WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 07-12

DYNAMICS OF WORKER FLOWS AND VACANCIES: EVIDENCE FROM THE AGNOSTIC IDENTIFICATION APPROACH

Shigeru Fujita
Federal Reserve Bank of Philadelphia

March 2007

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

Dynamics of Worker Flows and Vacancies: Evidence from the Agnostic Identification Approach*

Shigeru Fujita†
Federal Reserve Bank of Philadelphia

March, 2007

Abstract

This paper establishes *robust* cyclical features of the U.S. labor market by estimating VAR models of the job loss rate, job finding rate, and vacancies. To identify the "aggregate business cycle shock," I adopt the agnostic Bayesian identification approach developed by Uhlig (2005) and others. My approach traces not only responses of transition rates and vacancies but also those of gross job losses and hires and thereby the stock of unemployment in one unified framework. I find that when a negative shock occurs, (i) both the job loss rate and gross job losses rise quickly and remain persistently high, (ii) the job finding rate and vacancies drop in a hump-shaped manner, and (iii) gross hires respond little initially, but eventually *rise*. I argue that these results point to the importance of job loss in understanding U.S. labor market dynamics. The paper also considers the "disaggregate model" which uses data disaggregated by six demographic groups and incorporates transitions into and out of the labor force. I find that job loss continues to play a dominant role among prime-age male workers, while for other groups, changes in the job finding rate are more important.

JEL codes: C32, J63, J64
Keywords: Agnostic identification, hiring, job loss, vacancies, VAR

---

*This paper is a significantly revised version of the previously circulated paper "Vacancy Persistence." I am grateful for useful discussions with Hal Cole, Giorgio Primiceri, Garey Ramey, Frank Schorfheide, and Keith Sill. I also thank Craig Burnside, Wouter den Haan, Chris Foote, Jim Hamilton, and Harald Uhlig for their guidance and/or comments on the previous version of the paper. All remaining errors are mine. The views expressed here are mine and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/econ/wps/index.html.

†E-mail: shigeru.fujita@phil.frb.org.

1 Introduction

The purpose of this paper is to provide *robust* business cycle features of the U.S. labor market. In particular, I pay close attention to the variables useful for evaluating the quantitative abilities of the search/matching models which are widely used in macro/labor economics. Although a number of recent papers serve a similar purpose (e.g. Elsby et al. (2007), Fujita and Ramey (2006*a*, 2007), Fujita et al. (2007), Hall (2005*b*), Shimer (2005*b*), and Yashiv (2006)), all of these papers rely on descriptive measures for their evaluations of the data. This paper instead adopts a structural approach by the use of identified VAR models estimated with the job loss rate, the job finding rate, and vacancies. I make use of the CPS monthly hazard rates recently constructed by Fujita and Ramey (2006*a*) that are corrected for margin error (Abowd and Zellner (1985)) and time-aggregation error (Shimer (2005*b*)).

In seeking robust evidence, I apply the agnostic Bayesian identification approach developed by Uhlig (2005) and others.[1] This approach is very useful for this purpose because it identifies a structural shock by imposing only minimal inequality restrictions on the pattern of impulse response functions and considers all possible responses consistent with those restrictions. Specifically, I identify the "aggregate business cycle shock" by imposing very weak restrictions on the sign of responses of unemployment and vacancies. In particular, I assume that in response to a negative aggregate shock , unemployment rises for at least several months while vacancies drop in the impact month. These restrictions are consistent with a wide range of Mortensen-Pissarides style search/matching models with and without endogenous job loss.[2]

Furthermore, my framework allows me to trace not only the responses of the transition rates and vacancies but also those of gross job losses and hires, and thereby the stock of unemployment. The data generating process of transition rates (together with vacancies) is specified in the form of a VAR. I can then compute implied gross job losses and hires whose difference corresponds to changes in unemployment. Using the stock-flow identity of unemployment evolution makes it further possible to trace the stock of unemployment. One of the nice features of this approach is that I can identify the structural shock by imposing the inequality restrictions at the level of

[1] The other papers include Faust (1998), Canova and De Nicolo (2002), and Rubio-Ramirez et al. (2005).

[2] See, for example, Pissarides (1985), Pissarides (2000, Chapter 1) and Shimer (2005*b*) for dynamic features of search/matching models without endogenous job loss. The dynamic features of the models with endogenous job loss are examined, for example, by Pissarides (2000, Chapter 3), Mortensen and Pissarides (1994), and Fujita (2004).

unemployment behavior without imposing any restrictions on the behavior of transition rates, making it possible to assess the dynamic features of transition rates. Although this paper builds on the important contributions by Blanchard and Diamond (1989, 1990) and Davis and Haltiwanger (1999), who also evaluate worker (or job) flows over the business cycles in a VAR framework with the use of inequality restrictions, neither of these papers provides such a comprehensive and robust assessment.[3]

My approach uncovers the following features of the U.S. labor market: when a recessionary aggregate shock occurs, (i) both the job loss rate and job loss flows rise quickly and remain persistently high in the subsequent periods, (ii) the job finding rate and vacancies drop in a gradual and hump-shaped manner, and (iii) gross hires respond little initially, but eventually *rise* in later periods. Findings (i) and (ii) indicate that fluctuations in both the job loss rate and the job finding rate play an important role in shaping unemployment fluctuations over business cycles. The third finding, however, suggests that job loss is more dominant.[4] To support this argument, I conduct a simple, illustrative counterfactual experiment using the estimated VAR, where I fix either the job loss rate or job finding rate at the steady state level and examine the hypothetical paths of the remaining variables. Consider first the case where the job loss rate follows estimated paths, while the job finding rate is held fixed. In this case, unemployment increases because of more job losses. Because job finding takes place at the same fixed rate, the increases in unemployment result in more hires, which is consistent with the actual paths. On the other hand, when the job loss rate is fixed while letting the job finding rate follow the actual estimated paths, gross hires go *down*, not*up*, because of the direct consequence of slower job finding.[5] Although the experiment is only illustrative, it indicates that ignoring fluctuations in the job loss rate paints a misleading picture about U.S. labor market dynamics.

The extension section estimates a "disaggregate model" where the VAR model is formulated with hazard rates for six demographic groups (instead of aggregate hazard rates) together with the aggregate vacancy series. The analysis in this section incorporates transitions into and out of the labor force. The specification is motivated by the findings by Fujita and Ramey (2006*a*), who emphasize differences in job loss and hiring activity

[3]More recent work in a similar vein includes Braun (2005) and Braun et al. (2006).

[4]Fujita and Ramey (2006*a*) adopt the "job loss driven view" based on their examinations of unconditional second moments of the band-pass filtered data, contrasting it with the "hiring driven view" of Shimer (2005*b*) and Hall (2005*b*,*a*).

[5]The other counterfactual implication of this latter scenario is that combining lower employment and the fixed job loss rate implies declines in job losses in the face of the negative aggregate shock.

across different demographic groups, when transitions into and out of the labor force are taken into consideration. This system allows me to trace both disaggregate- and aggregate-level behavior of gross flows and the stock of unemployment. The aggregate shock is again identified by restricting aggregate-level unemployment and vacancy behavior. I show that the pattern of worker reallocation found in the aggregate model continues to hold among prime-age male workers, thus suggesting an important role of job loss. Among other groups of workers, on the other hand, countercyclicality of job loss becomes unclear while the job finding rate continues to respond procyclically, which indicates dominance of the hiring activity.

This paper proceeds as follows. Section 2 reviews the agnostic identification scheme in a general setting. In Section 3, I apply the method to the labor market data. The inequality restrictions used to identify the aggregate business cycle shock are discussed in reference to existing search/matching models. Section 4 presents the main results. Section 5 extends the aggregate model to a larger, disaggregate model that includes the data for the six demographic groups. Section 6 concludes the paper by discussing the implications of this paper's findings for the development of macroeconomic models with labor market frictions.

2 Review of the Method

Let Y_t be a vector of n endogenous variables containing time-t values whose dynamic relationships are described by the following vector autoregression of order k (VAR(k)):

$$\Phi(L)Y_t = \nu_t, \tag{1}$$

where ν_t is an $n \times 1$ vector containing time-t values of reduced-form disturbances whose variance-covariance matrix is written as $E\nu_t\nu_t' = \Sigma$, and $\Phi(L) = I - \Phi_1 L - \Phi_2 L^2 - \cdots - \Phi_k L^k$. For later reference, I stack the coefficient matices Φ_i into one coefficient matrix as $\Phi \equiv [\Phi_1, ..., \Phi_k]'$.Assuming that $\Phi(L)$ is invertible, VAR(k) has a Wold moving-average representation:

$$Y_t = \Psi(L)\nu_t, \tag{2}$$

where $\Psi(L) = \Phi(L)^{-1} = \Sigma_{j=0}^{\infty}\Psi_j L^j$. Let ω_t be an $n \times 1$ vector containing time-t values of structural disturbances. The reduced-form residuals and structural disturbances are linked through

$$\nu_t = A\omega_t, \tag{3}$$

where it is assumed that the structural disturbances are mutually independent as is standard in the literature. Also, I adopt the normalization that

$E\omega_t\omega_t' = I$. Using Equation (3) in Equation (2) implies that

$$Y_t = \Psi(L)A\omega_t. \tag{4}$$

Thus, Ψ_j can be constructed from Φ_j, which can be estimated by ordinary least squares, and knowledge about A allows one to fully characterize the process of Y_t in terms of the structural disturbances ω_t. The variance-covariance structure of the reduced-form residuals puts constraints on the matrix A:

$$AA' = \Sigma. \tag{5}$$

With an estimate of Σ at hand, the identification problem amounts to uncovering the $\frac{n(n-1)}{2}$ free elements in A by imposing identifying restrictions.

An important result in Uhlig (2005)'s paper is that the matrix A can always be written as

$$A = X\Lambda^{1/2}Q, \tag{6}$$

where X is an orthogonal matrix whose columns are the orthonormal eigenvectors of Σ, Λ denotes a diagonal matrix with the eigenvalues of Σ on its principal diagonal, and Q denotes some orthogonal matrix (i.e., $QQ' = I$). Equation (6) shows that determining the free elements in A can be conveniently transformed into the problem of choosing elements in an orthonormal set. Furthermore, if one is interested only in responses to one particular shock, say, an aggregate shock, then the problem amounts to determining an orthonormal vector q in the following expression:

$$a = X\Lambda^{1/2}q, \tag{7}$$

where a is a column of A (which Uhlig calls an *impulse vector*) containing the contemporaneous responses of n endogenous variables to the structural shock of interest to us, and q is a column of Q in the corresponding location. The main idea of the identification scheme is to impose a set of inequality constraints on $\Psi_j a$. This, of course, does not uniquely identify a but gives us ranges of possible responses consistent with the inequality constraints.

Computations. For each set of the estimates for (Φ, Σ), we can compute impulse vectors and hence impulse response functions corresponding to different unit vectors in an m-dimensional sphere. To uniformly cover the points on the m-dimensional space, I make use of the following algorithm: I generate m numbers from a normal distribution with mean zero and standard deviation one, treat them as coordinates, and normalize the resulting vector into a unit vector. The normalized m-dimensional vector corresponds to each point on the sphere. We can repeatedly generate m-dimensional vectors to uniformly cover the sphere.

I deal with the sampling uncertainty about the VAR parameters (Φ, Σ) in a Bayesian manner. As in Uhlig, I assume that prior and posterior distributions for (Φ, Σ) belong to the Normal-Wishart family. Let $\hat{\Phi}$ and $\hat{\Sigma}$ be the MLE for Φ and Σ, respectively. Under the use of a non-informative prior, the Normal-Wishart posterior distribution is characterized by (i) Σ^{-1} follows a Wishart distribution $\mathcal{W}(\hat{\Sigma}^{-1}/T, T)$ with $E[\Sigma^{-1}]$ where T is the sample size, and (ii) conditional on Σ, the coefficient matrix Φ in its column-wise vectorized form, $vec(\Phi)$ follows a multivariate Normal distribution $\mathcal{N}(vec(\hat{\Phi}), \Sigma \otimes (X'X)^{-1})$ where $X = [Y_1, ..., Y_T]$ with $X_t = [Y'_{t-1}, ...Y'_{t-m}]$. I use the Matlab routines wishrnd and mvnrnd to simulate the Normal-Wishart posterior distribution.

I simulate 1,000 pairs of Σ and Φ. For each pair, I evaluate 1,000 unit vectors on the m-dimensional sphere. Thus a total of 1,000,000 q's and impulse vectors are evaluated. After computing each set of the impulse response functions corresponding to each unit vector, I check if the inequality restrictions are satisfied. I store only the impulse vectors that meet the restrictions.

3 Application to Labor Market Dynamics

This section applies the agnostic identification scheme explained above to U.S. labor market data and argues that the method is particularly suitable for this application. I specify a VAR model to characterize the dynamics of the U.S. labor market with three variables: the job finding rate, the job loss rate, and vacancies. Although I could alternatively use gross hires and job losses instead of the two hazard rates, the literature's growing interest in cyclical behavior of hazard rates motivates me to utilize hazard rates.[6] But, as will be shown later in this section, responses of gross flows, which are implied by the behavior of hazard rates, are also computed. The system also includes the vacancy series as well because vacancies are the key endogenous variable in search/matching models, capturing firms' recruitment efforts.[7]

[6] In fact, Blanchard and Diamond (1990) and Davis et al. (1996) use gross flows in their VAR. More precisely, Davis and Haltiwanger (1999) use job creation and destruction rates, both of which are normalized by the size of aggregate employment.

[7] More specifically, in the search/matching literature, vacancies enter the matching function, which directly generates hiring and affects the job finding rate. The other argument for the aggregate matching function is the number of job seekers, which is equivalent to unemployment under the assumption of no on-the-job search. I trace unemployment behavior internally.

3.1 Data

Hazard rates. I adopt two hazard rate series from my previous paper with Garey Ramey (Fujita and Ramey (2006*a*)). The series are available at monthly frequency and cover the period January 1976 through December 2005. Fujita and Ramey (2006*a*) correct for so-called margin error in the CPS, building on the method developed by Abowd and Zellner (1985). The margin error refers to inconsistency in the stock-flow identities. In the CPS, labor market transition information can be computed for at most 75 percent of all the individuals included in the stock calculations. If the information is missing at random, the missing observations per se should not cause important inconsistencies, given that the sample size is large.[8] However, it is known that the missing individuals, amounting to at least 25 percent of the sample size, create systematic biases in the flow calculations. Fujita and Ramey therefore parameterize true flows as flexible nonlinear functions of the missing-at-random flows and minimize the distance between the stocks implied by the parameterized flows and the official CPS stocks by the use of nonlinear regressions. Fujita and Ramey (2006*a*)'s model nests the missing at random model, and the data strongly reject the latter model.[9]

Fujita and Ramey's series are also corrected for time aggregation error pointed out by Shimer (2005*b*). The error arises due to the fact that the CPS records workers' labor market status at one point (more precisely, during the reference week) in a month and thus misses the within-month spells. However, one can compute continuous-time hazard rates implied by the discrete-time observations under the assumptions that the stock variables evolve in continuous time, and that hazard rates are fixed over each sampling period, namely, a month in the CPS. In particular, by focusing on employment and unemployment (leaving out the not-in-the-labor-force state), one can calculate the U-to-E and E-to-U hazard rates as:

$$\lambda_t = -\log(1 - \hat{\lambda}_t - \hat{p}_t)\frac{\hat{\lambda}_t}{\hat{p}_t + \hat{\lambda}_t}, \tag{8}$$

$$p_t = -\log(1 - \hat{\lambda}_t - \hat{p}_t)\frac{\hat{p}_t}{\hat{p}_t + \hat{\lambda}_t}, \tag{9}$$

where $\hat{\lambda}_t$ and $\hat{p}_t$ are the average job loss rate and average job finding rate, respectively, measured by the CPS's discrete-time observations at time t, and

[8] Under the missing-at-random assumption, one can focus on the individuals whose labor market status is known for two consecutive months. In fact, many previous papers including Bleakley et al. (1999) and Shimer (2005*b*) rely on this assumption to compute gross flows and transition rates.

[9] The rejection of the missing-at-random model may pertain to the so-called rotation group bias in the CPS. See Solon (1986) for this.

where variables without hats are corresponding hazard rates. Given that I focus on transitions between the two states, employment and unemployment, average transition rates are computed as:

$$\hat{\lambda}_t = \frac{eu_t}{e_{t-1}}, \qquad \hat{p}_t = \frac{ue_t}{u_{t-1}},$$

where eu_t and ue_t indicate month-t margin-error-adjusted flows from employment to unemployment and from unemployment to employment, respectively. The denominators e_t and u_t denote the numbers of employment and unemployment, respectively.

Treatment of NILF flows. It is well-known that there are large flows into and out of the not-in-the-labor-force (NILF) state (e.g., Abowd and Zellner (1985), Blanchard and Diamond (1989)). Fujita and Ramey (2006*a*) propose a sensible way of incorporating such flows into the two-state framework above; however, they show that when the NILF flows are incorporated, the aggregate behavior of the hazard rates and flows paints misleading pictures of U.S. labor market dynamics. In particular, they show that breaking down the aggregate data into demographic groups reveals important differences in the cyclical behavior of young workers and prime-age workers. In Section 5 where I estimate the disaggregate model, I incorporate NILF flows into the above framework while first focusing on transitions between employment and unemployment here.[10]

Vacancies. I use the index of help-wanted advertisements released by the Conference Board as an approximation for vacancies. Because this series simply represents the index of the aggregate number of newspaper help-wanted advertisements in 51 major newspapers in the U.S., the approximation may be crude. However, there are several pieces of evidence that this series closely tracks the cyclicality of actual job vacancies in the U.S.[11] The series is available at monthly frequency starting the January 1951, but I use the observations starting in January 1976, because the hazard rate series are available only from then on.

[10] Note that as far as we concentrate on flows between E and U states, there are no large, systematic differences in the pattern of worker flows across different demographic groups.

[11] Abraham (1987) compares the index with actual vacancies in Minnesota where both series are available through two business cycles from 1972 to 1981 and finds that the index closely tracks actual vacancies. More recently, the BLS started a comprehensive survey on job vacancies (Job Openings and Labor Turnover Survey; JOLTS). Shimer (2005*a*) compares the help-wanted index with this series over the recent three-year period after 2000 and finds again that they move closely with each other.

Detrending. The first two panels of Figure 1 plot the seasonally adjusted data for the job finding rate p_t and job loss rate λ_t. While the job finding rate does not seem to have noticeable trending behavior, the job loss rate has been drifting down since the early 1980's.

The last panel shows the vacancy series, which also appear to exhibit low frequency movements. In particular, one can observe a significant downward trend in recent years. This could be due to the shift of recruitment methods toward the Internet.[12]

To remove such low frequency movements from the analysis, I detrend the data by using deterministic polynomial time trends of up to second order. The detrending is necessary given that one of the purposes of this paper is to provide useful guidance in evaluating widely used search/matching models, which do not feature such low frequency trends.[13] The trend components which are plotted in Figure 1 also show that as expected, the job finding rate has a flat trend, whereas the job loss rate and vacancies exhibit trends that initially rise and then gradually decline over time. The trend of the last two variables is well captured by the quadratic trend.

Tracing gross flows and unemployment. Suppose that I have obtained impulse response functions of hazard rates. I can then trace behavior of (i) gross flows, (ii) changes in unemployment, and (iii) the stock of unemployment as follows. Fujita and Ramey (2006*a*) show that hazard rates and gross flows are related by:

$$l_t = \lambda_t \left[-\left(u_{t-1} - \frac{\lambda_t}{\lambda_t + p_t} \right) \frac{1 - e^{-(\lambda_t + p_t)}}{\lambda_t + p_t} + \frac{p_t}{\lambda_t + p_t} \right], \tag{10}$$

$$h_t = p_t \left[\left(u_{t-1} - \frac{\lambda_t}{\lambda_t + p_t} \right) \frac{1 - e^{-(\lambda_t + p_t)}}{\lambda_t + p_t} + \frac{\lambda_t}{\lambda_t + p_t} \right], \tag{11}$$

where l_t and h_t stand for period-t gross job losses and hires, respectively, and $\tilde{u}_{t-1}$ is the unemployment rate in the previous period. Using the responses of λ_t and p_t computed through Equation (4), I can trace gross worker flows conditional on the initial value of $\tilde{u}_t$, which is chosen to be $\frac{\bar{\lambda}}{\bar{\lambda}+\bar{p}}$, where $\bar{\lambda}$ and $\bar{p}$ are historical averages.[14] Note again that λ_t and p_t are hazard rates in continuous time obtained based on the CPS's discrete time observations

[12]Note, however, that there is no a priori reason that such a shift of recruitment methods affects the cyclicality of the data.

[13]To my knowledge, there is no paper in the literature that explicitly models such trend and business cycle components of labor market flows within one framework.

[14]Evolution of the unemployment rate in continuous time is $\dot{u} = \lambda(1-u) + pu$ and the steady state unemployment rate is $\frac{\bar{\lambda}}{\bar{\lambda}+\bar{p}}$.

through Equations (8) and (9), and that l_t and h_t therefore capture all flows that occur over the month under the assumption that hazard rates are constant over the period. See Appendix B of Fujita and Ramey (2006*a*) for details. Fujita and Ramey further note that

$$\tilde{u}_t - \tilde{u}_{t-1} = l_t - h_t, \tag{12}$$

which simply states that net changes in unemployment equal differences in gross flows. This identity allows me to trace the stock of unemployment conditional on its initial value.

I argue that my approach above, which tracks in a unified framework not only responses of hazard rates but also gross flows and thereby the stock of unemployment, provides cleaner and more comprehensive analysis than the previous literature. On the one hand, papers in the recent literature, such as Shimer (2005*b*), Braun et al. (2006), Elsby et al. (2007), focus exclusively on behavior of hazard rates without any reference to gross flows. On the other hand, the earlier literature tends to put more emphasis on gross flows. For example, Blanchard and Diamond (1990) estimate the VAR with CPS worker flows constructed by Abowd and Zellner (1985) for answering questions similar to those in this paper. Davis and Haltiwanger (1999) use their job flow series in their VAR, similar to the one used in this paper. But they do not consider the behavior of transition rates, while this paper builds on the important contributions of these previous papers, none of them provide such a comprehensive and robust assessment of labor market dynamics.

3.2 Identification: Inequality Restrictions

A nice thing about my approach is that I can impose inequality restrictions on the behavior of unemployment rather than that of hazard rates, which are directly used in the estimation. This allows me to examine how hazard rates respond to the shock identified by the behavior of unemployment. Specifically, my benchmark identification relies on the following two restrictions on the behavior of unemployment and vacancies:[15]

Restriction 1: The negative aggregate shock causes changes in unemployment to be non-negative for at least 6 months.

Restriction 2: The negative aggregate shock does not raise vacancies in the impact month.

The first restriction is very intuitive, and it is hard to make the opposite case that the adverse aggregate shock causes unemployment to decrease.

[15] I describe the pattern of responses for the case of the negative aggregate shock. One can identify the positive shock symmetrically.

In fact, these restrictions are consistent with various of specifications of search/matching models. For example, the productivity shock in the simplest possible form of the labor matching model with exogenous job loss (like the one in Pissarides (2000, Chapter 1)) induces the above pattern of impulse responses. DSGE versions of this model also generate the pattern in response to the negative TFP shock as shown by Merz (1995) and Andolfatto (1996).

The productivity shock in the model with endogenous job loss, such as the one of Mortensen and Pissarides (1994), is also consistent with the above pattern under plausible calibrations. Note however that it is a priori unclear that Restriction 2 is satisfied in such a model. That is, vacancies could possibly go up in the face of an adverse aggregate shock in the model with endogenous job loss, even though lower output induced by the negative shock discourages vacancy postings. This possibility arises due to the fact that the adverse shock immediately raises the job loss rate and thereby unemployment, which in turn encourages vacancy postings, because the increased number of job seekers raises the job filling rate for firms. This second channel counters the first output channel, thus making it a priori difficult to qualitatively predict the effect of the negative shock. However, as is shown by Fujita (2004), in the calibrated model of Mortensen and Pissarides (1994) the negative shock in fact causes an initial decline in vacancies, although vacancies quickly bounce back to a level higher than the steady state value.[16]

Blanchard and Diamond (1989), who estimate a VAR with the stock variables including unemployment and vacancies, impose similar inequality restrictions on their identification of the aggregate shock; they assume that the negative aggregate shock causes unemployment to go up and vacancies to go down for at least 9 months. Compared with theirs, my identifying assumptions are weaker. In particular, Restriction 2 constrains vacancy behavior only in the impact period. Although Blanchard and Diamond's restrictions are intuitive, it seems important to be conservative here, especially because the theory makes an ambiguous prediction with respect to vacancy behavior as explained in the previous paragraph.

It should be further noted that there are different kinds of shocks that produce the same pattern as suggested by the two inequality restrictions. For instance, Cooley and Quadrini (1999) show that the technology shock and monetary shock produce qualitatively similar responses in vacancies and unemployment.[17] This implies that my approach is unable to disentangle these different types of shocks. To distinguish different types of shocks requires

[16] Shimer (2005*a*) shows that exogenous shocks to the job loss rate produce a strong positive correlation between unemployment and vacancies, since the output channel is absent in this case. .

[17] See also Trigari (2005).

expanding the empirical model by including more variables in the system. It will then necessarily introduce more controversial issues into identification of the shocks. Relative to this broader analysis, this paper's approach is much simpler and less controversial, which is an important element in my analysis seeking robust features of the data.[18]

Yet another nice feature of my approach is related to the fact that the data on the job finding and job loss rates exhibit significant high frequency variations. One can observe them in the top two panels of Figure 1. In Fujita and Ramey (2006*a*), we utilize the band-pass filter to eliminate those high frequency variations and focus on the business cycle frequency movements of the data. This paper, on the other hand, deals with the same problem by way of explicitly identifying shocks that are associated with the *systematic* business cycle movements of unemployment and vacancies defined by the inequality restrictions above, thereby leaving out such high frequency variations from the analysis.[19]

Tighter identification. In the analysis below, I also present results from the "tighter" case, where I impose the following two additional restrictions on the responses of hazard rates in the impact month of the negative aggregate shock:

Restriction 3: The response of the job loss rate is non-negative.

Restriction 4: The response of the job finding rate is non-positive.

Restriction 3 is consistent with search/matching models with and without endogenous job loss, while Restriction 4 is consistent with the models with endogenous job loss. These restrictions serve to sharpen the predictions from the benchmark case. Note that although imposing both of the restrictions at the same time exclude the model with exogenous job loss, one can still examine persistence and magnification of these variables because they qualitatively restrict the responses only in the first month.

4 Results

Figure 2 displays responses of hazard rates and worker flows. The impulse responses of changes in unemployment, the stock of unemployment, and

[18] I can make a similar argument in comparison with the conventional exact-identification scheme, which typically takes the form of either short-run zero restrictions and/or long-run zero restrictions (e.g., Blanchard and Quah (1989) and Shapiro and Watson (1988)). These zero restrictions require more restrictive assertions about the nature of the shocks.

[19] The sources of such high frequency movements are beyond of the scope of this paper. One possible explanation includes classification error in the CPS. This possibility is, for example, mentioned by Blanchard and Diamond (1990).

vacancies are separately plotted in Figure 3.

The panels in the left column of Figure 2 show that the negative aggregate shock identified by Restrictions 1 and 2 leads to increases in the job loss rate and declines in the job finding rate. More specifically, observe first that the error bands for the two hazard rates initially include zeros, but after the first month, both series significantly deviate from their steady state levels. Further notice that (i) the job loss rate quickly reaches its highest level, whereas the job finding rate follows a familiar hump-shaped pattern, reaching its lowest level after about a year, and (ii) the largest deviations from their steady state levels are of similar magnitude, suggesting that the two margins contribute roughly equally to unemployment fluctuations. These results are highly in line with Fujita and Ramey (2006*a*)'s results, which are based on the unconditional second moments of the band-pass-filtered series.

The right column of Figure 2 presents responses of gross job losses and hires. Not surprisingly, gross job losses behave similarly to the job loss rate because the pool size – the employment rate – is always close to one.[20] The response of gross hires is not distinguishable from zero in the first few months. However, gross hires subsequently *rise* to a level higher than the steady state level. Again, this result is consistent with the finding by Fujita and Ramey (2006*a*), who emphasize countercyclicality of gross hires. Countercyclicality may sound somewhat counterintuitive, given that the job finding rate is procyclical. Note, however, that the number of job seekers (unemployment) rises in the face of the adverse shock, and therefore it is a priori unclear whether the negative shock increases or decreases gross hires. The countercyclicality result thus suggests that the "pool size effect" outweighs the effect from slower job finding.

One can see in the top panel of Figure 3 that changes in unemployment are restricted to be non-negative for the first 6 months. Accordingly, the stock of unemployment, plotted in the middle panel, keeps rising for the same period, generating hump-shaped responses. The last panel shows that the initial response of vacancies is restricted to be non-positive, but responses in the subsequent periods exhibit again the hump-shaped pattern. Recall that I impose the restriction that the unemployment response be hump shaped since the changes in unemployment are constrained to be positive for at least 6 months. However, such restrictions are not imposed on the vacancy response.

[20]Remember that in the current framework, the employment rate equals $1 - \tilde{u}_t$, and the steady state levels of the unemployment and employment rates are set to 5.5% and 94.5%, respectively. As can be seen in the middle panel of Figure 3, the maximum percentage deviation of the unemployment rate from its steady state level is around 3-4 percent when facing a one-standard-deviation aggregate shock. The magnitude is thus much smaller with respect to $1 - \tilde{u}_t$.

Combining the responses of unemployment and vacancies forms the familiar Beveridge curve relationship. As emphasized by Fujita (2004) and Fujita and Ramey (2006*b*), the search/matching model in its standard form is unable to generate such hump-shaped patterns, especially in vacancies.[21] In addition to the negatively correlated hump-shaped responses, observe that vacancies lead unemployment (in Figure 3, note that the trough of vacancies comes before the peak of unemployment). This relationship is also emphasized in the literature (e.g., Pissarides (1985)).

Variance decomposition. Although the inequality restrictions seem plausible in capturing cyclical fluctuations of the data, it is important to make sure that the identified shock accounts for significant portions of variations of the variables of interest. Figure 4 displays three panels showing the portion of the variances explained by the aggregate shock for each horizon. Three lines correspond to the 10th percentile, 50th percentile, and 90th percentile of the posterior distribution.

The results show that the shock initially accounts for relatively small portions of the variations of hazard rates. The explanatory power of the shock gradually increases over the horizon, and the median estimates reach about 50 percent for both variables. The fact that the shock explains only a little of the variations of hazard rates is neither surprising nor problematic for the current analysis. As noted above, the data for hazard rates exhibit important high frequency variations, which appear unrelated to business cycle variations. My approach is supposed to filter out such high frequency variations that exist in short horizons, and the results here in fact suggest that the identifying assumptions successfully remove those variations. Observe in Figure 1 that the vacancy series does not exhibit such large high frequency movements, and hence more than half of the variations in vacancies are accounted for by the identified shock, even in the short run.

Illustrative counterfactual experiment. Recently, there has been a debate in the literature as to which margin – job finding rate or job loss rate – is more important in explaining unemployment fluctuations.[22] The results above have already indicated that *both* margins are important. But the counterfactual exercises in this section shed some more light on which

[21] Fujita and Ramey (2006*b*) show that the model does generate a hump-shaped response in unemployment, but the response is too quick relative to the data, and that this problem also originates from the counterfactual behavior of vacancies in the model.

[22] For example, papers that favor the "hiring driven view" include Shimer (2005*b*) and Hall (2005*b*) whereas papers that emphasize the importance of the job loss side include Fujita and Ramey (2006*a*, 2007), Elsby et al. (2007), Fujita et al. (2007), and Yashiv (2006).

view is quantitatively more plausible. The idea is to fix either λ_t or p_t at the steady state level and examine the hypothetical paths of the remaining variables. Specifically, by plugging a constant path for either one of the two variables and the estimated MA representation of the other into (10) and (11), the hypothetical paths for gross job losses and hires, thereby changes of unemployment and the stock of unemployment under each scenario can be computed. I can then compare the hypothetical paths of these variables with the estimated responses.

First, consider the case where the job finding rate is held fixed at the steady state level. The results are shown in Figure 5. Note that even though the job finding rate is fixed, it is entirely possible that gross hires move considerably. This is because changes in the job loss rate result in changes in the number of job seekers (the stock of unemployment) and therefore gross hires can move even under the fixed job finding rate. Further, the path of gross job losses under this exercise is different from the estimated paths of job losses.[23] However, the right top panel in the figure shows that gross job losses behave very similarly to the estimated paths under the fixed rate of job finding. The second row presents the hypothetical path of the job finding rate (which is constant by construction) and gross hires. Somewhat surprisingly, the behavior of gross hires is very similar to the estimated actual paths. The countercyclicality of gross hires is preserved because of the feedback effect mentioned above: higher job losses raise the stock of unemployment, thereby increasing gross hires as well. The right panel in the last row displays the unemployment response, showing that its variations explained by the fluctuations in the job loss rate alone amount to roughly one-half.[24]

Figure 6 considers what happens if I fix the job loss rate at the steady state level while letting the job finding rate take the estimated actual paths. Note again that it is entirely possible that job losses actually move due to the feedback effect from changes in the pool size ($= 1 - \tilde{u}_t$) driven by changes in the job finding rate. The top right panel shows that this is actually not the case; while job losses do move due to the feedback effect, the effect is quantitatively minuscule. Moreover, when the job loss rate is constant, declines in the job finding rate induce gross hires to *decrease*. This is highly counterfactual. The right-hand side of the last row presents the unemployment response. Comparing this with the actual estimated paths shows that this second case accounts for roughly one-half of total unemployment variations. Recall, however, that variations in unemployment

[23] Note also that p_t and λ_t directly enter the formula for l_t (Equation (10)) and h_t (Equation (11)), respectively. But the effects of p_t on l_t and of λ_t on h_t are quantitatively minimal.

[24] Fujita and Ramey (2006*b*) reach a similar conclusion using a different approach.

in this case are produced by the counterfactual behavior of *both* gross job losses and hires.

Although the exercise here is purely illustrative, it nonetheless sheds some light on the quantitative and qualitative importance of fluctuations in the two hazard rates. In particular, it suggests that ignoring movements of the job loss rate produces highly counterfactual implications on the remaining variables.

Imposing tighter restrictions. The discussion so far entails the case with imposing only two benchmark inequality restrictions on the behavior of unemployment and vacancies. To sharpen predictions from the VAR, I now add Restrictions 3 and 4. Figures 7 and 8 plot the graphs that correspond to Figures 2 and 3 for the benchmark case. It turns out that the two additional restrictions do not change the assessment of the labor market dynamics, although it does narrow the error bands considerably.

5 Extension: Disaggregate Model

5.1 Motivation and the Method

This section extends the above benchmark model to incorporate differences in the cyclicality of worker flows across demographic groups. This extension is motivated by the results reported by Fujita and Ramey (2006*a*), who find that the aggregate job loss rate becomes far less countercyclical when NILF flows are incorporated into the analysis. Behind this is the composition effect that young workers' job loss rate becomes essentially acyclical when we treat the E-to-NILF flow as part of job losses, whereas that of prime-age (25-54) workers, especially prime-age male workers, is strongly countercyclical regardless of the inclusion of the E-to-NILF flow. Motivated by these results, I estimate a VAR using disaggregated data across age and gender with inclusion of NILF flows.[25] I consider three age groups, young (16-24), prime-age (25-54) and old (55 or above), and therefore a total of six demographic groups are included in the analysis.

Incorporating the NILF flows. Before discussing the estimation issues of the VAR, I briefly describe the way Fujita and Ramey (2006*a*) incorporate the NILF flows into their framework. Because I am applying the same procedure for each group, I use subscript i in order to be explicit about

[25] Importantly, when I run the disaggregate model with the same six demographic groups focusing on employment and unemployment transitions (as I did in the previous section), the impulse responses of those demographic groups are similar to each other and to the behavior in the aggregate model in the previous section.

reference to the demographic groups. The definition of the average job loss rate is expanded to include flows from employment into NILF, denoted as en_t below:

$$\hat{\lambda}_{it} = \frac{eu_{it} + en_{it}}{e_{it-1}}.$$

Defining the average job finding rate is less straightforward, since it is difficult to know the number of job seekers that are out of the labor force. To impute the size of the pool, Fujita and Ramey adopt the assumption that workers flowing into the employment relationships from NILF have faced the same average job finding rate as those officially unemployed.[26] Under this assumption, the average job finding rate can be computed as:

$$\hat{p}_{it} = \frac{ue_{it} + ne_{it}}{(1 + \frac{ne_{it}}{ue_{it}})u_{it-1}},$$

where ne_t represents the flow into employment from the NILF pool. Having obtained the two average rates, I simply use the same formulas (8) and (9) to convert them into continuous time hazard rates.

Aggregation and the inequality restrictions. I estimate the VAR model with those hazard rates for 6 demographic groups and vacancies. There are therefore a total of 13 variables (i.e., two hazard rates for each of the six demographic groups plus the aggregate vacancy series). All series are pre-detrended in the same way as in the aggregate model. I do not impose any restrictions on the cross effects among those demographic groups. Lag length is set to 3 months.[27]

Once the paths of the hazard rates are obtained, we can apply the formulas for gross job losses (10) and hires (11) for each demographic group. This procedure gives gross flows for each group i.[28] I aggregate gross flows across six demographic groups by using the fixed average labor force weights computed by the data over the entire sample.[29]

$$l_t = \sum_i w_i l_{it},\ h_t = \sum_i w_i h_{it}, \text{ with } i = 1, ..., 6,$$

[26]This approach is not uncontroversial, yet reasonable. See page 14 in Fujita and Ramey for justification.

[27]In the aggregate model, I chose a lag length of 6 months. I have chosen the shorter lag length mainly to ease computational burden. However, the results are unchanged, even when the lag length is set to 6 months. See footnote 30.

[28]Note that in applying the formulas, we need to expand the definition of the unemployment rate $\tilde{u}_{it}$ by including the imputed job seekers out of the labor force in addition to the officially unemployed.

[29]In reality, the weights are changing over time. An assumption is that movements of the weights are orthogonal to the identified shock. The next question then is whether my results are robust with respect to different sets of weights. In fact, using different sets of weights changes the results little.

where l_{it} and h_{it} are gross job losses and hires, respectively, of group i, and w_i is the associated weight. The change in aggregate unemployment is then computed by taking the difference between job losses and hires at the aggregate level. The *aggregate* shock is then identified by looking at the *aggregate* level behavior of unemployment and vacancies. Specifically, Restrictions 1 and 2 are again used to identify the shock. Further, I also restrict the size of the initial response of vacancies. Specifically, I assume that the negative aggregate shock lowers vacancies by more than 1 percent. This is motivated by the results in the aggregate model. As Figure 3 shows, the 10th percentile of the vacancy response roughly corresponds to a 1 percent decline. Note also that Restrictions 3 and 4 are not relevant here in that we are not interested in the behavior of the aggregate hazard rates.

The disaggregate system is much larger than the previous aggregate model with only three variables and therefore incurs a large computational burden.[30] As in the aggregate model, I simulate 1,000 pairs of Σ and Φ and evaluate 1,000 unit vectors for each of the pairs.

5.2 Results

Figure 9 presents the aggregate level behavior of changes in unemployment, the stock of unemployment, and vacancies. Although the magnitude of deviations from the steady state levels are somewhat smaller in this disaggregate model (compared to corresponding results from the aggregate model), overall patterns of the responses are quite similar: both unemployment and vacancies exhibit hump-shaped responses with vacancies leading unemployment.[31]

Figure 10 plots unemployment responses for four (out of six) demographic groups.[32] The four panels in the figure clearly indicate that the recessionary shock induces gradual positive responses in unemployment. Further, observe that the response of the prime-age males is most noticeable among them.

[30] The size of the unit vector q and of draws from the Normal-Wishart posterior distribution are much larger than in the previous trivariate VAR model. Furthermore, since I need to keep all the responses of the hazard rates and gross flows for all demographic groups that satisfy the inequality restrictions, computation requires a large amount of memory.

[31] Note that "unemployment" here does not correspond to the definition of the official data because the pool of job seekers is expanded to include those who are outside of the labor force but looking for jobs, as described in the previous subsection.

[32] In the following discussion, I present the results for only four demographic groups, (i) young males, (ii) young females, (iii) prime-age males, and (iv) prime-age females. This is because the responses of workers older than 55 are relatively small and not cyclical. The estimation and identification are conducted with all six demographic groups, however.

Each of Figures 11 through 14 displays responses of the four variables – the job loss rate, gross job losses, the job finding rate, and gross hires – for each of the four demographic groups. First, consider responses of prime-age male workers, presented in Figure 13. The pattern strongly conforms to the findings of the aggregate model that focuses on E/U transitions; the job loss hazard rate and gross job losses are higher in the face of the adverse aggregate shock; while the job finding adjusts more slowly, the number of gross hires increases because the pool of job seekers expands as more workers leave employment.

Consider the other three groups of workers (Figures 11, 12, and 14). First, responses of the job loss rates for young male and prime-age female workers are not distinguishable from zero, while that for young-female workers is actually slightly negative. Responses of the job finding rates show procyclicality across the board. Reflecting declines in the job finding rates, gross hires tend to become lower within a year after the shock hits. Similarly, gross job losses are more likely to be lower over the same periods.

The results for these groups of workers indicate that slower job finding during recessionary periods drives the worker reallocation process. To be more specific, consider a hypothetical case similar to one of the two counterfactual scenarios before, where the job finding rate is lower while the job loss rate is simply constant. In this case, as studied in Figure 6, gross hires decline as a direct consequence of slower job finding, eventually causing declines in employment. The constant job loss rate then implies lower gross job losses.

When I estimate the disaggregate model with E/U transitions only, I find that all demographic groups show the same pattern as in the aggregate model. The results in this section therefore suggest that the participation decision plays an important role in understanding the cyclical adjustments among young workers and prime-age female workers. On the other hand, robustness of the results among prime-age male workers with respect to inclusion of NILF flows carries a large weight in thinking about the stylized facts of U.S. labor market dynamics from macroeconomic perspectives, as those workers tend to be in long-term, high-wage jobs. Further discussions on this issue can be found in Section 9 in Fujita and Ramey (2006*a*), who find similar results.

6 Conclusion

This paper has applied the agnostic identification scheme of Uhlig (2005) and others to uncover the robust features of U.S. labor market dynamics. In line with the evidence in papers such as Elsby et al. (2007), Fujita and

Ramey (2006*a*, 2007), Fujita et al. (2007), and Yashiv (2006). I have shown that countercyclicality of job loss is a quite robust feature of the data and is important in accounting for unemployment dynamics. In particular, I have also shown through the illustrative counterfactual experiment that the countercyclicality of gross hires cannot be replicated under the case where the job loss rate is acyclical. The importance of job loss continues to hold among prime-age male workers whether or not transitions into and out of the labor force are incorporated into the analysis.

In addition to the importance of the countercyclicality of the job loss rate, the results of this paper point to the important avenues for future research in macroeconomic models with labor market frictions. As shown throughout the paper, vacancies exhibit hump-shaped responses while leading unemployment. While this may not be surprising from an empirical point of view, the Mortensen-Pissarides style models, especially coupled with endogenous job loss, fail to replicate this feature, and instead introduce positive comovements between unemployment and vacancies.[33] In Fujita and Ramey (2006*b*), we extend the standard seach/matching model (with a fixed job loss rate) by introducing sunk job creation costs which are incurred when new jobs are created. This extension makes vacancies a predetermined variable (instead of a jump variable as in a standard model), generating highly realistic dynamics in vacancies and unemployment. Furthermore, it also introduces an important distinction between job and worker turnover because jobs exist until they become obsolete while worker-firm relationships can be severed independently of job obsolescence. The model in Fujita (2003) features a similar costly job creation process together with endogenous job loss. This model generates both gradual negative relationships between vacancies and unemployment and countercyclical job loss at the same time. There seem to be possibilities for further economically interesting work along these lines.

[33] See Mortensen and Pissarides (1994) and Fujita (2004). See footnote 16 as well.

References

Abowd, J. and Zellner, A. (1985), 'Estimating gross labor force flows', *Journal of Business and Economic Statistics* **3**, 254–83.

Abraham, K. (1987), 'Help-wanted advertising, job vacancies, and unemployment', *Brookings Papers on Economic Activity* **82**(1), 207–43.

Andolfatto, D. (1996), 'Business cycles and labor-market search', *American Economic Review* **86**(1), 112–132.

Blanchard, O. and Diamond, P. (1989), 'The Beveridge curve', *Brookings Papers on Economic Activity* **1**, 1–60.

Blanchard, O. and Diamond, P. (1990), 'The cyclical behavior of the gross flows of U.S. workers', *Brookings Papers on Economic Activity* **2**, 85–155.

Blanchard, O. and Quah, D. (1989), 'The dynamic effects of aggregate demand and supply disturbances', *American Economic Review* **79**(4), 655–673.

Bleakley, H., Ferris, A. and Fuhrer, J. (1999), 'New data on worker flows during business cycles', *New England Economic Review* **July/August**, 49–76.

Braun, H. (2005), 'Unemployment dymamics: The case of monetary polocy shocks'. Draft, Northwestern University.

Braun, H., DeBock, R. and DiCecio, R. (2006), 'Aggregate shocks and labor market fluctuations', *Federal Reserve Bank of St. Louis Working Paper 2006-004A* .

Canova, F. and De Nicolo, G. (2002), 'Monetary disturbances matter for business fluctuations in the G-7', *Journal of Monetary Economics* **49**(6), 1131–1159.

Cooley, T. and Quadrini, V. (1999), 'A neoclassical model of the phillips curve relation', *Journal of Monetary Economics* **44**, 165–193.

Davis, S. and Haltiwanger, J. (1999), 'On the driving forces behind cyclical movements in employment and job reallocation', *American Economic Review* **89**(5), 1234–1258.

Davis, S., Haltiwanger, J. and Schuh, S. (1996), *Job Creation and Destruction*, MIT Press.

Elsby, M., Solon, G. and Michaels, R. (2007), 'The ins and outs of cyclical unemployment', *NBER Working Paper No.12853* .

Faust, J. (1998), 'The robustness of identified VAR conclusions about money', *Carnegie-Rochester Conference Series in Public Policy* **49**, 207–244.

Fujita, S. (2003), 'The Beveridge curve, job creation and the propagation of shocks'. Draft, UC San Diego.

Fujita, S. (2004), 'Vacancy persistence', *Federal Reserve Bank of Philadelphia Working Paper* **04**(23).

Fujita, S., Nekarda, C. and Ramey, G. (2007), 'The cyclicality of worker flows: New evidence from the SIPP', *Federal Reserve Bank of Philadelphia Working Paper 07-05* .

Fujita, S. and Ramey, G. (2006*a*), 'The cyclicality of job loss and hiring', *Federal Reserve Bank of Philadelphia Working Paper 06-17, UC San Diego Discussion Paper 06-08* .

Fujita, S. and Ramey, G. (2006*b*), 'Job matching and propagation', *Journal of Economic Dynamics and Control* **forthcoming**.

Fujita, S. and Ramey, G. (2007), 'Reassessing the Shimer facts', *Federal Reserve Bank of Philadelphia Working Paper 07-02* .

Hall, R. (2005*a*), 'Employment efficiency and sticky wages: Evidence from flows in the labor market', *Review of Economics and Statistics* **87**(3), 397–407.

Hall, R. (2005*b*), 'Job loss, job finding, and unemployment in the U.S. economy over the past fifty years', pp. 275–310.

Merz, M. (1995), 'Search in the labor market and the real business cycle', *Journal of Monetary Economics* **36**, 269–300.

Mortensen, D. and Pissarides, C. (1994), 'Job creation and job destruction in the theory of unemployment', *Review of Economic Studies* **61**, 397–415.

Pissarides, C. (1985), 'Short-run equilibrium dynamics of unemployment, vacancies and real wages', *American Economic Review* **75**(4), 676–690.

Pissarides, C. (2000), *Equilibrium Unemployment Theory*, second edn, MIT Press.

Rubio-Ramirez, J., Waggoner, D. and Zha, T. (2005), 'Markov-switching structural vector regressions: Theory and application'. Draft, Federal Reserve Bank of Atlanta.

Shapiro, M. and Watson, M. (1988), 'Sources of business cycle fluctuations', *NBER Macroeconomics Annual* pp. 111–47. Olivier J. Blanchard and Stanley Fischer, eds.

Shimer, R. (2005*a*), 'The cyclical behavior of equilibrium unemployment and vacancies', *American Economic Review* **95**(1), 25–49.

Shimer, R. (2005*b*), 'Reassessing the ins and outs of unemployment'. Draft, University of Chicago.

Solon, G. (1986), 'Effects of rotation group bias on estimation of unemployment', *Journal of Business and Economic Statistics* **4**(1), 105–109.

Trigari, A. (2005), 'Equilibrium unemployment, job flows and inflation dynamics', *ECB Working Paper No. 304* .

Uhlig, H. (2005), 'What are the effects of monetary policy on output? Results from an agnostic identification procedure', *Journal of Monetary Economis* **52**(2), 381–419.

Yashiv, E. (2006), 'U.S. labor market dynamics revisited', *IZA dicussion paper 2445* .

Figure 1: Data

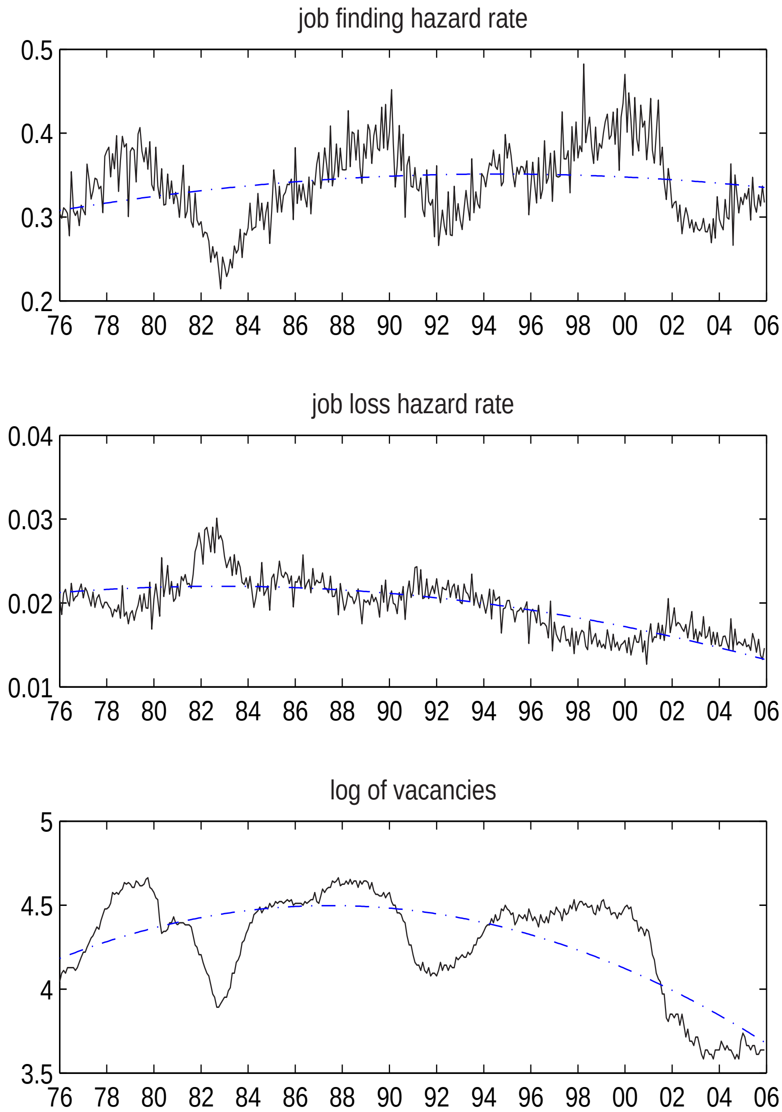

Notes: Hazard rate series are taken from Fujita and Ramey (2006*a*), and seasonally adjusted by Census X-12. The index of help-wanted advertisements represents the vacancy series. The seasonally adjusted series is released by the Conference Board. Trends are identified by regressing on time polynomials of up to second order.

Figure 2: Impulse response functions for hazard rates and worker flows: Restrictions 1 and 2

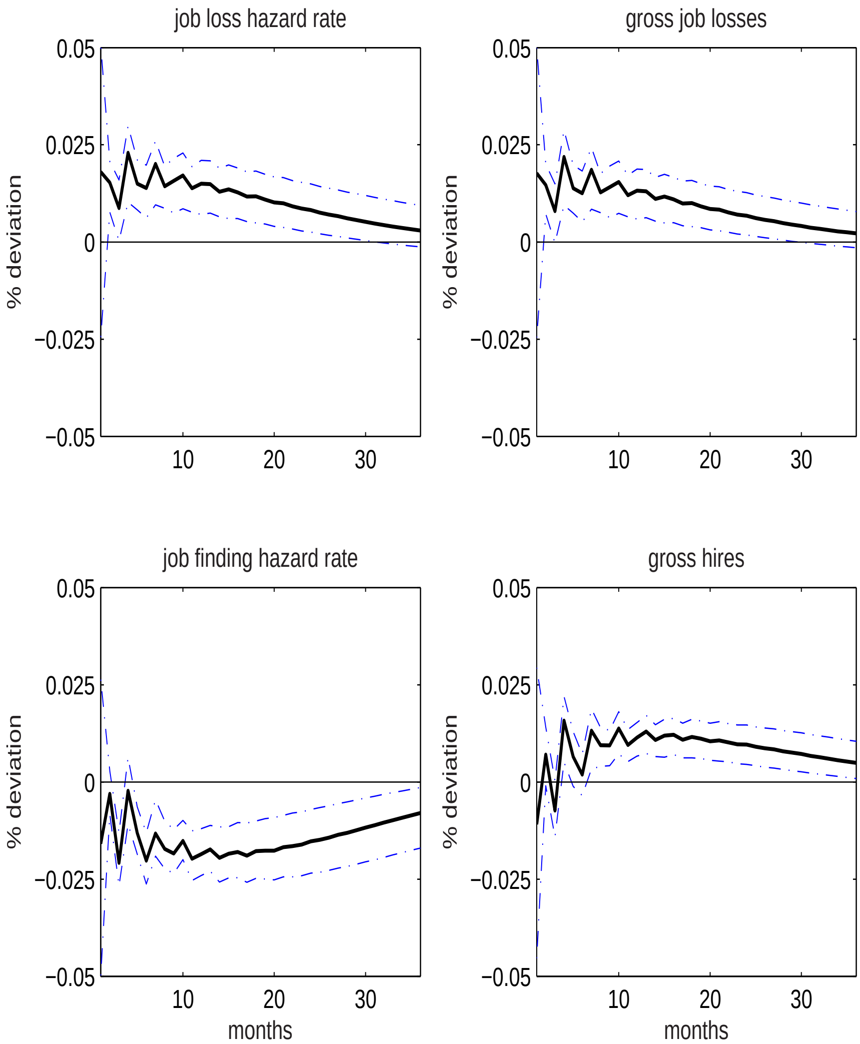

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 3: Impulse response functions for the change in unemployment, stock of unemployment and vacancies: Restrictions 1 and 2

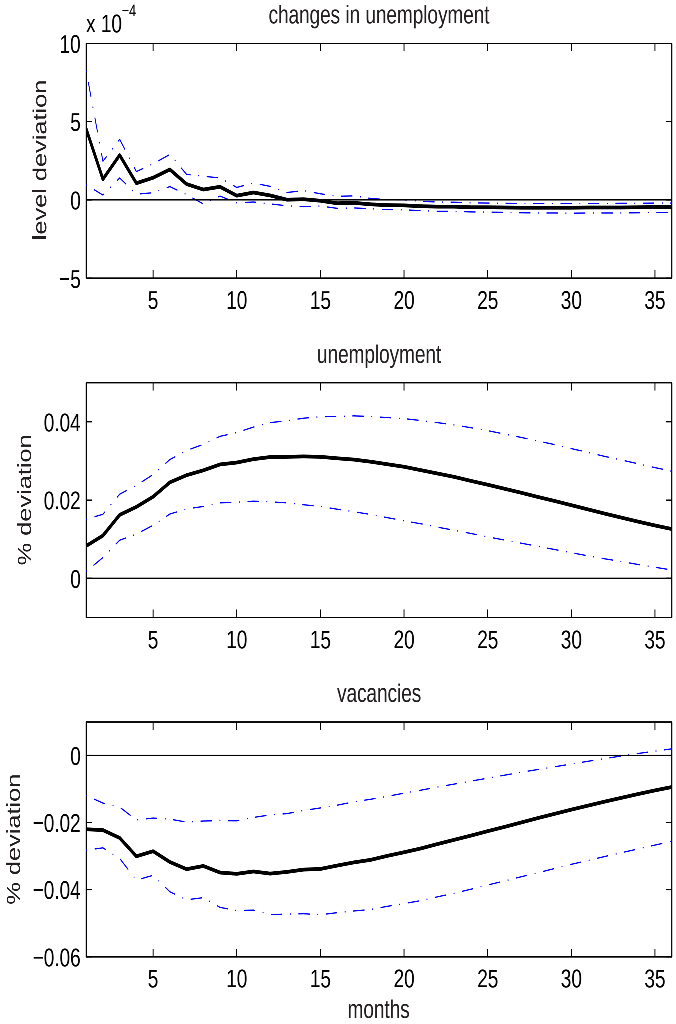

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses of the stock of unemployment and vacancies are expressed as percentage deviations from the steady state levels.

Figure 4: Variance decomposition: Restrictions 1 and 2

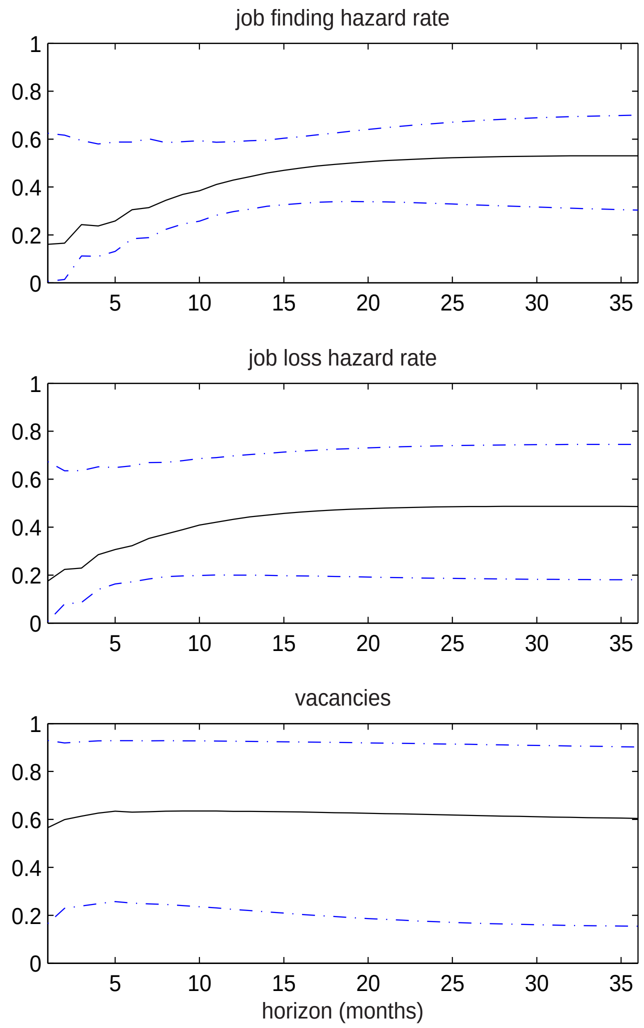

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution.

Figure 5: Counterfactual experiment: fixed job finding rate and variable job loss rate

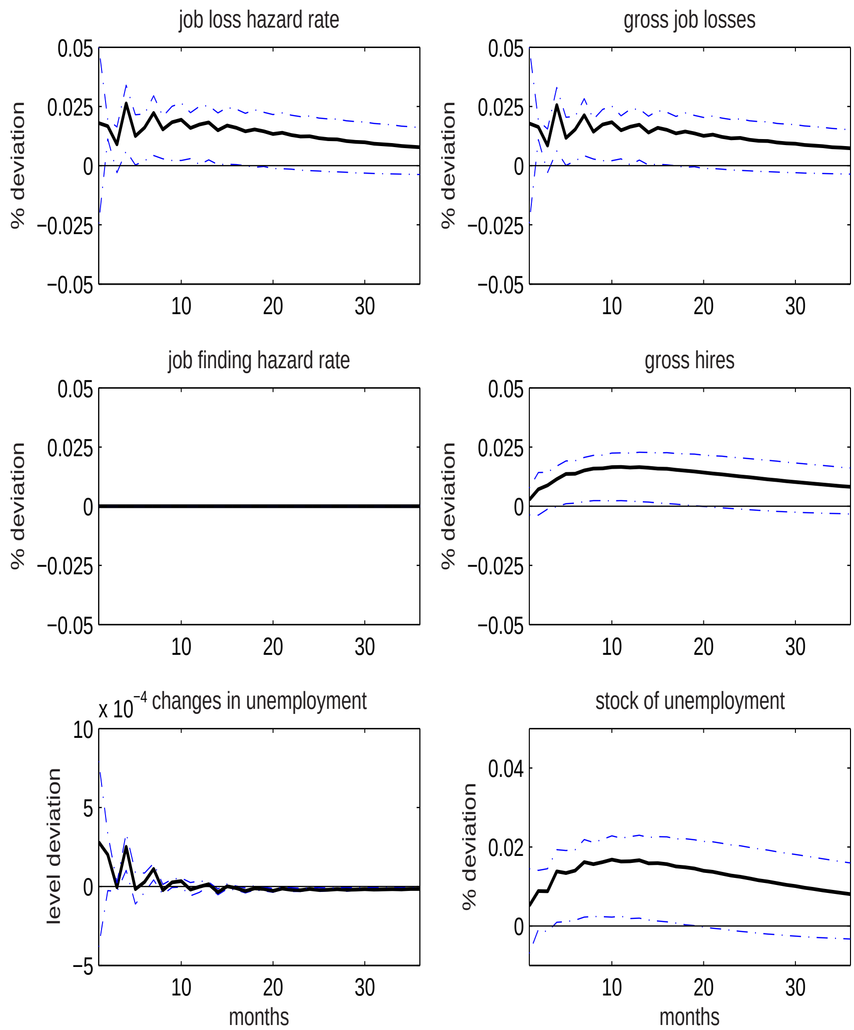

Notes: The shock is identified by imposing Restrictions 1 and 2.The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution.

Figure 6: Counterfactual experiment: fixed job joss rate and variable job finding rate

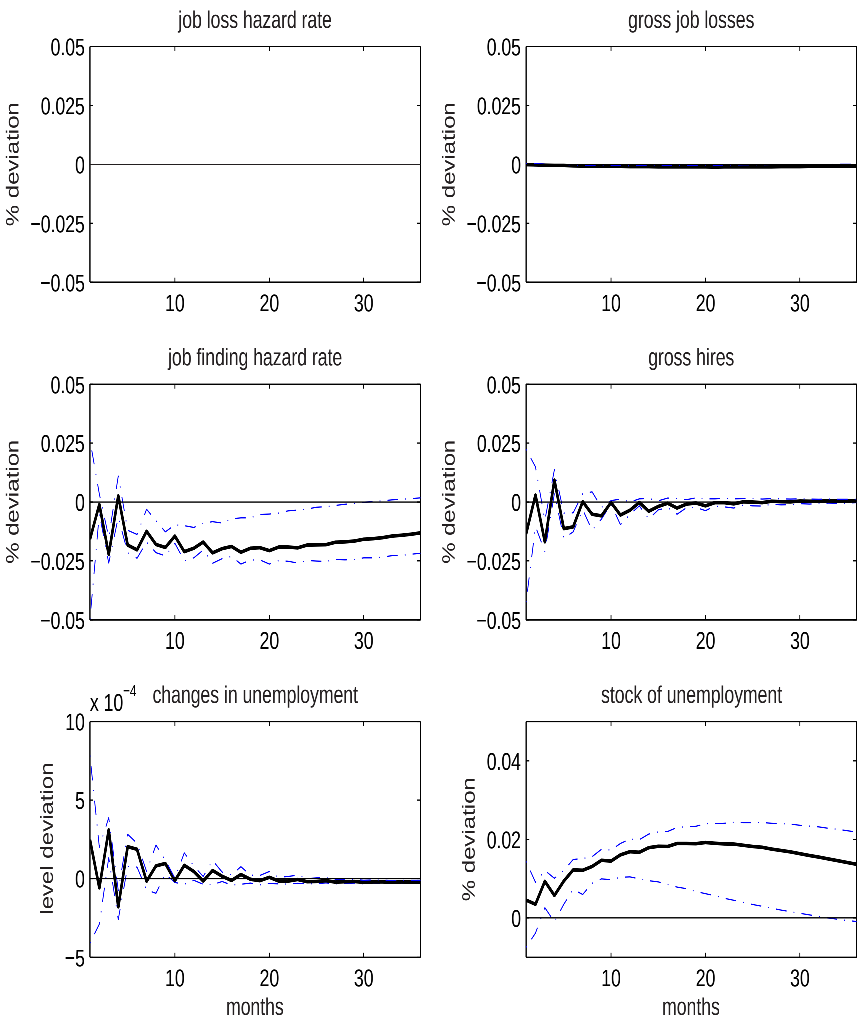

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution.

Figure 7: Impulse response functions for hazard rates and worker flows: Restrictions 1 and 2

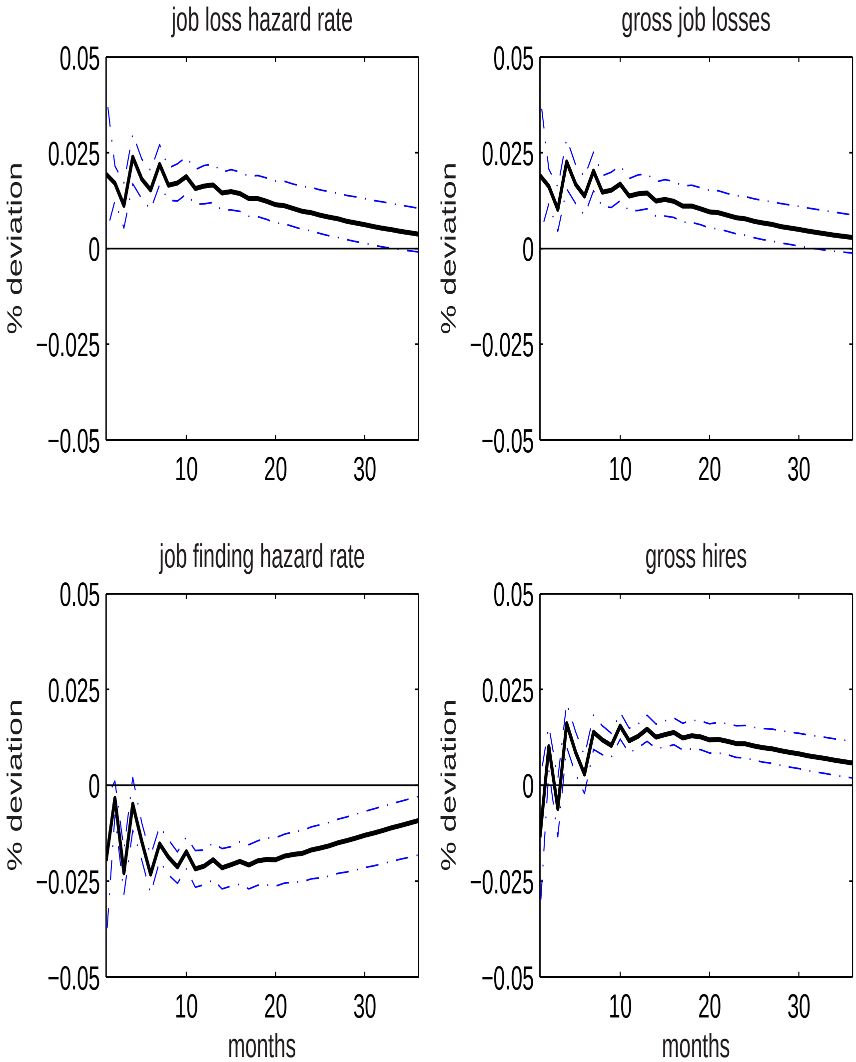

Notes: The shock is identified by imposing Restrictions 1 through 4. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 8: Impulse response functions for the change in unemployment, stock of unemployment and vacancies: Restrictions 1 through 4

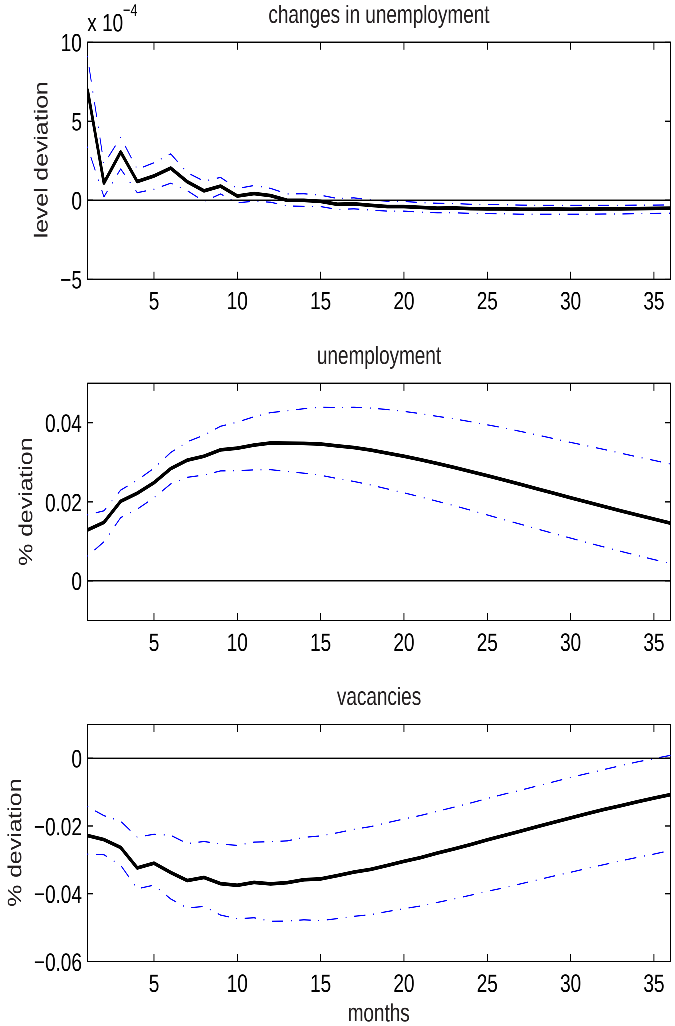

Notes: The shock is identified by imposing Restrictions 1 through 4. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses of the stock of unemployment and vacancies are expressed as percentage deviations from the steady state levels.

Figure 9: Aggregate level responses of unemployment and vacancies from the disaggregate model

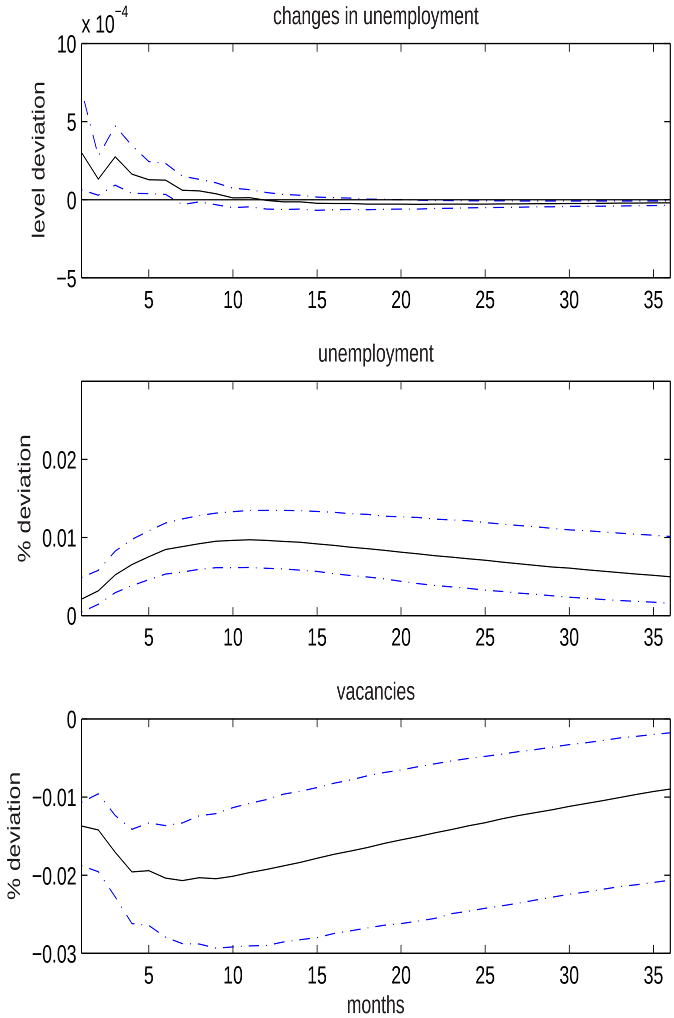

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 10: Unemployment responses for demographic groups

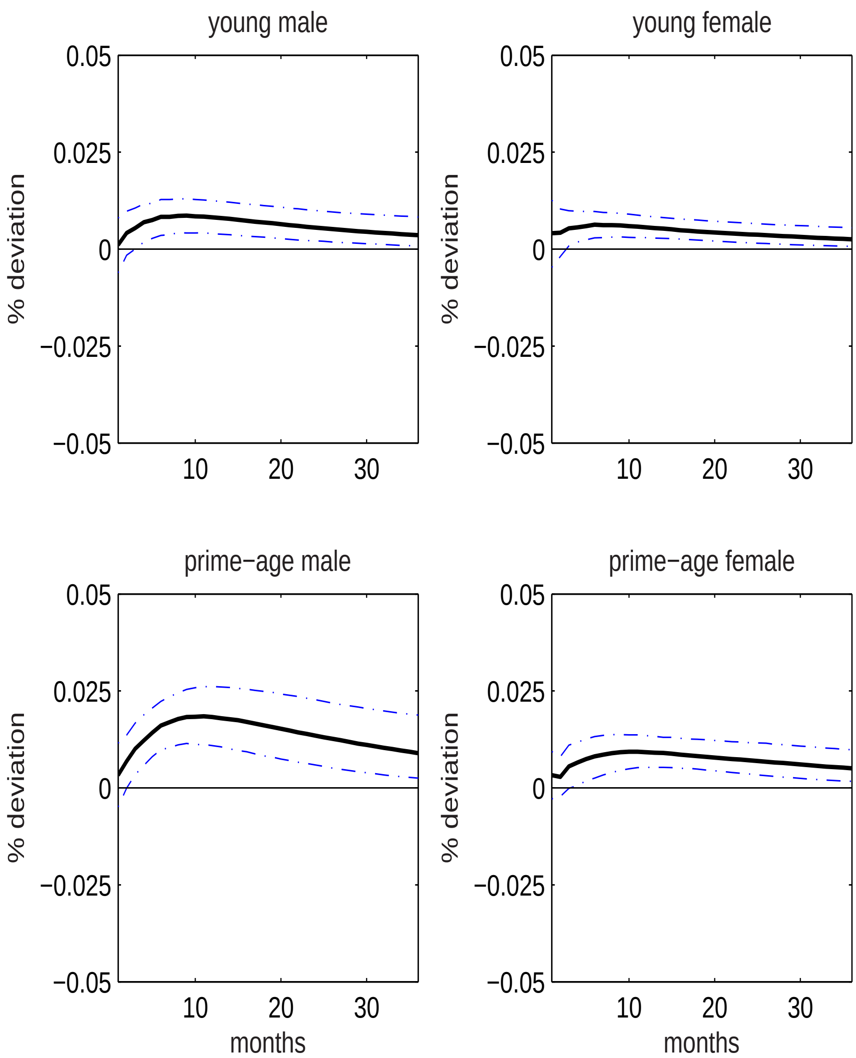

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 11: Responses of young male workers to the aggregate shock

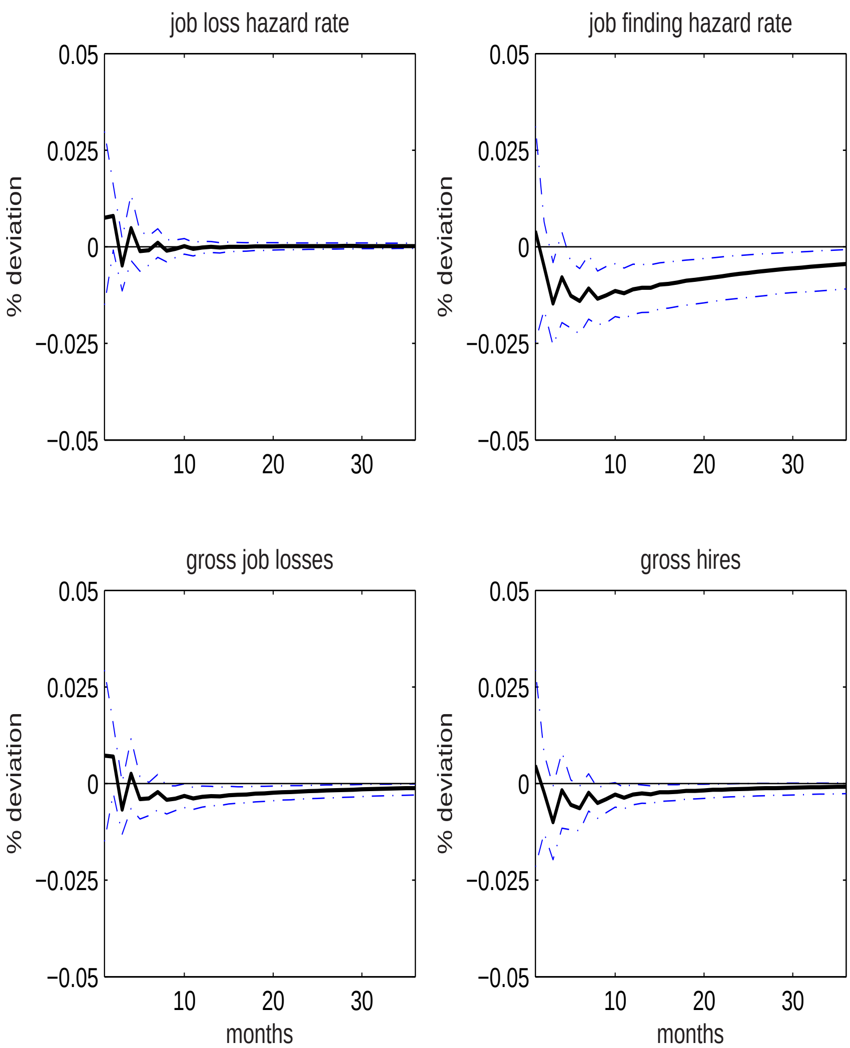

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 12: Responses of young female workers to the aggregate shock

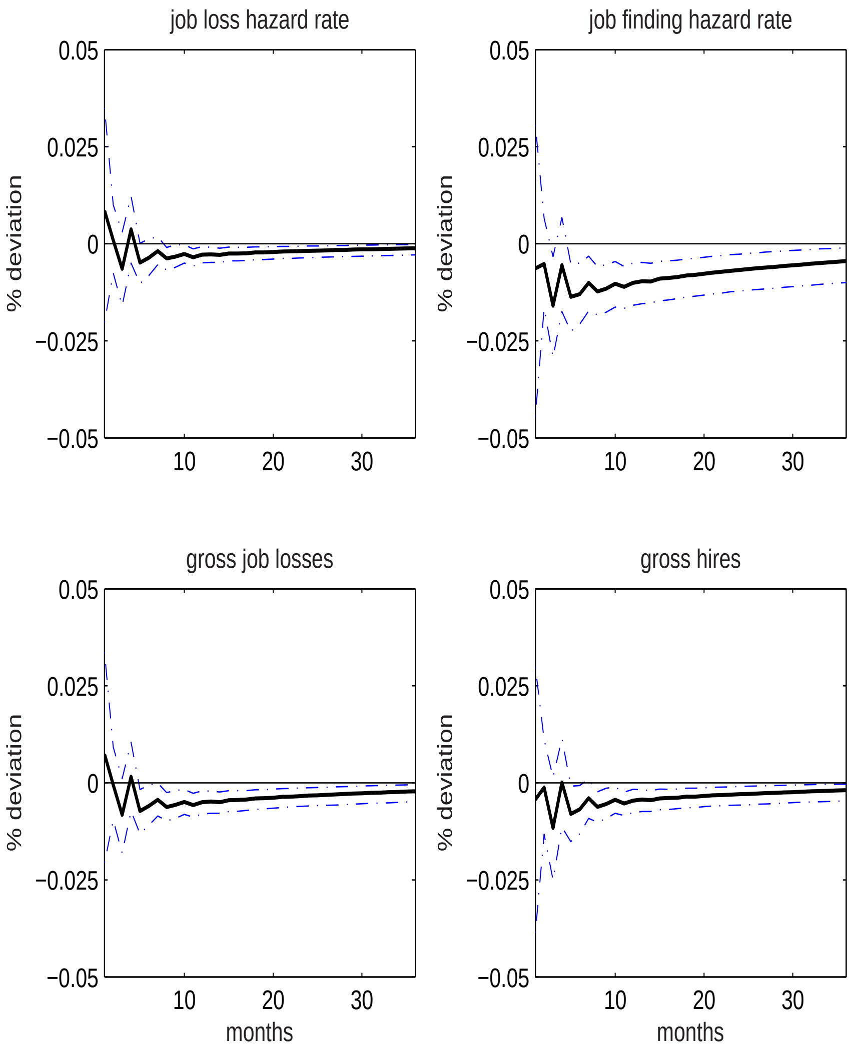

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 13: Responses of prime-age male workers to the aggregate shock

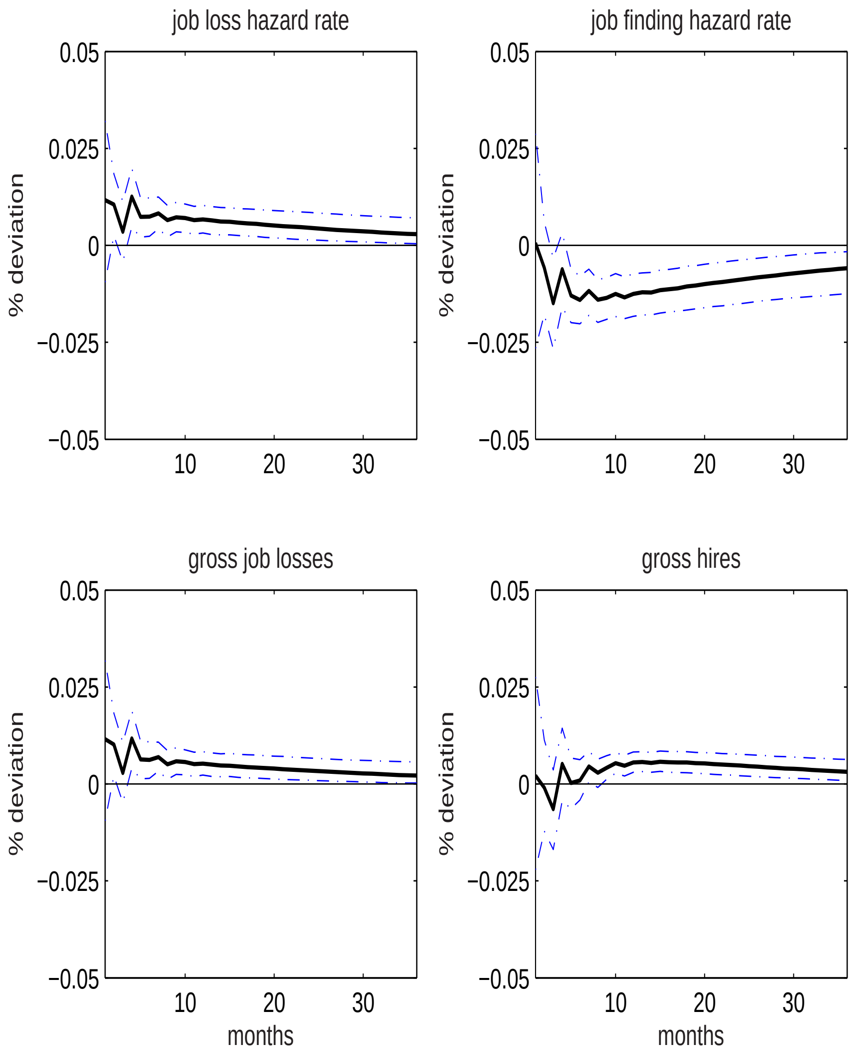

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.

Figure 14: Responses of prime-age female workers to the aggregate shock

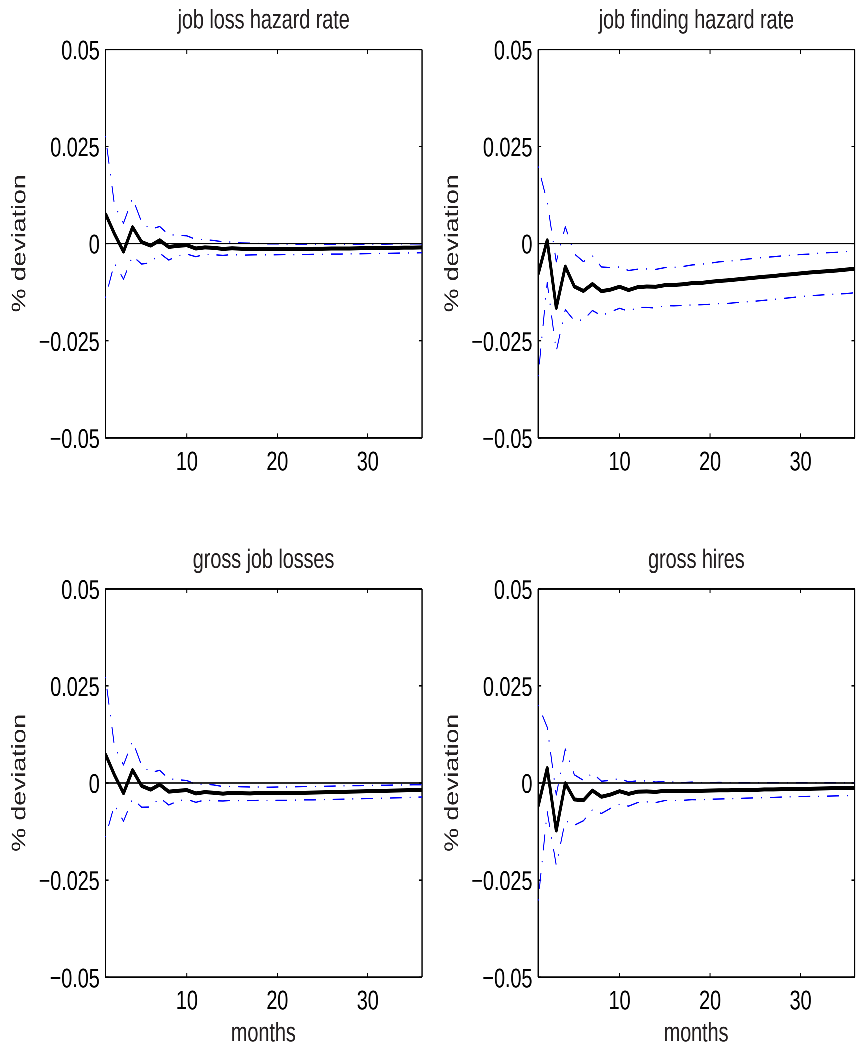

Notes: The shock is identified by imposing Restrictions 1 and 2. The black solid line represents the median of the posterior distribution. Error band represents the 10th and 90th percentiles of the posterior distribution. Responses are expressed as percentage deviations from the steady state levels.