The Ins and Outs of Unemployment: An Analysis Conditional on Technology Shocks*

Fabio Canova
ICREA-UPF

David Lopez-Salido
Federal Reserve Board

Claudio Michelacci†
CEMFI

This version: January, 2012

Abstract

We analyze how unemployment, job finding and job separation rates react to neutral and investment-specific technology shocks. Neutral shocks increase unemployment and explain a substantial portion of it volatility; investment-specific shocks expand employment and hours worked and contribute to hours worked volatility. Movements in the job separation rates are responsible for the impact response of unemployment while job finding rates for movements along its adjustment path. The evidence warns against using models with exogenous separation rates and challenges the conventional way of modelling technology shocks in search and sticky price models.

JEL classification: E00, J60, O33.
Key words: Unemployment, technological progress, labor market flows, business cycle models.

*We thank Andrew Scott (the Editor), four anonymous referees, 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, Universitat Pompeu Fabra, New York Fed, Philadelphia Fed, Richmond Fed, and Atlanta Fed for comments and suggestions. We also thank Jason Cummins, Gianluca Violante, Robert Shimer, Pau Rabanal, Sergio Rebelo, Gary Solon and Ryan Michaels for kindly making their data available to us. 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. Previous versions of the paper have ciculated under the title "The ins and out of unemployment: a conditional analysis " and "Schumpeterian Technology Shocks".

†Authors are also affiliated with CREI, BSGE, CREMeD, 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.

1 Introduction

Since the pioneering contributions of Darby et al. (1985, 1986), Jackman et al. (1989), and Blanchard and Diamond (1990), the literature has recognized the importance of characterizing cyclical employment adjustment in terms of workers flows in and out of unemployment. The conventional wisdom has generally been that recessions, defined as periods of sharply rising unemployment, typically begin with a wave of layoffs and persist over time because unemployed workers have hard time to find new jobs. Hall (2005) and Shimer (2007) have challenged this view by showing that, in the US over the business cycle, there are substantial fluctuations in the job finding rate (the rate at which unemployed workers find a job), while the job separation rate (the rate at which employed workers lose their job) is comparatively acyclical. However, Yashiv (2007), Fujita and Ramey (2009), and Elsby et al. (2009) looking at the same evidence, attribute to the separation rate a larger role in characterizing US unemployment fluctuations. Since the conclusions these authors reach are based on unconditional correlation analysis, the interpretation of the evidence is difficult. First, it could be that the response of the unemployment rate differs depending on the source of the shock. An unconditional analysis, lumping different responses together, may mask such differences. Second, it does not tell us how unemployment responds when important business shocks occur. Third, it leaves open the question of what drives fluctuations in finding and separation rates and the contribution of technology shocks relative to other source of fluctuations. Fourth, it could be that the contribution of the ins and outs of unemployment on impact and over the adjustment path are different.

To address these issues, this paper analyzes the dynamics of the ins and outs of unemployment during technology induced recessions. Since the pioneering work of Kydland and Prescott (1982), many authors have suggested that technology shocks are responsible for a large portion of the fluctuations in macroeconomic variables and, following recent literature (see Fisher 2006, and Michelacci and López-Salido 2007), we focus attention on investment-neutral and investment-specific technology shocks. Technology shocks are identified in a VAR by imposing 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. These restrictions follow directly from the neoclassical growth model. We analyze the induced labor market dynamics along the intensive margin (hours per-employee) and the extensive margin (number of employed workers) and characterize unemployment dynamics in terms of the job separation rate and the job finding rate.

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 the growth rate of labor productivity and of the relative price of investment goods exhibit significant long run swings which have gone together with important changes in unemployment, and labor market flows. Many authors have documented important changes in the pace of US technological progress over the post WWII period, see for example, Greenwood and Yorokoglu (1997), Gordon (2000) and Jorgenson and Stiroh (2000). These patterns have been greatly emphasized in the literature on growth and wage inequality and growth and unemployment (see Aghion and Howitt, 1994, Mortensen and Pissarides 1998, Violante, 2002, and Hornstein et al., 2007, among others). Here we emphasize that it is important to account for these patterns to correctly identify the effects of technology shocks on labor market variables.

Neutral technology shocks having positive long run effects on labour productivity, substantially increase unemployment in the short run and affect labor markets primarily through the extensive margin. Positive investment specific technology shocks, on the other hand, expand aggregate hours worked, both because hours per-worker increase and because unemployment falls, but the intensive margin contributes most to the adjustments. For neutral shocks, the impact response of unemployment is almost entirely due to the instantaneous response of the separation rate while movements in the finding rate account for the subsequent unemployment dynamics. Thus, positive neutral shocks may cause recessions and the workers flows they induce are 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.

The practical relevance of these findings depends on how important technology shocks are for labor market fluctuations and how accurately they represent important historical episodes. On average, technology shocks explain more than 50 per cent of the cyclical fluctuations in labor market variables with neutral technology shocks mattering

slightly more for the volatility of unemployment and investment specific technology shocks more for hours per-worker volatility. In addition, neutral technology shocks explain the recession of the late 80's and the subsequent recovery of the early 90's. They initially rise the job separation and the unemployment rate; subsequently output builds up until it reaches its new higher long run value, but over the transition path employment remains below normal levels because the job finding rate is persistently below its long run level —generating a "jobless"recovery, which is a distinctive feature of this episode. The contribution of technology shocks to the jobless recoveries of the 2000's is instead significantly smaller.

These results are important for the literature analyzing business cycle in the labor market through the lenses of search models. In fact, they provide a healthy warning to the ongoing tendency to analyze the effects of technology shocks in search models with exogenous separation rates, a point also made by Ramey (2008). They also challenge the conventional way to model technology shocks in search models (see, e.g., Shimer, 2007), both because the labor market responses to neutral and investment specific technology shocks are different and because neutral technology shocks are contractionary rather than expansionary on employment (see also Balleer, 2010).

Our evidence also challenges the standard sticky-price explanation for why hours fall in response to neutral technology shocks, 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. While there is very little empirical evidence comparing the cost of firing and the costs of changing prices, the mechanism naturally applies to the intensive margin since displacing workers is likely to be 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.

Finally, our findings that neutral technology shocks are contractionary on employment and the fall in hours is related to the time consuming process of reallocation of workers across jobs,while investment specific technology shocks are expansionary, challenge the wisdom conventional models provide. The finding however are fully 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. As shown by Caballero and Hammour (1994, 1996), and Michelacci and López-Salido (2007) when the labor market is characterized by search frictions, these adjustments can cause unemployment. Thus, our findings support the view by Michelacci and López-Salido (2007) that neutral technological progress prompts waves of Schumpeterian creative destruction, where outdated, technologically obsolete productive units are pruned out of the productive system, while investment-specific technology shocks lead to an expansion in economic activity, because a substantial proportion of old jobs upgrade their capital equipment and reap the benefits of the most recent advancements in capital equipment.

We are not the first to analyze labor market dynamics using conditional responses. Michelacci and López-Salido (2007), Ravn and Simonelli (2007), Barnichon (2010, 2011) and Balleer (2010) have also used long-run restrictions in SVARs to identify the effects of technology shocks on job and worker flows, vacancies and unemployment.[1] Braun et al. (2009) and Fujita (2011), on the other hand, have used sign restricted SVARs to study the labor market response to structural shocks but the restrictions they use are not necessarily satisfied if technology shocks have Schumpeterian features. In addition, we emphasize that to identify the effects of technology shocks using long run restrictions it is important to take care of the low frequency movements in the growth rate of productivity and of the price of investment observed in the data and that the effects of neutral shocks on labor market variables are strongly at variance with the dynamics implied by technology shocks, as modelled in the conventional search model.

The rest of the paper is structured as follows. Section 2 discusses the data, the empirical model, and the consequences of low frequency comovements in the variables. Section 3 presents basic results. Section 4 quantifies the relative contribution of job separation rates to the dynamics of technological unemployment. Section 5 measures the contribution of technology shocks to labor market fluctuations. Section 6 interprets the results in light of existing work. Section 7 examines robustness. Section 8 concludes.

[1] This paper bears similarities with Michelacci and López-Salido (2007). We complements their work in three ways: we consider workers flow data rather than using job creation and job destruction rates; the labor market flows we use are representative of the whole US economy rather then just the manufacturing sector; the datset is a longer and more informative.

2 The empirical model

Let X_t be a $n \times 1$ vector of variables and let X_{1t} and X_{2t} be the first difference of the price of investment, q_t, and labor productivity y_{nt}, respectively. Let $X_t = D(L)\eta_t$ be the (linear) Wold representation of X_t, where $D(L)$ has all its roots inside the unit circle and $E(\eta_t\eta_t') = \Sigma_\eta$. We assume that the reduced form shocks η_t and the structural shocks ϵ_t are related via $\eta_t = S\epsilon_t$, where S is a full rank matrix, that the structural shocks ϵ_t are uncorrelated and let $E(\epsilon_t\epsilon_t') = I$. Thus, impulse responses represent the effects of one-standard deviation shocks. To identify the shocks of interest, we use the restrictions that the non-stationarities in q_t originate exclusively from investment specific technology shocks and that the non-stationarities in y_{nt} are entirely produced by investment specific and neutral technology shocks. Hence, a neutral technology shock (a y_n-shock) is the shock 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; and no other shock has long-run effects on these two variables. This implies that the first row of $D(1)S$ has zeros everywhere except in the first position and the second row has zeros everywhere except in the first and second positions.

These restrictions can be derived from a simple neoclassical growth model where technological progress is non-stationary (see Fisher, 2006 and Michelacci and López-Salido, 2007). In 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 change in investment demand. Since we only restrict the price of investment in the long run, our identification strategy is robust to the existence of short run increasing marginal costs to produce investment goods.

There is controversy on how the price of investment and GDP should be deflated. Fisher (2006) and Michelacci and López-Salido (2007) 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, excluding indirect taxes, and discounting the fact that the CPI only includes a subset of the consumption goods and that its weights measures the prices paid by *urban* consumers, the CPI and the output deflator are

similar. In an open economy, differences may arise because some consumption goods are produced abroad. In our baseline specification, we deflate both variables with a output deflator.[2] Later, we show that the choice of deflator does not matter.

2.1 The data

Our model has six variables $X = (\Delta q,\ \Delta y_n, h,\ u,\ s,\ f)'$, where Δ denotes the first difference operator. The last four variables are in logs and all are multiplied by one hundred: q equals to the inverse of the relative price of a quality-adjusted unit of new equipment, y_n is labor productivity, measured as output per hours, 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, respectively. The dynamics of hours per-worker can be obtained from the responses of hours and unemployment; those of output per-capita can be derived from the responses of labor productivity and hours. Rather than using a standard lag selection criteria, we use a generous lag length (8 lags) and reduce overparametrization with a prior that stochastically restrict the lag decay toward zero assuming that the prior variance of lag j is proportional to j^{-2}. We choose this approach in order to avoid the problems emphasized by Giordani (2004), Chari et al. (2008), 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. The sensitivity of the results to the choice of lags is discussed in the robustness section.

The series for labor productivity, unemployment, and hours are from the USECON database commercialized by Estima and are all seasonally adjusted; q until 2000 is from Cummins and Violante (2002), who extend the Gordon (1990) measure of the quality of new equipment, see their paper for further details. The original q series is annual; we use Galí and Rabanal (2004) quarterly interpolated values up to 2000:IV. We extrapolate this quarterly series up to 2010:I using NIPA account. [3] Real output (mnemonics `LXNFO`) and the aggregate number of hours worked (`LXNFH`) correspond to

[2] One can show that this approach is consistent with the balanced growth conditions of a well defined open economy, see Kehoe and Ruhl (2007) for a similar point and Section 7 for further discussion.

[3] Several authors have constructed price of investment series directly from NIPA accounts (see for example Justiniano et al. (2010). These series are not very highly correlated with the original Cummins and Violante series making the comparison with the earlier literature difficult.

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)). CN and CS are nominal consumption of non-durable and services while CNH and CSH are the analogous values in real terms. The aggregate number of hours worked per capita is the ratio of LXNFH to the working age population (P16). The unemployment rate corresponds to the civilian unemployment rate. Starting from 1967:II, the monthly Current Population Survey public microdata can be used to 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). Following Shimer (2007), we calculate exact instantaneous rates at which workers move from employment to unemployment (finding rate) and viceversa (separation rate). The availability of workers' flows data restricts the analysis to 1967:II-2010:I period.

2.2 The low frequency comovements on the VAR

The first graph of Figure 1 plots the log of total hours (continuous line) and the log of the unemployment rate (dashed line). Hours display a clear cyclical pattern, highly negatively correlated (-0.8) with the one of unemployment. Whether the two series are stationary or exhibit persistent low frequency movements, is matter of controversy; see, for example, Francis and Ramey (2005) and Fernald (2007).

The next two graphs of Figure 1 plot the growth rate of y_n and of the relative price of investment (equal to minus q), measured in either output units (continuous line) or consumption units (dashed line). There is a dramatic fall in the value of q in 1974, reversed in the following years. Cummins and Violante (2002) attribute this drop 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 last graph of Figure 1 displays the log of the job finding rate and the log of the job separation rate. The job finding rate is relatively more persistent than the separation rate (AR1 coefficient is 0.86 vs. 0.73). Given that recessions are typically associated with a persistent fall in the job finding rate, the higher persistence of the job finding rate is consistent with Hall (2005) and Shimer (2007) observation that cyclical fluctuations

in the unemployment rate are highly correlated with those in the job finding rate.

The low frequency co-movements of the series of interest are highlighted in Figure 2. In the first row we follow the growth literature and choose 1973:I and 1997:I 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. 2007). Switching these dates forward by one or two years has no consequence on the points we make here. The rate of growth of the relative price of investment goods was minus 1.0 per cent per quarter over the period 1967:I to 1973:I and moved to minus 1.3 per cent per quarter in the period 1973:II-1997:I. This difference is statistically significant. During the productivity revival of the late 1990's the price of investment goods was falling at even a faster rate and before 2007 the growth rate of the price of investment exceeded minus 1.6 per cent per quarter. The rate of growth of labor productivity exhibits an opposite trend. It was higher in the 1967:I- 1973:I period than in the 1973:II-1997:I period, strongly recovered in the late 1990's to drop after 2007. Also in this case, subperiod 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.

The graphs in the second row of Figure 2 plot the trend component of labor productivity growth, the log of hours and the log unemployment obtained by using a Hodrick Prescott filter with smoothing coefficient equal to 1600. The trends are related: there are negative comovements between productivity growth and the log of the unemployment rate and positive comovements between productivity growth and the log of hours. The third row of Figure 2 shows that both the separation rate and the finding rate exhibit low frequency movements that mimic those of the unemployment rate.

2.3 Low-frequencies comovements and impulse responses

To show why these comovements are problematic when analyzing the responses to technology shocks, we plot the point estimates of the responses obtained for three different samples: 1967:I-2010:I, 1973:II-1997:I and 1997:II-2010:I. 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. Figure 4 presents the responses to an investment specific shock.

It is apparent that estimated responses to neutral shocks in the two subsamples are similar. Yet, they look quite different from the responses for the full sample. In particular, the movements in the relative price of investment are stronger and those in the separation rate weaker. Moreover the fall in hours and in 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. Differences in the estimates 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. Similar logic applies when comparing the responses to investment specific technology shocks in the full sample to those in the two subsamples. [4]

2.4 Discussion

Commentators have sometimes questioned our choice of break points. Some have suggested that taking a break point as known (when in fact it is not) may bias results, while others have suggested that a perhaps more relevant break point would be, as in the Great Moderation literature, somewhere around the beginning of 1980. As shown in the on-line appendices, moving forward by one or two years the break dates does not change the conclusion that, over subsamples, the responses of the variables are similar and different from those of the full sample. Furthermore, visual inspection of Figure 1 indicates that none of the series displays unusual behavior in the early 1980s.

The evidence contained in figure 3 and 4 indicates that the dynamic responses of the variables of the VAR to the two shocks are sufficiently homogeneous over subsamples. Therefore, the low frequency variations we have highlighted can be effectively taken care by adjusting the constant of the VAR and this is what we do in this paper.

[4] This evidence is robust to a number of standard modifications, including the use of the population-adjusted hours produced by Francis and Ramey (2005) and the choice of the price deflator.

In Canova et al. (2010), we elaborate on this issue and present cases where unaccounted level breaks within a sample produce sign switches or an extreme pattern of persistence in the responses, see also Fernald (2007). A simple way to eliminate the low frequency comovements is to estimate the VAR over sub-samples. While feasible, such an approach would be inefficient, since the dynamics are roughly unchanged, and it may cause biases, since imposing long run restrictions in a system estimated over a small sample may distort structural estimates (see Erceg et al. 2005).

Low frequency movements in the data are the object of controversy and our choice of eliminating them could be criticized in, at least, two ways. It could be argued, for example, that after a prolonged period of low productivity growth and in anticipation that productivity will pick up, labor input could be lower in the period of low productivity, making low frequency movements informative about business cycle fluctuations. One way to rationalize our decision of removing low frequency fluctuations is that breaks can not be forecasted so anticipatory effects are not present. It could also be argued that changes in productivity growth also affect agents decisions rule, not just the intercept. Since the dynamics in response to the shocks are very similar in the two subsamples (and different from the full sample where no adjustment for low frequency movements is made), agent's decision rules appear to be unaffected by the breaks. Furthermore, once breaks in the intercepts are considered, the full sample evidence roughly replicates the one the sub-samples, see next section.

3 The full sample results

Figure 5 plots the response of the variables of interest to a neutral technology shock in the VAR when the intercept is deterministically broken at 1973:I and 1997:I. The reported bands correspond to 90 percent confidence intervals. We also plot the median of the band rather than the point estimate, since the latter is known to be biased in small samples. 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 smaller but still statistically significant. The impact increase 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 quickly returns to normal levels while

the job finding rate takes up to twenty-five 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.

Figure 6 plots responses to an investment specific shock. 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 is in part due to a sharp drop in the separation rate and in part due to an increase in the job finding rate. In the impact period the unemployment response is, however, small and insignificant. Thus, the increase in hours is primarily explained by the sharp and persistent increase in the number of hours worked per employee. Hence, while labor market adjustments to neutral technology shocks occur mainly along the extensive margin, those in response to an investment specific technology shock mainly occur along the intensive margin.

3.1 Omitted variables

Our VAR has enough lags to make the residuals white noises. 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 policy variables, therefore making responses to Solow residuals shocks uninterpretable, while Forni and Gambetti (2010) show that many puzzles in responses to monetary shocks could be due to omitted variables. To check for this possibility we have correlated, up to 6 leads and 6 lags, 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, such as 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 7.

As the figure indicates, no variable is significantly correlated with the neutral technology shocks we extract. On the other hand, they are indications that the consumption to output ratio and inflation could be correlated with investment specific shocks at some large lag. For this reason we have repeated estimation, enlarging the VAR system to include these three new variables. Since, the dynamics of labor market variables are affected (see the on-line appendix), we continue to work with a six variable VAR.

4 The role of separation rates

Hall (2005) and Shimer (2007) 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 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 of Fujita and Ramey (2009) that the separation rate leads the cycle (by about one quarter) while the finding rate lags it (by about two months), see also Yashiv (2007) and Elsby et al. (2009) for similar considerations.

To further evaluate the role of the separation rate for unemployment fluctuations, we use a simple two state model of the labor market (see Jackman et al., 1989 and Shimer, 2007 and 2008) and assume that the stock of unemployment evolves as:

$$\dot{u}_t = S(l_t - u_t) - Fu_t \tag{1}$$

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:

$$\tilde{u} = \frac{S}{S+F} \equiv \frac{\exp(s)}{\exp(s)+\exp(f)}. \tag{2}$$

Shimer (2007) shows that the fictional unemployment rate $\tilde{u}$ closely tracks the actual unemployment rate series, so that one can fully characterize the evolution of the stock of unemployment just with the dynamics of labor market flows. After linearizing the log of $\tilde{u}$, we can also calculate its response using the response of (the log of) the separation rate s and the finding rate f. Thus, we can measure the contribution of finding and separation rates to the cyclical fluctuations of fictional unemployment $\tilde{u}$ and evaluate how accurately fictional unemployment approximates actual unemployment.

Figure 8 reports the results. 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) $\tilde{u}$ obtained if the job finding rate had remained unchanged at its average level. It therefore represents the contribution of the separation rate to fluctuations in fictional unemployment.

The dynamics of fictional unemployment after a neutral shock are explained to a large extent by fluctuations in the separation rate. Consistent with previous analysis, the separation rate explains almost 90 per cent of the impact effect on fictional unemployment. However, its contribution falls to 40 per cent after one quarter and drops to 20 per cent one year after the shock. Following an investment specific shock, unemployment falls on impact because of the fall in the separation rate.

Differences between the response of fictional and actual unemployment to investment-specific technology shocks are minimal, but there are some differences in the impact response of actual and fictional unemployment to neutral technology shocks. There are several reasons that could explain these differences in impact responses. First there could be responses in the flows at which workers move in and out of the labor force that affect the unemployment rate and that are not properly taken into account by our series for the finding and the separation rate.[5] Second one has to bear in mind that the equation (2) is just an approximation and that the fictional and the actual unemployment rates coincide only when labor market transitions are sufficiently fast and time periods are large enough: on impact the two series do not have to coincide. To check our conclusions we also performed the decomposition proposed by Elsby et al. (2009), which allows to calculate the contribution of the separation rate also out of the steady state. The results changes very little, because the discrepancies between the impulse responses of fictional and actual unemployment rate are large just on impact, and by definition impact effects are just deviations from the steady state.

5 The contribution of technology shocks

To put our findings in the right perspective, we study whether the contribution of technology shocks to fluctuations in the variables of interest is non-negligible. Otherwise, what we uncover is an interesting intellectual curiosity without practical implications. Table 1 reports the forecast error variance decomposition.

The combined effect of neutral and investment specific shocks for output fluctua-

[5] An earlier version of the paper showed that the participation rate is procyclical, and that the flows in and out of the labor force were generally little significant except on impact. This indicates that workers movements in and out of the labor force can play some role in characterizing the response of the unemployment rate.

tions is considerable: at the 2-8 years horizon the proportion of the variance of output jointly explained by the two shocks is between 40 and 60 percent. The two shocks also explain a substantial portion of the volatility of unemployment and hours per-worker: neutral technology shocks explain about 30-40 per cent of unemployment fluctuations and about 25-35 of the hours per-worker fluctuations at time horizons between 2 and 8 years. Investment specific technology shocks are also important for the volatility of these two variables: at time horizons between 2 and 8 years, they explain between 16 and 36 per cent of the variance of unemployment and between 22 and 44 percent of the variance of hours per-worker. Interestingly, while neutral shocks are major source of volatility in the short run, the importance of investment specific shocks picks up with time. Taken together, the shocks we consider explain 40-60 per cent of the volatility of unemployment and hours at horizons between 2 and 8 years.

The message is similar when considering the two alternative sub-samples. If we stop estimation prior to the productivity revival of the 1990s (see Panel B), the importance of neutral shocks for output fluctuations increases and while the proportion jointly explained by the two technology shocks falls, technology shocks still explain a considerable share of the fluctuations in unemployment and hours per-worker. If we stop estimation prior to the last crisis, the importance of investment specific shocks for output fluctuations increases, but their contribution to the fluctuations in unemployment and hour per-worker falls significantly (see panel C).

Technology shocks are also important for labor markets flows. Neutral shocks explain between 25 and 35 percent of the variability of the finding and separation rate at horizons between 2 and 8 years. The contribution of investment specific shocks is somewhat smaller and changes with the estimation period, but in the full sample they explain about 20-30 percent of the variability of the finding rate. In sum, the technology shocks we have identified represent major sources of fluctuations in both the goods and labor markets and induce significant movements in labor market flows.

Further evidence on the role of technology shocks can be obtained examining their historical contribution to fluctuations in unemployment, output, job finding and job separation. In Figure 9 we present the original series (solid line) and its component due to either neutral or investment specific shocks (dotted line), as recovered from the VAR. All series are detrended with a Hodrick Prescott filter with smoothing parameter

equal to 1600; grey area correspond to the NBER recessions.

The figure confirms that technology shocks account for several important business cycle episodes, including the recession of the late 80's and the subsequent remarkably slow labor market recovery of the early 90's. This episode have been extensively investigated in the literature, yet its causes are still unexplained; see for example Bernanke (2003). Two key features of the episode are that the downturn in employment was severe and that the peak in unemployment occurred two years later than the trough in output. This is a remarkable exception relative to other episodes, see McKay and Reis (2007). Figure 10 zooms in on the content of figure 9 and presents output, unemployment, finding and separation rates (solid lines) and their component due just to technology shocks (dotted lines) in that period. The technology component and the raw data tracks are quite close. This is mainly due to the evolution of neutral shocks that naturally tend to induce jobless recoveries: following the initial rise in job separation and unemployment, output increases to its new higher long run value, while unemployment remains above trend because of the low job finding rate, which induces a remarkably slow recovery in the labor market, see right column. It is worth mentioning that the relationship between technology shocks and the jobless recoveries in the 2000's is weaker.

6 Interpreting the evidence

Our findings indicate that the separation rate is important in characterizing the labor market response to technology shocks and that labor market adjustments to different technology shocks are different. Neutral shocks exercise their effects primarily along the extensive margin of the labor market and are contractionary on employment; investment specific shocks mainly affect labor along the intensive margin and they are expansionary on hours and employment. These results have important implications for both empirical analysis and theoretical models of business cycles.

First, failure to empirically distinguish between the two types of disturbances may lead to nonsensical representation of the dynamics following unexpected technological improvements. Second, our results qualify the conclusions of Hall (2005) and Shimer (2007) and show the importance to use search models with endogenous separation for

business cycle analysis, a point emphasized also by Ramey (2008). The difference in conclusions is partly due to our focus on correlations, conditional on technology shocks, rather than on unconditional correlations at generic business cycle frequencies. Third, for interpretation purposes, it is very important to decompose the response of total hours into the response along the extensive and the intensive margin. For example, since Galí (1999) it is common to interpret the evidence that hours fall in response to neutral technology shocks using sticky prices models. In sticky-price models, when technology improves and monetary policy is not accommodating enough, demand is sluggish to respond to the shocks and firms take advantage of technology improvements to economize on labor input. While this mechanism has its own appeal, it should most naturally apply to the intensive margin of the labor market since changing prices is arguably less costly than displacing workers—whose cost includes 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 firing (see e.g. Mankiw, 1985 and Hamermesh, 1993 for a review of the literature). Admittedly, no formal model analyzing the trade-off between changing prices and displacing workers exists in the literature and empirical evidence on the issue is scant but one can conjecture that when the decision of changing prices is endogenous and menu cost a-la Caballero and Engel (2007) are used, this is the expected outcome.[6] Our evidence indicates that labor market adjustments to neutral shocks occur primarily at the extensive margin and the fall in hours is mostly caused by the time consuming process of reallocation of workers across productive units. This challenges the sticky prices interpretation of the phenomenon.

As stressed by Balleer (2010), our findings also challenge the conventional way of modelling technology shocks in search models (see for example Pissarides 2000 and Shimer, 2008), both because the labor market responses to neutral and investment specific technology shocks are substantially different and because neutral technology shocks are contractionary rather than expansionary on employment. A possible interpretation of the finding is that investment specific technological progress has standard neoclassical features, while neutral shocks have Schumpeterian features so that the

[6] Of course this is just a conjecture, since in standard sticky price models based on Calvo pricing firms do not explicitly pay a menu cost to change prices and the probability of price adjustment is exogenous, see Barnichon (2010), for a model of this type.

introduction of new neutral technologies causes the destruction of technologically obsolete productive units and the creation of new technologically advanced ones. When the labor market features search frictions, this process leads to a temporary rise in unemployment. 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. (2007). Apart from Michelacci and López-Salido (2007), we are not aware of papers using this paradigm to interpret the response of worker flows to technology shocks. In their model, newly created jobs embody the most advanced technologies (both neutral and investment specific) available at the time of their creation while existing jobs may fail to upgrade their previously installed technologies. The idea is that the adoption of new technologies requires the performance of some new worker-specific tasks, so workers initially hired to operate specific technologies may not be suitable for their upgrading. In their model the short run response of the economy to a technology shock may be expansionary or contractionary on employment depending on how easily old jobs upgrade their technology and reap the benefits of technological advancements. When old-jobs' technology can not be easily upgraded, technological progress causes a wave of Schumpeterian creative destruction characterized by a simultaneous increase in the destruction of obsolete productive units and in the creation of new technologically advanced ones. Based on micro evidence on technology dynamics at the firm level, Michelacci and López-Salido (2007) argue that only technologies embodied in capital can be readily upgraded. As a result, a neutral technology shock prompts a wave of Schumpeterian creative destruction where job creation, job destruction and unemployment simultaneously increase, while the adoption of investment specific technologies operate essentially as in the standard neoclassical growth model, leading to an expansion in economic activity.

7 Robustness

This section describes some robustness exercises whose outcomes are documented in the on-line appendix. The main conclusions of these exercises is that our technology shocks are unlikely to stand in for other sources of disturbances and that the results

stand when we change i) the lag length of the VAR, ii) the way we remove low frequency fluctuations, iii) the timing of identifying restrictions, iv) the price deflator, v) the labor market data and the relative price of investment series.

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 disturbances. To check for this possibility, we have correlated the estimated technology shocks with oil prices shocks (mnemonics `PZTEXP`) deflated by the consumption deflator and federal fund rate shocks (`FFED`), both computed filtering these two series with an AR(1). Correlations are insignificant up to 6 leads and 6 lags.

VAR lag length The issue of the length of VAR has been recently brought back to the attention of applied researchers by Giordani (2004) and Chari et al. (2008), who show that the aggregate decision rules of a subset of the variables of a model may have not always be representable with a finite order VAR. This issue is unlikely to be important in our context since we have checked that the residuals of a VAR(8) are white noise and do not stand-in for other potential sources of shocks. To confirm this we have reestimated our VAR using 4 lags. The pattern of responses is unchanged because the prior effectively removes the noise that longer lags produce in the VAR.

Alternative treatments of trends We have considered two alternatives to the dummy approach we employ to remove low frequency movements in the variables of the VAR: we have allowed up to a third order polynomial in time as intercept in the VAR; we filtered all the variables, before entering them in the VAR, with the Hodrick Prescott filter with a smoothing parameter $\lambda = 12800$. In both cases responses have the same shape and approximately the same size as in the benchmark specification.

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. Literally, this implies the neutral shocks we have extracted may not be structural. To take care

of this problem, Uhlig suggests to check if conclusions change when medium term restrictions are used. We have recomputed our responses imposing 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 restrictions are imposed at any horizon of at least one year.

Price deflators In our benchmark VAR, labor productivity and the relative price of investment are deflated with the output deflator - this is equivalent to use domestic consumption as a numeraire. As mentioned, other authors have either used the CPI deflator or combination of output and CPI deflators. Such choice is potentially problematic since the identifying assumptions are no longer valid with such deflators.

To see this recall that when the price of investment and total factor productivity have unit roots, a standard Solow economy evolves around the (stochastic) trend

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

where Z is the TFP component and Q the price of investment component (see e.g. Michelacci and López-Salido (2007)) and that the quantities $Y \equiv \tilde{Y}/\left(XN\right)$, and $K \equiv \tilde{K}/\left(XQN\right)$ converge to $Y^* = (s/\delta)^{\frac{\alpha}{1-\alpha}}$ and $K^* = (s/\delta)^{\frac{1}{1-\alpha}}$, respectively. Consequently, the logged level of aggregate productivity, $y_n \equiv \ln \tilde{Y}/N$, evolves according to

$$y_n = y^* + v + x = y^* + v + \frac{1}{1-\alpha} z + \frac{\alpha}{1-\alpha} q \tag{3}$$

where small letters denote the log of the corresponding quantities in capital letters and v is a stationary term, accounting for transitional dynamics. Hence, in the long run labour productivity is due to the neutral and the investment specific shocks.

Now, 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 Consumer Price Index , P_c, defined as $P_c = \left(\frac{P_c^H}{a}\right)^a \left(\frac{P_c^F}{1-a}\right)^{1-a}$, where P_c^H and P_c^F are the prices of consumption goods produced in the US and abroad; and a represents the share of domestic consumption goods. Tedious calculations show that

$$y_n^c = cte + \frac{1}{1-\alpha-\beta} z + \frac{\alpha+\beta}{1-\alpha-\beta} q^c + \frac{1}{1-\alpha-\beta} \left(1-a\right) \left(p_c^H - p_c^F\right) \tag{4}$$

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 could affect long run labor productivity and confused with "neutral" technology shocks (see also Kehoe and Ruhl, 2008). Since the real exchange rate exhibits remarkable persistence, one should worry about mixing neutral and real exchange rate shocks.

When we deflate the relative price of investment with the CPI index and output with the GDP deflator we obtain

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

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

Our choice of deflator is the right one, in the sense that it implies a well defined balanced growth path in an open economy version of the Solow model, and does not suffer from misspecification issues. Figure 1 indicates that the difference due to the use of an incorrect price deflator are small. Nevertheless, it is worth investigating whether the results we obtain could be altered if the CPI deflator is used. Responses are roughly similar to our benchmark ones. The main difference concerns the response of the price of investment to a neutral technology shock, which is now more pronounced.

Alternative data sets Elsby et al. (2009) have recently calculated an alternative series for the job finding and job separation rates, by slightly modifying the methodology of Shimer (2007). Our results are unchanged when this alternative series for labor market flows is used in the VAR.

An earlier version of the paper (see Canova, et al., 2009) also used Shimer's (2007) approximate finding and separation rates, which are available from 1948 and results were broadly similar. These rates are constructed from the Bureau of Labor Statistics data 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 measures, which are quarterly averages of monthly rates, are calculated under the assumption that employment and unemployment are the only two possible states of the labor market. Since they abstract from workers' labor force participation decisions, they approximate the true labor market rates at which unemployed workers find a new job and employed workers lose their job.

It is also worth mentioning that if one uses the q series directly obtained from the NIPA table (as e.g. in Justiniano et al., 2010), the responses to investment specific shocks change somewhat. The series we use and the one obtained from the NIPA table are positively but not perfectly correlated. In addition, the average growth rate of the variable is different. Which series should be used in both structural and semi-structural estimation exercises is unclear and future work should explore in more details the construction of the appropriate price of investment series.

8 Conclusions

We have analyzed the effects of neutral and investment specific technology shocks on unemployment, job finding, job separation rates and other labor market variables. We show that positive neutral technology shocks affect labor market variables primarily along the extensive margin and substantially increase unemployment. Positive investment specific technology shocks, on the other hand, expand aggregate hours worked, both because hours per-worker increase and because unemployment falls, but the intensive margin contributes most to the adjustments. For both shocks, the short run response of unemployment is almost entirely due to the instantaneous response of the separation rate while movements in the finding rate account for the dynamic adjustments of unemployment. Thus, positive neutral shocks can cause recessions and the induced flows in and out of unemployment are 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. Technology shocks explain over 50 per cent of the cyclical fluctuations in labor market variables and around 30 percent of fluctuations in workers flows and accurately characterize the "jobless" recovery of the early 90's. Our findings are robust to a number of specification choices, to the selection of price deflators and to changes in auxiliary assumptions.

The evidence we uncover casts doubts on the recent tendency to use search models with exogenous separation rates to analyze the effects of technology shocks. It also challenges the standard sticky price explanation for why hours fall in response to neutral technology shocks. The evidence may instead be consistent with the idea that investment specific technological progress has standard neoclassical features, while neutral

technological progress is Schumpeterian. According to this view the introduction of new neutral technologies causes the destruction of technologically obsolete productive units and the creation of new technologically advanced ones. When the labor market is characterized by search frictions, these adjustments lead to a temporary rise in unemployment (see e.g. Canova et al., 2007). If correct, this interpretation questions the conventional way of modelling technology shocks in search models.

References

[1] Aghion, P. and Howitt, P. (1994): "Growth and Unemployment," *Review of Economic Studies* 61, 477-494.

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

[3] Balleer, A. (2010): "New Evidence, Old Puzzles: Technology Shocks and Labor Market Dynamics," Mimeo Institute for International Economic Studies Stockholm.

[4] Barnichon, R., (2010): "Productivity and Unemployment over the Business Cycle," *Journal of Monetary Economics* 57, 1013-1025.

[5] Barnichon, R., (2011): "The Shimer Puzzle and the Identification of Productivity Shocks," Mimeo Universitat Pompeu Fabra.

[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 Diamond, P. (1990): "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity* 2, 85–143.

[8] Blanchard, O. and Quah, D. (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review* 79, 655-673.

[9] Braun, H., De Bock R. and DiCecio R. (2009): "Supply Shocks, Demand Shocks, and Labor Market Fluctuations," *Federal Reserve Bank of St.Louis Review*, May, 155-176.

[10] Caballero, R. and Hammour, M. (1994): "The Cleansing Effect of Recessions," *American Economic Review* 84, 1350-1368.

[11] Caballero, R. and Hammour, M. (1996): "On the Timing and Efficiency of Creative Destruction," *Quarterly Journal of Economics* 111, 805-852.

[12] Caballero, R. and Engel, E. (2007): "Price Stickiness in Ss Models: New Interpretations of Old Results," *Journal of Monetary Economics* 54, 1001-1121.

[13] Canova, F. López-Salido, D. and Michelacci, C. (2007): "Shumpeterian technology shocks," CEPR working paper.

[14] Canova, F. López-Salido, D. and Michelacci, C. (2009): The Ins and Out of Unemployment: A conditional Analysis, UPF working paper 1213.

[15] Canova, F. López-Salido, D. and Michelacci, C. (2010): "On the Robust Effects of Technology Shocks on Hours worked and Output," *Journal of Applied Econometrics* 25, 755-773.

[16] Chari, V. , Kehoe, P. and McGrattan, E. (2008): "Are structural VARs with long-run restrictions useful in developing business cycle theory?," *Journal of Monetary Economics* 55, 1337-1352.

[17] 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.

[18] Darby, M., J. Haltiwanger, and Plant, M. (1985): "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations," *American Economic Review* 75, 614–637.

[19] Darby, M., J. Haltiwanger, and Plant, M. (1986): "The Ins and Outs of Unemployment: The Ins Win," NBER Working Paper.

[20] Elsby, M., Michaels, R. and Solon, G. (2009): "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics* 1, 84–110.

[21] Erceg, C., Guerrieri, L. and Gust, C. (2005), "Can Long Run Restrictions Identify technology shocks?", *Journal of the European Economic Association* 3, 1237-1278.

[22] Evans, C. L. (1992): "Productivity Shocks and Real Business Cycles," *Journal of Monetary Economics* 29, 191–208.

[23] Fernald, J. (2007): "Trend Breaks, Long Run Restrictions, and Contractionary Technology Improvements," *Journal of Monetary Economics* 8, 2467-2485.

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

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

[26] Forni, M. and Gambetti, L. (2010) "The Dynamic Effects of Monetary Policy: A structural Factor Model Approach", *Journal of Monetary Economics* 54, *203-213.*

[27] Foster, L., Haltiwanger, J. and Krizan, C. (2001): "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in *New Developments in Productivity Analysis*, University of Chicago Press.

[28] 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.

[29] Fujita, S. (2011): "Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach," *Journal of Applied Econometrics* 26, 89-121.

[30] Fujita, S. and Ramey, G. (2009): "The Cyclicality of Separation and Job Finding Rates," *International Economic Review* 50, 415-430.

[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. and Yorokoglu, M. (1997): "1974," *Carnegie-Rochester Conference Series on Public Policy* 46, 49-95.

[37] 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).

[38] Hamermesh, D. (1993): *Labour Demand,* Princeton University Press.

[39] Hornstein, A. Krusell, P. and Violante, G. (2007): "Technology-Policy Interaction in Frictional Labor Markets," *Review of Economic Studies* 74, 1089-1124

[40] Jackman, R., Layard, R., and Pissarides, C (1989): "On Vacancies," *Oxford Bulletin of Economics and Statistics,* 51, 377-394.

[41] 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.

[42] Justiniano, A., Primiceri, G. and Tambalotti, A. (2010) Investment shock and business cycle, *Journal of Monetary Economics* 57, 132-145.

[43] Kehoe, T. and Ruhl, K. (2008): "Are Shocks to the Terms of Trade Shocks to Productivity?," *Review of Economic Dynamics* 11, 804-819.

[44] Kydland F. and Prescott E. (1982): "Time to Build and Aggregate Fluctuations," *Econometrica* 50, *1345-1370.*

[45] Mankiw, G. (1985), "Small Menu Costs and Large Business Cycles: A Macreoconomic Model of Monopoly", *Quarterly Journal of Economics* 100, 529-539.

[46] McKay, A. and Reis, R. (2007): "The Brevity and Violence of Contractions and Expansions," *Journal of Monetary Economics* 55, 738-751.

[47] Michelacci, C. and López-Salido, D. (2007): "Technology Shocks and Job Flows," *Review of Economic Studies* 74, 1195-1227.

[48] Mortensen, D. and Pissarides, C. (1998): "Technological Progress, Job Creation, and Job Destruction," *Review of Economic Dynamics* 1, 733-753.

[49] Pissarides, C. (2000): *Equilibrium Unemployment Theory*, 2nd Edition, MIT Press.

[50] Ramey, G. (2008): "Exogenous vs. Endogenous Separation," Mimeo University of California San Diego.

[51] Ravn, M. and Simonelli S. (2009): "Labor Market Dynamics and the Business Cycle: Structural Evidence for the United States," *Scandinavian Journal of Economics* 109, 743-77.

[52] Shimer, R. (2007): "Reassessing the Ins and Outs of Unemployment," NBER Working Paper 13421.

[53] Shimer, R. (2008): "Labor Markets and Business Cycles," CREI Lectures in Macroeconomics, available at http://robert.shimer.googlepages.com/workingpapers

[54] Uhlig, H. (2004): "Do Technology Shocks Lead to a Fall in total Hours Worked?," *Journal of the European Economic Association* 2, 361-371.

[55] Violante, G. (2002): "Technological Acceleration, Skill Transferability and the Rise of Residual Inequality," *Quarterly Journal of Economics* 117, 297-338.

[56] Yashiv, E. (2007): "U.S. labor market dynamics revisited," *Scandinavian Journal of Economics* 109, 749-806.

Variable	Neutral				Investment specific			
	Horizon (quarters)				Horizon (quarters)			
	1	8	16	32	1	8	16	32
A:**Full sample**								
Output	1	3	15	32	16	34	34	27
Hours	31	21	16	12	28	49	58	63
Hours per Worker	52	41	33	26	2	22	34	44
Unemployment	57	43	37	30	1	16	27	36
Finding Rate	22	35	31	26	1	11	22	33
Separation Rate	37	35	34	33	3	11	14	16
B. 1967:II-1997:IV sample								
Output	1	6	27	48	4	31	35	25
Hours	24	17	10	6	16	45	62	70
Hours per Worker	44	31	25	21	3	11	26	37
Unemployment	43	33	29	26	8	7	18	27
Finding Rate	17	30	28	27	9	4	9	13
Separation Rate	28	26	25	23	1	7	12	16
C. 1967:II-2007:I sample								
Output	2	10	25	47	7	35	40	38
Hours	14	12	8	7	18	42	58	67
Hours per Worker	35	32	28	26	1	7	17	25
Unemployment	36	35	33	31	5	4	9	14
Finding Rate	8	24	24	23	5	3	7	10
Separation Rate	21	28	28	28	1	2	4	5

Table 1: Forecast Error Variance Decomposition: percentage of the forecast error variance explained by neutral and investment-specific technology shocks. The VAR has siz variables and intercept deterministically broken at 1973:II and 1997:I. The variables of the VAR 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.

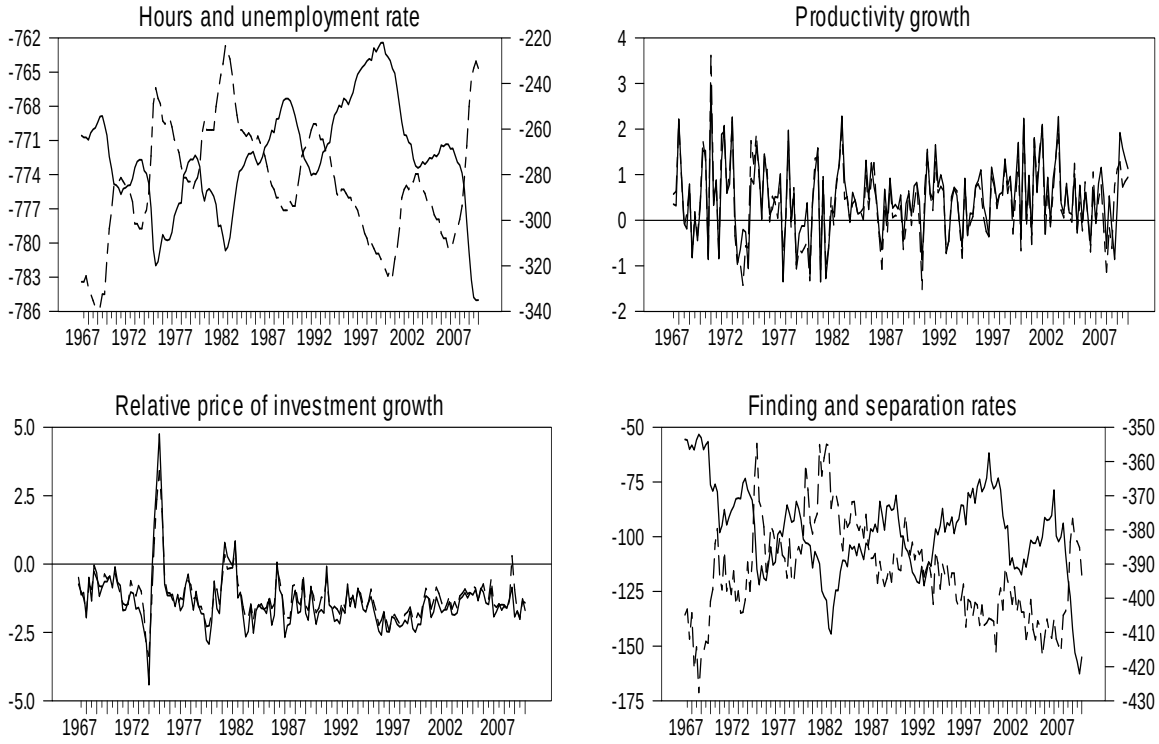

Figure 1. First graph: the continuous line is the log of the aggregate number of hours worked per-capita; the dashed line is the log of civilian unemployment. Second graph: the continuous line is the growth rate of labor productivity in the non-farm business sector, measured in output units; the dashed line the same variable, measured in consumption units. Third graph: the continuous line is the growth rate of the relative price of investment goods measured in output units; the dashed line the same variable measured in consumption units. Fourth graph: the continuous line is the log of job finding rate and the dashed line is the log of job separation rate.

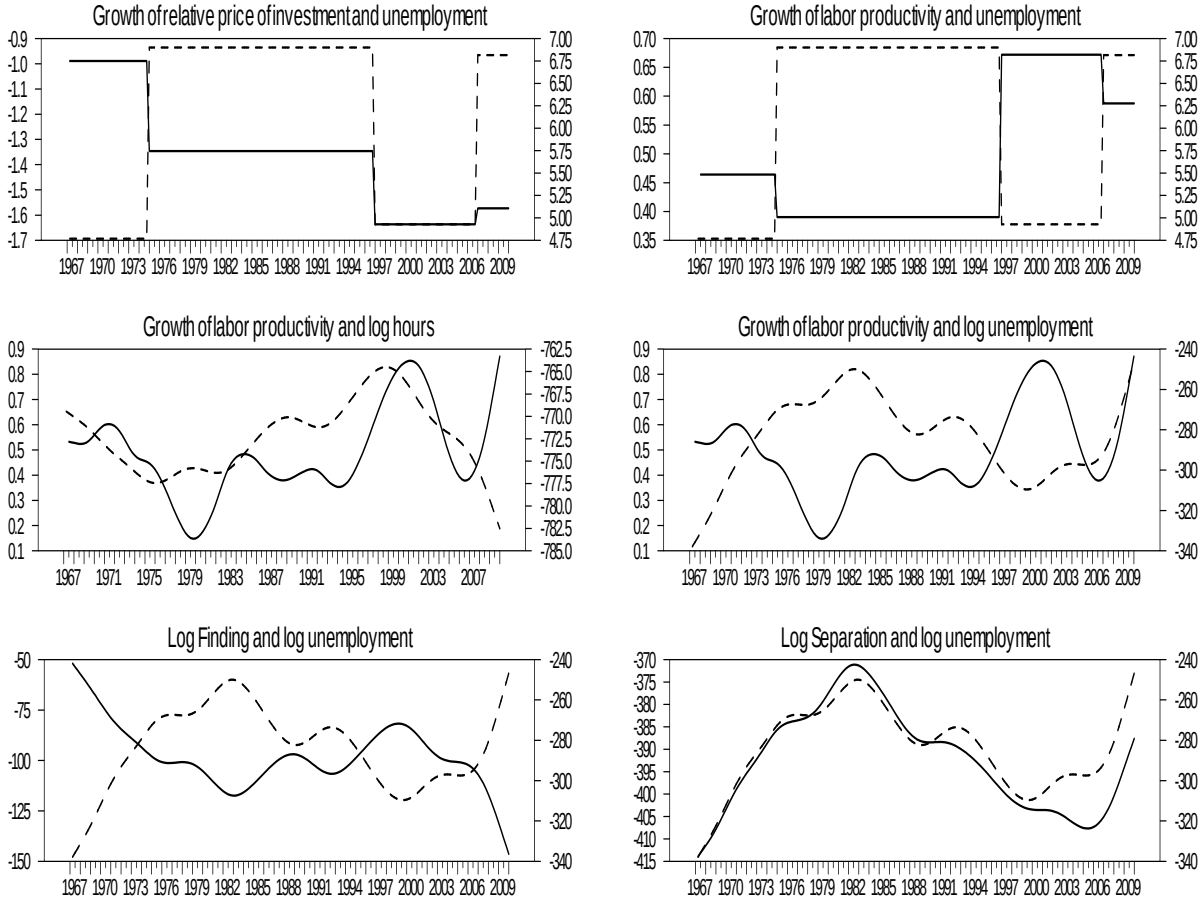

Figure 2.First graph: the continuous line is the average quarterly growth rate of the relative price of investment and the dashed line the unemployment rate. Second graph: the continuous line is the average quarterly growth rate of labour productivity and the dashed line the unemployment rate. Third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

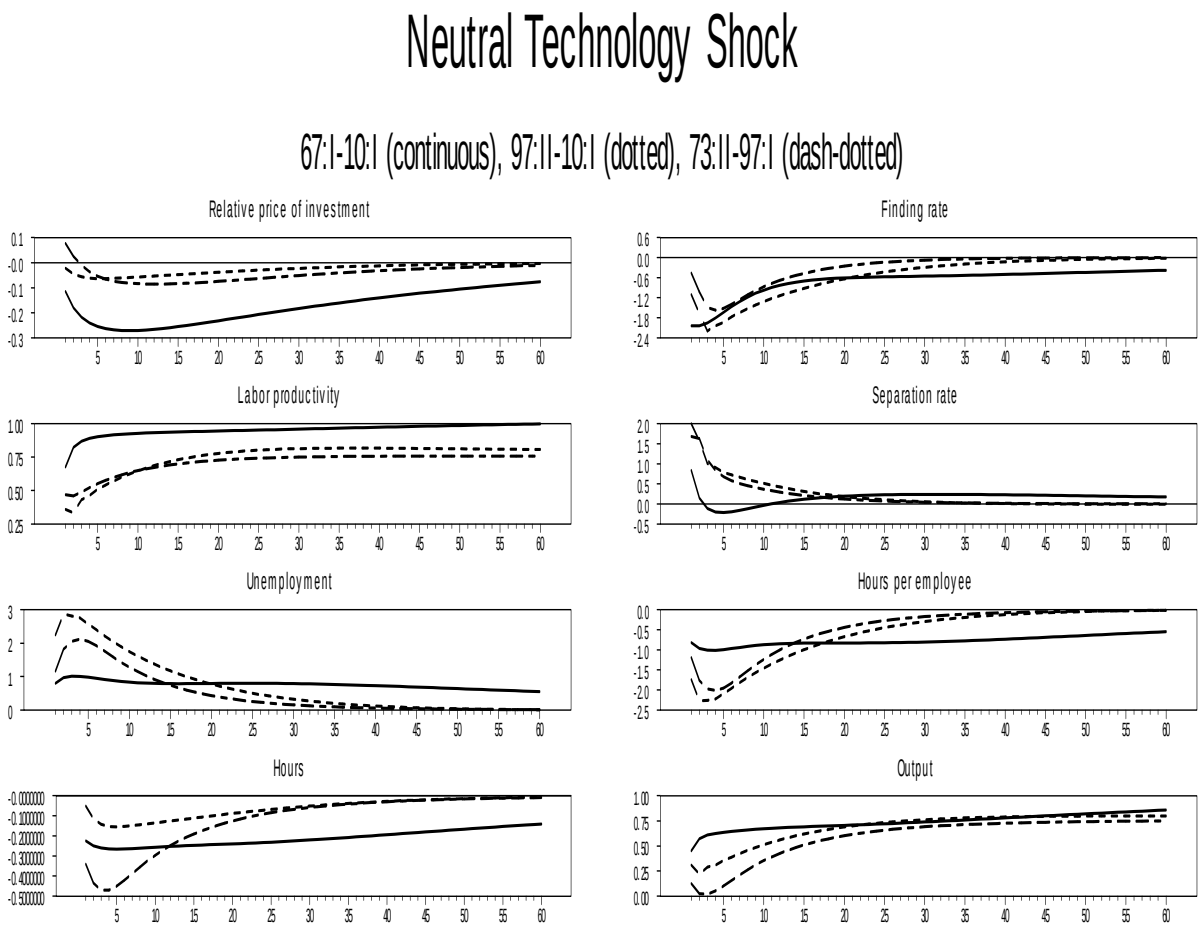

Figure 3. Responses to a one-standard deviation neutral shock. Each line corresponds to a six variable VAR(8) with the rate of growth of the relative price of investment, the rate of growth of labour productivity, the (logged) unemployment rate, and the (logged) aggregate number of hours worked per capita,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

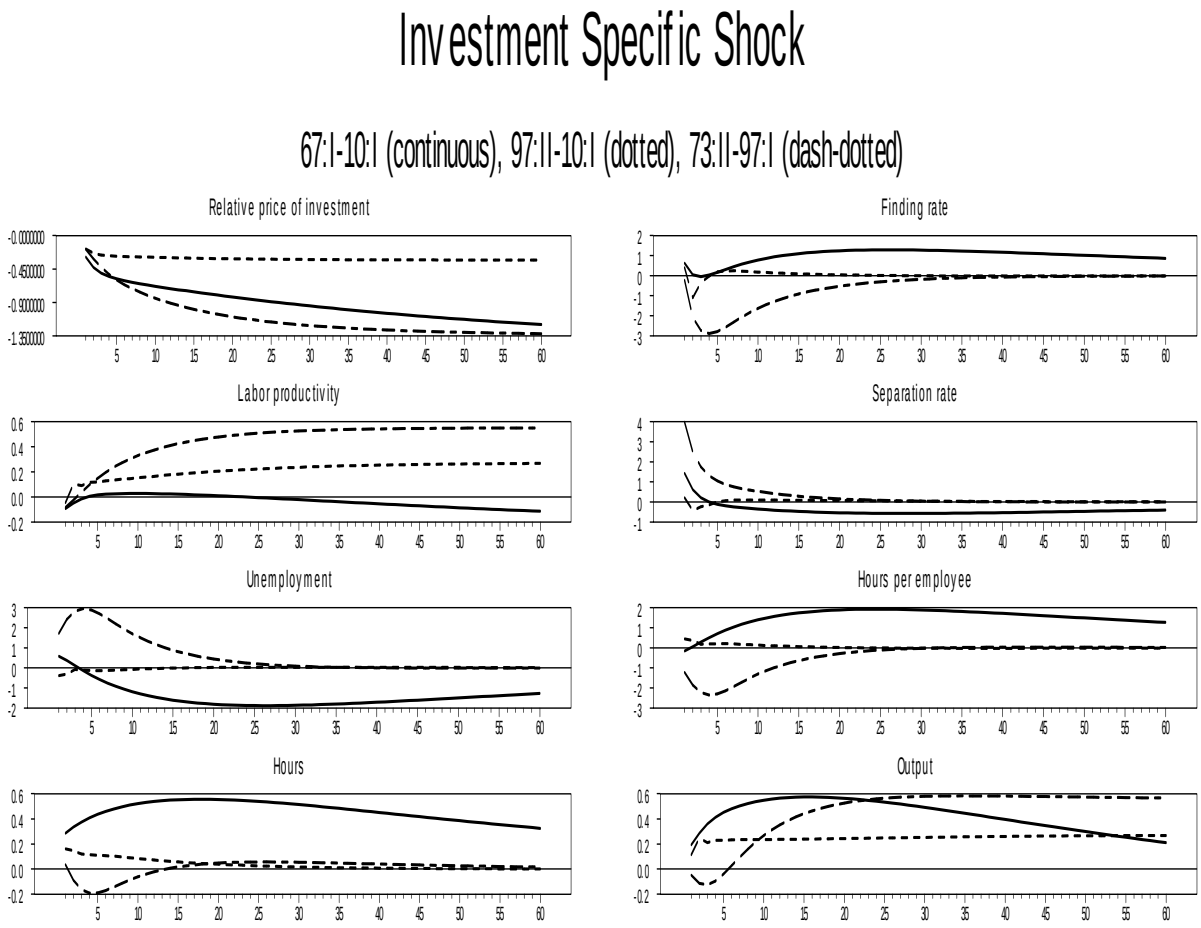

Figure 4. Responses to a one-standard deviation investment specific shock. Each line corresponds to a six variable VAR(8) with the rate of growth of the relative price of investment, the rate of growth of labour productivity, the (logged) unemployment rate, and the (logged) aggregate number of hours worked per-capita, the log of separation and finding rates.

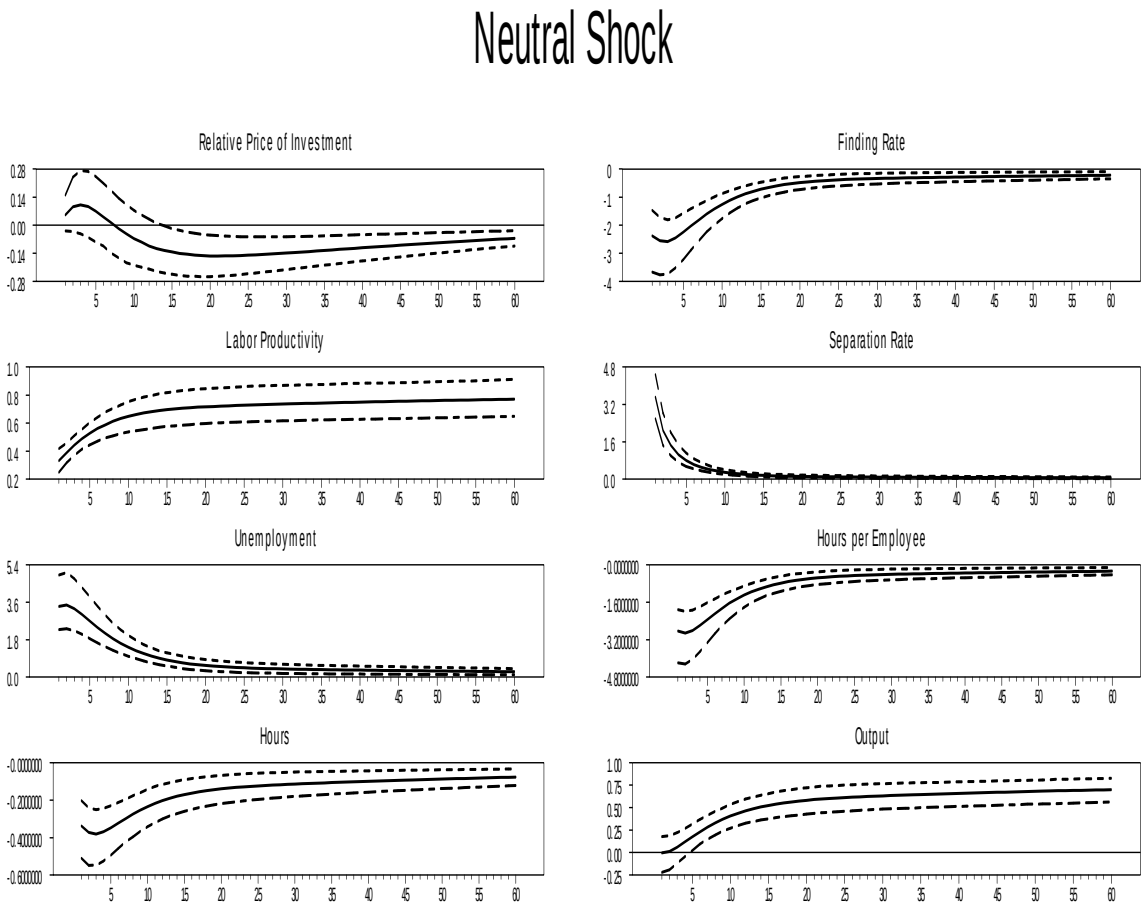

Figure 5. Responses to a one-standard deviation shock. Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I. Six variables VAR(8). Dotted lines are 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line is the median estimate.

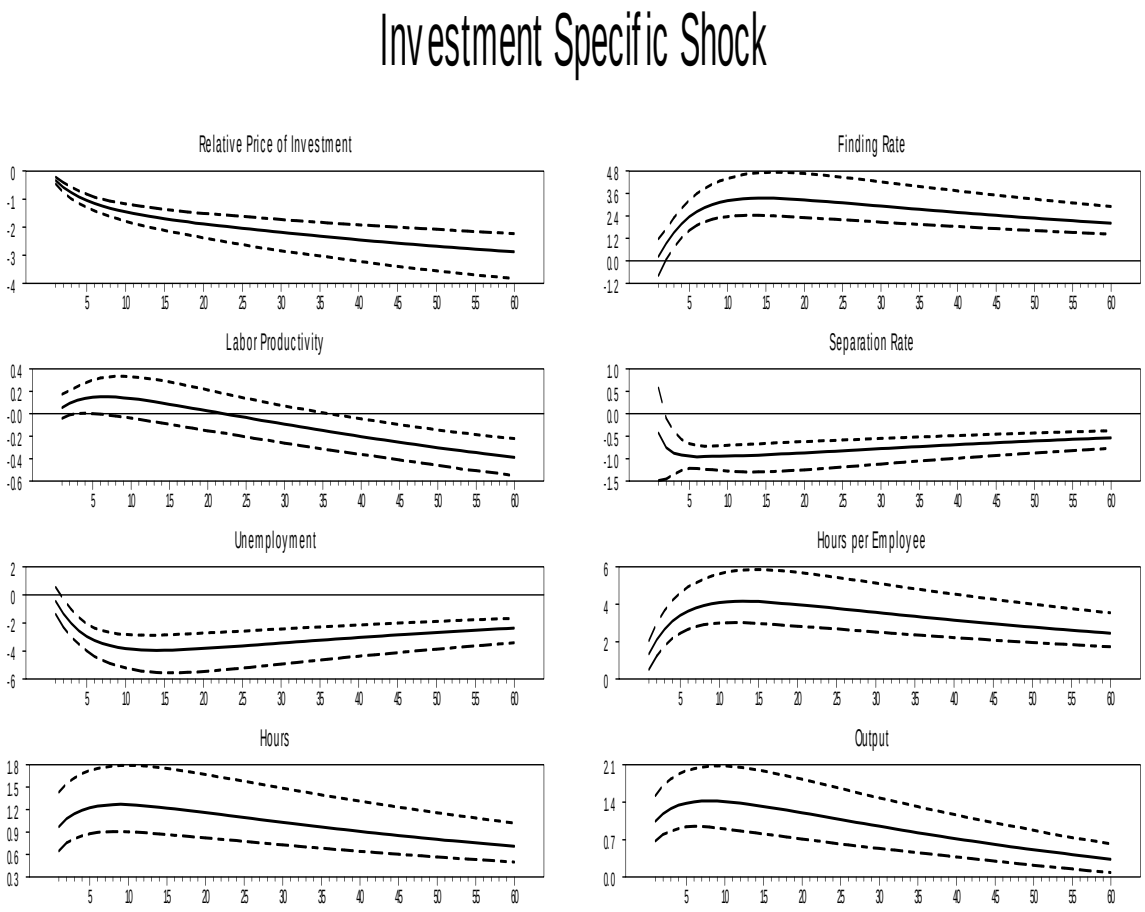

Figure 6. Responses to a one-standard deviation shock. Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I.Six variables VAR(8). Dotted lines are 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line is the median estimate.

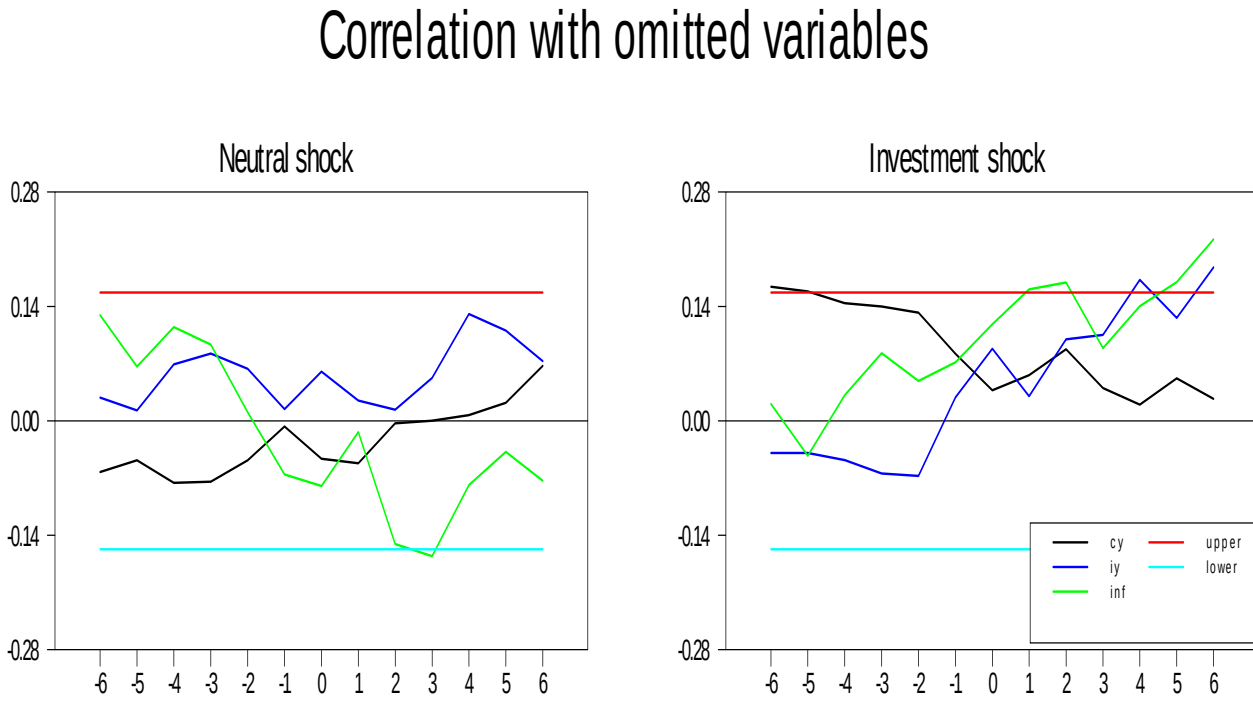

Figure 7. Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I. Plotted are the correlations of the neutral shocks (left column) and investment specific shock (right column) with the consumption-output ratio, the investment-output ratio andthe inflation rate. The shocks are estimated from the six variables VAR(8). The two horizontal lines correspond to an asymptotic 95 percent confidence interval for the null of zero correlation.

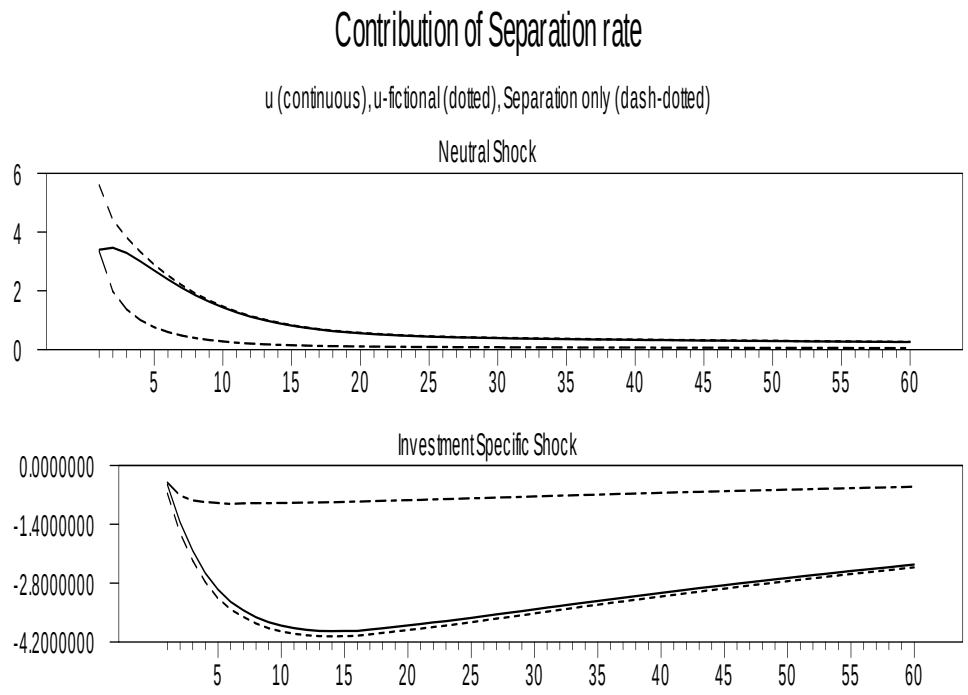

Figure 8. Six variables VAR(8). Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I. Reported are median estimates from 500 bootstrap replications.

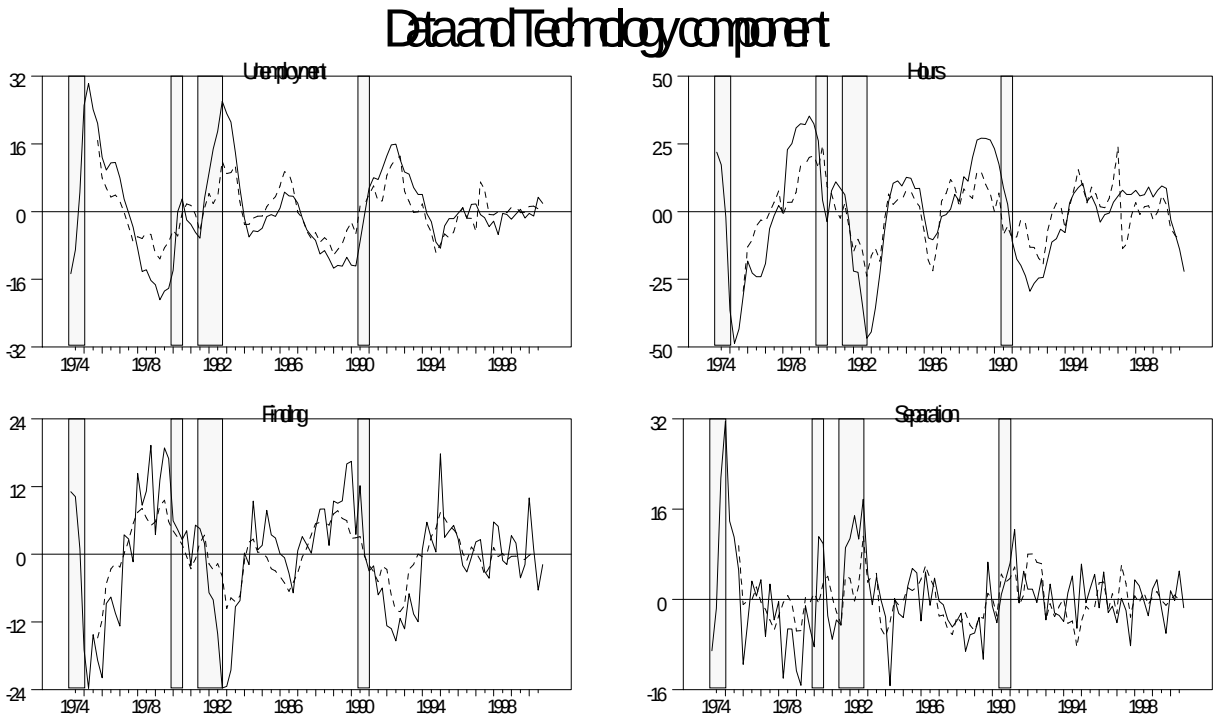

Figure 9. Technology shocks and labor market fluctuations. The solid line refers to raw data, the dashed lines to the component due to technology shock (either neutral or investment specific) as recovered from the six variables VAR(8). Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I. All series are detrended with a Hodrick Prescott filter with smoothing parameter equal to 1600.The areas in grey correspond to the NBER recessions.

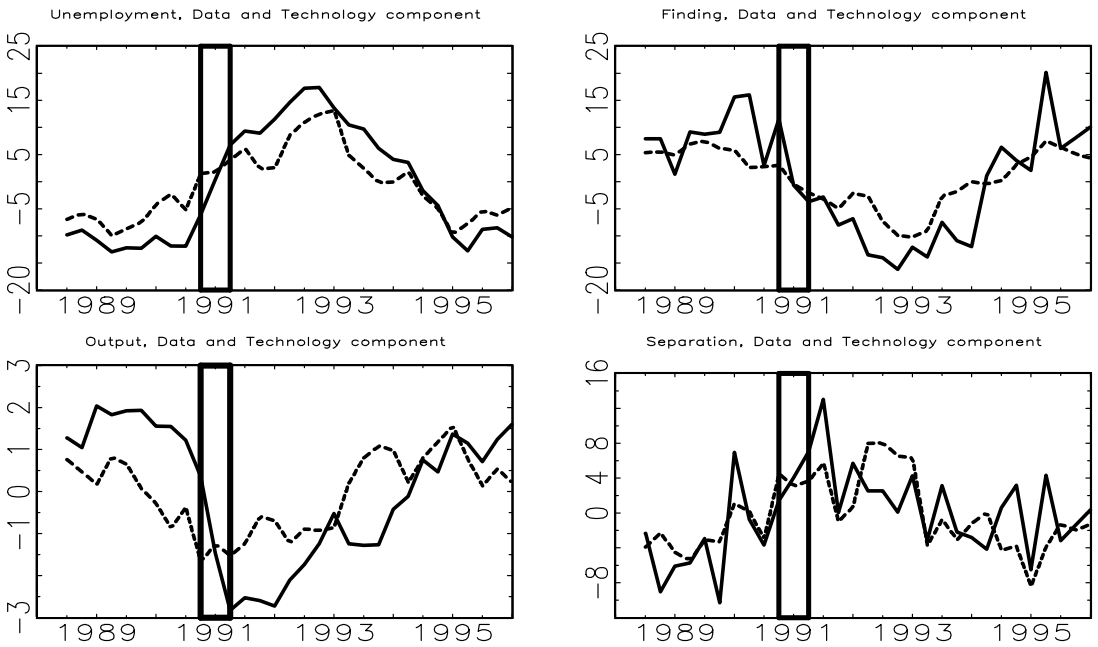

Figure 10. The jobless recovery of the 90s. Solid lines are raw data (either unemployment, finding rates, separation rates or output), the dashed lines the component due to technology shocks (either neutral or investment specific) as recovered from the six variables VAR(8) Sample 1967:II-2010:I with intercept deterministically broken at 1973:II and 1997:I. All series are detrended with a Hodrick Prescott filter with smoothing parameter equal to 1600. The vertical lines identifies the NBER recession.

Appendix

THIS APPENDIX CONTAINS ADDITIONAL EMPIRICAL RESULTS. IT IS PROVIDED FOR BACKING UP STATEMENTS MADE IN THE PAPER AND IT IS NOT INTENDED FOR PUBLICATION.

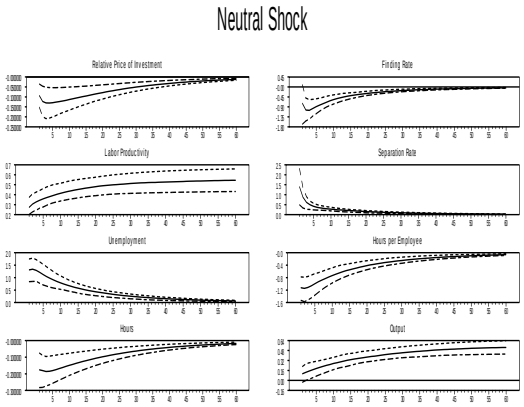

(a) Neutral technology shock

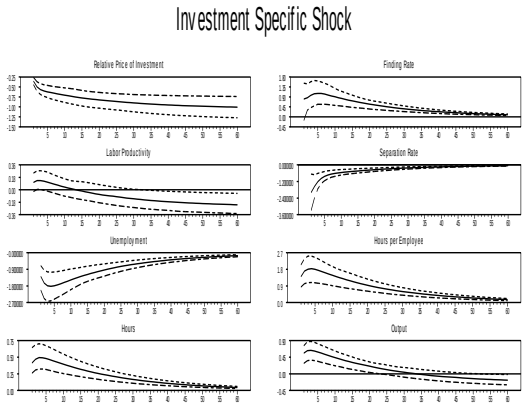

(b) investment specific technology shock

Figure A-1: The sample period is 1973:I-1997:I. The VAR has eight lags and contains: the rate of growth of the relative price of investment, the rate of growth of labour productivity, the (logged) job finding rate, the (logged) job separation rate, the (logged), unemployment rate (logged), and the (logged) aggregate number of hours worked per capita. Dotted lines represent the 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line corresponds to median estimate from bootstrap replications.

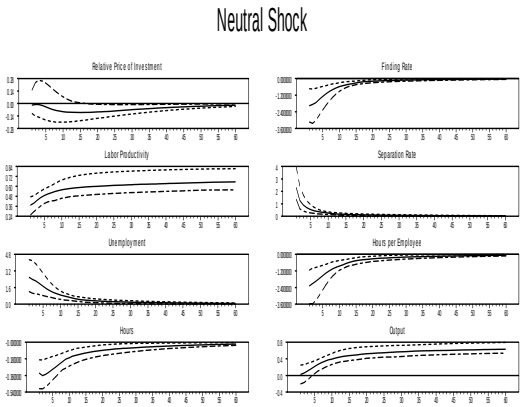

(a) Neutral technology shock

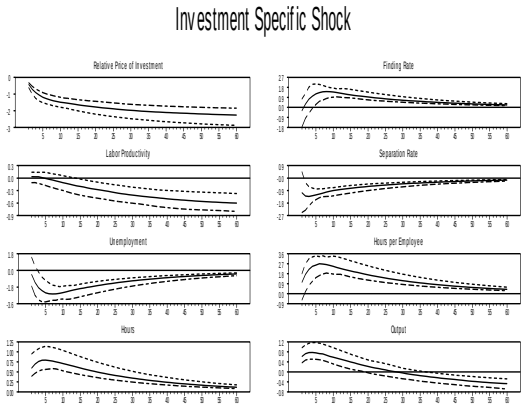

(b) Investment specific technology shock

Figure A-2: The sample period is 1967:II-1997:I. The VAR has eight lags and contains: the rate of growth of the relative price of investment, the rate of growth of labour productivity, the (logged) job finding rate, the (logged) job separation rate, the (logged), unemployment rate (logged), and the (logged) aggregate number of hours worked per capita. Dotted lines represent the 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line corresponds to median estimate from bootstrap replications.

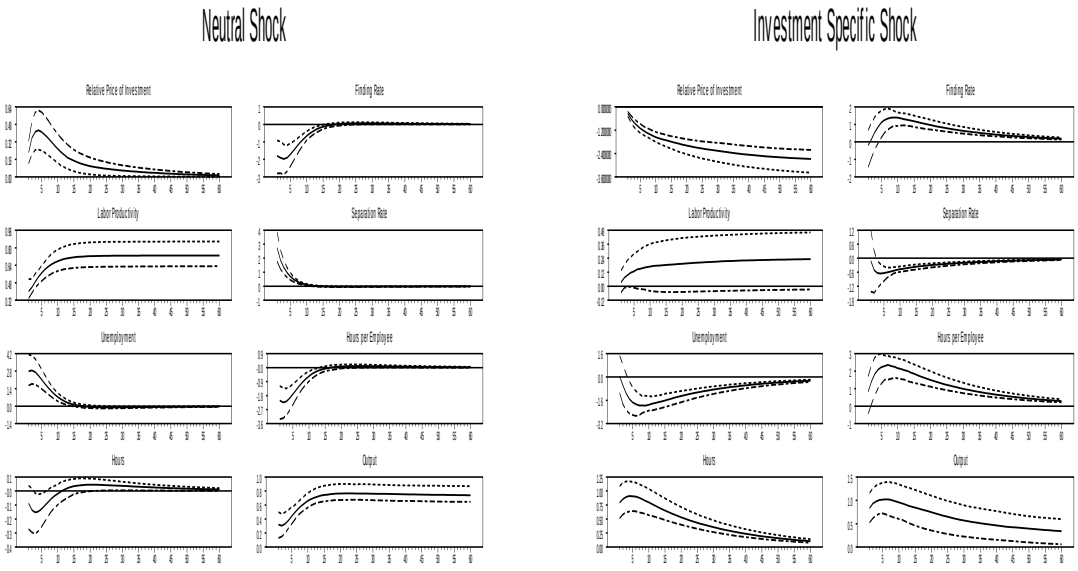

(a) Neutral technology shock

(b) Investment specific technology shock

Figure A-3: The sample period is 1967:II-2007:I. The VAR has 8 lags and contaians: 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. Dotted lines represent the 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line corresponds to median estimate from bootstrap replications.

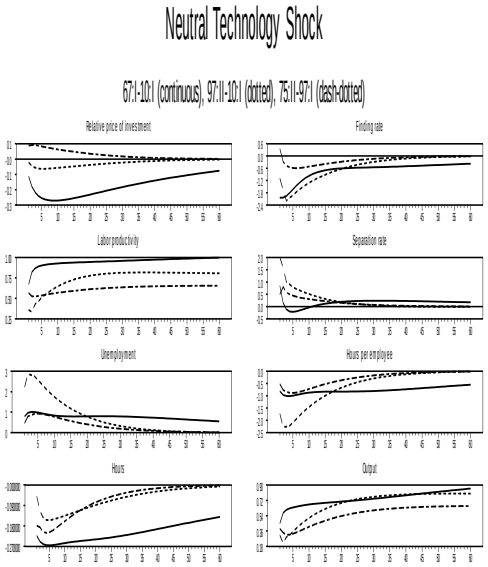

(a) Neutral technology shock

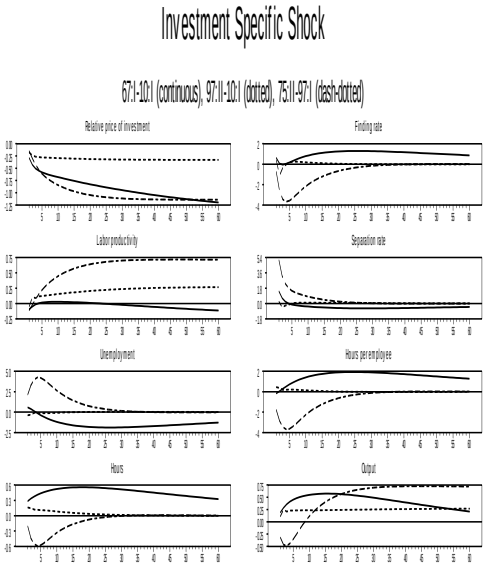

(b) Investment specific technology shock

Figure A-4: Responses to a one-standard deviation shocks in the samples: 1967:II-2010:I, 1975:II-1997:I, and 1997:II-2010:I. 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.

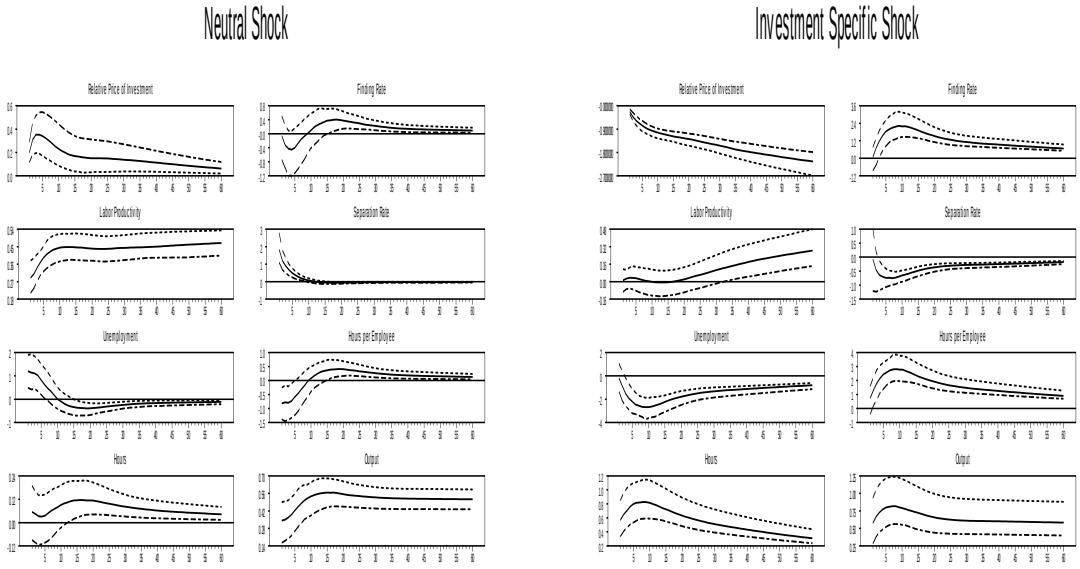

(a) Neutral technology shock

(b) Investment specific technology shock

Figure A-5: Response to a neutral or an investment-specific technology shock in a six variables VAR(8), 1967:II-2010:I sample with intercept deterministically broken at 1973:I and 1997:I. Dotted lines represent the 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR. The continuous line corresponds to median estimate.

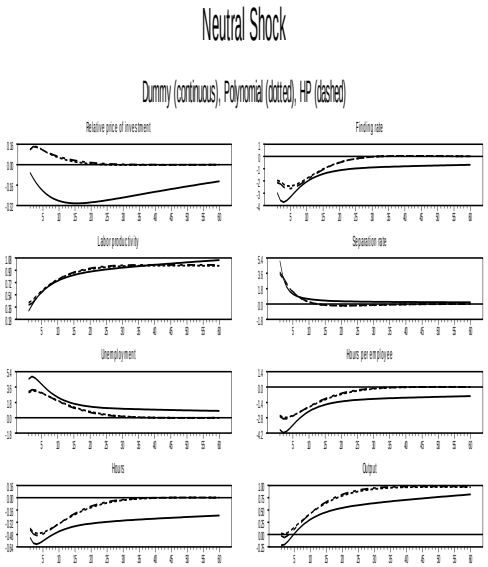

(a) Neutral technology shock

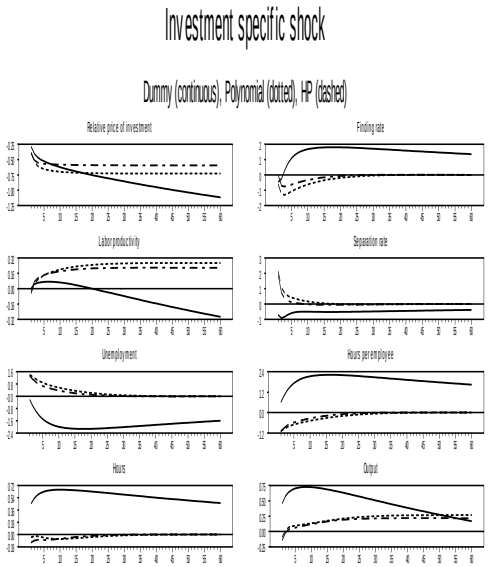

(b) Investment specific technology shock

Figure A-6: Six variable VAR with 8 lags, sample 1967:II-2010:I with intercept deterministically broken at 1973:1 and 1997:I. The continous line has responses for the dummy specification, the dotted line responses where the intercept is a 3rd order polynomial in time, the dashed lines are responses after filtering with an Hodrick Prescott filter and smoothing parameter $\lambda = 12800$.

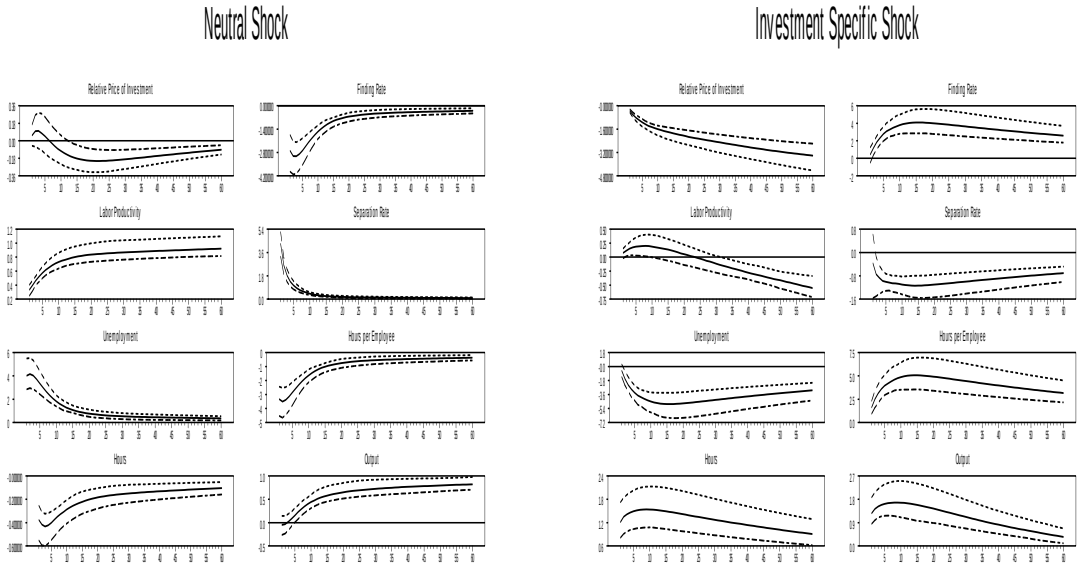

(a) Neutral technology shock

(b) Investment specific technology shock

Figure A-7: Six variable VAR with 4 lags, sample 1967:II-2010:I with intercept deterministically broken at 1973:1 and 1997:I. Dotted lines represent the 5% and 95% quantiles of the distribution of the responses simulated by bootstrapping 500 times the residuals of the VAR; the solid line is the median of the distribution.

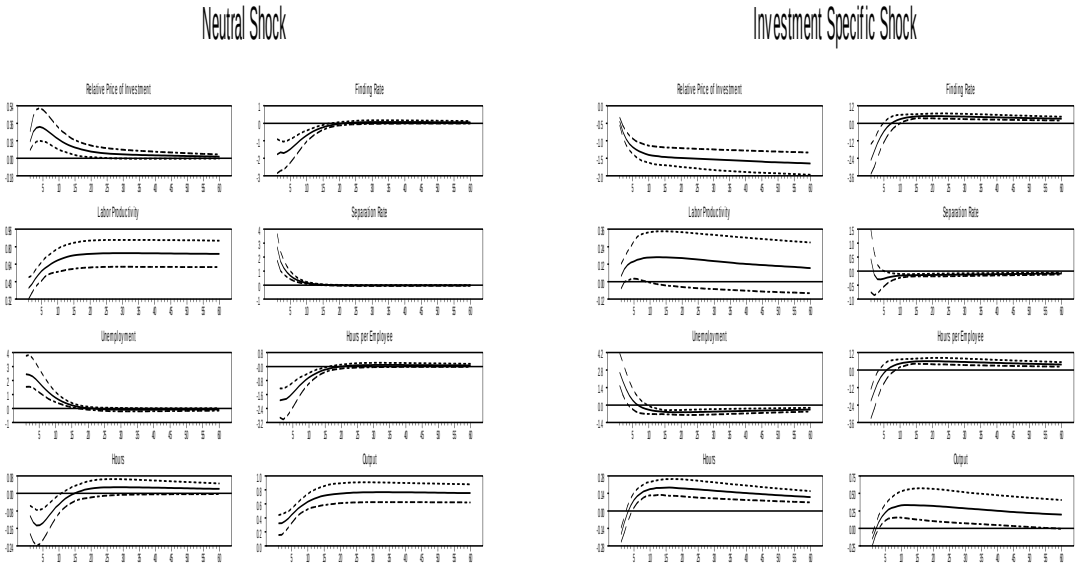

(a) Neutral technology shock

(b) Investment specific technology shock

Figure A-8: Six variable VAR with 8 lags, sample 1967:II-2010:I with intercept deterministically broken at 1973:1 and 1997:I. Identifycation restrictions imposed at medium horizon (12 quarters). Dotted lines represent the 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 of the distribution.

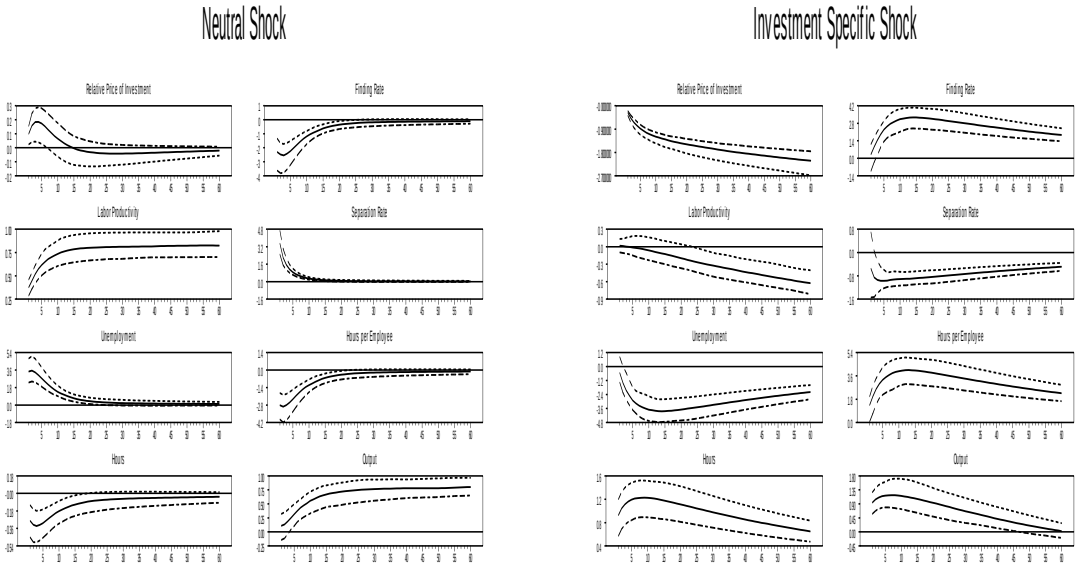

(a) Neutral technology shock

(b) Investment specific technology shock

Figure A-9: Six variable VAR with 8 lags, sample 1967:II-2010:I with intercept deterministically broken at 1973:1 and 1997:I. The variables in VAR are deflated with a consumer price index. Dotted lines represent the 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 of the distribution.