, Lopez-Salido D, Michelacci C. 2006. Schumpeterian technology shocks. www.crei.cat/people/canova/ [14 May 2009].

Canova F, Gambetti L, Pappa E. 2007. The structural dynamics of output Growth and inflation: some international evidence. *Economic Journal* **117**: C167–C191.

Canova F, Lopez-Salido D, Michelacci C. 2008. The effects of technology shocks on hours and output: a robustness analysis. *CEPR working paper 6720*.

Chari V, Kehoe P, McGrattan E. 2005. A critique of SVAR using Real Business Cycle theory. *Working Paper 631*, Federal Reserve Bank of Minneapolis.

Christiano L, Eichenbaum M, Vigfusson R. 2003. What happens after a technology shock. http://www.faculty.econ.northwestern.edu/faculty/christiano/research.htm [14 May 2009].

Comin D, Gertler M. 2006. Medium term business cycles. *American Economic Review* **96**: 523–551.

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.

Erceg C, Guerrieri L, Gust C. 2005. Can long run restrictions identify technology shocks? *Journal of the European Economic Association* **3**: 1237–1278.

Evans C. 1992. Productivity shocks and real business cycles. *Journal of Monetary Economics* **29**: 191–208.

Faust J, Leeper E. 1997. Do long run restrictions give reliable results? *Journal of Business and Economic Statistics* **15**: 343–353.

DOI: 10.1002/jae

Fernald J. 2007. Trend breaks, long run restrictions, and the contractionary effects of technology shocks. *Journal of Monetary Economics* **54**: 2467–2485.
Fernandez-Villaverde J, Rubio-Ramirez J, Sargent T, Watson M. 2007. The ABC (and D) for understanding VARs. *American Economic Review* **97**: 1021–1026.
Fisher J. 2006. The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy* **114**: 413–451.
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.
Francis N, Ramey V. 2006. Measures of per-capita hours and their implications for the technology-hours debate. Working paper, University of North Carolina.
Galí J. 1999. Technology, employment, and the business cycle: do technology shocks explain aggregate fluctuations?, *American Economic Review* **89**: 249–271.
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.
Gambetti L. 2005. Technology shocks and the response of hours worked: time varying dynamics matter. http://www.pareto.uab/gambetti.
Giordani P. 2004. An alternative explanation of the price puzzle. *Journal of Monetary Economics* **51**: 1271–1296.
Greenwood J, Yorokoglu M. 1997. 1974. *Carnegie–Rochester Conference Series on Public Policy*, Vol. 46; 49–95.
Kehoe T, Ruhl K. 2007. Are shocks to the terms of trade shocks to productivity? *NBER Working Paper 13111*.
Michelacci C, Lopez-Salido D. 2007. Technology shocks and job flows. *Review of Economic Studies* **74**: 1195–1227.
Uhlig H. 2004. Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association* **2**: 361–371.
Vigfusson R. 2004. The delayed response to technology shock: a flexible price explanation. *International Finance Discussion Paper 810*, Federal Reserve Board.

DOI: 10.1002/jae