

Schumpeterian Technology Shocks*

Fabio Canova
ICREA-UPF

David Lopez-Salido
Federal Reserve Board

Claudio Michelacci[†]
CEMFI

This version: November, 2007

Abstract

We analyze the effects of neutral and investment-specific technology shocks on hours worked and unemployment. We characterize the response of unemployment in terms of job separation and job finding rates. Job separation rates mainly account for the impact response of unemployment while job finding rates for movements along its adjustment path. Neutral shocks increase unemployment and explain a substantial portion of unemployment and output volatility; investment-specific shocks expand employment and hours worked and mostly contribute to hours worked volatility. This evidence is consistent with the view that neutral technological progress prompts Schumpeterian creative destruction, while investment specific technological progress has standard neoclassical features.

JEL classification: E00, J60, O33.

Key words: Search frictions, technological progress, creative destruction.

*We thank Jason Cummins, Gianluca Violante, Robert Shimer, Pau Rabanal, Sergio Rebelo, Gary Solon and Ryan Michaels for kindly making their data available to us. We also appreciate comments from Marios Angeletos, Roc Armenter, Robert Barro, Olivier Blanchard, Jesus Fernandez-Villaverde, Robert Hall, Jim Nason, Ivan Werning, Tao Zha, and participants at the 2006 CEPR-ESSIM conference, and seminar audiences at MIT, New York Fed, Philadelphia Fed, and Atlanta Fed. The opinions expressed here are solely those of the authors and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

[†]Authors are also affiliated with CREI, AMeN, and CEPR; CEPR; and CEPR, respectively. Address for correspondence: CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. Tel: +34-91-4290551. Fax: +34-91-4291056. Email: c.michelacci@cemfi.es, fabio.canova@upf.edu, david.j.lopez-salido@frb.gov.

1 Introduction

There has been a renewed interest in examining how labor market variables respond to technology shocks. The analysis has generally focused on the dynamics of total or per-capita hours worked—see, among others, Galí (1999), Uhlig (2004), Francis and Ramey (2005), and Fernald (2007). This focus is partly motivated by having as reference the basic neoclassical growth model, where a representative household offers his labor services in a competitive market. However, such a focus obscures whether fluctuations in labor input are due to fluctuations in hours per employee (the intensive margin of labor market adjustment) or in the number of employed workers (the extensive margin) and whether employment adjustments arise because of changes in the hiring or in the firing policies of firms. Analyzing these different margins can convey useful information for at least two reasons. First, hours and employment have different volatility and their correlation is far from perfect (see, for example, Cooley, 1995). Second, worker flows provide key insights into employment adjustments. For example, the conventional wisdom has generally been that recessions—periods of sharply rising unemployment—begin with a wave of layoffs and persist over time because unemployed workers have hard time to find a new job. Shimer (2005b) and Hall (2005) have challenged this view by arguing that the flow of workers out of jobs hardly increases in recessions. But are all the recessions alike? Can we safely neglect the role of the separation rate in characterizing unemployment dynamics?

In this paper we address these issues by analyzing how labor market variables respond to technology shocks along the extensive and the intensive margin. We characterize employment dynamics in terms of the job separation rate (the rate at which workers move from employment to unemployment) and the job finding rate (the rate at which unemployed workers find a job). Our analysis focuses on the response to investment-neutral and investment-specific technology shocks. The identification restrictions we use are taken directly from Solow (1960) growth model and require that investment specific technological progress is the unique driving force for the secular trend in the relative price of investment goods, while neutral and investment specific technological progress explain long-run movements in labor productivity (see also Altig et. al. (2005), Fisher (2006) and Michelacci and Lopez Salido (2007)).

As in Blanchard and Quah (1989) and in Fernald (2007), we recognize that low frequency movements could give a misleading representation of the effects of shocks. This is a relevant concern since in the sample the growth rate of both labor productivity and the relative price of investment goods exhibit significant long run swings which have gone together with important changes in labor market conditions. These patterns have been greatly emphasized in the literature on growth and wage inequality (see Violante, 2002 and Greenwood and Yorokoglu, 1997, among others). The productivity revival of the late 90's has also been heralded as the beginning of a new era in productivity growth and it has been a matter of extensive independent research, see for example Gordon (2000), Jorgenson and Stiroh (2000). Once we efficiently take care of the low frequency movements in the variables entering the VAR we find that:

1. Labor market adjustment mainly occurs along the extensive margin in response to neutral technology shocks and the intensive margin in response to investment specific technology shocks.
2. Neutral technology shocks increase unemployment. The separation rate accounts for the impact response of unemployment, the finding rate for its dynamic adjustment. Thus, the response to a neutral technology shock is in line with the conventional wisdom: unemployment initially rises because of a wave of layoffs and remains high because the job finding rate takes time to recover.
3. Investment specific technology shocks expand aggregate hours worked both because hours per worker increase and because unemployment falls. Again, the job separation rate accounts for a major portion of the impact response of unemployment, and the job finding rate for its dynamic path.
4. Neutral technology shocks explain a substantial proportion of the volatility of unemployment and output while investment specific technology shocks mainly account for the volatility of hours worked. Taken together, technology shocks explain around 30 per cent of the cyclical fluctuations of key labor market variables at time horizons between 2 and 8 years.
5. Our estimated technology shocks accurately characterize certain historical business cycle episodes. In particular, the recession of the early 90's and the subse-

quent remarkably slow labor market recovery appear to be driven almost entirely by advancements in the neutral technology. Neutral technology shocks initially cause a rise in job separation and unemployment; output builds up until it reaches its new higher long run value, but over the transition path employment remains below normal levels because of the low job finding rate. This makes the output recovery appear to be “jobless”.

These findings are robust to the choice of the lag length, to the presence of omitted variables, to the identification scheme, to the measurement of the labor variables, and to other auxiliary assumptions needed in specifying the VAR.

We shows that this evidence is consistent with the Schumpeterian view that the introduction of new neutral technologies causes the destruction of technologically obsolete productive units and the creation of new technologically advanced ones. In the presence of search frictions in the labor market, this leads to a temporary rise in unemployment. Investment specific technological progress has instead standard neoclassical features. Schumpeterian creative destruction matters for productivity dynamics at the micro level, see Foster et al. (2001) and it is a prominent paradigm in the growth literature, see Aghion and Howitt (1994), Mortensen and Pissarides (1998), Violante (2002) and Hornstein et al. (2005). Yet it has generally been overlooked in business cycle analysis—notable exceptions are Caballero and Hammour (1994, 1996) and Michelacci and Lopez-Salido (2007). We show that Shumpeterian creative destruction also plays a key role in qualitatively and quantitatively explaining key business cycle episodes, including the jobless recovery of the early 90’s.

The standard explanation for the fall in hours in response to neutral technology shocks is based on sticky-prices, see for example Galí (1999). In sticky-price models, when technology improves and monetary policy is not accommodating enough, demand is sluggish to respond and firms take advantage of technology improvements to economize on labor input. This mechanism applies most naturally to the intensive margin, since displacing workers is typically more costly than changing prices—due to both the direct cost of firing and the value of the sunk investment in training and in job specific human capital that is lost with workers displacement(see e.g. Mankiw, 1985

and Hamermesh, 1993 for a review of the literature).¹ The Schumpeterian mechanism has distinctive features and different policy implications. For example, the extensive margin plays a key role in the adjustment since the fall in hours caused by creative destruction is due to the time consuming process of reallocation of workers across productive units. The policy implications also differ. The fall in hours in sticky-price models is due to an inefficient response of monetary policy. In our model, it is the result of the (possibly) socially optimal process of technology adoption in the presence of creative destruction and search frictions in the labor market.

This paper extends in a number of ways the analysis of Michelacci and Lopez-Salido (2007) who consider a vintage model with search frictions in the labor market to analyze the effects of technology shocks on job flows. First, the model economy we use to interpret the evidence is a streamlined version of theirs. Second, we focus the analysis on the social planner problem to provide evidence that the labor market response to technology shocks may be socially efficient. Third, the labor market flow data here are representative of the whole US economy rather than just of the manufacturing sector. Fourth, instead of using job creation and job destruction rates, we consider workers flow data, which have recently received much attention in the business cycle literature. This allows us to revisit the conclusions by Shimer (2005b) and Hall (2005) on the importance of the separation rate over the business cycle. Finally, our analysis of the episode of the recession of the early 90's and the subsequent remarkably slow labor market recovery provides novel evidence that Schumpeterian technology shocks are important in explaining jobless recoveries.

There are two reasons why our conclusions are apparently at odds with those in Shimer (2005b) and Hall (2005). First, our analysis is *conditional* on technology shocks rather than unconditional. This allows to separately quantify the contribution of the separation rate to unemployment dynamics on impact and over the adjustment path. Second, our analysis shows that part of the response of the finding rate is due to the initial response of the separation rate: since both finding and separation rates are

¹It should be pointed out, however, that this is only a conjecture since we are aware of no paper that has formally analyzed the trade-off between changing prices and displacing workers. This would require replacing the staggered price setting a la Calvo (1983) of standard sticky price models with an endogenous decision to change prices subject to menu-costs (as in Caballero and Engel, 2007) and then allowing firms to choose whether to adjust hours or employment at a cost.

endogenous, the separation rate can indirectly contribute to cyclical unemployment through the effects it exerts on the finding rate. This may provide a warning to the recent tendency to use search models with exogenous separation rates for business cycle analysis.

The rest of the paper is structured as follows. Section 2 discusses the identification of shocks. Section 3 describes the data and shows the biases caused by low frequency movements. Section 4 presents impulse responses. Section 5 examines the role of potentially omitted variables. Section 6 quantifies the contribution of the separation and finding rates to unemployment dynamics. Section 7 presents a model which helps to interpret the evidence. Section 8 analyzes cyclical fluctuations induced by technology shocks. Section 9 deals with robustness. Section 10 concludes.

2 Identification of technology shocks

We use a version of Solow (1960) model to decompose aggregate productivity into the sum of a stationary component and a component driven by neutral and investment specific technology shocks. This decomposition holds in several versions of the model (including the one in Section 7), and justifies its use for identification purposes.

Solow model Assume technological progress is exogenous and the rate of saving and capital depreciation are stationary. The production function is:

$$\tilde{Y} = Z\tilde{K}^\alpha N^{1-\alpha}, \quad 0 < \alpha < 1,$$

where \tilde{Y} is final output, \tilde{K} is capital, N is labor and Z is the investment-neutral technology. Final output can be used for either consumption \tilde{C} , or investment \tilde{I} . A stationary fraction of output s is invested, $\tilde{I} = s\tilde{Y}$. Next period capital is

$$\tilde{K}' = (1 - \delta)\tilde{K} + Q\tilde{I},$$

where $0 < \delta < 1$ is a stationary depreciation rate. The variable Q formalizes the notion of investment specific technological change. A higher Q implies a fall in the cost of producing a new unit of capital in terms of output or an improvement in the quality of new capital produced with a given amount of resources. If the sector producing new capital is competitive, the inverse of its relative price is an exact measure of Q .

One can check that this economy evolves around a (stochastic) trend given by

$$X \equiv Z^{\frac{1}{1-\alpha}} Q^{\frac{\alpha}{1-\alpha}}$$

and that the quantities $Y \equiv \tilde{Y}/(XN)$, and $K \equiv \tilde{K}/(XQN)$ converge to $Y^* = (s/\delta)^{\frac{\alpha}{1-\alpha}}$ and $K^* = (s/\delta)^{\frac{1}{1-\alpha}}$, respectively. As a result the logged level of aggregate productivity, $y_n \equiv \ln \tilde{Y}/N$, evolves according to

$$y_n = y^* + v + x = y^* + v + \frac{1}{1-\alpha}z + \frac{\alpha}{1-\alpha}q \quad (1)$$

where small letters denote the log of the corresponding quantities in capital letters and v is a stationary term which accounts for transitional dynamics. Equation (1) decomposes aggregate productivity into the sum of a stationary term plus a trend induced by the evolution of the neutral and the investment specific technologies. This result can be used to identify technology shocks from a VAR: a neutral technology shock (a z -shock) is the disturbance having zero long-run effects on the relative price of investment goods and non-negligible long-run effects on labor productivity; an investment specific technology shock (a q -shock) affects the long-run level of both labor productivity and the price of investment. No other shock has long-run effects on the price of investment or labor productivity. In the simple model above, the price of investment goods is exogenous and there is a one-to-one mapping between the price and q . But in more general models with variable capital utilization and adjustment costs, the short run marginal cost of producing capital is increasing and the price of investment goods responds in the short run to any change in investment demand. Since in the long run investment specific technological progress is the only determinant of the price of investment, our identification strategy is robust to the existence of short run increasing marginal costs to produce investment goods.

Choice of deflator There is some controversy on how the price of investment and GDP should be deflated. We show below that our results are not sensitive to the choice of deflator. In our baseline specification we deflate them both by using the output deflator. Fisher (2006) and Michelacci and Lopez-Salido (2007) instead deflate both of them by the CPI index. Altig et al. (2005) appear to deflate the relative price of investment with the CPI index, and output with the output deflator (although

they are not entirely clear about the issue). In a closed economy, and if we exclude indirect taxes, the CPI and the output deflator are similar, but in an open economy important differences arise because some consumption goods are produced abroad.² In the appendix we show that our approach is consistent with the balanced growth conditions of a well defined *open* economy, while other approaches may imply that the decomposition (1) no longer holds exactly and that the real exchange rate, in addition to the z and the q shocks, determines long run productivity (see also Kehoe and Ruhl (2007) for a similar point). Using the GDP deflator is equivalent to use as a numeraire *domestic* consumption—i.e. the consumption goods produced in the US. The Consumer Price Index, P_c , is $P_c = \left(\frac{P_c^H}{a}\right)^a \left(\frac{P_c^F}{1-a}\right)^{1-a}$, where P_c^H and P_c^F are the prices of consumption goods produced in the US and abroad, respectively; and a represents the share of domestic consumption goods. Let q^c and y_n^c denote the inverse of the relative price of investment and labor productivity (both in logs), when deflated with the CPI index. In appendix A we show that

$$y_n^c = cte + \frac{1}{1 - \alpha - \beta} z + \frac{\alpha + \beta}{1 - \alpha - \beta} q^c + \frac{1}{1 - \alpha - \beta} (1 - a) (p_c^H - p_c^F) \quad (2)$$

where α and β are the output elasticities to domestic and foreign capital, respectively. Hence, with this choice of numeraire, a permanent change in the real exchange rate affects long run labor productivity measured in CPI units and could be confused with “neutral” technology shocks. This may be a relevant concern since the real exchange rate is known to exhibit remarkable persistence. Similarly, when we deflate the relative price of investment with the CPI index and output with the GDP deflator we obtain that

$$y_n = cte + \frac{1}{1 - \alpha - \beta} z + \frac{\alpha + \beta}{1 - \alpha - \beta} q^c + \frac{\alpha + \beta}{1 - \alpha - \beta} (1 - a) (p_c^H - p_c^F),$$

and again a permanent change in $p_c^H - p_c^F$ has long run effects on productivity.

Empirical implementation Let X_t be a $n \times 1$ vector of variables and let X_{1t} and X_{2t} be the first difference of q_t and y_{nt} , respectively. The Wold representation of

²There could be differences between the output deflator and the CPI even in closed economies because the CPI only includes a subset of the consumption goods while the output deflator includes all of them. Moreover, the weights in the CPI may differ from those in the output deflator because the CPI measures the prices paid by *urban* consumers for a market basket of consumer goods and services.

$X_t = (X_{1t}, \dots, X_{nt})$ is $X_t = D(L)\eta_t$, where $D(L)$ has all its roots inside the unit circle and $E(\eta_t\eta_t') = \Sigma_\eta$. In general, η_t is a combination of several structural shocks ϵ_t . We assume a linear relationship between η_t and ϵ_t , $\eta = S\epsilon$, where, by convention, the first element of ϵ_t is taken to be the q -shock and the second the z -shock. We also assume that the structural shocks ϵ_t are uncorrelated and normalize their variance so that $E(\epsilon_t\epsilon_t') = I$. Under this normalization, impulse responses represent the effects of shocks of one-standard deviation of magnitude. The restrictions that the nonstationarities in q_t and y_{nt} originate exclusively from technology shocks imply that the first row of $G = D(1)S$ is a zero vector except in the first position, while the second row is a zero vector except in the first and second position.³ With the assumed orthogonality of structural shocks, these restrictions are sufficient to separate the two technology shocks and to analyze the dynamic responses to each disturbance.

3 Effects of low frequency comovements on the VAR

Our benchmark model has six variables $X = (\Delta q, \Delta y_n, h, u, s, f)'$, where Δ denotes the first difference operator. All variables are in logs (and multiplied by one hundred): q is equal to the inverse of the relative price of a quality-adjusted unit of new equipment, y_n is labor productivity, h is the number of per-capita hours worked (thereafter simply hours), u is the unemployment rate and s and f are the job separation rate and the job finding rate. The dynamics of hours per worker in response to shocks can be obtained assuming that labor force participation is insensitive to shocks. We show below that this is a reasonable assumption. The dynamics of output per-capita can be obtained from those of labor productivity and hours. We use 8 lags and stochastically restrict their decay toward zero.

The series for labor productivity, unemployment, and hours are from the USECON database commercialized by Estima and are all seasonally adjusted; q is from Cummins and Violante (2002), who extend the Gordon (1990) measure of the quality of new equipment till 2000:4. The availability of data for q restricts the sample period to 1955:1-2000:4. The original series for q is annual and it is converted into quarters as

³Equation (1) implies that G_{12} , the long run effect of a q -shock on labor productivity is $\frac{\alpha}{1-\alpha}$. We leave this coefficient unrestricted since its exact magnitude depends on the production function and the details of the law of motion of the capital stock.

in Galí and Rabanal (2004).⁴

The series for the job separation and the job finding rates are from Shimer (2005b). They are quarterly averages of monthly rates. Shimer calculates two different series for the job separation and job finding rates. The first two are available from 1948 up to 2004. Their construction uses data from the Bureau of Labor Statistics for employment, unemployment, and unemployment duration to obtain the *instantaneous* (continuous time) rate at which workers move from employment to unemployment and viceversa. The two rates are calculated under the assumption that workers move between employment to unemployment and viceversa. Since they abstract from workers' labor force participation decisions, they are an approximation to the true labor market rates. Starting from 1967:2, the monthly Current Population Survey public microdata can be used to directly calculate the flow of workers that move in and out of the three possible labor market states (employment, unemployment, and out of the labor force). With this information Shimer calculates an exact instantaneous rates at which workers move from employment to unemployment and viceversa. We analyze both measures: the first two are termed *approximated* rates, the others *exact* rates.

The first graph in the first row of Figure 1 plots hours and the unemployment rate together with NBER recessions (the grey areas). Hours display a clear U-shaped pattern and are highly negatively correlated with unemployment (-0.8). Whether the two series are stationary or exhibit persistent low frequency movements, is matter of controversy in the literature, see for example Fernald (2007) and Francis and Ramey (2005). The second graph plots hours worked per employee (measured as hours over aggregate employment). Clearly, the series exhibits some low frequency changes, primarily at the beginning of the 1970s.

The two graphs in the second row of Figure 1 plot the first difference of y_n and of the relative price of investment (equal to minus q), respectively. One can notice the existence of a dramatic fall in the value of q in 1974 and its immediate recovery in the

⁴Real output (LXNFO) and the aggregate number of hours worked (LXNFH) correspond to the non-farm business sector. The relative price of investment is expressed in output units by subtracting to the (log of the) original Cummings and Violante series the (log of) the output deflator (LXNFI) and then adding the log of the consumption deflator $\ln((CN+CS)/(CNH+CSH))$. Here CN and CS denotes nominal consumption of non-durable and services while CNH and CSH are the analogous values of consumption in real terms. The aggregate number of hours worked per capita is calculated as the ratio of LXNFH to the working age population (P16), i.e. $h \equiv \ln(LXNFH/P16)$.

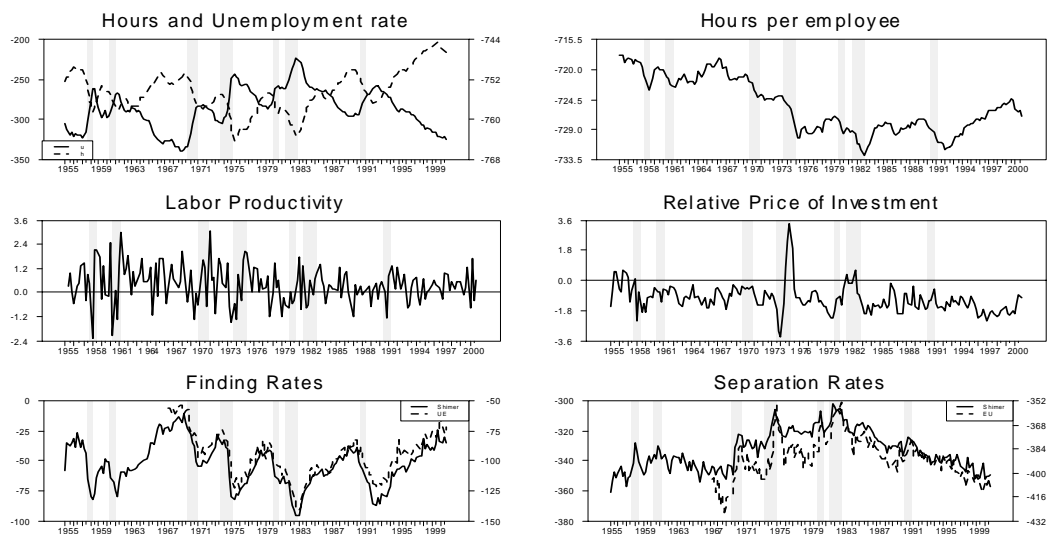


Figure 1: First graph: the dashed line is the aggregate number of hours worked per capita; the continuous line is civilian unemployment both series in logs. Second graph: (logged) hours per employee. Third graph: rate of growth of labor productivity in the non-farm business sector. Fourth graph: growth rate of the relative price of investment goods. Fifth and sixth graph: job finding rate and job separation rate (both in logs), respectively. The solid line corresponds to the approximated rate, the dashed to the exact rate. Shaded areas are NBER recessions.

following years. Cummins and Violante (2002) attribute this to the introduction of price controls during the Nixon era. Since price controls were transitory, they do not affect the identification of investment specific shocks, provided that the sample includes both the initial fall in q and its subsequent recovery. The two panels in the third row of Figure 1 display the job finding rate and the job separation rate. Each graph plots approximated and exact rates. The two job finding rate series move quite closely. The exact job separation rate has a lower mean in the 1968-1980 period, higher volatility but tracks the approximated series well. The job finding rate is relatively more persistent than the separation rate (AR1 coefficient is 0.86 vs. 0.73). Recessions are typically associated with a persistent fall in the job finding rate. This has motivated Shimer (2005b) and Hall (2005) to claim that cyclical fluctuations in the unemployment rate are driven mainly by fluctuations in the job finding rate.

The low frequency co-movements of the series are highlighted in Figure 2. We follow

the growth literature and choose 1973:2 and 1997:1 as a break points, two dates that many consider critical to understand the dynamics of technological progress and of the US labor market (see Greenwood and Yorokoglu, 1997, Violante, 2002, Hornstein et al. 2002). The rate of growth of the relative price of investment goods was minus 0.8 per cent per quarter over the period 55:1 to 73:1 and moved to minus 1.2 per cent per quarter in the period 73:2-97:1. This difference is statistically significant. During the productivity revival of the late 90's the price of investment goods was falling at even a faster rate. The rate of growth of labor productivity exhibits an opposite trend. It was higher in the 55:1 to 73:1 period than in the 73:2-97:1 period, and recovered in the late 90's. Also in this case, differences are statistically significant. Shifts in technological progress occurred together with changes in the average value of the unemployment rate, see the first row of Figure 2.

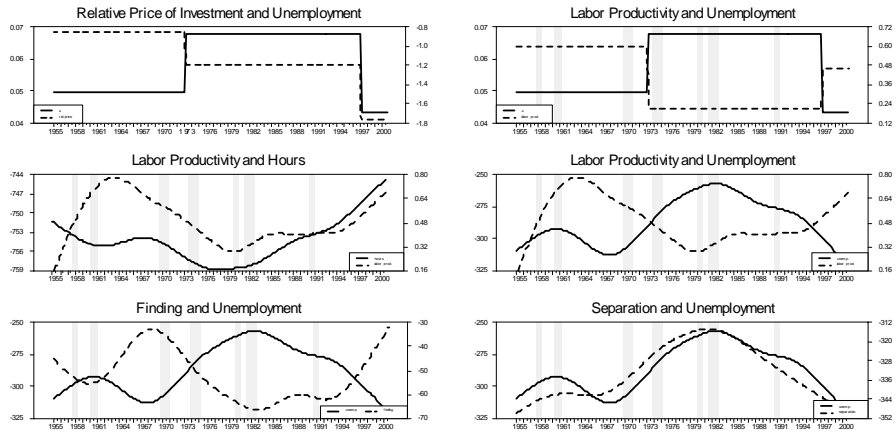


Figure 2: First graph: average quarterly growth rate of the relative price of investment (dotted line) and unemployment rate (solid line). Second graph: average quarterly growth rate of labour productivity (dotted line) and unemployment rate (solid line). Third graph: Hodrick Prescott trend of labor productivity growth (dotted line) and hours per capita (solid line). Fourth graph: Hodrick Prescott trend of labor productivity growth (dotted line) and unemployment rate (solid line). Fifth and sixth graph: Hodrick Prescott trend of finding and separation rates (dotted lines) and unemployment rate (solid line). The smoothing coefficient is $\lambda = 1600$.

The graphs in the second row of Figure 2 plot the trend component of labor productivity growth, hours and unemployment obtained by using a Hodrick Prescott filter

with smoothing coefficient equal to 1600. The trends are related: there appears to be a negative comovement between productivity growth and the unemployment rate and a positive comovement between productivity growth and hours. The third row of Figure 2 shows that the separation rate exhibits low frequency movements that closely mimic those present in the unemployment rate. The opposite is true for the finding rate. Next, we show why these comovements are problematic.

The effects of low-frequencies comovements on impulse responses Panel (a) in Figure 3 displays the responses of labor productivity, the relative price of investment, unemployment, hours, hours per employee, the separation rate, and the finding rate to a neutral shock. Panel (b) present responses to an investment specific shock. We plot together the point estimates obtained for three different samples: 1955:I-2000:IV, 1955:I-1973:I, and 1973:II-1997:I. The responses of labor productivity and output to either shock in the full sample specification are similar to those found in Fisher (2006). However, we have a slight initial fall in hours and in the price of investment in response to a neutral shock that Fisher does not have. As shown in Canova et al. (2006), the inclusion of the additional labor market variables (unemployment, job finding rate and job separation rate) in the empirical model explain the differences.

When considering panel (a), it is apparent that the estimated responses to neutral shocks in the two subsample are similar. Yet, they look quite different from the responses for the full sample. In the full sample, the relative price of investment and the separation rate fall, while they increase in the two subsamples. Moreover the fall in hours and the job finding rate and the increase in unemployment are much less pronounced in the full sample than in each sub-sample. Finally, output and labor productivity respond faster in the full sample. The potential bias present in the estimated responses for the full sample can be related to the low frequency correlations previously discussed. In the full sample, a permanent change in the rate of productivity growth is at least partly identified as a series of neutral technology shocks. Thus, over the period 1973:II-1997:I when productivity growth is on average lower, the full sample specification finds a series of negative neutral technology shocks. Since in this period the unemployment rate and the separation rate are above their full sample average, while hours and the finding rate are below, biases emerge leading, for example, to a lower

response of the unemployment rate and of the separation rate, and a higher response of hours and the job finding rate.

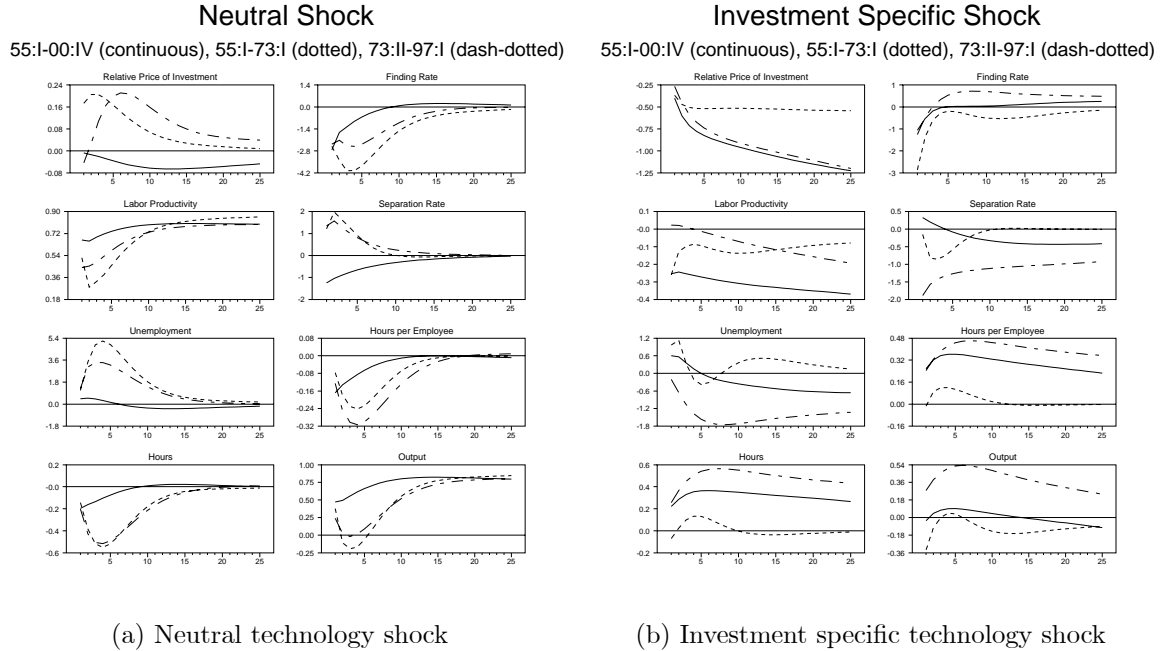


Figure 3: Responses to a one-standard deviation shocks. Each line corresponds to a six variable VAR(8) with the rate of growth of the relative price of investment, the rate of growth of labour productivity, the (logged) unemployment rate, and the (logged) aggregate number of hours worked per capita, the log of separation and finding rates, estimated over a different sample period.

Panel (b) deals with the effects of investment specific shocks for the same three samples. In comparing the results, one should bear in mind two important facts (see Figures 11 and 12 in Appendix C): i) the estimated responses in the first subsample are almost never significant (with the exception of the response of the relative price of investment) and ii) investment specific technology shocks contribute little to the volatility of all variables in the first subsample (again leaving aside the price of investment). In the second sub-period the contribution of investment specific shocks instead becomes important. Hence, it is appropriate to compare estimates for the full sample and the 1973:2-1997:1 sub-period. The bias in the estimated responses for the full sample is in line with the low frequency correlations previously discussed. In the full sample, a permanent change in the rate of growth of the relative price of investment is at least partly identified as a series of investment specific technology shocks. Thus,

over the period 1973:II-1997:I when the price of investment falls at a faster rate on average, the full sample specification tends to identify a series of positive investment specific technology shocks. Since over the period, the unemployment rate and the separation rate are also higher than their full sample average, while hours, the job finding rate, and productivity growth are lower, the full sample specification biases estimates towards a higher response of the unemployment rate and of the separation rate, and a lower response of hours, the job finding rate, and productivity.

These results are robust to a number of modifications: they are unaffected if the second subsample is 1973:II-2000:IV (see panels (a) and (b) in Figures 13 in Appendix C) or if we use the population-adjusted hours produced by Francis and Ramey (2005). In fact, as shown in Canova et al. (2006), this adjusted hours series exhibits the same low frequency variations as the one used here. In sum, sub-sample instabilities maybe minor and difference with the full sample estimates are due to the low frequency comovements exhibited by the variables of the VAR.

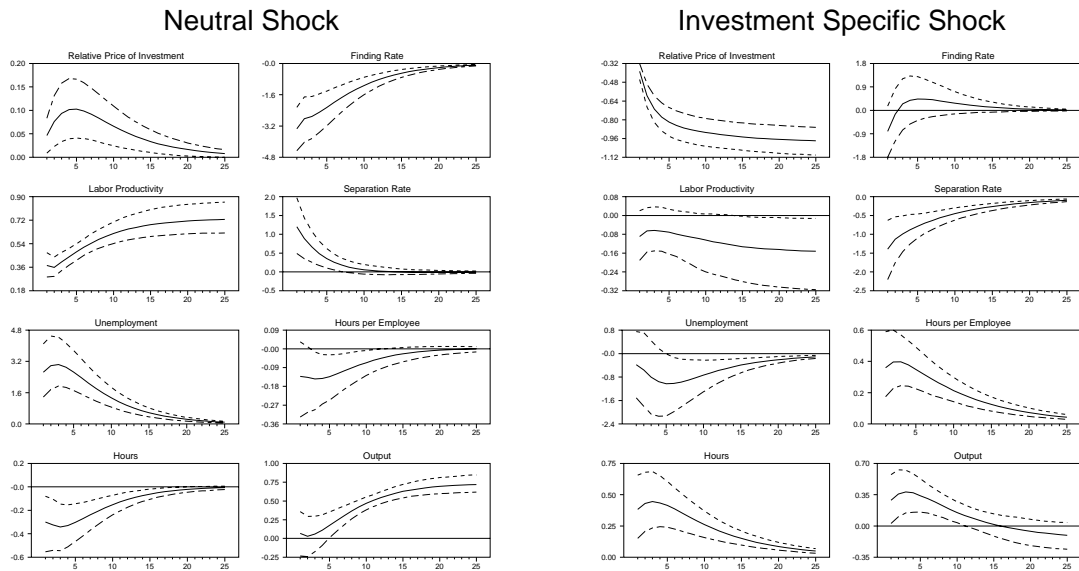
4 The full sample results after dealing with trends

To tackle the issue of the low frequency comovements one could estimate the VAR in each sub-sample. Splitting the sample is however inefficient, since the dynamics are roughly unchanged over the sub-samples. Moreover, imposing as identifying long run restrictions in a system estimated over a small sample may induce biases in the structural estimates (see Erceg et al. 2005). As an alternative, we allow the intercept of all VAR equations to vary over time but restrict the slopes to be time invariant. We have considered several options: in the baseline specification (the “dummy” specification) the intercept is deterministically broken at 1973:2 and 1997:1. We show below that conclusions are robust to several alternative low frequency removal approaches.

4.1 Evidence using the approximated rates

Panel (a) in Figure 4 plots the response of the variables of interest to a neutral technology shock for the full sample using the approximated job finding and job separation rates. The reported bands correspond to 90 percent confidence intervals. A neutral shock leads to an increase in unemployment and to a fall in the aggregate number of

hours. The effects on hours worked per employee are small and generally statistically insignificant. The impact rise in unemployment is the result of a sharp rise in the separation rate and of a significant fall in the job finding rate. In the quarters following the shock, the separation rate returns to normal levels while the job finding rate takes up to fifteen quarters to recover. Hence, the dynamics of the job finding rate explains why unemployment responses are persistent. Output takes about 5 quarters to significantly respond but then gradually increases until it reaches its new higher long-run value. Interestingly, once low frequency movements are taken into account, the dynamic responses for the full sample look like those of the two subsamples.



(a) Neutral technology shock

(b) Investment specific technology shock

Figure 4: Responses to a one-standard deviation shocks. Full sample with intercept deterministically broken at 1973:II and 1997:I. Six variables VAR(8). Dotted lines are 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line is the median estimate.

Panel (b) in Figure 4 plots responses to an investment specific shock. The estimated responses are very similar to those obtained in the 1973:2-1997:1 sub-sample. An investment specific technology shock leads to a short run increase in output and hours per capita and a fall in unemployment. The fall of unemployment on impact is due to a sharp drop in the separation rate. Since this effect is partly compensated by a fall

in the job finding rate, the initial fall in unemployment rate is small in absolute terms and statistically insignificant. Hence, the increase in the number of hours is primarily explained by the sharp and persistent increase in the number of hours worked per employee. Thus, labor market adjustment to an investment specific technology shock mainly occurs along the intensive margin.

4.2 Evidence using the exact rates

We next analyze the effects of technology shocks when considering exact job finding and separation rates. Again, we report results obtained with the dummy specification. Panel (a) in Figure 5 presents the responses to a neutral technology shock with the exact rate (dotted line) together with the previously discussed responses obtained with the approximated rates (solid line). Both specifications agree on the sign and shape of the responses. There are however two important quantitative differences. When considering the exact rates, the separation rate rises on impact twice as much, while the finding rate falls slightly less. Furthermore, over the adjustment path the separation rate exhibits more persistence when exact rates are used.

Panel (b) in Figure 5 reports responses to an investment specific technology shock when exact and approximated rates are used. Also in this case, the two specifications agree on the sign and shape of the responses, but there are two significant quantitative differences. When the exact rates are used, the response of the separation rate is more pronounced and falls on impact twice as much. Instead, the job finding rate is now unaffected on impact and remains above normal levels all along the adjustment path. As a result, the fall in the unemployment rate is more pronounced both on impact and during the transition suggesting that the extensive margin plays a more important role in accounting for the rise in hours when exact rates are used. Nevertheless, the increase in hours per employee remains predominant.

5 Omitted variables

Our specification has allowed for enough lags, so that the residuals are clearly white noise processes. Yet, it is possible that omitted variables play a role in the results. For example, Evans (1992) showed that Solow residuals are correlated with a number of

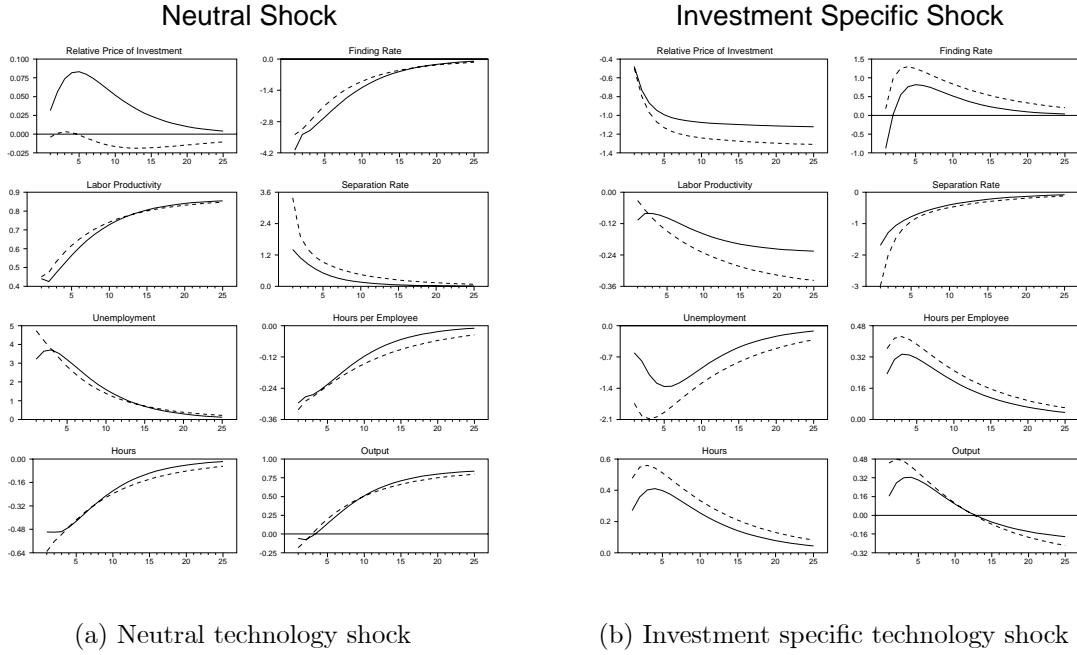


Figure 5: Exact rates (dotted lines) and approximated rates (solid lines). Both VAR includes dummies corresponding to the breaks in technology growth. Each VAR has 8 lags and six variables. Reported are point estimates of the responses.

policy variables, therefore making responses to Solow residuals shocks uninterpretable. To check for this possibility we have correlated our two estimated technology shocks with variables which a large class of general equilibrium models suggest as being jointly generated with neutral and investment specific shocks. In particular, we compute correlations up to 6 leads and lags between each of our technology shocks and the consumption to output ratio, the investment to output ratio, and the inflation rate. The point estimates of these correlations together with an asymptotic 95 percent confidence tunnel around zero are in Figure 6. The shocks are obtained in the dummy specification with the approximated rates (similar results are obtained with the exact rates).

There is some evidence that the consumption to output and the investment to output ratios help to predict neutral technology shocks, while none of the three potentially omitted variables significantly correlate with investment specific shocks. Hence, we investigate what happens when we enlarge the system to include these three new variables. Panels (a) and (b) in Figures 14 in Appendix C present the responses when considering a VAR which includes the original six variables plus the consumption to

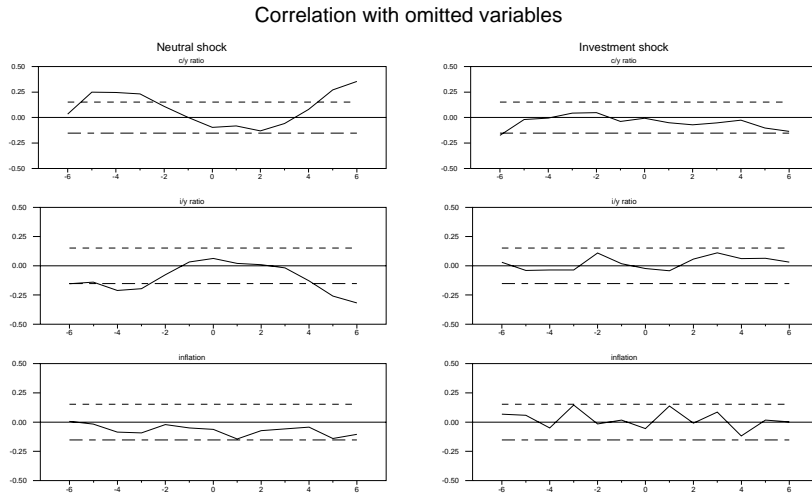


Figure 6: Left column corresponds to neutral technology shocks; right column to investment specific technology shocks. The first row plots the correlation with the consumption-output ratio, the second with the investment-output ratio, the third with the inflation rate. The shocks are estimated from the six variables VAR with approximated rates in the dummy specification. The horizontal lines correspond to an asymptotic 95 percent confidence interval for the null of zero correlation.

output and the investment to output ratios and the inflation rate in the dummy specification, when approximated rates are used. None of our previous conclusions is affected and this is still the case when exact rates are used. Only the volatility of technology shocks falls somewhat when considering the extended VAR.

6 The dynamics of fictional unemployment rates

Shimer (2005b) and Hall (2005) have challenged the conventional view that recessions—defined as periods of sharply rising unemployment—are the result of higher job-loss rates. They argue that recessions are mainly explained by a fall in the job finding rate. Our impulse responses suggest instead that the separation rate plays a major role in determining the impact effect of technology shocks on unemployment. This is consistent with the evidence by Fujita and Ramey (2006) that the separation rate leads the cycle (by about one quarter) while the finding rate lags it (by about two months).

To further evaluate the role of the separation rate, we use a simple two state model

of the labor market (see Jackman et al., 1989 and recently Shimer, 2005b) and we assume that the stock of unemployment evolves as:

$$\dot{u}_t = S(l_t - u_t) - Fu_t \quad (3)$$

where l_t and u_t are the size of the labor force and the stock of unemployment, respectively; while S and F are the separation and finding rates in levels, respectively. The unemployment rate tends to converge to the following *fictional* unemployment rate:

$$\tilde{u} = \frac{S}{S + F} \equiv \frac{\exp(s)}{\exp(s) + \exp(f)}.$$

Shimer (2005b) shows that the fictional unemployment rate \tilde{u} tracks quite closely the actual unemployment rate series, so that one can fully characterize the evolution of the stock of unemployment just by characterizing the dynamics of labor market flows. After linearizing the log of \tilde{u} , we can calculate its response using the information contained in the response of (the log of) the separation rate s and the finding rate f . This allows to measure the contribution of finding and separation rates to the cyclical fluctuations of fictional unemployment \tilde{u} ; and to evaluate how accurately fictional unemployment approximates actual unemployment, that is, whether workers movements in and out of the labor force play a role in determining unemployment.

Panel (a) in Figure 7 reports results for the specification with approximated rates, panel (b) with the exact rates. In both cases, the same nine variable VAR employed in section 5 is used. In each panel, the response of the true unemployment rate appears with a solid line and the response of (logged) \tilde{u} appears with a dotted line. The dash-dotted line corresponds to the response of (logged) \tilde{u} that would be obtained if the job finding rate had remained unchanged at its average level in the sample. It therefore represents the contribution of the separation rate to fluctuations in fictional unemployment.

There are several important features of figure 7. First, the dynamics of fictional unemployment after a neutral shock are explained to a large extent by fluctuations in the separation rate, especially when considering the specification with exact rates. In agreement with previous results, the separation rate explains almost 90 per cent of the impact effect on fictional unemployment. However, after only one quarter, its contribution falls to 40 per cent and drops to 20 per cent one year after the shock.

There are some differences in the impact response of actual and fictional unemployment. This suggests that workers movements in and out of the labor force play some role in characterizing the response of the unemployment rate, at least on impact.

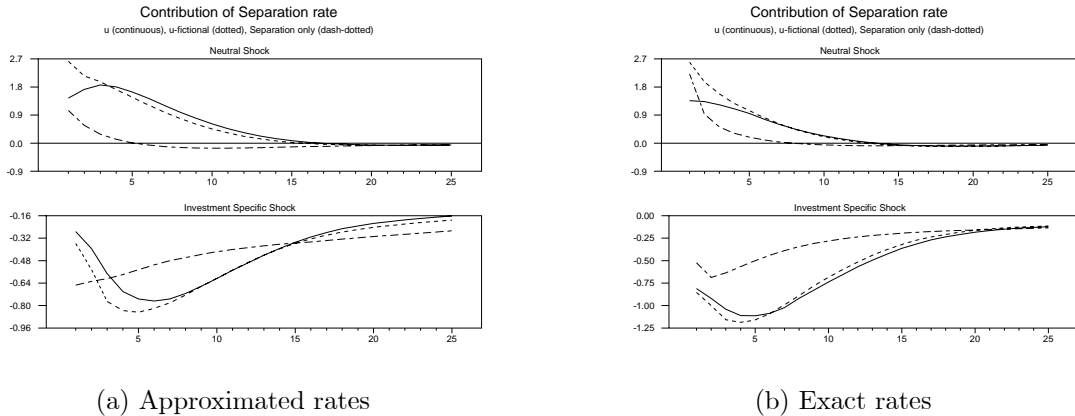


Figure 7: Nine variables VAR with approximated or exact rates. Full sample with deterministic time dummies. Reported are median estimates from 500 bootstrap replications.

Following an investment specific shock, in the specification with approximated rates, unemployment falls little on impact because the fall in the separation rate makes unemployment decrease while the fall in the job finding rate makes unemployment increase. When considering the specification with exact rates, unemployment falls substantially on impact and this is mainly due to the fall in the separation rate. The differences between the response of fictional and actual unemployment are minimal both with approximated and with exact rates. Hence, other labor market flows are likely to play a minor role in determining the unemployment responses to investment specific shocks. This reinforces the conclusion that labor market adjustments to investment specific shocks mainly occur along the intensive margin.

7 Interpretation

Next, we present a model which can be used to interpret the evidence we have uncovered. We assume there are no frictions in the adoption of the investment specific technology, while we impose a vintage structure on the neutral technology. As reviewed by Brynjolfsson and Hitt (2000), the idea is that the adoption of new neutral technologies, such as more effective managerial practices, better organization of production

or the introduction of new products and services, alter the entire job structure of the firm and require the performance of new tasks for which previously hired workers may not be suitable. This implies that the firm has to replace part of its current workforce to upgrade its technology. In the analysis we focus on the social planner problem to stress that the observed responses could result from the optimal process of technology adoption in the presence of Schumpeterian creative destruction and labor market frictions.⁵ We first describe the economy and then discuss impulse responses. Appendix B contains the derivation of equilibrium conditions.

7.1 Assumptions

There is one consumption good, the numeraire. Output is produced according to

$$\tilde{Y} = F(\tilde{K}, \tilde{H}) = \tilde{K}^\alpha \tilde{H}^{1-\alpha},$$

where \tilde{K} is the capital stock and \tilde{H} the amount of labor intensive intermediate goods used in production. Labor intensive intermediate goods are produced in *jobs* which consist of firm-worker pairs. A worker can be employed in at most one job where he supplies one unit of labor at an effort cost (in utility terms) c_w . A job with *neutral technology* z produces an amount of intermediate goods equal to $\exp\left(\frac{z}{1-\alpha}\right)$. As in standard vintage models (see for example Jovanovic and Lach, 1989, Caballero and Hammour, 1996, and Aghion and Howitt, 1994), newly created jobs always embody leading-edge technologies while *old jobs* do not upgrade their previously installed technologies. Specifically, a job which starts producing at time t operates with a neutral technology z_{it} equal to the economy *leading technology* z_t of that time, while the current period neutral technology of old jobs, z_{it} , remains (in expected value) unchanged:

$$z_{it} = z_{it-1} + \epsilon_{it} \tag{4}$$

⁵Michelacci and Lopez Salido (2007) analyze the decentralized equilibrium of a related economy. Caballero (2007) discusses the inefficiencies involved in the Schumpeterian process of worker reallocation across productive units. Generally the inefficiencies due to search frictions are solved when workers can direct their search as in the models by Moen (1997) and Acemoglu and Shimer (1999).

where ϵ_{it} is an idiosyncratic shock which is iid normal with standard deviation σ_ϵ .⁶ The leading edge neutral technology evolves as:

$$z_t = \mu_z + z_{t-1} + \epsilon_{zt} \quad (5)$$

where ϵ_{zt} is iid normal with standard deviation σ_z . Hereafter, we will refer to the difference between the leading technology z_t and the job's neutral technology z_{it} as the job *technological gap*, $\tau_{it} \equiv z_t - z_{it}$.

The law of accumulation of capital is $\tilde{K}' = (1 - \delta)\tilde{K} + e^q \tilde{I}$ where \tilde{I} is the amount of investment expenditures measured in final output and q is the investment specific technology, which evolves according to

$$q_t = \mu_q + q_{t-1} + \epsilon_{qt} \quad (6)$$

where ϵ_{qt} is iid normal with standard deviation σ_q .

At every point in time jobs are exogenously destroyed with probability λ . Jobs can also be destroyed when their technological gap is too large relative to an endogenously determined critical threshold τ_t^* . Jobs created at time t starts producing at time $t + 1$. Creating new jobs requires the services of recruiters. The cost of creating n new jobs involves a cost in utility terms to recruiters equal to:

$$C(u, n) = cu^{-\eta_0} n^{\eta_1}, \quad \eta_0, \eta_1 > 0 \quad (7)$$

so that unemployment reduces the cost of creating new jobs, as it is standard in search models, see e.g. Pissarides (2000). This formulation embeds others present in the literature. For example, if the matching function has constant returns to scale and the utility cost of posting a vacancy is constant, then $\eta_1 - \eta_0 = 1$. If instead the cost of posting vacancies is increasing in the number of posted vacancies or in the number of newly created jobs, as in Caballero and Hammour (1996) and Michelacci and Lopez-Salido (2007), $\eta_1 - \eta_0 > 1$.

The population of workers is constant and normalized to one. We assume that a representative household exists so that workers and recruiters pool their income at the

⁶The idiosyncratic shocks ϵ guarantee that the cross-sectional distribution of job technology has no mass points. In turn, this property ensures a smooth transitional dynamics by ruling out the possibility that persistent oscillations occur over the transition path —i.e. the “echo effects” that typically arise in vintage models, see for example Benhabib and Rustichini (1991).

end of the period and choose consumption and effort costs to maximize the sum of the expected utility of the household's members. The instantaneous utility is:

$$\ln \tilde{C} - c_w(1 - u) - C(u, n) \quad (8)$$

where \tilde{C} is aggregate consumption, while u and n denote the unemployment rate and the flow of newly created jobs, respectively. The last two terms in (8) account for the effort cost of working for workers and recruiters, respectively. The household's discount factor is β . The aggregate resource constraint is: $\tilde{Y} = \tilde{I} + \tilde{C}$.

We adopt the following convention about the timing of events within a period t :

- i. Aggregate technology shocks ε_{zt} and ε_{qt} are realized;
- ii. Old jobs realize whether their job is exogenously destroyed (which occurs with probability λ) and their idiosyncratic shocks ε_{it} . New jobs (resulting from matches at time $t - 1$) start with neutral technology z_t ;
- iii. Decisions about job destruction, job creation, and investment are taken;
- iv. Output is produced, income pooled and consumed. Next period begins.

7.2 Equilibrium conditions

Let $f_t(\tau)$ denote the time- t measure of old jobs which, in case they are kept in operation, would produce with technological gap τ . In the described sequence of events, this is the distribution resulting after the events in ii). Then unemployment is

$$u_t = 1 - \int_{-\infty}^{\tau_t^*} f_t(\tau) d\tau - n_{t-1} \quad (9)$$

since jobs are destroyed when their technology gap is greater than the critical technological gap τ_t^* , while all newly created jobs are productive. The fraction of jobs destroyed between time $t - 1$ and time t (i.e. the job separation rate) is

$$S_t = \lambda + \frac{\int_{\tau_t^*}^{\infty} f_t(\tau) d\tau}{1 - u_{t-1}}$$

while the job finding probability for workers searching between time $t - 1$ and time t is $F_t = \frac{n_t}{u_{t-1}}$ so that unemployment evolves as

$$u_t = u_{t-1} + S_t(1 - u_{t-1}) - F_t u_{t-1}$$

which is the discrete time analogue of equation (3). Jobs are created up to the point that the marginal cost of job creation is equal to its expected future net value, so that

$$c\eta_1 n_t^{\eta_1-1} u_t^{-\eta_0} = \beta E_t(V_{t+1}(0)) \quad (10)$$

where $V_{t+1}(0)$ is the next period (utility) value of a job with technological gap equal to zero (see Appendix B for further details).

This economy fluctuates around the stochastic trend given by $X_t \equiv e^{x_t}$, where

$$x_t = \frac{1}{1-\alpha} z_t + \frac{\alpha}{1-\alpha} q_t$$

Hence, we scale quantities by X_t and solve log-linearizing the first order conditions around the steady state version of the model when $\varepsilon_{z,t} = 0$ and $\varepsilon_{q,t} = 0$. To characterize the beginning-of-period distribution, f_t , we follow Campbell (1998) and Michelacci and Lopez-Salido (2007) and use values at a fixed grid of technological gaps.⁷

The logged level of unscaled aggregate productivity, $y_{nt} \equiv \ln(Y_t/(1-u_t)) + x_t$, evolves as in (1), where $Y_t \equiv \tilde{Y}_t/X_t$ denotes scaled aggregate output. Specifically, let Y and $1-u$ denote the constant level of scaled output and employment around which the economy fluctuates. Then (1) holds for $y^* = \ln Y - \ln(1-u)$ and v accounts for the stationary fluctuations of Y_t and $1-u_t$ around their mean. Therefore, our identification approach is fully consistent with the structure of this model.

7.3 The response to technology shocks

We calibrate the model at the quarterly frequency and derive the implied average monthly rate of the associated labor market flows to make results comparable with the empirical analysis.⁸ The values of the parameters used are in Table 1. Most of the choices are standard, including the value of the discount factor β and of the output elasticity to capital, α , (see for example Cooley (1995)). Following Greenwood et al. (1997), μ_z , μ_q and δ are chosen so as to yield, at the yearly level, a growth rate of z of 0.39 percent, a growth rate of q of 3.21 per cent, and a capital depreciation

⁷A Computational Appendix (available at <http://www.cemfi.es/~michela>) describes in more detail the procedure used.

⁸Alternatively we could calibrate the model at the monthly frequency and aggregate the results at the quarterly level. This alternative approach, however, would force us into specifying when the shock has occurred within a given quarter, an issue that can be sidestepped here.

Parameter Values						
β :0.99	α :0.36	μ_z :0.0975%	μ_q : 0.8025%	δ : 3.2%	σ_z :0.56%	σ_q :1.3%
λ :3%	η_0 :0.66	η_1 :4	σ_ϵ : 4.9%	c_w : 0.62	c :408.85	

Table 1: Parameters values used in the baseline specification.

rate of capital of 12.4 per cent, respectively. The standard deviation of the shocks is obtained from the previous analysis by noticing that a one-standard-deviation neutral technology shock leads to a long run increase in labor productivity of approximately 0.85 percentage points, while a one-standard-deviation investment specific shock leads to a long-run fall in the relative price of investment of 1.3 percentage points, see Figure 5. The resulting volatility values are just smaller than those typically used in the real business cycle literature (see for example Cooley (1995) and Greenwood et al. (2000)). The parameter λ is obtained as in den Haan et al. (2000), assuming that exogenous separation accounts for about one half of total separation and η_0 is set by assuming that it exist a constant return to scale matching function where the matching elasticity to unemployment is 0.4, which is the estimated value by Blanchard and Diamond (1990). To calibrate η_1 we assume that the cost of posting vacancy is increasing in the number of newly created jobs, say, because recruits require some training to be productive in new jobs and recruiters have decreasing marginal utility to leisure. If we assume that these services are exchanged in a competitive labor market, we can use standard estimates for the Frisch elasticity of labor supply—which is typically slightly greater than one half, see Blundell et al. (1993) and Lee (2001)—together with the reported estimates for the matching elasticity to vacancies to estimate η_1 . The remaining parameters σ_ϵ , c_w , and c are set to match, in the steady state version of the model without aggregate shocks, i) that the fraction of existing jobs more productive than a newly created job is around 60 percent, ii) that the job finding probability is 80 per cent, and iii) that the separation rate is 6 percent. The first condition is in line with Baily et al. (1992). The last two are the quarterly counterpart of a monthly job finding rate of 40 percent and separation rate of 2 per cent, which are the averages in our sample.

Panel (a) in Figure 8 characterizes the response of the economy to a one-standard deviation z -shock, (i.e. an increase in ε_{zt} of σ_z). These responses are obtained inte-

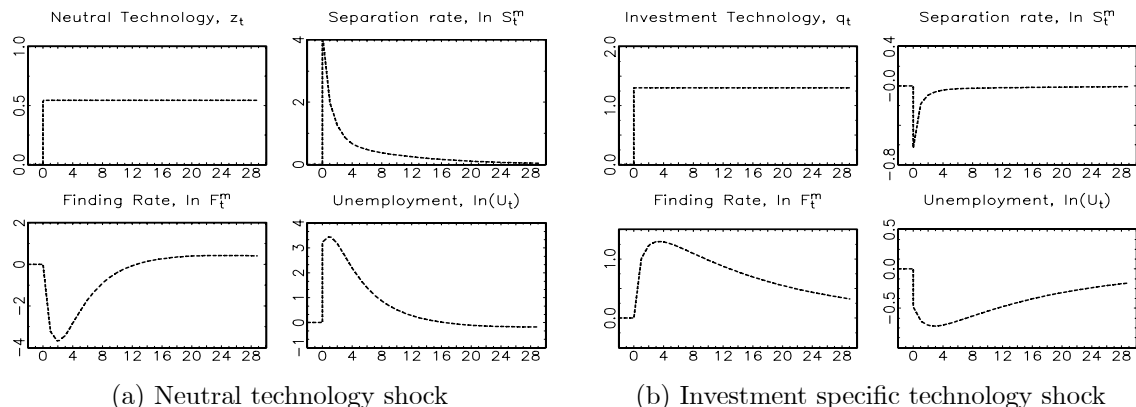


Figure 8: Impulse responses to a one-standard deviation z -shock at time zero in the model, Panel (a), and to to one-standard deviation q -shock, Panel (b). The separation and finding rates are implied average monthly rate. All responses are multiplied by 100.

grating out of the aggregate decision rules variables which are not used in the VAR. As discussed in Canova (2007) this implies that the model responses are fully compatible with those obtained in the empirical analysis. The implied monthly separation rate S_t^m and finding rate F_t^m , are obtained by using the relations $1 - S_t = (1 - S_t^m)^3$ and $1 - F_t = (1 - F_t^m)^3$. Neutral technology shocks bring about a simultaneous increase in the destruction of technologically obsolete jobs and in the creation of new highly productive units which prompts a contractionary period during which employment temporarily falls. More formally, as z_t increases, jobs with a given technology become obsolete relative to the technological frontier so the distribution of old jobs $f_t(\tau)$ shifts to the right on impact. This leads to an initial cleansing of technologically outdated jobs which makes the separation rate and the unemployment rate increase. Quantitatively, the shock leads to an increase of about four per cent in the monthly separation rate and in the unemployment rate, which are close to what we obtained in the VAR with exact rates, see Panel (a) in Figure 5.

In the quarters after the shock, more jobs are created both because the pool of unemployed workers has increased and because the value of new jobs $V_t(0)$ has increased. Thus, the initial upsurge in unemployment is gradually absorbed and, as new jobs embody the more advanced technology, output, investment and consumption reach their permanently higher new long-run value. Unemployment takes around 4 years to return to normal levels—which is in line with the empirical evidence. The dynamics of the

job finding rate, that remains below its steady state level over the whole adjustment path, explains these persistent effects. The job finding probability falls because the increase in reallocation pushes up the costs of job creation, which slows the pace of job creation. Quantitatively, the maximal fall in the job finding rate is of about four percentage points, which is similar to the effects obtained in the VAR with exact rates, see Panel (a) in Figure 5.⁹

Panel (b) in Figures 8 presents responses to a one-standard deviation fall in the price of capital (i.e. an increase in ε_{qt} of σ_q). As q_t rises, it is optimal to accumulate more capital. Since capital accumulation is costly, \tilde{C}_t falls below its state value. This reduces the value of the effort cost of working which increases the value of jobs with a given technological gap τ . This pushes up the critical technological gap τ_t^* and makes the separation rate fall. In other words, the desire to smooth consumption makes the economy spread over time the pruning of relatively outdated technologies, so more obsolete technologies are temporarily kept in operation. Quantitatively, the job separation rate falls by less than a percentage point, which is slightly smaller than the effect observed in the data. In the quarters following the shock, job creation falls due to the reduction in the pool of searching workers. The initial fall in unemployment is gradually absorbed and, after about seven years, employment returns to its pre-shock level while output, consumption and productivity reach their new long-run values. The persistent effects on unemployment are driven by the response of the job finding rate, that remains above its steady state level over the adjustment path. This is due both to the increase in the value of new jobs $V_t(0)$, and to the fall in reallocation that reduces the cost of job creation. Quantitatively, the shock leads to a maximal increase in the job finding rate of more than one per cent, in line with the empirical findings.

⁹Shimer (2005a) has recently questioned the ability of the standard Mortensen and Pissarides (1994) model to reproduce the right volatility of key labor market variables. Hagedorn and Manovskii (2006) show that this is possible only if the difference between job output and the income forgone by employed workers is low enough. In our baseline calibration the difference between new jobs output and the value of the effort cost of working is around to 0.191, which is somewhat closer to the favorite value by Hagedorn and Manovskii of 0.057 than to the value of 0.6 chosen by Shimer (2005a).

7.4 The intensive margin

It is interesting to analyze the response of the intensive margin to technology shocks and contrast it with the response of the extensive margin. For this purpose, assume that a job with neutral technology z , produces an amount of intermediate goods equal to $\exp\left(\frac{z}{1-\alpha}\right) e$, where e denotes the number of hours worked in the job. Assume also that the utility cost of working e hours is: $c_w = \bar{c} + c_e \frac{e^{1+\frac{1}{\phi}}}{1+\frac{1}{\phi}}$, where ϕ is the elasticity of the disutility of working with respect to the number of hours worked. At any point in time and for any job, the social planner chooses e so that

$$(1 - \alpha) \left(\frac{\tilde{K}_t}{\tilde{H}_t} \right)^\alpha \exp\left(\frac{z}{1-\alpha}\right) \frac{1}{\tilde{C}_t} = c_e e^{\frac{1}{\phi}},$$

which can be solved to obtain the equilibrium number of hours worked, and to trace out how they respond to shocks. Since the adoption of new technologies requires time and investment in capital, \tilde{C}_t falls below its long run value in response to either technology shock, so the marginal disutility of working falls, and a worker in a job with a given technological gap works longer hours. As a result the average number of hours worked per employee increases. Thus, in response to a z -shock, the number of employed workers fall, but the average number of hours worked per employee tend to increase. This composition effect is such that neutral shocks contribute relatively less to the volatility of aggregate hours worked than to the volatility of unemployment while the opposite is true for investment specific shocks which is precisely what we find in the data (see next section).

8 The contribution of technology shocks

Here we analyze the contribution of technology shocks to business cycle fluctuations. Table 1 reports the forecast error variance decomposition using either the approximated rates or the exact rates. We focus on the VAR(8) with nine variables (given by the growth in the investment relative price and in labor productivity, hours, the unemployment rate, job finding and job separation rates, the consumption and the investment to output ratio, and the inflation rate). In the six variables VAR, the contribution of technology shocks is slightly larger.

Variable	Neutral				Investment specific			
	Horizon (quarters)				Horizon (quarters)			
	1	8	16	32	1	8	16	32
A. Approximated rates, full sample								
Investment Relative Price	16	13	12	12	42	45	46	46
Labor Productivity	23	21	21	21	3	4	4	4
Output	1	6	30	55	3	5	5	4
Hours	8	9	8	7	14	16	21	22
Hours per Worker	5	5	4	4	17	23	29	29
Unemployment	23	21	21	21	3	3	6	6
Finding Rate	17	17	17	17	0	1	2	2
Separation Rate	10	8	7	6	5	8	12	14
B. Approximated rates, 1973:II-2000:IV sample								
Investment Relative Price	4	3	4	3	38	36	34	35
Labor Productivity	18	18	18	18	0	1	1	1
Output	1	4	24	43	22	11	10	9
Hours	12	14	12	11	37	18	20	21
Hours per Worker	10	10	8	9	44	30	31	32
Unemployment	12	18	16	14	13	2	2	3
Finding Rate	7	13	12	12	4	1	2	2
Separation Rate	28	28	12	14	2	4	8	12
C. Exact rates								
Investment Relative Price	3	2	3	3	35	35	34	34
Labor Productivity	7	11	11	11	1	1	2	2
Output	8	4	17	37	14	8	6	6
Hours	22	19	18	16	24	15	14	14
Hours per Worker	14	12	11	10	35	27	28	28
Unemployment	34	30	29	27	3	1	1	1
Finding Rate	1	25	24	24	0	1	2	3
Separation Rate	34	34	30	26	0	1	1	1

Table 2: Forecast Error Variance Decomposition: percentage of variance explained by neutral or investment-specific technology shocks at different time horizons for the selected variables. All VARs have nine variables with intercept deterministically broken at 1973:II and 1997:I. The variables are the growth in the relative price of investment and in labor productivity, hours per capita, the unemployment rate, the job separation and the job finding rate, the consumption to output ratio, the investment to output ratio, and the inflation rate. Panel A deals with a VAR with approximated rates, Panel B restrict the analysis to the 1973:II-2000:IV sub-sample, Panel C deals with the exact rates.

Neutral technology shocks explain a substantial proportion of the volatility of unemployment and output. In the specification with approximated rates, neutral technology shocks explain between 30 and 50 per cent of output fluctuations at time horizons between 4 and 8 years and 20 percent of unemployment volatility (see panel A). The contribution of neutral technology shocks to fluctuations in hours per worker is however small. Investment specific technology shocks instead account for a substantial proportion of the volatility of hours worked: around 20 per cent of the volatility of hours per capita and 30 per cent of the volatility of hours per worker. The contribution of investment specific technology shocks to output and unemployment volatility is instead small (generally smaller than 10 per cent). Taken together, technology shocks explain a relevant proportion of the business cycle volatility: at horizons between 2 and 8 years they explain around 40 per cent of the volatility of output, and about 30 per cent of the volatility of unemployment and hours. The importance of technology shocks is generally greater when exact rates are used (see panel C). This is however due to the greater importance of technology shocks in the 1973:II-2000:IV sample period. When we estimate the VAR with approximated rates in the 1973:II-2000:IV sample, we find that technology shocks explain roughly the same amount with approximated and exact rates (see panel B). The only exception is in the contribution of neutral technology shocks to the volatility of the separation rate, which is three times larger with exact rates.

To further examine whether technology shocks are an important source of cyclical fluctuations, we analyze the historical contribution of technology shocks to fluctuations in logged unemployment, job separation and job finding. The graphs in the left column of Figure 9 represent as a solid line the original series and as a dotted line its component due to technology shocks (either neutral or investment specific), as recovered from the nine variables VAR in the dummy specification with the exact rates. All series are detrended with a Hodrick Prescott filter with smoothing parameter equal to 1600. It is apparent that technology shocks are an important driving force of business cycles. They explain several business cycle episodes including the recession of the early 80's and of the early 90's and the subsequent recovery. The graphs in the right column permit us to evaluate how accurately the model replicates the fluctuations due to technology in the corresponding variable. Each graph contains the previously

calculated technology component of the relevant series (again represented as a dotted line) together with the model generated series obtained by feeding the z -shocks and the q -shocks recovered by the VAR into the model. Hence, the two series share the shocks while the transmission mechanism is independently obtained. If the series look alike, there is evidence that the model closely replicates the transmission mechanism of the data. Overall, the model is quite successful in quantitatively reproducing the technology component of unemployment and job separation of the data and reproduces well the dynamics of the finding rate, although fluctuations are slightly larger in the model than in the data.

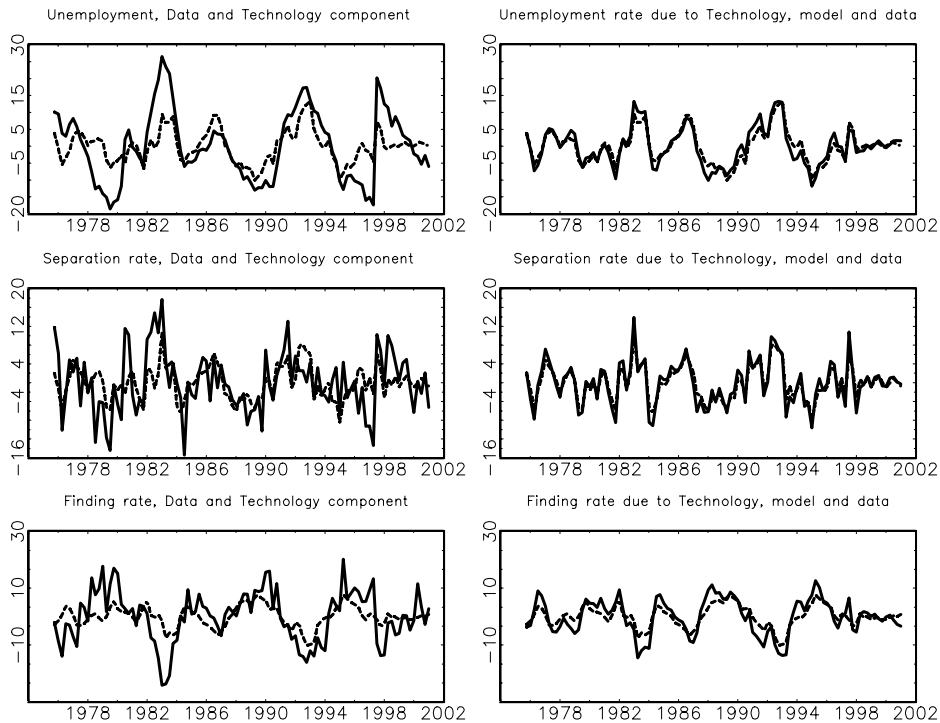


Figure 9: Effects of technology shocks in data and model. Left column: solid line is the raw data, the dotted line the component due to technology shocks (either neutral or investment specific) as recovered from the nine variables VAR with the exact rates. Right column: dotted line is the component due to technology shocks in the data, solid line is the series obtained after feeding the shocks obtained from the VAR into the model. The separation and finding rates correspond to the implied average monthly rate. All series are detrended with a Hodrick Prescott filter with smoothing parameter equal to 1600.

Finally, we study the recession of the early 1990s and the subsequent recovery. This episode have been extensively investigated in the literature, yet its causes are still un-

explained; see for example Bernanke (2003). A key feature of the episode is that the downturn in employment was severe. Another is that the peak in unemployment occurred about two years later than the trough in output. This is a remarkable exception relative to other business cycle episodes, see McKay and Reis (2007). The graphs in the left column of Figure 10 presents the original output and unemployment series (solid lines) and their component due just to technology shocks (dotted lines), again obtained from the nine variables VAR(8) with the exact rates. All series are detrended with the Hodrick Prescott filter. The vertical lines capture the NBER recession. Technology shocks explain well the recession of the early 90's and the subsequent remarkably slow recovery in the labor market. This is due to the contribution of neutral technology shocks that naturally tend to induce jobless recoveries since, following the initial rise in job separation and unemployment, output increase to their new higher long run value, while unemployment remains above trend because of the low job finding rate. Hence, the rise in output appears to be jobless. The graphs in the right column show how the model can account for the jobless recovery of the early 90's. It plots as dotted line the technology component of the original series and as a solid line the model generated series obtained by feeding the technology shocks recovered by the VAR into the model. Again all series are detrended with the Hodrick Prescott filter. The model accurately reproduces both the magnitude of the effects and the faster recovery of output relative to employment. Only the fall in output at the start of the recession is somewhat more pronounced in the data than in the model.

9 Robustness

This section briefly describes some robustness exercises we have undertaken. The conclusions are that our technology shocks are unlikely to stand in for other sources of disturbances and that our results persist when we change i) the method to remove low frequency fluctuations, ii) the lag length, iii) the identifying restrictions, iv) the price deflator and v) the labor market data and the series for the relative price of investment.

Other disturbances Despite the fact that our technology shocks do not proxy for omitted variables, it is still possible that they stand in for other sources of dis-

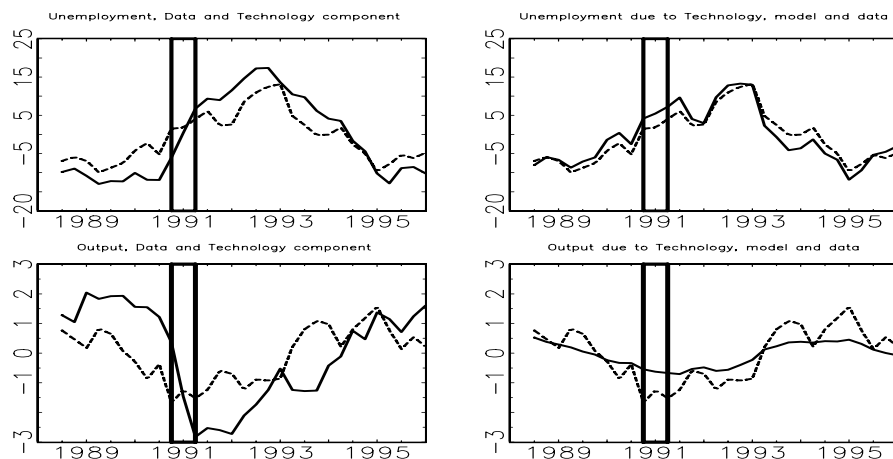


Figure 10: The jobless recovery of the 90s. Left Column: Solid lines are raw data (either unemployment or output), the dotted lines the component due to technology shocks (either neutral or investment specific) as recovered from the nine variables VAR with the exact rates. Right Column: the dotted line is again the component due to technology shocks in the data, the solid line is the technology component generated by the model after feeding the technology shocks from the VAR into the model. All series are detrended with a Hodrick Prescott filter with smoothing parameter equal to 1600. The vertical lines identifies the NBER recession.

turbances. To check for this possibility, we have correlated the estimated technology shocks obtained from the nine variables VAR with the approximated rates with oil price and federal fund rate shocks.¹⁰ Figure 15 in Appendix C shows that correlations are insignificant.

Alternative treatments of trends We have considered two alternatives to remove low frequency movements: we have allowed up to a fifth order polynomial in time in the intercept; we filtered all the variables, before entering them in the VAR, with the Hodrick Prescott filter with a smoothing parameter $\lambda = 10000$. Figure 16 in Appendix C show that responses have the same shape and approximately the same size as with the dummy specification.

VAR lag length The issue of the length of VAR has been recently brought back to the attention of applied researchers by Giordani (2003) and Chari et al. (2005), who show that the aggregate decision rules of a subset of the variables of a model may

¹⁰The mnemonics for the corresponding variables are PZTEXP and FFED, respectively. Technology shocks are correlated with $\ln(\text{FFED})$ and $\ln(\text{PZTEXP}) - \ln((\text{CN}+\text{CS})/(\text{CNH}+\text{CSH}))$, the last term being the consumption deflator.

have not always be representable with a finite order VAR. This issue is unlikely to be important in our context since we checked that VAR and model based responses are fully compatible. To further investigate whether this is an issue, we have reestimated our VAR using 4, 8 and 12 lags. The results using approximated rates and the dummy specification are in Figure 17 in Appendix C. Responses are unchanged.

Medium versus long-run identifying restrictions Uhlig (2004) has argued that disturbances other than neutral technology shocks may have long run effects on labor productivity and that, in theory, there is no horizon at which neutral (and investment specific) shocks fully account for the variability of labor productivity. To take care of this problem Uhlig suggests to check if conclusions change when medium term restrictions are used. In Panel (a) and (b) in Figure 18 in Appendix C we report the responses obtained when the restrictions that the two shocks are the sole source of fluctuations in labor productivity and the price of investment is imposed at the time horizon of 3 years rather than in the long-run. The sign and the shape of responses are almost unchanged. Similar results are obtained if the restriction is imposed at any horizon of at least one year.

Relative price effects So far labor productivity and the relative price of investment are deflated by using the output deflator. To investigate whether this choice matters for our results we have computed responses for the VAR with approximated rates in the dummy specification deflating output and the price of investment by the CPI (see Figures 19 in Appendix C). Responses are unaffected by this choice except for the response of the price of investment to a neutral technology shock, which is more pronounced when the price of investment is deflated with the CPI index.

Alternative data sets Elsby et al. (2007) have recently calculated an alternative series for the job finding and job separation rates, by slightly modifying the methodology of Shimer (2005b). Jaimovich and Rebelo (2006) have also extended the series for the investment specific technology up to the mid 2000's. Our results are unaffected by the use of these alternative series for labor market flows and for q .

10 Conclusions

We analyzed the labor market effects of neutral and investment specific technology shocks on unemployment, hours worked and other labor market variables. We characterized the dynamic response of unemployment in terms of job separation and job finding rates. After efficiently taking care of the low frequency movements in the variables entering the VAR we found that the job separation rate accounts for a major portion of the impact response of unemployment. Later unemployment is mainly explained by fluctuations in the job finding rate. Neutral shocks prompt an increase in unemployment while investment specific shocks rise employment and hours worked. Neutral technology shocks are an important source of cyclical variability. They almost entirely explain the recession of the early 90's and the subsequent jobless recovery, a recession typically hard to interpret with conventional models.

We show that the evidence is consistent with the view that neutral technological progress is Schumpeterian, while investment specific progress operates essentially as in a neoclassical growth model. Neutral technology shocks leads to a simultaneous increase in the destruction of technologically obsolete productive units and in the creation of new technologically advanced ones. But since labor market frictions make reallocation sluggish, employment temporarily falls. Contrary to what happens in sticky price models, the rise in unemployment is not ascribed to an inefficient response of monetary policy to technology shocks, but it results from a process of technological adoption in the presence of creative destruction and search frictions in the labor market. We evaluate the quantitative performance of the model by feeding the technology shocks recovered by the VAR into the model and we compare model generated data with their technology component observed in the data. We find that the model is quite successful in quantitatively reproducing the technology component of unemployment, job separation and job finding and accounts well for the business cycle experience of the early 90's.

References

- [1] Acemoglu, D. and Shimer, R. (1999), “Holdups and Efficiency with Search Frictions,” *International Economic Review*, 40, 827-851.
- [2] Aghion, P. and Howitt, P. (1994), “Growth and Unemployment,” *Review of Economic Studies*, 61, 477-494.
- [3] Altig, D. Christiano, J. Eichenbaum, M. and Linde, J. (2005), “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” NBER Working Paper 11034.
- [4] Baily, M. Hulten, C. and Campbell, D. (1992), “Productivity Dynamics in Manufacturing Plants,” *Brookings Papers on Economic Activity: Microeconomics*, 0, 187-249.
- [5] Benhabib, J. and Rustichini, A. (1991), “Vintage Capital, Investment, and Growth,” *Journal of Economic Theory*, 55, 323-339.
- [6] Bernanke, B. (2003), “The Jobless Recovery. Remarks by Governor Ben Bernanke,” speech at the Global Economic and Investment Outlook Conference. Available at www.federalreserve.gov/boarddocs/speeches.
- [7] Blanchard, O. and Quah, D. (1989), “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, 79, 655-673.
- [8] Blanchard, O. and Diamond, P. (1990), “The Aggregate Matching Function,” in Diamond, P. (Eds), *Growth, Productivity, Unemployment, Essays to Celebrate Bob Solow’s Birthday*, MIT Press, Cambridge MA, 159-201.
- [9] Blundell, R., Meghir, C., and Neves, P. (1993), “Labour Supply and Intertemporal Substitution,” *Journal of Econometrics*, 59, 137-160.
- [10] Brynjolfsson, E. and Hitt, L. (2000), “Beyond Computation: Information Technology, Organizational Transformation and Business Performance”, *Journal of Economic Perspectives*, 14, 23-48.
- [11] Caballero, R. and Hammour, M. (1994), “The Cleansing Effect of Recessions,” *American Economic Review*, 84, 1350-1368.
- [12] Caballero, R. and Hammour, M. (1996), “On the Timing and Efficiency of Creative Destruction,” *Quarterly Journal of Economics*, 111, 805-852.
- [13] Caballero, R. (2007), *Specificity and the Macroeconomics of Restructuring*, MIT Press.
- [14] Caballero, R. and Engel, E. (2007), “Price Stickiness in Ss Models: New Interpretations of Old Results,” Mimeo, MIT and Yale University.

- [15] Calvo, G. (1983), “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 983-998.
- [16] Campbell, J. (1998), “Entry, Exit, Embodied Technology, and Business Cycles,” *Review of Economic Dynamics*, 1, 371-408.
- [17] Canova, F. López-Salido, D. and Michelacci, C. (2006), “On the Robust Effects of Technology Shocks on Hours worked and Output,” mimeo UPF.
- [18] Canova, F. (2007), *Methods for Applied Macroeconomic Research*, Princeton University Press.
- [19] Chari, V. , Kehoe, P. and McGrattan, E. (2005), “A critique of SVAR using Real Business Cycle theory,” Fed of Minneapolis, working paper 631.
- [20] Cooley, T.F. (1995) *Frontiers of Business Cycle Research*, Princeton University Press.
- [21] Cummins, J. and Violante, G. (2002), “Investment Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, 5, 243-284.
- [22] Den Haan, W. Ramey, G. and Watson, J. (2000), “Job Destruction and Propagation of Shocks,” *American Economic Review*, 90, 482-498.
- [23] Elsby, M., Michaels, R. and Solon, G. (2007), “The Ins and Outs of Cyclical Unemployment,” NBER Working Paper 12853.
- [24] Erceg, C., Guerrieri, L. and Gust, C. (2005), “Can Long Run Restrictions Identify Technology shocks?,” *Journal of the European Economic Association*, 3, 1237-1278.
- [25] Evans, C. (1992), “Productivity Shocks and Real Business Cycles,” *Journal of Monetary Economics*, 29, 191-208.
- [26] Fernald, J. (2007), “Trend Breaks, Long Run Restrictions, and Contractionary Technology Improvements,” *Journal of Monetary Economics*, forthcoming.
- [27] Fisher, J. (2006), “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114, 413-451.
- [28] Foster, L., Haltiwanger, J. and Krizan C. (2001), “Aggregate Productivity Growth: Lessons from Microeconomic Evidence,” in *New Developments in Productivity Analysis*, NBER / University of Chicago Press.
- [29] Francis, N. and Ramey, V. (2005), “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *Journal of Monetary Economics*, 52, 1379-1399.

- [30] Fujita S. and Ramey, G. (2006), “The Cyclicalities of Job Loss and Hiring,” Mimeo, Federal Reserve Bank of Philadelphia.
- [31] Galí, J. (1999), “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?,” *American Economic Review*, 89, 249-271.
- [32] Galí, J. and Rabanal, P., (2004), “Technology Shocks and Aggregate Fluctuations: How Well Does the RBC Model Fit Postwar U.S. Data?,” *NBER Macroeconomics Annual*, 19, 225-288.
- [33] Giordani, P. (2004), “An Alternative Explanation of the Price Puzzle,” *Journal of Monetary Economics*, 51, 1271-1296.
- [34] Gordon, R. (1990), *The Measurement of Durable Good Prices*, NBER Monograph Series, University of Chicago Press.
- [35] Gordon, R. (2000), “Does the ‘New Economy’ Measure Up To The Great Inventions of the Past?,” *Journal of Economic Perspectives*, 14, 49-74.
- [36] Greenwood, J. Hercowitz, Z. and Krusell, P. (1997) “Long-Run Implications of Investment-Specific Technological Change,” *American Economic Review*, 87, 342-362.
- [37] Greenwood, J. Hercowitz, Z. and Krusell, P. (2000) “The Role of Investment-Specific Technological Change in the Business Cycle,” *European Economic Review*, 44, 91-115.
- [38] Greenwood, J. and Yorokoglu, M. (1997), “1974,” *Carnegie-Rochester Conference Series on Public Policy* 46, 49-95.
- [39] Hagedorn, M. and Manovskii (2006), “The Cyclical Behaviour of Equilibrium Unemployment and Vacancies Revisited,” Mimeo, University of Pennsylvania.
- [40] Hall, R. (2005), “Job Loss, Job Finding, and Unemployment in the US Economy over the past Fifty Years,” *NBER Macroeconomic Annual*, Vol. 20, Mark Gertler and Kenneth Rogoff (editors).
- [41] Hamermesh, D. (1993), *Labour Demand*, Princeton University Press.
- [42] Hornstein, A. Krusell, P. and Violante, G. (2005), “The Replacement Problem in Frictional Economies: An Equivalence Result,” *Journal of the European Economic Association*, 3, 1007-1057.
- [43] Jackman, R., Layard, R., and Pissarides, C (1989), “On Vacancies,” *Oxford Bulletin of Economics and Statistics*, 51, 377-394.

- [44] Jaimovich, N. and Rebelo S. (2006), “Can News about the Future Drive the Business Cycle?,” Mimeo, Kellogg School of Management.
- [45] Jorgenson, D. and Stiroh, K. (2000) “Raising the Speed Limit: U.S. Economic Growth in the Information Age,” *Brookings Papers on Economic Activity*, 31:1, 125-210.
- [46] Jovanovic, B. and Lach, S. (1989), “Entry, Exit, and Diffusion with Learning by Doing,” *American Economic Review*, 79, 690-699.
- [47] Kehoe, T. and Ruhl, K. (2007), “Are Shocks to the Terms of Trade Shocks to Productivity?,” NBER Working Paper 13111.
- [48] Lee, C., (2001), “Finite Sample Bias in IV Estimation of Intertemporal Labor Supply Models: Is the Intertemporal Substitution Elasticity Really Small?,” *Review of Economics and Statistics*, 83, 638-646.
- [49] Mankiw, G. (1985), “Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly,” *Quarterly Journal of Economics*, 100, 529-539.
- [50] McKay, A. and Reis, R. (2007), “The Brevity and Violence of Contractions and Expansions,” Mimeo, Princeton University.
- [51] Michelacci, C. and Lopez-Salido, D. (2007), “Technology Shocks and Job Flows,” *Review of Economic Studies*, 74, 1195-1227.
- [52] Moen, E. (1997), “Competitive Search Equilibrium,” *Journal of Political Economy*, 105, 385-411.
- [53] Mortensen, D. and Pissarides, C. (1994), “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61, 397-415.
- [54] Mortensen, D. and Pissarides, C. (1998), “Technological Progress, Job Creation, and Job Destruction,” *Review of Economic Dynamics*, 1, 733-753.
- [55] Pissarides, C. (2000), *Equilibrium Unemployment Theory*, 2nd Edition, MIT Press.
- [56] Shimer, R. (2005a), “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95(1): 25-49.
- [57] Shimer, R. (2005b), “Reassessing the Ins and Outs of Unemployment,” Mimeo, University of Chicago.
- [58] Solow, R. (1960), “Investment and Technical Progress,” in *Mathematical Methods in the Social Sciences, 1959*, edited by Arrow, K., Karlin, S. and Suppes, P., Stanford University Press.

- [59] Uhlig, H. (2004), "Do Technology Shocks Lead to a Fall in total Hours Worked?," *Journal of the European Economic Association*, 2, 361-371.
- [60] Violante, G. (2002), "Technological Acceleration, Skill Transferability and the Rise of Residual Inequality," *Quarterly Journal of Economics*, 117, 297-338.