

Labor Market Dynamics: A Time-Varying Analysis*

HAROON MUMTAZ[†], FRANCESCO ZANETTI[‡]

[†]*Queen Mary University, Mile End Road, London, E1 4NS, UK
(e-mail: h.mumtaz@qmul.ac.uk)*

[‡]*Department of Economics, University of Oxford, Manor Road Oxford, OX1 3UQ, UK
(e-mail: francesco.zanetti@economics.ox.ac.uk)*

Abstract

This paper studies how key labour market stylized facts and the responses of labour market variables to technology shocks vary over the US postwar period. It uses a benchmark dynamic, stochastic, general equilibrium model enriched with labour market frictions and investment-specific technological progress that enables a novel identification scheme based on sign restrictions on a SVAR with time-varying coefficients and stochastic volatility. Key findings are: (i) the volatility in job finding and separation rates has declined over time, while their correlation varies across time; (ii) the job finding rate plays an important role for unemployment, and the two series are strongly negatively correlated over the sample period; (iii) the magnitude of the response of labour market variables to technology shocks varies across the sample period.

I. Introduction

The dynamics of the labour market have been a subject of intense empirical and theoretical research over the past three decades. This paper contributes to this realm of research by studying how the statistical relation among key labour market aggregates and their response to technology shocks change over the US postwar period. To this aim, it develops an estimated time-varying parameter vector autoregression (TVP-VAR) model with stochastic volatility, whose variables' reaction to technology shocks is identified using the cyclical properties of a dynamic, stochastic, general equilibrium (DSGE) model of the business cycle characterized by labour market frictions.

We use a model with labour market frictions to investigate the dynamics of the labour market due to their empirical relevance and theoretical appeal. Empirically, Rogerson and Shimer (2011) summarize evidence showing that labour markets are characterized by

JEL Classification numbers: E32, C32.

*We would like to thank Anindya Banerjee and three anonymous referees for extremely helpful suggestions and Sophocles Mavroeidis, Lydia Silver and seminar participants at the Bank of England and the University of Oxford for very useful comments.

frictions that prevent the competitive market mechanism from determining labour market equilibrium allocations, thereby suggesting that their presence is important for a realistic description of the functioning of the labour market. Theoretically, labour market frictions introduce the extensive margin of labour (i.e. (un)employment) into the model, whose dynamics depend on the flows of workers in and out of unemployment. Importantly for the analysis of the paper, we make the rate at which jobs are destroyed endogenous, so that flows in and out of unemployment result from the incentives that workers and firms have to engage in production or terminate their relation. In this way, the theoretical framework details the dynamics of unemployment, job finding and job separation rates, whose reaction to shocks enables a new identification scheme.

The analysis establishes important stylized facts of the US labour market. In particular, it uncovers the following findings:

- (i) The volatility of job finding and job separation rates declines over the sample period, after reaching a peak around the mid-1980s. However, changes in the volatility of the job finding and job separation rates display different patterns over sub-periods.
- (ii) The correlation of the labour market variables with GDP growth is relatively stable over time, whereas the correlation between the job finding and job separation rates and the unemployment rate shows significant time variation. In particular, the job finding rate plays an important role for the unemployment rate, and the two series are strongly negatively correlated over the sample period.
- (iii) The magnitude of the response of labour market variables to neutral technology shocks varies over time, whereas the response to investment-specific technology shocks is substantially constant. Neutral technology shocks decrease unemployment.
- (iv) Across the sample period, neutral technology shocks explain approximately half of the movements in unemployment and the job finding rate, whereas they explain only 20% of fluctuations in the job separation rate. Investment-specific technology shocks contribute to, on average, < 20% of fluctuations in unemployment, job finding and job separation rates.
- (v) The time-varying trends of the job separation rate and the unemployment rate show a similar pattern, and they peak in the early 1980s.

The contribution of this paper is threefold. First, to the best of our knowledge no papers have yet investigated how the statistical relation among key labour market variables and their response to technology shocks changes over time. A few studies, detailed below, investigate the time-varying response of macroeconomic variables to shocks, but none of them focuses on the dynamics of the labour market.

Second, the theoretical framework embeds endogenous job destruction in a model in which technology shocks are distinguished between neutral and investment-specific technological processes, since the latter are key to study the dynamics of the technological progress, as shown in Greenwood, Hercowitz and Krusell (1997), Fisher (2006) and Justiniano and Michelacci (2011). Therefore, the model provides new insights into the time-varying effect of investment-specific technology shocks on labour market variables, and, importantly for our analysis, enables a novel identification scheme.

Finally, we propose a novel identification scheme based on sign restrictions that use information from real activity, job finding and separation rates and unemployment. In this way, our high-frequency identification scheme imposes a minimal set of constraints on the model compared to low- or medium-frequency identification restrictions. This is particularly important in the context of labour market variables, since Fernald (2007) points out that any procedure that includes low- or medium frequencies generates an artificial positive comovement between labour input and neutral technology shocks that disappear once controlling for long cycles. In general, Faust and Leeper (1997) and Chari, Kehoe and McGrattan (2008) show that long-run restrictions may generate unreliable estimates as they are unable to accurately recover true underlying impulse response functions when estimated using data generated from a structural model.

Before proceeding, we relate this study to the literature. This paper contributes to two strands of the literature. First, it contributes to the literature that studies the cyclical properties of the labour market. Shimer (2012) identifies important labour market stylized facts using business cycle statistics. We enrich this realm of research by allowing for time variation in the analysis, thereby uncovering changes in labour market dynamics across time. In addition, we also allow for both neutral and investment-specific technology shocks, so as to provide a more comprehensive assessment of the influence of technological process on the dynamics of the labour market. Michelacci and Lopez-Salido (2007) and Ravn and Simonelli (2008) use theoretical models with labour market frictions to identify the effect of technology shocks using a SVAR. Compared to these studies, we use worker flows and allow for time variation in coefficients and stochastic volatility in the estimation. Similarly to us, studies by Gali and Gambetti (2009), Primiceri (2005), Cogley, Primiceri and Sargent (2010) and Benati and Surico (2008) investigate the time-varying response of macroeconomic variables such as output, inflation, monetary aggregates and the nominal interest rate to shocks. However, our focus is on the dynamics of the labour market, and we identify shocks using short-run restrictions.

Second, this paper contributes to the literature on the identification of technology shocks. Similarly to us, Uhlig (2004), Dedola and Neri (2007), Paustian (2007), Pappa (2009), Canova and Paustian (2011) and Mumtaz and Zanetti (2012), use short-run identification schemes to investigate the reaction of labour input to technology shocks. However, none of these studies, with the exception of Mumtaz and Zanetti (2012), uses restrictions based on labour market variables. In particular, this last study identifies technology shocks using labour market variables such as hiring and labour market tightness. Instead, we identify technology shocks based on their effect on real activity, job finding and separation rates and unemployment, whose interaction provides a theoretically consistent characterization of the relation among labour market variables. Finally, all these mentioned studies focus on the impact of technology shocks, whereas our paper focuses on how the statistical relations among key labour market aggregates and their responses to technology shocks change over the US postwar period.

The remainder of the paper is organized as follows. Section II describes the theoretical model and details the sign restrictions from the theoretical model. Section III describes the TVP-VAR model with stochastic volatility and the implementation of the identification scheme. Section IV presents the results. Section V concludes.

II. The theoretical model and theoretical restrictions

We use a model with search and matching frictions that resembles Mumtaz and Zanetti (2012). However, we enrich the theoretical framework with endogenous job destruction, as in Mortensen and Pissarides (1994) and Thomas and Zanetti (2009), which relates movements in both job creation and finding rates to changes in the endogenous variables and exogenous changes to technology. In addition, we also embed the investment-specific technology progress as in Greenwood *et al.* (1997) and the cost of hiring as in Mandelman and Zanetti (2014).

The model economy is populated by three types of agents: households, firms and a fiscal authority. Households consist of a large number of members, a fraction of which are unemployed and searching for jobs. On the other side of the labour market, firms hire workers. The fiscal authority simply balances the budget in every period. The theoretical model and its calibration are described in Mumtaz and Zanetti (2014).

To derive the sign restrictions to impose on the empirical TVP-VAR model we use the theoretical model to determine how each variable reacts to shocks. To produce robust responses to one positive percentage point neutral and investment-specific technology shocks we simulate the theoretical model by drawing 10,000 times from a range of plausible parameters described in Mumtaz and Zanetti (2014). As in Dedola and Neri (2007), Pappa (2009) and Mumtaz and Zanetti (2012), to eliminate extreme responses, we discard the regions of the two distributions below and above 2.5 and 97.5 percentiles respectively. To illustrate how the variables of the theoretical model react to each shock, Figure 1 plots the impulse responses of variables to a one positive percentage point deviation of the neutral and investment-specific technology shocks respectively. Independently from the shock considered, capital and investment show similar dynamics, as they both rise. In addition, the impact response of output growth could be either positive or negative for both shocks, depending on the model's calibration. However, the impact response is more pronounced in the case of a neutral technology shock, which corroborates the findings in Greenwood *et al.* (1997) and Fisher (2006). The impact reactions of consumption, hiring and the job finding rate to a neutral technology shock are positive, whereas the response of the job destruction rate is negative. The intuition of these results is straightforward. In response to a positive technology shock, hiring increases as firms expand production by increasing labour input, while job destruction falls since the productivity of the marginal job increases. Consequently, the unemployment rate falls, which combined with the increase in hiring generates a rise in the job finding rate. On the other hand, in the face of an investment-specific technology shock, the unemployment rate rises since capital is more productive and, as described, firms respond to this by expanding production. As a consequence, hiring and the number of workers both decrease, thereby decreasing the job finding rate. Importantly for the analysis of this paper, since the job finding rate and the unemployment rate have opposite reactions to neutral or investment-specific technology shocks we are able to disentangle the effects of these two shocks in the data. To implement the identification scheme, we impose the described sign restrictions, as summarized in Table 1, on the first-period reaction of the TVP-VAR model. Note that we include GDP growth in the empirical model, but leave its reaction unconstrained as the impact reaction in the theoretical model could be either positive or negative, depending on the model's calibration.

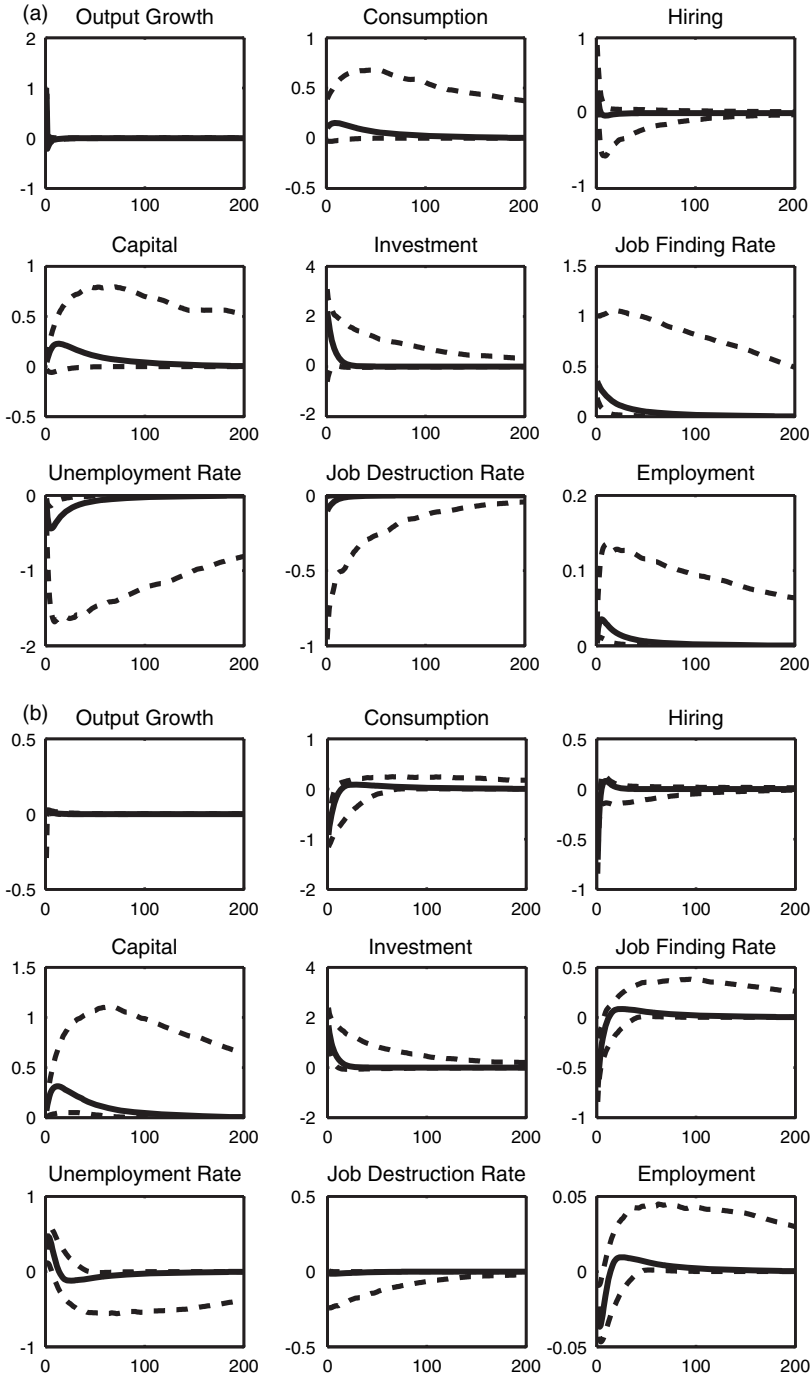


Figure 1. Theoretical impulse-response functions. (a) Neutral technology shock. (b) Investment-specific technology shock

Notes: Panel a (panel b) shows the percentage-point response of one of the model's variables to a one-percentage-deviation neutral (investment specific) technology shock. The solid line reports the median responses and the dashed lines report the 2.5 and 97.5 percentiles of the responses.

TABLE 1
Sign restrictions on the first-period TVP-VAR variables

Variable	Neutral technology shock	Investment-specific technology shock
GDP growth	Free	Free
Job finding rate	+	-
Job destruction rate	-	-
Unemployment rate	-	+

Notes: Entries show sign restrictions on the first period TVP-VAR variables to neutral and investment-specific technology shocks.

III. The empirical model

In this section, we describe the empirical TVP-VAR model with stochastic volatility, the identification scheme based on sign restrictions, the Bayesian estimation and the data.

We consider the following VAR model

$$Z_t = c_t + \sum_{l=1}^L \phi_{l,t} Z_{t-l} + v_t, \tag{1}$$

where Z_t contains the job finding rate, the job separation rate, GDP growth and the unemployment rate. Our benchmark model allows for time variation in the parameters. We postulate the following law of motion for the coefficients $\tilde{\phi}_{l,t} = \tilde{\phi}_{l,t-1} + \eta_t$, where $\tilde{\phi}_{l,t} = \{c_t, \phi_{l,t}\}$. As in Cogley and Sargent (2005), the covariance matrix of the innovations v_t is factored as $\text{VAR}(v_t) \equiv \Omega_t = A_t^{-1} H_t (A_t^{-1})'$. The time-varying matrices H_t and A_t are defined as:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad A_t \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix} \tag{2}$$

with the $h_{i,t}$ evolving as geometric random walks, $\ln h_{i,t} = \ln h_{i,t-1} + \tilde{v}_t$. Following Primiceri (2005), we postulate the non-zero and non-one elements of the matrix A_t to evolve as driftless random walks, $\alpha_t = \alpha_{t-1} + \tau_t$. The time-varying VAR model can be written compactly as

$$y_t = x_t' \tilde{B}_t + A_t^{-1} H_t^{1/2} \varepsilon_t \tag{3}$$

where $y_t = \text{vec}(Z_t)$, $x_t = I \otimes [1, Z_{t-1}, Z_{t-2}, \dots]$, $\tilde{B}_t = \text{vec}([c_t, \phi_{1,t}, \phi_{2,t}, \dots])$ and $\text{VAR}(\varepsilon_t) = I$.

The time-varying VAR model in equation (3) represents a flexible framework which is particularly suited to our analysis that considers changes in the role and transmission of technology shocks. Consider re-writing equation (3) as

$$y_t = x_t \tilde{B}_t + \tilde{A}_{0,t} \varepsilon_t, \tag{4}$$

where $\tilde{A}_{0,t}$ represents a time-varying structural impact matrix such that: $\Omega_t = \tilde{A}_{0,t} \tilde{A}_{0,t}'$.

The structural VAR in equation (4) allows flexibility along several dimensions. First, the magnitude of the contemporaneous relationships amongst v_t is allowed to be different across time. This seems particularly appropriate for the US which has experienced several structural shifts. Within a simple economic model, these structural changes imply a change in the contemporaneous reaction of macroeconomic variables to structural shocks. Therefore, an empirical model with fixed impact matrix \tilde{A}_0 is unable to account for this feature of the data. Moreover, structural changes in the economy may have occurred along several dimensions, implying independent shifts in different (structural) equations of the model. By allowing for independent time variation in each contemporaneous and lagged coefficient, it is likely that the model is a good proxy for structural change with these features. In a similar vein, the time-varying VAR has a flexible formulation for volatility allowing shifts in shock volatility that are independent from changes in the coefficients B_t .

Identification of structural shocks

As mentioned, the structural analysis using the VAR model is based on the identification of the signs of the responses of the endogenous variables to neutral and investment-specific technology shocks, as described in the previous section and summarized in Table 1.

In the empirical analysis below, we impose these restrictions on the contemporaneous impulse responses estimated using the VAR model in equation (1). The sign restrictions are imposed using the procedure described in Rubio-Ramírez, Waggoner and Zha (2010), and the identification scheme is implemented as follows. Let $\Omega_t = P_t P_t'$ be an arbitrary decomposition of the VAR covariance matrix Ω_t , and let $\tilde{A}_{0,t} \equiv P_t$. We draw an $N \times N$ matrix, J , from the $N(0, 1)$ distribution. We take the QR decomposition of J . That is, we compute the matrices Q and R such that $J = QR$. This gives us a candidate structural impact matrix as $A_{0,t} = \tilde{A}_{0,t} Q$. We check if the rows of the A_0 matrix are consistent with the restrictions in Table 1. If this is the case we store $A_{0,t}$. If the sign restrictions are not satisfied, we draw another J and repeat the above.

With time-varying coefficients, the calculation of impulse responses is complicated by the possibility of coefficient variation over the impulse response horizon. To tackle this issue, we follow the procedure in Koop, Pesaran and Potter (1996) and use Monte Carlo integration to account for future coefficient uncertainty. In particular, the impulse response functions at each point in time are defined as:

$$IRF_t = E(Y_{t+k} \setminus \Psi_t, Y_{t-1}, \mu) - E(Y_{t+k} \setminus \Psi_t, Y_{t-1}) \quad (5)$$

where Ψ_t denotes all the parameters and hyperparameters of the VAR, k is the horizon under consideration and μ denotes the shock. Equation (5) states that the impulse response functions are calculated as the difference between two conditional expectations. The first term in equation (5) denotes a forecast of the endogenous variables conditioned on one of the structural shocks, μ . The second term is the baseline forecast, i.e. conditioned on the scenario where the shock equals zero. Koop *et al.* (1996) describe how to approximate these conditional expectations via a stochastic simulation of the VAR model.

Estimation and data

The TVP-VAR model is estimated using the Bayesian methods described in Kim and Nelson (1999). In particular, we employ a Gibbs sampling algorithm that approximates the posterior distribution. A detailed description of the prior distributions, the sampling method and evidence of convergence are provided in an appendix available upon request from the authors. Here, we summarize the basic algorithm which involves the following steps:

1. The VAR coefficients \tilde{B}_t and the off-diagonal elements of the covariance matrix A_t are simulated by using the methods described in Carter and Kohn (2004). As is common practice in this literature (see Cogley and Sargent, 2005), we impose the constraint that \tilde{B}_t should be stable at each point in time.
2. The volatilities of the reduced form shocks, H_t , are drawn using the date by date blocking scheme introduced in Jacquier, Polson and Rossi (2004).
3. The hyperparameters Q and S are drawn from an inverse Wishart distribution while the elements of G are simulated from an inverse gamma distribution.

The lag length is set at 2. We compare the relative fit of a TVP-VAR(1) and TVP-VAR(2) using deviance information criterion (*DIC*).¹ The estimated *DIC* is virtually identical for the VAR with these two lag lengths ($-7,255.74$ for the TVP-VAR(1) and $-7,254.35$ for the TVP-VAR(2)). Therefore, we select two lags but check our results using a TVP-VAR(1). The results from the TVP-VAR(1) are very similar to the benchmark estimates presented below.²

The data are quarterly, seasonally adjusted and cover the period 1948:Q2–2007:Q1. Real GDP and the unemployment rate are obtained from the FRED database (mnemonics GDP96 and UNRATE respectively). Data on job creation and job separation probabilities are from Shimer (2012), and are quarterly averages of monthly transition probabilities, corrected for time-aggregation. As in Cogley and Sargent (2005) we obtain starting values and set priors using a training sample of 10 years with the estimation carried out from 1959:Q1.

IV. Results

This section documents the results. In particular, it shows time-varying statistics for the volatility and correlation among variables, the impulse response functions, the forecast error variance decomposition, the trends and an index of persistence and predictability of the variables.

Volatility and correlation

Figure 2 plots the time-varying unconditional variance (solid line) of each endogenous variable and the estimated variance when either the neutral technology shock (dashed line),

¹The DIC can be thought of as a generalization of the Akaike information criterion (see Berg, Meyer and Yu, 2004 for details). The calculation of the DIC requires the evaluation of the likelihood function of the TVP-VAR. We accomplish this via a particle filter.

²These results are available upon request from the authors.

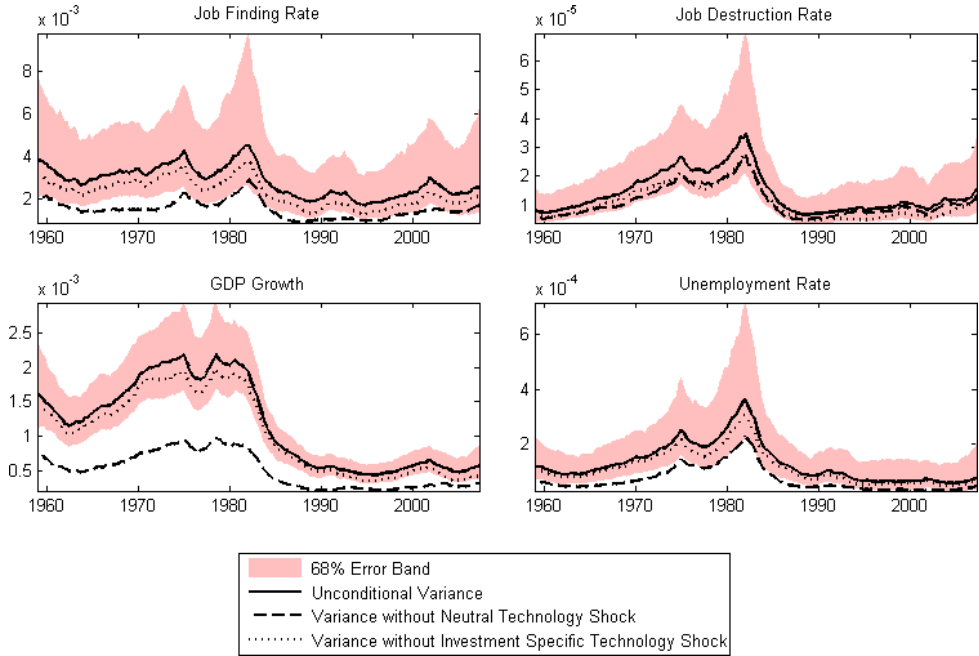


Figure 2. Time-varying unconditional volatilities

Notes: Each entry shows the unconditional variance (solid line) of each endogenous variable and the estimated variance when either the neutral technology shock (dashed line), or the investment-specific technology shock (dotted line) is set to zero.

or the investment-specific technology shock (dotted line) is set to zero. The unconditional volatility is calculated as $\text{vec}[\text{VAR}(Z_t)] = (I - \bar{\phi}_{l,t} \otimes \bar{\phi}_{l,t})^{-1} \text{vec}(\Omega_t)$. Given the structural impact matrix $A_{0,t}$ (with diagonal elements normalized to 1), the variance of structural shocks can be recovered as $\tilde{H}_t = A_{0,t}^{-1} \Omega_t A_{0,t}^{-1}$. We set the variance of either the neutral technology shock or the investment-specific technology shock in \tilde{H}_t to zero. We use this counterfactual estimate (denoted \tilde{H}_t^*) to build a counterfactual covariance matrix $\Omega_t^* = A'_{0,t} \tilde{H}_t^* A_{0,t}$ and recalculate $\text{VAR}(Z_t)$.

The volatility of the job finding rate is fairly constant over the sample period, although there is a slight decrease in the median after the early 1980s. It is clear from the figure that the unconditional volatility of the job finding rate is largely driven by the neutral technology shock. This shock explains about 60% of the variance during the 1960s until the mid-1970s and from the mid-1980s onwards. The peak in the job finding volatility during the late 1970s-early 1980s displays a decline in the contribution of the neutral technology shock to around 40%. The contribution of the investment-specific technology shock is also fairly constant, at approximately 10% throughout the sample period.

The volatility of the job destruction rate increases substantially during the early 1980s and rapidly decreases afterwards to return to its 1960s' level. Interestingly, the unconditional variance of the job destruction rate displays a pattern that mimics the unconditional volatility of the job finding rate for the post-1980s period, but it differs in the pre-1970s and early 2000s periods. In the pre-1970s the unconditional volatility of the job destruction rate increases, whereas it remains largely unchanged for the job finding rate. In the early 2000s

the unconditional volatility of the job destruction rate increases, whereas it increases for the job finding rate. Finally, the contribution of neutral and investment-specific technology shocks is similar throughout the sample period.

The estimated volatility of GDP growth is high during the 1970s and the first half of the 1980s and shows a sharp decline around 1984, which corroborates the findings in Stock and Watson (2003), who detect a significant fall in the volatility of output around the early 1980s. The volatility of GDP growth is substantially smaller throughout the sample period if the variance of the neutral technology shock is set to zero. In fact, this shock explains about 60% of the variance of GDP growth. The contribution of the investment-specific technology shock is smaller, ranging at around 10% throughout the sample period.

Finally, the volatility of the unemployment rate increases until the mid-1980s and then shows a sharp decline. The neutral technology shock contributes 50% to this variance over most of the sample period except during the early 1980s when this contribution falls to 40%. The contribution of the investment-specific technology shock is fairly constant at 10% throughout the sample period.

Overall, these results show that the volatility of labour market aggregates remarkably declines after the early 1980s, and then stabilizes at a level similar to that of the 1960s. This observed pattern is consistent with the numerous studies that investigate the changes in the volatility of output and other macroeconomic aggregates, as in Stock and Watson (2003). However, we find that movements in the volatility of job finding and job destruction rates display different patterns over sub-periods. Also, the time-varying volatility of the job finding rate is high, roughly twice the volatility of GDP growth. Moreover, the volatility of the job finding rate is higher than the volatility of the job destruction rate throughout the sample period, which supports the evidence based on summary statistics in Shimer (2012).

Figure 3 shows the time-varying correlations among the endogenous variables as implied by the TVP-VAR model (solid line), the 68% error band (shaded area), and the correlations conditional on setting the variance of either the neutral technology shock (dashed line), or the investment-specific technology shock (dotted line) to zero. We calculate the unconditional correlation matrix at each point in time by using the unconditional covariance matrix of the variables $\text{vec}[\text{VAR}(Z_t)] = (I - \bar{\phi}_{1,t} \otimes \bar{\phi}_{1,t})^{-1} \text{vec}(\Omega_t)$. There are large changes in the correlation between the job finding rate and the job destruction rate across time. In particular, the correlation is significantly negative until the mid-1980s, and it then decreases and remains insignificantly different from zero after the mid-1990s. Removing the influence of the investment-specific technology shock has limited influence on the evolution of this correlation. Without the neutral technology shock, however, this correlation goes to zero (and becomes positive) quicker in the post-1985 period. In general, a negative correlation between the job finding and destruction rates has been detected by several studies. However, focusing on its variation across time shows that the correlation weakens over time, and virtually vanishes from the late 1990s onwards.³ This change in the sign of the correlation is likely linked with changes in labour market institutions during the early-1990s that might have altered the response of job finding and destruction to shocks,

³ Note that the fact that the unconditional and conditional correlation is similar across investment-specific and neutral technology shocks is not necessarily evidence that the destruction rate is well approximated by an endogenous rate, since for both scenarios the job destruction is endogenous. Such a similarity simply suggests that both margins have a similar response throughout the sample period to investment-specific and neutral technology shocks.

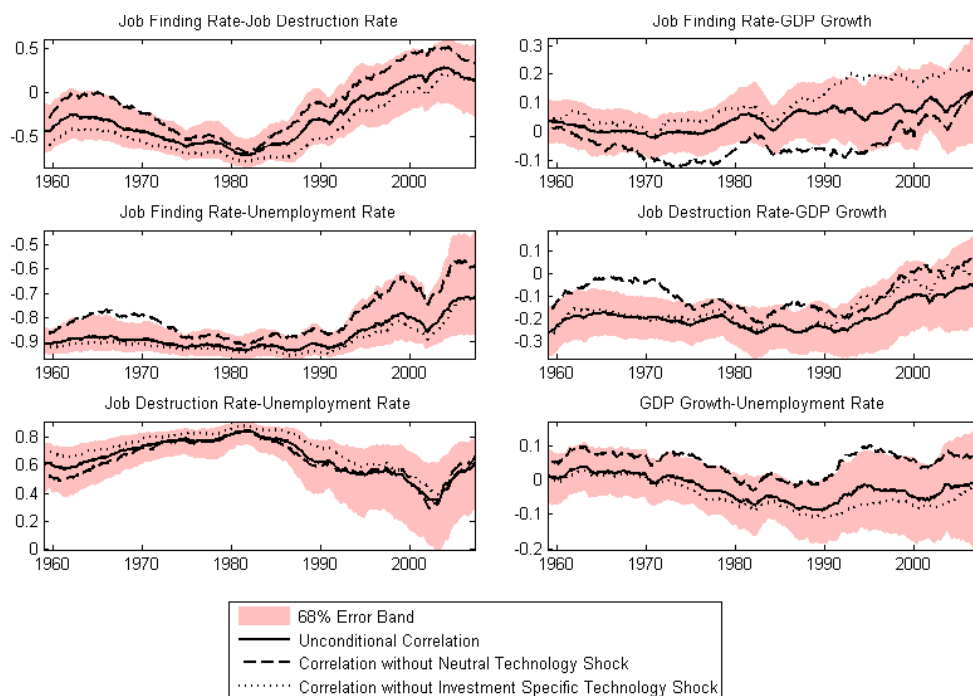


Figure 3. Time-varying correlations

Notes: The figure shows the time-varying correlations among the endogenous variables as implied by the TVP-VAR model (solid line), the 68% error band (shaded area) and the correlations conditional on setting the variance of either the neutral technology shock (dashed line) or the investment-specific technology shock (dotted line) to zero.

as detailed in Zanetti (2011) and Zanetti (2014). Extending the analysis to consider the role of labour market institutions explicitly and their effect of the job finding and destruction rates would certainly be a useful extension for future research.

The correlation between the job finding rate and the unemployment rate is strongly negative throughout the sample period and it weakens after 1990. Without the neutral technology shock, the correlation is smaller in magnitude and declines at a higher rate after 1990. The post-1985 period is also associated with a decrease in the correlation between the job destruction rate and the unemployment rate (from a peak of 0.8 in the mid-1980s to a low of 0.4 in 2003/04). However, the identified technology shock has limited influence on the time-path of this correlation. A comparison of the correlations between the unemployment rate and the job finding and job destruction rates enables us to evaluate which variable is more strongly correlated with the unemployment rate. Interestingly, the job finding rate displays a sample average correlation with the unemployment rate around -0.85% , whereas the job destruction rate has a correlation with the unemployment rate around 0.65% throughout the sample period. This supports the view that although the job finding and job destruction rates jointly determine unemployment, the flow out of unemployment (i.e. the job finding rate) plays a relatively more important role for the dynamics of unemployment. However, the correlations of both the job finding and job destruction rates with the unemployment rate decline substantially from 2000 onwards.

The correlation between the job finding rate and GDP growth is estimated to be insignificantly different from zero over the sample period with the neutral technology shock playing an important part. In contrast, the correlation between job destruction and GDP growth is negative, and it becomes insignificantly different from zero from the early 2000s onwards.⁴ Finally, the correlation between GDP growth and the unemployment rate is close to zero across the sample period.

Overall, it is worth noting that the correlation of the variables with GDP growth is relatively stable over time, whereas the correlation between the job finding rate, the job destruction rate and the unemployment rate shows significant variation across time. This evidence of changes in the joint dynamics of labour market variables, underlines that the statistical relation between the job finding and destruction rates with the unemployment rate is time-varying.

Impulse response functions

Figures 4 and 5 plot the cumulated impulse response to neutral and investment-specific technology shocks. Both shocks are normalized to increase GDP growth by 1% on impact at each point in the sample period. This normalization allows us to focus on possible changes in the responses of the labour market variables.⁵ The left panels in each figure present the median impulse response at each point in time. The *X*-axis in these panels represents the time periods while the *Y*-axis is the impulse response horizon. The remaining panels show the average cumulated impulse response in each decade of the sample for the TVP-VAR (solid line) together with the response from a standard fixed coefficient VAR (dashed line) estimated over the entire sample period. Before considering the variables' response to each shock, looking across the different impulse response functions provides a few useful insights. First, allowing for time variation in the VAR is important to describe the responses of variables to shocks, as most of the variables' responses from the fixed coefficient VAR significantly differ from those of the TVP-VAR. For instance, the reaction of the job destruction rate from the fixed coefficient VAR is systematically higher compared to its TVP-VAR counterpart. Second, neutral and investment-specific technology shocks increase GDP, and they have an opposite effect on labour market variables. The neutral technology shock decreases the unemployment rate by increasing the job finding rate and decreasing the job destruction rate, whereas the investment-specific technology shock leaves the unemployment rate substantially unchanged. As discussed, Michelacci and Lopez-Salido (2007) have also analysed the effect of technology shocks using a SVAR model and find that the neutral technology shock increases job destruction thereby increasing unemployment, whereas the impact of the investment-specific technology shock on unemployment is contractionary. Our analysis has a number of important differences from this study. First, we use a different identification strategy. While these authors use

⁴Note that the unconditional correlation and the correlation without investment-specific technology shocks are around zero during the sample period, with a wide error band surrounding the estimates. Instead the correlation without neutral technology shock is positive and increases slightly from the 1980s onwards, in line with the estimates of the investment-specific technological process in Fisher (2006).

⁵Note that the time-varying nature of the VAR means that we take low-frequency movements in the data into account when estimating impulse responses.

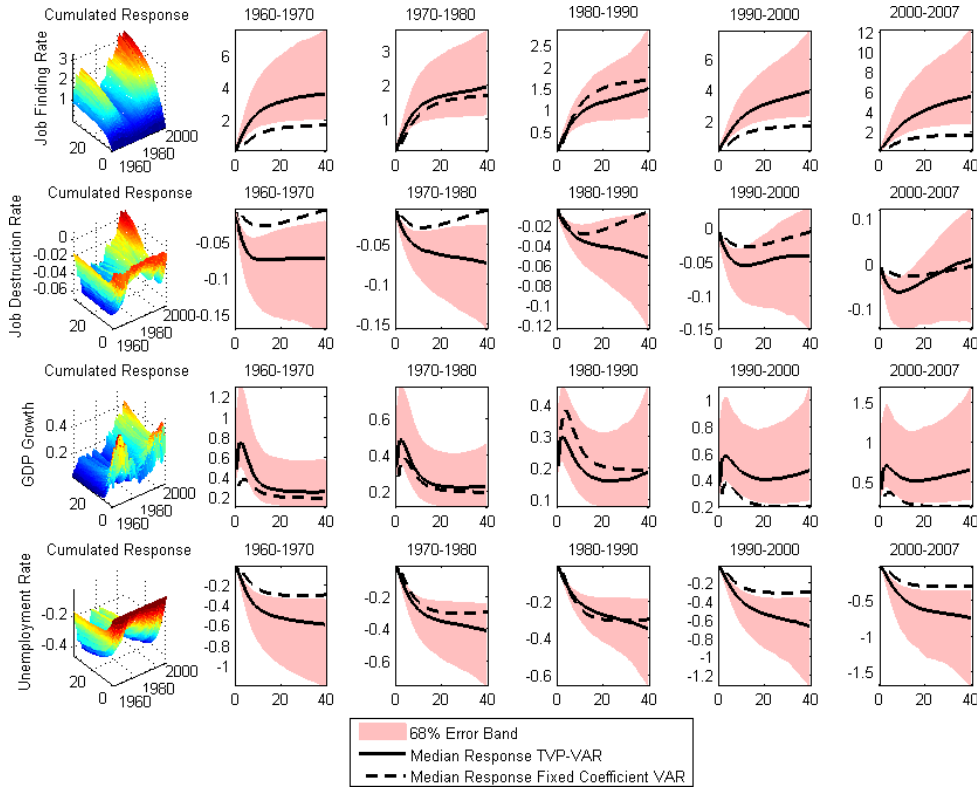


Figure 4. Empirical impulse-response functions to a neutral technology shock

Notes: Cumulated impulse responses to a neutral technology shock. The shock is normalized to decrease the unemployment rate by 1% on impact at each point in the sample period. The left panels present the median impulse response at each point in time. The X -axis in these panels represents the time periods, while the Y -axis is the impulse response horizon. The remaining panels show the average cumulated impulse response in each decade of the sample for the TVP-VAR (solid line), its 68% error band (shaded area), together with the response from a standard fixed coefficient VAR (dashed line) estimated over the entire sample period.

long-run restrictions to identify technology shocks, we use high-frequency restrictions. Our approach has a number of advantages, as detailed in the outset. Second, we allow for time variation in the analysis. Third, the data series differ. While these authors use quarterly data series for job flows from Davis, Haltiwanger and Schuh (1998) that cover the period 1972–93, we use job flows data from Shimer (2012) that cover the period 1948–2007.

Figure 4 shows the response to the neutral technology shock. The top panel shows that the response of the job finding rate to the shock is significant and positive throughout the sample period, with the cumulated effect ranging from around 2% to 6%. The response of the job finding rate to the neutral technology shock displays significant time variation. The top left panel shows a large increase in the cumulated response of the job finding rate at the end of the 1980s, especially at longer horizons. This can be seen quite clearly from the average impulse response functions with the estimates over the period 1990–2007, which are larger than the previous years. Note that over this period the fixed coefficient VAR tends to underestimate the response of the job finding rate on average. There is limited evidence of time variation in the response of the job destruction rate to the neutral technology shock

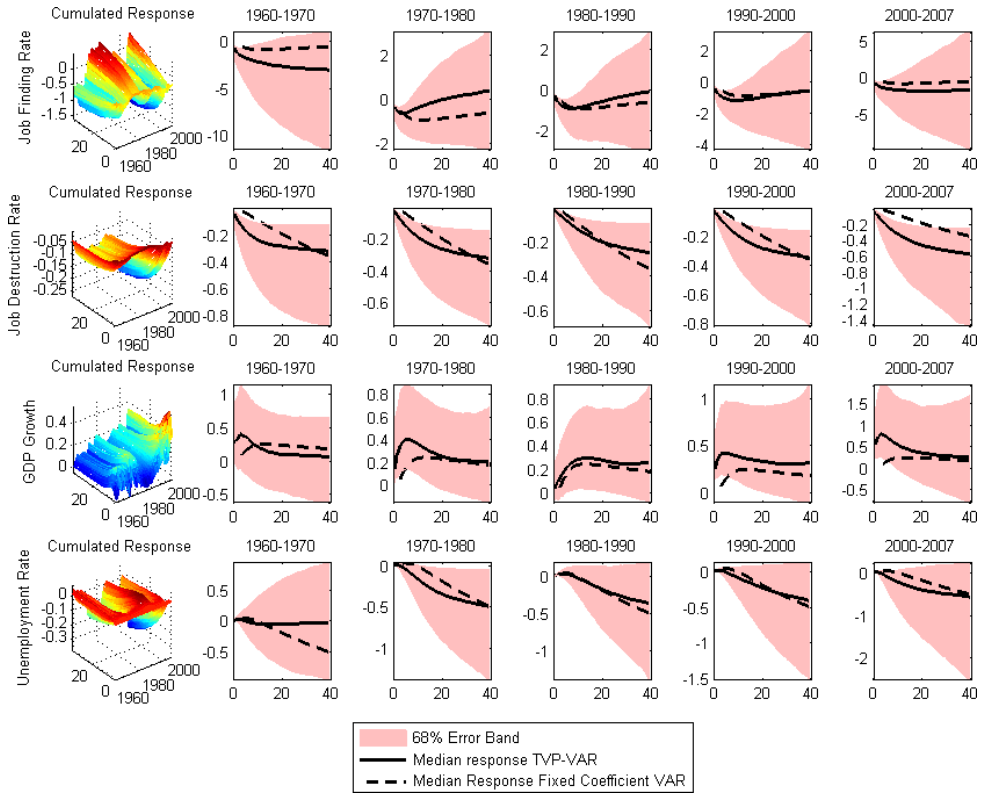


Figure 5. Empirical impulse-response functions to an investment-specific technology shock

Notes: Cumulated impulse responses to an investment-specific technology shock. The shock is normalized to decrease the unemployment rate by 1% on impact at each point in the sample period. The left panels present the median impulse response at each point in time. The X -axis in these panels represents the time periods, while the Y -axis is the impulse response horizon. The remaining panels show the average cumulated impulse response in each decade of the sample for the TVP-VAR (solid line), its 68% error band (shaded area), together with the response from a standard fixed coefficient VAR (dashed line) estimated over the entire sample period.

until the period 1990–2000. Similarly, the response of the unemployment rate to the neutral technology shock is similar across periods until the late 1980s, and increases afterwards.

To investigate further the time-varying relations between job finding and destruction rates and the unemployment rate, the top panel of Figure 6 shows the cumulated responses of these variables to the neutral technology shock after two years. The entries point out that the response of the job destruction rate is remarkably constant from the early 1970s, whereas the responses of the job finding rate and the unemployment rate are stronger after 1980. The response of the job creation rate almost doubles, whereas the response of the unemployment rate decreases by approximately the same amount. This reinforces the results based on time-varying cross correlations between unemployment and job finding and destruction rates, and it is in line with the evidence in Shimer (2012), who, using descriptive statistics, points out that movements in the job finding rate are key drivers of unemployment fluctuations. Interestingly, the figure also shows that the response of output to the neutral technology shock increases from 1980 onwards, in line with the findings in Galí and Gambetti (2009).

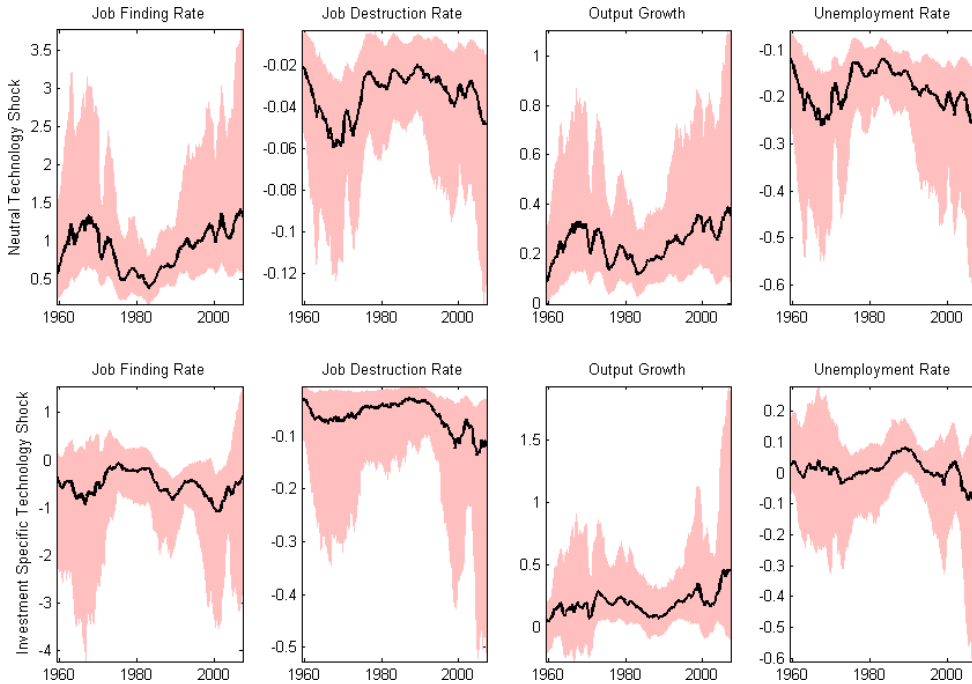


Figure 6. Empirical cumulated impulse-response functions after 2 years

Notes: The top row shows the cumulated responses after 2 years from the TVP-VAR model to a neutral technology shock. The bottom row shows the cumulated responses after 2 years from the TVP-VAR model to an investment-specific technology shock. Each plot shows the 68% error band.

Figure 5 shows the response to the investment-specific technology shock. Entries show that the response of the job destruction rate to the shock displays significant time variation. Over the period 1960–2000, the job destruction rate declines by around 0.3% in response to this shock at the two-year horizon. However, from 2000 onwards the magnitude of the response increases to around 0.5%. The bottom panel of Figure 6 shows the cumulated responses of these variables after two years. The reactions of the variables to an investment-specific technology shock are substantially constant over time, and often around zero, suggesting that changes in labour market dynamics are more likely explained by the neutral technology shock.⁶

Forecast error variance decomposition

To understand the extent to which the movements of each variable are explained by the shocks, Figures 7 and 8 report the forecast error variance decompositions to a neutral and investment-specific technology shock respectively for the TVP-VAR model (solid line),

⁶In order to evaluate the empirical role of the endogeneity of the job destruction rate, we re-estimate a version of the model where no restriction is placed on job destruction rate. Removing this restriction has a large impact on the response to investment-specific shocks, suggesting that the restriction on job destruction is important for identification of the shocks in our model. A companion appendix that documents these results is available on request from the authors.

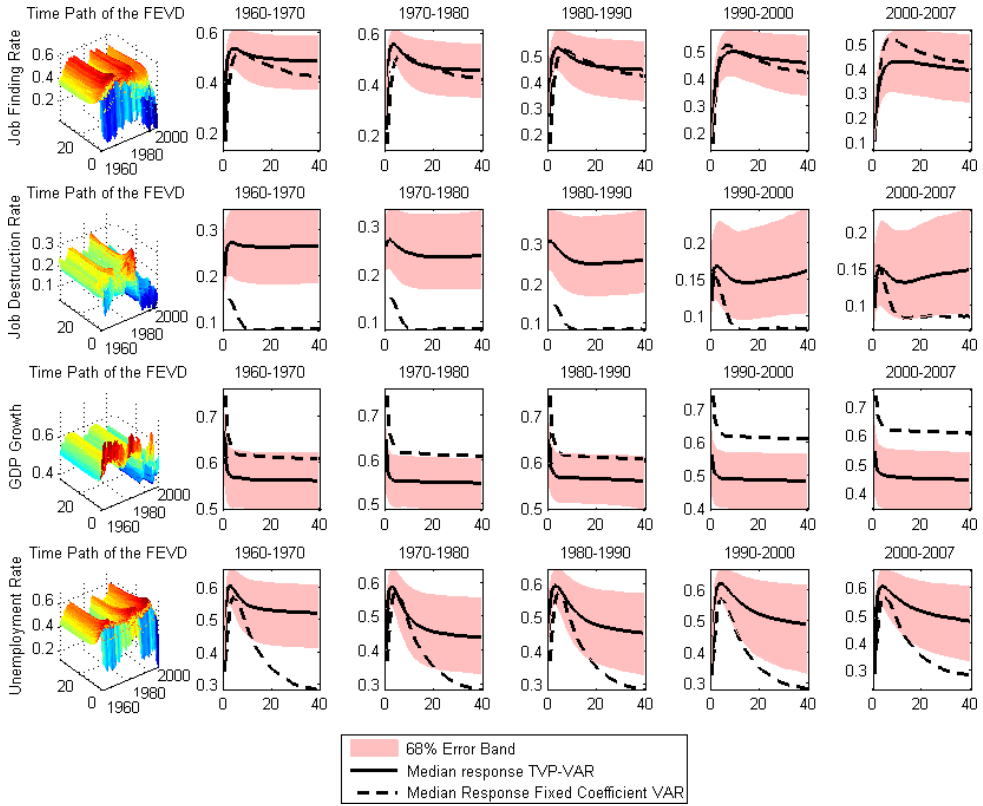


Figure 7. Forecast error variance decompositions, neutral technology shock

Notes: The left panels present the forecast error variance decompositions to a neutral technology shock at each point in time. The *X*-axis in these panels represents the time periods, while the *Y*-axis is the forecast error variance decomposition at different forecast horizons. The remaining panels show the forecast error variance decompositions to a neutral technology shock from the TVP-VAR model (solid line), its 68% error band (shaded area), and the median response from the fixed coefficient VAR (dashed line).

the 68% error band (shaded area), and the median response from the fixed coefficient VAR (dashed line). The left panels present the forecast error variance decompositions to the shock at each point in time. The *X*-axis in these panels represents the time periods, while the *Y*-axis is the forecast error variance decomposition at different forecast horizons. The remaining panels show the forecast error variance decompositions in each decade of the sample.

Looking across the entries shows that the contribution of each shock to movements in the data varies over time. As shown in Figure 7, the contribution of the neutral technology shock to the job finding and job destruction rates decreases over the sample period. In the period 1960–70 neutral technology shocks explain approximately 51% and 28% of movements of these series at low frequencies, whereas their contribution declines to approximately 40% and 15% over the period 2000–07. The contribution of the neutral technology shock to GDP growth displays a similar pattern, since this shock explains approximately 55% of GDP growth over the period 1960–90, but the contribution decreases to approximately 45% over the period 2000–07. Finally, the contribution of the neutral technology shock to

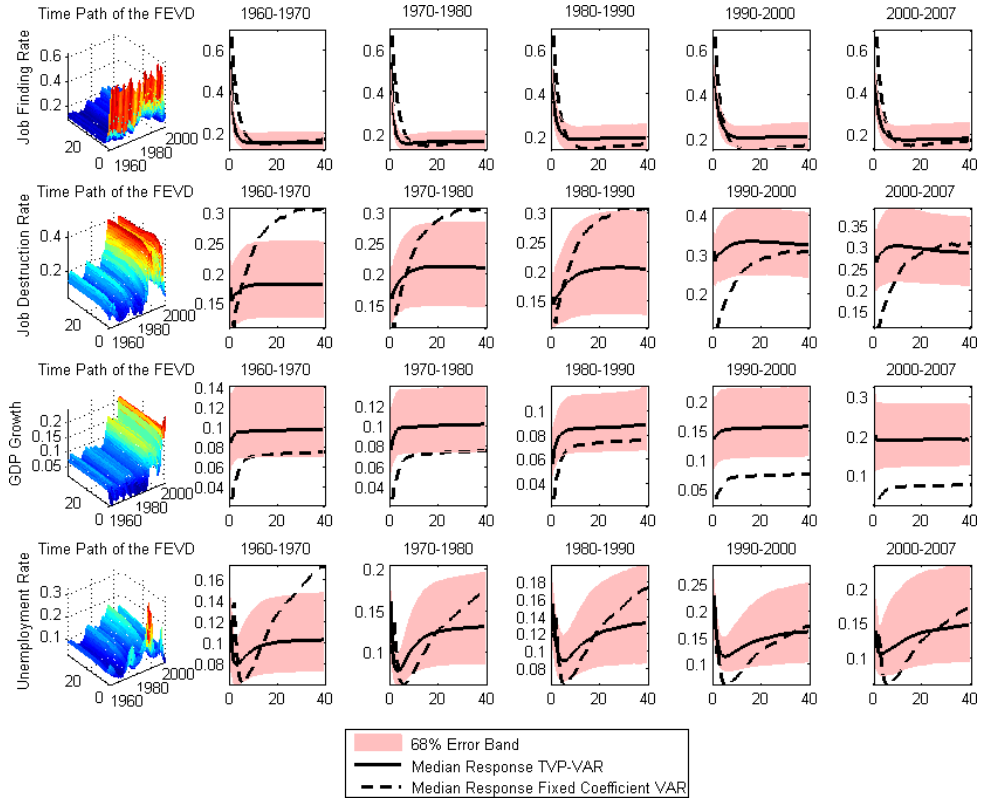


Figure 8. Forecast error variance decompositions, investment-specific technology shock

Notes: The left panels present the forecast error variance decompositions to an investment-specific technology shock at each point in time. The X-axis in these panels represents the time periods, while the Y-axis is the forecast error variance decomposition at different forecast horizons. The remaining panels show the forecast error variance decompositions to a neutral technology shock from the TVP-VAR model (solid line), its 68% error band (shaded area), and the median response from the fixed coefficient VAR (dashed line).

the unemployment rate shows a low degree of time variation, as it is at approximately 52% over the period 1970–90, and at around 47% on average over the subsequent periods.

As shown in Figure 8, the contribution of the investment-specific technology shock explains approximately 18% of the job creation and destruction rates at low frequencies over the period 1960–70, while the contribution almost doubles over the period 2000–07. The investment-specific technology shock explains on average 8% of GDP growth at high frequencies over the sample period 1960–80 and its contribution increases to approximately 18% on average over the period 1990–2007. The investment-specific technology shock contributes to around 18% on average to short-run fluctuation in the unemployment rate, although its contribution declines to approximately 14% at low frequencies. On average, the investment-specific technology shock contributes < 20% on average of fluctuations in the unemployment rate, job finding and job destruction rates.

In general, neutral and investment-specific technology shocks contribute significantly to explain the variance of the variables, although the explanatory power of the investment-specific technology shock is lower than the neutral technology shock, which

corroborates the findings in Zanetti (2008) obtained by estimating standard real business cycle.⁷ Both neutral and investment-specific technology shocks contribute to explain around 55% on average of unemployment fluctuations at low frequencies, in line with Fisher (2006) and Christiano, Eichenbaum and Vigfusson (2004). Moreover, although both neutral and investment-specific shocks explain the bulk of the variables' fluctuations, they are unable to explain the whole variance of the variables, therefore indicating that other shocks are important to describe the dynamics in the data. For both shocks, the forecast error variance decompositions are always statistically significant albeit a sizeable degree of uncertainty surrounds the estimates.

V. Conclusion

This paper has investigated how the statistical relations among key labour market aggregates and their responses to technology shocks have changed over the US postwar period. The analysis is conducted using an estimated VAR that allows for time-varying coefficients and stochastic volatility, whose variables' reactions to shocks are identified using a novel identification scheme based on the cyclical properties of a DSGE model of the business cycle characterized by labour market frictions. The results clearly point out that the dynamics of labour market variables significantly change over time, suggesting that allowing for time variation is an important dimension for a comprehensive assessment of labour market dynamics. We establish several results. The volatility of job finding and job destruction rates has declined over time, after reaching a peak around the mid-1980s. However, changes in the volatility of the job finding and job destruction rates display different patterns over sub-periods. The correlation of the labour market variables with GDP growth is relatively stable over time, whereas the correlation between the job finding and job destruction rates and the unemployment rate shows significant time variation. The job finding rate plays an important role for the unemployment rate, and the two series are strongly negatively correlated over the sample period. The response of labour market variables to neutral technology shocks varies over time. Moreover, neutral technology shocks decrease unemployment, whereas investment-specific technology shocks have a limited impact on unemployment. Finally, across the sample period, neutral technology shocks explain approximately half of the movements in the unemployment rate and the job finding rate, whereas they explain only 20% of fluctuations in the job destruction rate. Investment-specific technology shocks contribute to below 20% on average of fluctuations in unemployment, job finding and job destruction rates.

The analysis of this paper suggests that most of the variation in the unemployment rate over time is explained by job creation. It also points out that changes in labour market dynamics involve deep variation in the functioning of the labour market, suggesting that structural changes play a role to account for the time-varying dynamics of labour market

⁷ Note that our results are different from those of Justiniano and Primiceri (2008) who find that investment-specific technology shocks account for the bulk of the decline in the variability of output and hours at business cycle frequencies. Our approach and information set differ from those of this study since we use a time-varying VAR model that accounts for key labour market variables, therefore including data for unemployment, job creation and job separation probabilities. The abovementioned study instead uses a general equilibrium model that does not account for these labour market series. In addition, our model abstracts from nominal rigidities.

variables. Extending the model to investigate what structural changes could account for the observed variations in labour market dynamics remains an outstanding task for future research. This would prove to be a difficult task however, because it requires the estimation of fully fledged DSGE models with time-varying parameters, as recently outlined in Fernández-Villaverde and Rubio-Ramírez (2008).

Finally, the empirical model could be extended to allow for additional variables, so as to investigate the time-varying statistical relations of labour market variables with a broader set of macroeconomic aggregates. This extension is also open for future research.

Final Manuscript Received: January 2015

References

- Benati, L. Surico, P. (2008). 'Evolving US monetary policy and the decline of inflation predictability', *Journal of the European Economic Association*, Vol. 6, pp. 634–646.
- Berg, A. Meyer, R. Yu, J. (2004). 'Deviance information criterion for comparing stochastic volatility models', *Journal of Business & Economic Statistics*, Vol. 22, pp. 107–120.
- Canova, F. Paustian, M. (2011). 'Business cycle measurement with some theory', *Journal of Monetary Economics*, Vol. 58, pp. 345–361.
- Carter, C. Kohn, P. (2004). 'On Gibbs sampling for state space models', *Biometrika*, Vol. 81, pp. 541–553.
- Chari, V. Kehoe, P. J. McGrattan, E. R., (2008). 'Are structural VARs with long-run restrictions useful in developing business cycle theory?', *Journal of Monetary Economics*, Vol. 55, pp. 1337–1352.
- Christiano, L. J. Eichenbaum, M. Vigfusson, R. (2004). 'The response of hours to a technology shock: evidence based on direct measures of technology', *Journal of the European Economic Association*, Vol. 2, pp. 381–395.
- Cogley, T. Sargent, T. J. (2005). 'Drifts and volatilities: monetary policies and outcomes in the post WWII U.S.', *Review of Economic Dynamics*, Vol. 8, pp. 262–302.
- Cogley, T. Primiceri, G. E. Sargent, T. J. (2010). 'Inflation-gap persistence in the US.', *American Economic Journal: Macroeconomics*, 2, pp. 43–69.
- Davis, S. J. Haltiwanger, J. C. Schuh, S. (1998). *Job Creation and Destruction*, Vol. 1 of *MIT Press Books*, The MIT Press, Cambridge, MA.
- Dedola, L. Neri, S. (2007). 'What does a technology shock do? A VAR analysis with modelbased sign restrictions', *Journal of Monetary Economics*, Vol. 54, pp. 512–549.
- Faust, J. Leeper, E. M. (1997). 'When do long-run identifying restrictions give reliable results?', *Journal of Business & Economic Statistics*, Vol. 15, pp. 345–353.
- Fernald, J. G. (2007). 'Trend breaks, long-run restrictions, and contractionary technology improvements', *Journal of Monetary Economics*, Vol. 54, pp. 2467–2485.
- Fernández-Villaverde, J. Rubio-Ramírez, J. F. (2008). 'How structural are structural parameters?', in *NBER Macroeconomics Annual 2007*, Vol. 22, Cambridge, MA, NBER Books, pp. 83–137.
- Fisher, J. D. M. (2006). 'The dynamic effects of neutral and investment-specific technology shocks', *Journal of Political Economy*, Vol. 114, pp. 413–451.
- Gali, J. Gambetti, L. (2009). 'On the sources of the great moderation', *American Economic Journal: Macroeconomics*, Vol. 1, pp. 26–57.
- Greenwood, J. Hercowitz, Z. Krusell, P. (1997). 'Long-run implications of investment-specific technological change', *American Economic Review*, Vol. 87, pp. 342–362.
- Jacquier, E. Polson, N. Rossi, P. (2004). 'Bayesian analysis of stochastic volatility models', *Journal of Business and Economic Statistics*, Vol. 12, pp. 371–418.
- Justiniano, A. Michelacci, C. (2011). 'The cyclical behavior of equilibrium unemployment and vacancies in the U.S. and Europe', in *NBER International Seminar on Macroeconomics 2011*, Cambridge, MA, University of Chicago Press, pp. 65–114.

- Justiniano, A. Primiceri, G. E. (2008). 'The time-varying volatility of macroeconomic fluctuations', *American Economic Review*, Vol. 98, pp. 604–641.
- Kim, C. J. Nelson, C. R. (1999). *State-Space Models with Regime Switching*, MIT Press, Cambridge, MA.
- Koop, G. Pesaran, H. Potter, S. (1996). 'Impulse response analysis in nonlinear multivariate models', *Journal of Econometrics* Vol. 74, pp. 119–147.
- Mandelman, F. S. Zanetti, F. (2014). 'Flexible prices, labor market frictions and the response of employment to technology shocks', *Labour Economics*, Vol. 26, pp. 94–102.
- Michelacci, C. Lopez-Salido, D. (2007). 'Technology shocks and job flows', *Review of Economic Studies*, Vol. 74, pp. 1195–1227.
- Mortensen, D. T. Pissarides, C. A. (1994). 'Job creation and job destruction in the theory of unemployment', *Review of Economic Studies*, Vol. 61, pp. 397–415.
- Mumtaz, H. Zanetti, F. (2012). 'Neutral technology shocks and the dynamics of labor input: results from an agnostic identification', *International Economic Review* Vol. 53, pp. 237–257.
- Mumtaz, H. Zanetti, F. (2014). *Labor Market Dynamics: A Time-varying Analysis*, Economics Series Working Papers No. 728. University of Oxford, Department of Economics.
- Pappa, E. (2009). 'The effects of fiscal shocks on employment and the real wage', *International Economic Review*, Vol. 50, pp. 217–244.
- Paustian, M. (2007). 'Assessing sign restrictions', *The B.E. Journal of Macroeconomics*, Vol. 7, pp. 1–33.
- Primiceri, G. (2005). 'Time varying structural vector autoregressions and monetary policy', *Review of Economic Studies*, Vol. 72, pp. 821–852.
- Ravn, M. O. Simonelli, S. (2008). 'Labor market dynamics and the business cycle: structural evidence for the united states', *Scandinavian Journal of Economics*, Vol. 109, pp. 743–777.
- Rogerson, R. Shimer, R. (2011). 'Search in macroeconomic models of the labor market', in Ashenfelter O. Card D. (eds), *Handbook of Labor Economics*, chapter 4, North Holland: Elsevier, pp. 619–700.
- Rubio-Ramirez, J. F. Waggoner, D. F. Zha, T. (2010). 'Structural vector autoregressions: theory of identification and algorithms for inference', *Review of Economic Studies*, Vol. 77, pp. 665–696.
- Shimer, R. (2012). 'Reassessing the ins and outs of unemployment', *Review of Economic Dynamics*, Vol. 15, pp. 127–148.
- Stock, J. H. Watson, M. W. (2003). 'Has the business cycle changed and why?', in *NBER Macroeconomics Annual 2002*, Vol. 17, Cambridge, MA: National Bureau of Economic Research, Inc. NBER Chapters, pp. 159–230.
- Thomas, C. Zanetti, F. (2009). 'Labor market reform and price stability An application to the Euro Area', *Journal of Monetary Economics*, Vol. 56, pp. 885–899.
- Uhlig, H. (2004). 'Do technology shocks lead to a fall in total hours worked?', *Journal of the European Economic Association*, Vol. 2, pp. 361–371.
- Zanetti, F. (2008). 'Labor and investment frictions in a real business cycle model', *Journal of Economic Dynamics and Control*, Vol. 32, pp. 3294–3314.
- Zanetti, F. (2011). 'Labor market institutions and aggregate fluctuations in a search and matching model', *European Economic Review*, Vol. 55, pp. 644–658.
- Zanetti, F. (2014). *Labour Market and Monetary Policy Reforms in the UK: A Structural Interpretation of the Implications*, Economics Series Working Papers No. 702. University of Oxford, Department of Economics.