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

Fabio Canova
ICREA-UPF, CREI 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. '*News shocks*' are uncorrelated with the estimated technology shocks.

JEL classification: E00, J60, O33.

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

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 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 e.g. Galí (1999)). However, it is know 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 as 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) and 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 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 (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 an 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 the conclusions 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, 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 also 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 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 NBER peaks and throughs, 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 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), 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. These cycles are of longer duration than those considered in the business cycle literature and may reflect 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, efficiency losses, and potentially uninterpretable results. Our point of view therefore leads to conclusions which are different from those by Francis and Ramey (2005) who, in order to justify a specification in first differences for hours, argue that one should find economic reasons that lead to a unit root in per-capita hours.


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

Fernald (2007) has also stressed the importance of characterizing shifts 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 a robustness analysis - using various measures of hours, different specification choices, medium versus long run restrictions, and several alternative detrending procedure; 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 some interesting economic shocks.

4

Our evidence is hard to reconcile with simple flexible price and sticky price models of the type advocated by Galí (2005) and Altig et al. (2005). 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 discuss in that paper, 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. Section 4 examines the robustness of our findings. Section 5 discuses 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 VAR with labor productivity, the price of investment, both measured in consumption units and per-capita hours. All variables are in logs. 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 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 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}. Rather than using a standard lag selection criteria, we use a generous lag length and reduce overparametrization with a prior. We choose this approach 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 that are not always representable with a finite order VAR—a problem that may be severe in a three equation system. Unless otherwise stated, 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 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 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.

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. 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.

6

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 just two quarters, but remain insignificant at all horizons. Notice that if instead of the 1993-1997 sample we had considered the 1973-2000 sample, we would have found that hours instantaneously increase in both specifications even though the increase is significant only in the level system (see fourth row of figure 2).
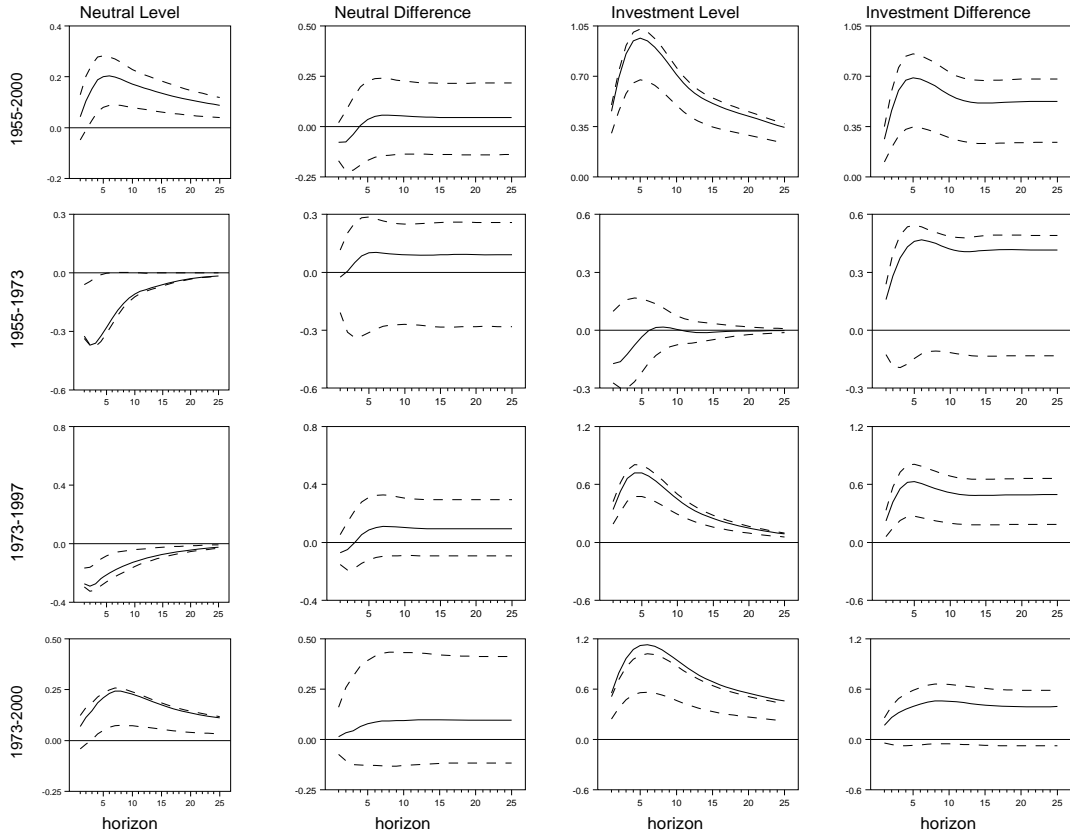

Figure 2: Responses of per-capita hours

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 insignificantly fall in the 1955-1973 sample, and significantly increase in the 1973-1997 and in the 1973-2000 samples. For the difference system, the point estimate is positive in all three sub-samples, but responses are significant only in the 1973-1997 sample.

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

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 particular type of data heterogeneity. In particular, they indicate the presence of level shifts and the need to properly account for them in order to pin down the sign of the conditional dynamics induced by technology shocks. To see why this may be the case 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 impact relationship between neutral shocks (on the x axis) and per-capita hours (y axis), respectively. Notice 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 sub-sample 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).


Figure 3: Interpretation of the evidence.

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 sub-samples 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 by the US during the productivity revival of the late 90's. If not properly accounted for, these shifts lead to biases in the estimated responses because average changes in the 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 little changes in the structural relationship, and may make small sample biases large. 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 to project estimated VAR coefficients infinitely far into the future.

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 also that the hump after 4-5 quarters that was present in figure 2—which has been greatly emphasized by Vigfusson (2004) —is no longer present.
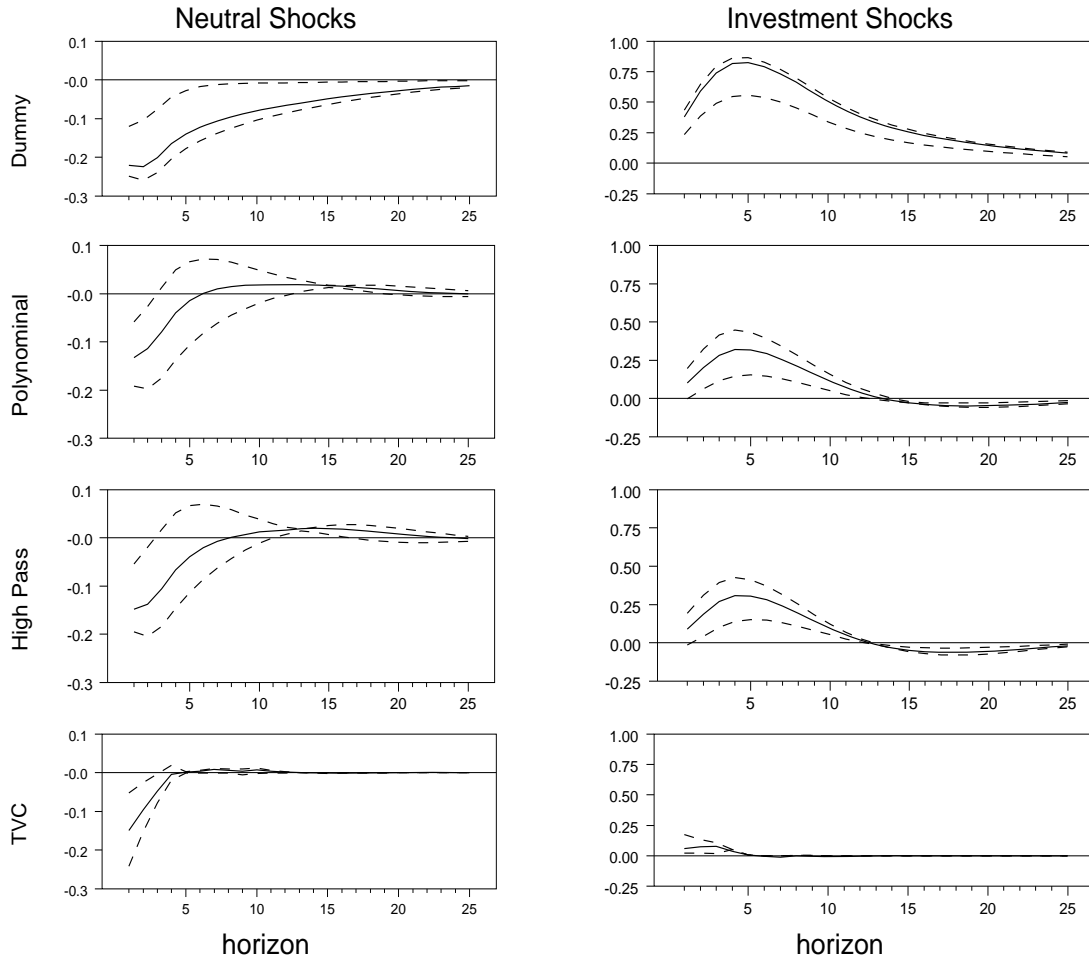

Figure 4: Responses of hours, sample 1955-2000

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). Interestingly, 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.


Figure 5: Variance decomposition, dummy specification, 90 percent bands

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 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 neglect long cycles in hours, which we show are key to understand 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 sub-sample evidence presented in figure 2 was uninteptetable because of the potential interactions between small sample biases and improper treatment of data heterogeneity. Clearly,

small sample biases will not disappear. Nevertheless, to the extent that data heterogeneity is now properly taken into account, sub-sample analysis should be more informative. For the sake of space, figure 6 reports results only for the dummy specification. The responses of per-capita hours to neutral shocks are now instantaneously negative in all sub-samples and responses 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, the relationship between per-capita hours and technology shocks appears to be stable over time.
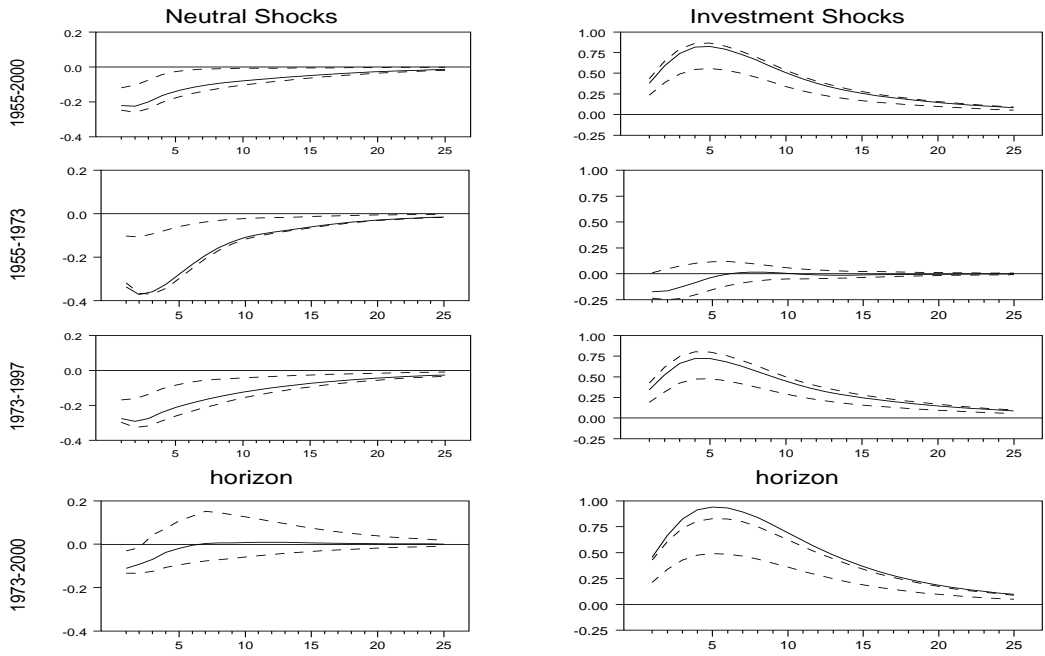

Figure 6: Responses of hours, different samples, dummy specification

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.
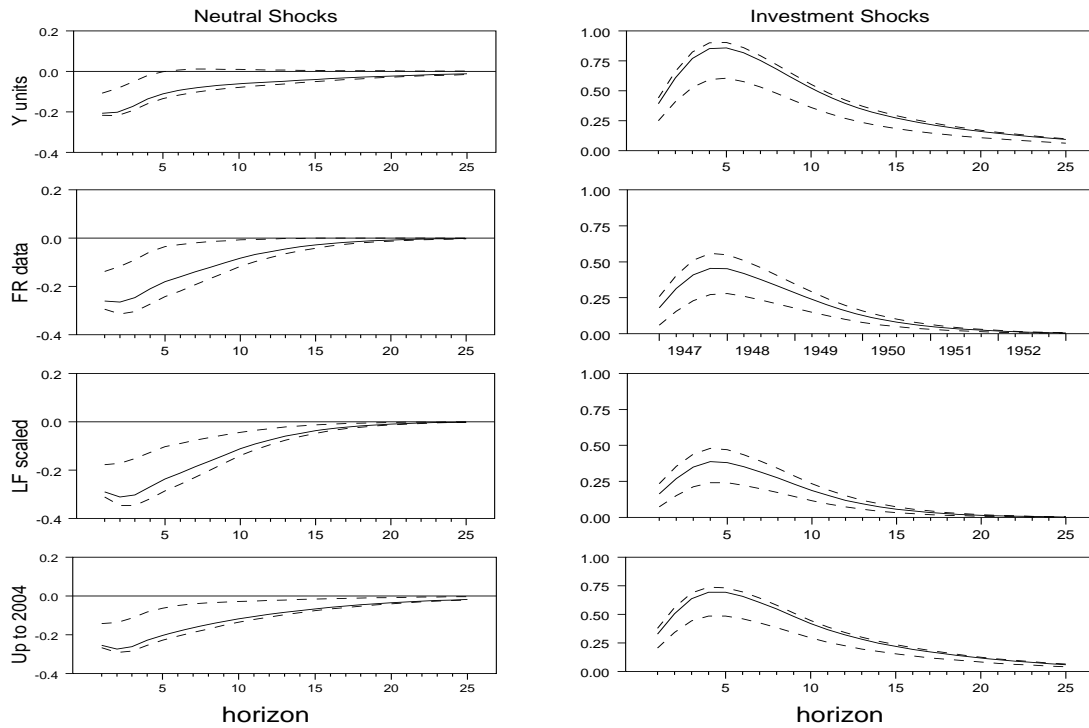

Figure 7: Responses of per-capita hours, Alternative measurements, Dummy specification, 1955-2000

In the systems we have run, labor productivity and the price of investment are

14

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 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 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. However, given the similarities in the paths of the two hours series, the bias in labor productivity introduced by the lack of adjustment is unlike 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 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. 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 with it, since, while free of trends, such a series still displays long cycles. As shown in the third row of figure 7, 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 enough 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, 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 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: Cross correlation structural shocks-omitted variables, dummy specification, 1955-2000

Figure 8 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

16

outcome is confirmed by the results of bivariate Granger causality tests between the three 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.

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 4 lags and 4 leads of the disturbances. Hence, estimated technological shocks do not stand-in for other sources of technological and non-technological disturbances.


Figure 9: Contemporaneous response of hours, dummy specification, 1955-2000

Another way to examine the potential effect of omitted variables on the relation-

ship 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 9. 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 both shocks becomes greater in magnitude.

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. 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 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 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 qualita-

tive 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 an event 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.

The first row of figure 10 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 cyclicality. In particular, the series displays throughs 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 throughs 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 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 these episodes we present in figure 10 an historical decomposition exercise. Each panel 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 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.
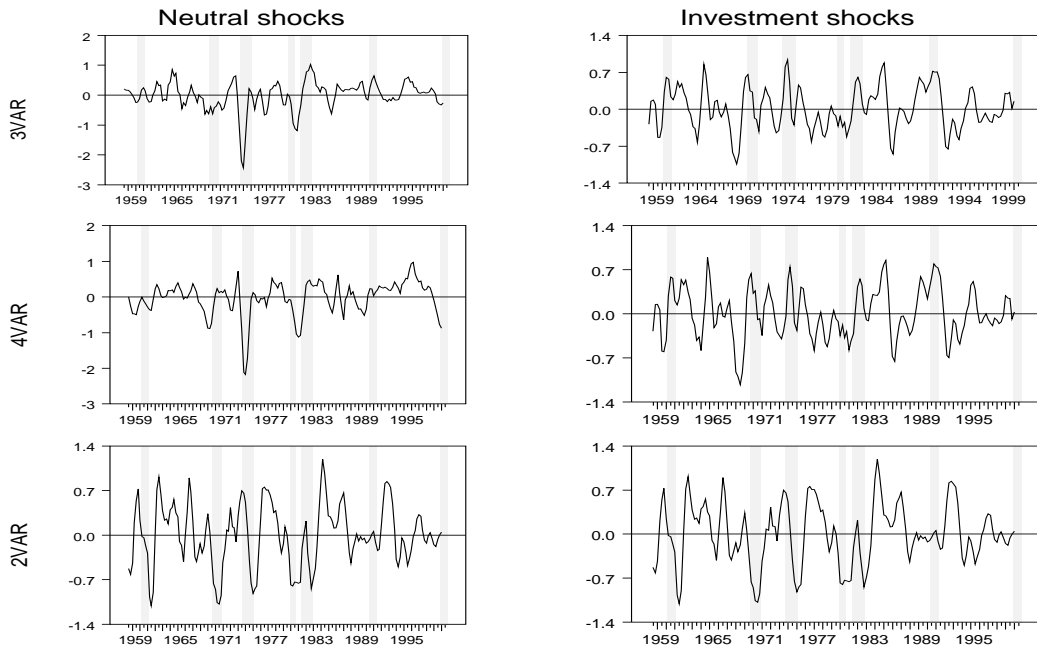

Figure 10: Technology shocks

Figure 11 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

20

of investment and labor productivity.

Figure 11: Historical decomposition

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 $\frac{y_t}{n_t} = (\frac{u_t k_t}{n_t})^\alpha A_t$, 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 12 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 in some NBER recessions. Similar conclusions are obtained if we sample annually our neutral shocks and compare them with the aggregate technol-

21

ogy 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 12: 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 12 plots the resulting series and the estimated investment specific shocks. As it was the case with 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 the relationships in the second row of figure 12. It is clear that the correlation between the two series 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 specific shocks are negative. 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-statisitic* 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, 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 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 the time series properties of technology shocks to vary in bivariate systems. Indeed, the time series properties of technology shocks do change substantially in bivariate systems (com-

pare 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 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.

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 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. It turns out that such a shock is negatively and significantly correlated with the news shock contemporaneously (estimate is -0.03 and t-statistic -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 throughs at major NBER recessions, and induces positive per-capita hours and output responses.

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 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 splitting the sample in pieces is the right approach to follow. The first choice induces overdifferencing and enhances the importance of measurement error; the second introduces important small sample biases (and inefficiencies) in the estimation. Allowing the intercept of the hours series in the VAR to vary over time is instead a much more natural 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, 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 in any reasonable way, the data robustly suggest a number of interesting facts which we believe are useful for empirically distinguishing models of the business cycle. 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.

25

- 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 may be 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 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 attempted to do this 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] Canova F, Lopez-Salido D, Michelacci C. 2006. Schumpeterian technology shocks. Available at www.crei.cat/people/canova/

[5] Canova F, Gambetti L, Pappa E. 2007. The Structural Dynamics of Output Growth and Inflation: Some international Evidence. *Economic Journal* **117**: C167-C191.

[6] Chari V, Kehoe P, McGrattan E. 2005. A critique of SVAR using Real Business Cycle theory. Fed of Minneapolis, working paper 631.

[7] 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.

[8] Comin D, Gertler M. 2006. Medium Term Business Cycles. *American Economic Review* **96**: 523-551.

[9] 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.

[10] Erceg C, Guerrieri L, Gust C. 2005. Can Long Run Restrictions Identify technology shocks. *Journal of the European Economic Association* **3**: 1237-1278.

[11] Evans, C. 1992. Productivity Shocks and Real Business Cycles. *Journal of Monetary Economics* **29**: 191-208.

[12] Faust J, Leeper E. 1997. Do Long Run Restrictions give reliable Results? *Journal of Business and Economic Statistics* **15**: 343-353.

[13] Fernald J. 2007. Trend Breaks, Long Run Restrictions, and the Contractionary Effects of Technology Shocks. *Journal of Monetary Economics* **54**: 2467-2485.

[14] Fernandez-Villaverde J, Rubio-Ramirez J, Sargent T, Watson M. 2007. The ABC (and D) for understanding VARs, *American Economic Review* **97**: 1021-1026.

[15] Fisher J. 2006. The Dynamic Effects of Neutral and Investment-Specific Technology Shocks. *Journal of Political Economy* **114**: 413-451.

[16] 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.

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

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

[19] 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.

[20] Gambetti L. 2005. Technology shocks and the response of hours worked: time varying dynamics matter. Available at http:// www.pareto.uab/gambetti.

[21] Giordani P. 2004. An alternative explanation of the price puzzle. *Journal of Monetary Economics* **51**: 1271-1296.

[22] Greenwood J, Yorokoglu M. 1997. 1974. *Carnegie-Rochester Conference Series on Public Policy* **46**: 49-95.

[23] Kehoe T, Ruhl K. 2007. Are Shocks to the Terms of Trade Shocks to Productivity?. NBER Working Paper 13111.

[24] Michelacci C, Lopez-Salido D. 2007. Technology Shocks and Job Flows. *Review of Economic Studies* **74**: 1195-1227.

[25] Uhlig H. 2004. Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association* **2**: 361-371.

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