Working Paper Series

The Time-Varying Beveridge Curve

WP 13-12

Luca Benati
University of Bern

Thomas A. Lubik
Federal Reserve Bank of Richmond

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

The Time-Varying Beveridge Curve*

Luca Benati
University of Bern†

Thomas A. Lubik
Federal Reserve Bank
of Richmond‡

August 2013

Working Paper No. 13-12

Abstract

We use a Bayesian time-varying parameter structural VAR with stochastic volatility to investigate changes in both the reduced-form relationship between vacancies and the unemployment rate, and in their relationship conditional on permanent and transitory output shocks, in the post-WWII United States. Evidence points towards similarities and differences between the Great Recession and the Volcker disinflation, and widespread time variation along two key dimensions. First, the slope of the Beveridge curve exhibits a large extent of variation from the mid-1960s on. It is also notably pro-cyclical, whereby the gain is positively correlated with the transitory component of output. The evolution of the slope of the Beveridge curve during the Great Recession is very similar to its evolution during the Volcker recession in terms of both its magnitude and its time profile. Second, both the Great Inflation episode and the subsequent Volcker disinflation are characterized by a significantly larger negative correlation between the reduced-form innovations to vacancies and the unemployment rate than the rest of the sample period. Those years also exhibit a greater cross-spectral coherence between the two series at business-cycle frequencies. This suggests that they are driven by common shocks.

JEL Classification: C11, C32, E20, E24,
Keywords: Bayesian VARs; stochastic volatility; time-varying parameters; Great Recession; Great Inflation; long-run restrictions

*The views in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond, the Board of Governors, or the Federal Reserve System. We are grateful to participants at the Applied Time Series Econometrics Workshop at the Federal Reserve Bank of St. Louis and the Midwest Macroeconomics Meetings at the University of Colorado Boulder for useful comments and suggestions.

†Department of Economics, University of Bern, Schanzeneckstrasse 1, CH-3001 Bern, Switzerland. Email: luca.benati@vwi.unibe.ch

‡Research Department, Federal Reserve Bank of Richmond, 701 E Byrd Street, Richmond, VA 23218, USA. E-mail: Thomas.Lubik@rich.frb.org

1 Introduction

The Beveridge curve describes the relationship between the unemployment rate and open positions, that is, vacancies, in the labor market. Plotting the former against the latter in a scatter diagram reveals a downward-sloping relationship that appears to be clustered around a concave curve (see Figure 1). The curve reflects the highly negative correlation between unemployment and vacancies that is a hallmark of labor markets in market economies.

Empirical work on the Beveridge curve has explored the relationship between vacancies and the unemployment rate under the maintained assumption that it can be regarded, as a first approximation, as time-invariant. The behavior of the two series during the Great Recession, with the unemployment rate seemingly stuck at high levels, even in the presence of a vacancy rate that has been progressively improving, has, however, raised doubts about the validity of the assumption of time-invariance. This suggests exploring the relationship between the two series allowing for the possibility that it may have evolved over time.

Our paper builds directly on the seminal contribution of Blanchard and Diamond (1989). These authors reintroduced the concept of the Beveridge curve as one of the key relationships in macroeconomic data. They conducted a vector autoregression (VAR) analysis of unemployment, vacancies, and the labor force in order to identify the driving forces behind movements in the Beveridge curve. We build upon their analysis by identifying both permanent and transitory structural shocks in a time-varying VAR context. By doing so, we are able to trace out the sources of movements, shifts, and tilts in the Beveridge curve over time.

The theoretical background for our study, and one that we use for identifying the structural shocks, is the simple search and matching approach to modeling labor markets (see Shimer, 2005). The Beveridge curve encapsulates the logic of this model. In times of economic expansions, unemployment is low and vacancies, that is, open positions offered by firms, are plentiful. Firms want to expand their workforce, but they are unable to do so since the pool of potential employees (that is, the unemployed) is small. As economic conditions slow down and demand slackens, firms post fewer vacancies and unemployment rises, consistent with a downward move along the Beveridge curve. At the trough of the business cycle, firms may have expectations of a future uptick in demand and start posting open positions. This decision is amplified by the large pool of unemployed, which guarantees firms high chances of finding suitable candidates and thus outweighs the incurred search costs. As the economy improves, unemployment falls and vacancy postings rise in an upward move along the Beveridge curve.

We approach the Beveridge curve empirically by specifying a time-varying parameter VAR with stochastic volatility. Our choice is informed by the observation that there are patterns in the Beveridge-curve relationship that are ill described by a linear framework. Specifically, the data suggest that the slope of the Beveridge curve is different for each business cycle episode, that the curve shifts over time, and that the pattern of driving forces change in a nonlinear fashion as well. Naturally, nonlinearity can take many forms, as do the framework to capture this. We utilize a time-varying parameter framework since it is a reasonably straightforward extension of linear VARs. Moreover, and perhaps more importantly, it can capture and approximate a wide range of underlying nonlinear behavior. To this point, we will introduce time variation in a nonlinear theoretical model of the labor market in order to relate it to the results from the VAR.

Our empirical analysis starts by documenting the presence of time variation in the relationship between vacancies and the unemployment rate by means of Stock and Watson's (1996, 1998) time-varying parameter median-unbiased estimation, which allows us to test for the presence of random-walk time variation in the data. Having detected evidence of time variation in the bivariate relationship between vacancies and the unemployment rate, we then use a Bayesian time-varying parameter structural VAR with stochastic volatility to characterize changes over time in the relationship. Evidence points towards both similarities and differences between the Great Recession and the Volcker disinflation, and widespread time variation along two key dimensions.

First, the slope of the Beveridge curve, which we capture by the average gain of the unemployment rate onto vacancies at business cycle frequencies is strongly negatively correlated with the Congressional Budget Office's (CBO) estimate of the output gap. The evolution of the slope of the Beveridge curve during the Great Recession is very similar to its evolution during the Volcker recession in terms of both its magnitude and its time profile. This suggests that the seemingly anomalous behavior of the Beveridge curve during the Great Recession, which has attracted much attention in the literature, may not have been that unusual. Second, both the Great Inflation episode and the subsequent Volcker disinflation, are characterized by a significantly larger (in absolute value) negative correlation between the reduced-form innovations to vacancies and the unemployment rate than the rest of the sample period. These years also show a greater cross-spectral coherence between the two series at business-cycle frequencies. This suggests that they are driven, to a larger extent than the rest of the sample, by common shocks.

Having characterized changes over time in the relationship between vacancies and the unemployment rate, we then proceed to interpret these stylized facts based on an estimated search and matching model. Specifically, we explore within a simple theoretical model how changes in individual parameter values affect the relationship between vacancies and the unemployment rate in order to gauge the origin of the variation in the Beveridge relationship.

The paper is organized as follows. The next section presents preliminary evidence on the presence of (random-walk) time variation in the bivariate relationship between the vacancy rate and the unemployment rate. Section 3 describes the Bayesian methodology we use to estimate the time-varying parameter VAR with stochastic volatility, whereas Section 4 discusses the evidence of changes over time in the Beveridge relationship. Section 5 details our structural identification procedure based on insights from a simple search and matching model, where we discuss implementation of both long-run and sign restrictions. We present the results of the structural identification procedure in Section 6. We also explore how changes in individual structural parameters of the search and matching model map into corresponding changes in the relationship between vacancies and the unemployment rate. Section 7 concludes. An appendix contains a detailed description of the econometric methods and the theoretical model we use for identification purposes.

2 Searching for Time Variation in the Beveridge Relationship

Figure 1 presents a time series plot of the unemployment rate and vacancies from 1949 to 2011. The negative relationship between the two series is readily apparent. At the peak of the business cycle unemployment is low and vacancies are high. Over the course of a downturn the former rises and the latter declines as fewer and fewer workers are employed and firms have fewer and fewer open positions. Volatility and serial correlation of both series appear of similar magnitude. The second panel in Figure 1 depicts the same series in a scatter plot of vacancies against unemployment, resulting in the well-known downward-sloping relationship that has come to be known as the Beveridge curve. In the graph, we plot individual Beveridge curves for each NBER business cycle. Each episode starts at the business cycle peak and ends with the period before the next peak. Visual inspection reveals two observations. First, all curves are downward-sloping, but with different slopes. Second, there is substantial lateral movement in the individual Beveridge curves, ranging from the innermost cycle, the 1953-1957 episode, to the outermost, 1982-1990. We take these observations as motivating evidence that the relationship between unemployment and vacancies exhibits substantial variation over time, which a focus on a single aggregate Beveridge curve obscures.

Time variation in data and in theoretical models can take many forms, from continuous variations in unit-root time-varying parameter models to discrete parameter shifts such as regime-switching. We regard both discrete and continuous changes as a priori plausible. In this paper, we focus on the latter. We thus provide evidence of time variation in the bivariate relationship between vacancies and unemployment. We apply the methodology developed by Stock and Watson (1996, 1998) to test for the presence of random-walk time-variation in the two-equation VAR representation for the two variables.[1]

The regression model we consider is:

$$x_t = \mu + \alpha(L)V_{t-1} + \beta(L)U_{t-1} + \epsilon_t \equiv \theta' Z_t + \epsilon_t, \tag{1}$$

where $x_t = V_t$, U_t, with V_t and U_t being the vacancy rate and the unemployment rate, respectively. $\alpha(L)$ and $\beta(L)$ are lag polynomials; $\theta = [\mu, \alpha(L), \beta(L)]'$ and $Z_t = [1, V_{t-1},$ ..., $U_{t-p}]'$. We select the lag order as the maximum of the lag orders individually chosen by the Akaike, Schwartz, and Hannan-Quinn criteria. Letting $\theta_t = [\mu_t, \alpha_t(L), \beta_t(L)]'$, the time-varying parameter version of (1) is given by:

$$x_t = \theta_t' Z_t + \epsilon_t, \tag{2}$$

$$\theta_t = \theta_{t-1} + \eta_t, \tag{3}$$

with $\eta_t \sim iid\ \mathcal{N}(0_{4p+1}, \lambda^2\sigma^2 Q)$, where 0_{4p+1} is a $(4p+1)$-dimensional vector of zeros. σ^2 is the variance of ϵ_t, Q a covariance matrix, and $E[\eta_t\epsilon_t] = 0$. Following Stock and Watson (1996, 1998), we set $Q = [E(Z_t Z_t')]^{-1}$. Under this normalization, the coefficients on the transformed regressors, $[E(Z_t Z_t')]^{-1/2} Z_t$, evolve according to a $(4p+1)$-dimensional standard random walk, where λ^2 is the ratio between the variance of each transformed innovation and the variance of ϵ_t. We estimate the matrix Q as $\hat{Q} = \left[T^{-1}\sum_{t=1}^{T} z_t z_t'\right]^{-1}$.

We estimate the specification (1) by OLS, from which we obtain an estimate of the innovation variance, $\hat{\sigma}^2$. We then perform an exp- and a sup-Wald joint test for a single unknown break in μ and in the sums of the α's and β's, using either the Newey and West

[1] From an empirical perspective, we prefer their methodology over, for instance, structural break tests for reasons of robustness to uncertainty regarding the specific form of time-variation present in the data. While time-varying parameter models can successfully track processes subject to structural breaks, Cogley and Sargent (2005) and Benati (2007) show that break tests possess low power when the true data-generating process (DGP) is characterized by random walk time variation. Generally speaking, break tests perform well if the DGP is subject to discrete structural breaks, while time-varying parameter models perform well under both scenarios.

(1987) or the Andrews (1991) HAC covariance matrix estimator to control for possible autocorrelation and/or heteroskedasticity in the residuals. Following Stock and Watson (1996), we compute the empirical distribution of the test statistic by considering a 100-point grid of values for λ over the interval [0, 0.1]. For each element of the grid we compute the corresponding estimate of the covariance matrix of η_t as $\hat{Q}_j = \lambda_j^2 \hat{\sigma}^2 \hat{Q}$; conditional on $\hat{Q}_j$ we simulate the model (2)-(3) 10,000 times, drawing the pseudo innovations from pseudo-random *iid* $N(0, \hat{\sigma}^2)$. We compute the median-unbiased estimate of λ as that particular value for which the median of the simulated empirical distribution of the test is closest to the test statistic previously computed based on the actual data. Finally, we compute the p-value based on the empirical distribution of the test conditional on $\lambda_j = 0$, which we compute based on Benati's (2007) extension of the Stock and Watson (1996, 1998) methodology.

We report the estimation results in Table 1. We find strong evidence of random-walk time variation in the equation for the vacancy rate. The p-values for the null of no time variation range from 0.0028 to 0.0195, depending on the specific test statistic. The median-unbiased estimates of λ are comparatively large, between 0.0235 and 0.0327. On the other hand, the corresponding p-values for the unemployment rate are much larger, ranging from 0.1661 to 0.2594, which suggests time-invariance. However, the densities of the median-unbiased estimates of λ in Figure 2 paint a more complex picture. A substantial fraction of the probability mass are clearly above zero, whereas median-unbiased estimates of λ range between 0.0122 and 0.0153. Although, strictly speaking, the null hypothesis of no time variation cannot be rejected at conventional significance levels in a frequentist sense, the evidence reported in Figure 2 suggests more caution. In what follows, we will proceed under the assumption that both equations feature random-walk time variation. We now investigate the changing relationship between the vacancy rate and the unemployment rate based on a Bayesian time-varying parameter VAR.

3 A Bayesian Time-Varying Parameter VAR with Stochastic Volatility

We define the data vector $Y_t \equiv [\Delta y_t, V_t, U_t]'$, where Δy_t is real GDP growth, computed as the log-difference of real GDP; V_t is the vacancy rate based the Conference Board's Help-Wanted Index and Barnichon's (2010) extension; and U_t is the unemployment rate. The data are all quarterly. The vacancies and unemployment series are both normalized by the labor force, seasonally adjusted, and converted from the original monthly series by simple averaging. The overall sample period is 1951Q1-2011Q4. We use the first 15 years of data to compute the Bayesian priors, which makes the effective sample period 1965Q1-2011Q4. Appendix A contains a complete description of the data and of their sources.

We specify the time-varying parameter $VAR(p)$ model:

$$Y_t = B_{0,t} + B_{1,t}Y_{t-1} + ... + B_{p,t}Y_{t-p} + \epsilon_t \equiv X_t^{'}\theta_t + \epsilon_t. \tag{4}$$

The notation is standard. As is customary in the literature on Bayesian time-varying parameter VARs, we set the lag order to $p = 2$. The time-varying lag coefficients, collected in the vector θ_t, are postulated to evolve according to:

$$p(\theta_t \mid \theta_{t-1}, Q) = I(\theta_t)\ f(\theta_t \mid \theta_{t-1}, Q), \tag{5}$$

where $I(\theta_t)$ is an indicator function that rejects unstable draws and thereby enforces stationarity on the VAR. The transition $f(\theta_t \mid \theta_{t-1}, Q)$ is given by:

$$\theta_t = \theta_{t-1} + \eta_t, \tag{6}$$

with $\eta_t \sim iid\ \mathcal{N}(0, Q)$.

We assume that the reduced-form innovations ϵ_t in (4) are normally distributed with zero mean, where we factor the time-varying covariance matrix Ω_t as:

$$Var(\epsilon_t) \equiv \Omega_t = A_t^{-1} H_t (A_t^{-1})'. \tag{7}$$

The time-varying matrices H_t and A_t are defined as:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 \\ 0 & h_{2,t} & 0 \\ 0 & 0 & h_{3,t} \end{bmatrix}, \qquad A_t \equiv \begin{bmatrix} 1 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 \end{bmatrix}. \tag{8}$$

We assume that the $h_{i,t}$ evolve as geometric random walks:

$$\ln h_{i,t} = \ln h_{i,t-1} + \nu_{i,t}, i = 1, 2, 3. \tag{9}$$

For future reference, we define $h_t \equiv [h_{1,t},\ h_{2,t},\ h_{3,t}]'$ and $\nu_t = [\nu_{1,t},\ \nu_{2,t},\ \nu_{3,t}]'$, with $\nu_t \sim iid\ \mathcal{N}(0, Z)$ and Z diagonal. We assume, as in Primiceri (2005), that the non-zero and non-unity elements of the matrix A_t, which we collect in the vector $\alpha_t \equiv [\alpha_{21,t},\ \alpha_{31,t},\ \alpha_{32,t}]'$, evolve as random walks:

$$\alpha_t = \alpha_{t-1} + \tau_t\ , \tag{10}$$

where $\tau_t \sim iid\ \mathcal{N}(0, S)$. Finally, we assume that the innovations vector $[u_t',\ \eta_t',\ \tau_t',\ \nu_t']'$ is distributed as:

$$\begin{bmatrix} u_t \\ \eta_t \\ \tau_t \\ \nu_t \end{bmatrix} \sim N\left(0, V\right), \text{ with } V = \begin{bmatrix} I_3 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & Z \end{bmatrix}, \tag{11}$$

where u_t is such that $\epsilon_t \equiv A_t^{-1} H_t^{\frac{1}{2}} u_t$.

We follow the literature in imposing a block-diagonal structure for V, mainly for parsimony, since the model is already quite heavily parameterized. Allowing for a completely generic correlation structure among different sources of uncertainty would also preclude any structural interpretation of the innovations. Finally, following Primiceri (2005) we adopt the additional simplifying assumption of a block-diagonal structure for S:

$$S \equiv Var(\tau_t) = Var(\tau_t) = \begin{bmatrix} S_1 & 0_{1\times 2} \\ 0_{2\times 1} & S_2 \end{bmatrix}, \tag{12}$$

with $S_1 \equiv Var(\tau_{21,t})$, and $S_2 \equiv Var([\tau_{31,t}, \tau_{32,t}]')$. This implies that the non-zero and non-one elements of A_t, which belong to different rows, evolve independently. This assumption simplifies inference substantially, since it allows Gibbs sampling on the non-zero and non-one elements of A_t equation by equation. We estimate (4)-(12) using standard Bayesian methods. Appendix B discusses our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data.

4 Reduced-Form Evidence

Figure 3 presents the first set of reduced-form results. It shows statistics of the estimated time-varying innovations in the VAR (4). The first panel depicts the median posterior estimate of the correlation coefficient of the innovations to vacancies and the unemployment rate and associated 68% and 90% coverage regions. The plot shows substantial time variation in this statistic. From the late 1960s to the early 1980s the correlation strengthens from -0.4 to -0.85 before rising (in absolute value) to a low of -0.25. Over the course of the last decade, the correlation has strengthened again, settling close to the average median value of -0.55. This suggests that the unemployment-vacancy correlation strengthens during periods of broad downturns and high volatility, whereas it weakens in general upswings with low economic turbulence. The evidence over the last decade also supports the impression that the U.S. economy is in a period of a prolonged downswing.[2]

The impression of substantial time variation is strongly supported by the second panel, which shows the fraction of draws from the posterior distribution for which the correlation coefficient is greater than the average median value over the sample. The fraction of draws

[2] We also note that at the same time the coverage regions are tightly clustered around the median estimate during the period of highest instability, namely the last 1970s and the Volcker disinflation, whereas they are more spread out in the beginning and towards the end of the sample.

sinks toward zero at the end of the Volcker disinflation, while it oscillates for much of the Great Moderation between 0.6 and 0.9. Similarly to the results in the first panel, the period since the beginning of the financial crisis in August 2007 is characterized by a substantial decrease in the fraction of draws.

In the third panel of Figure 3, we highlight the ratio of the estimated standard deviations of the unemployment and vacancy innovations. The graph shows substantial time variation in this ratio, although overall both innovation variances are of roughly equal size. While the innovation variance of the vacancy rate appears overall dominant, unemployment innovations play a relatively larger role at the end of the Great Inflation, the Volcker disinflation, and the Great Recession. All of these are periods during which the unemployment rate shot up sharply. This suggests a dominant role of specific shocks, namely those tied closely to reduced-form innovations to the unemployment rate, at the onset of an economic downturn. We attempt to identify the sources for this behavior in the following section.

We now narrow our focus to the behavior of unemployment and vacancies at the business-cycle frequencies between six quarters and eight years. We report these results using statistics from the frequency domain. Figure 4 shows median posterior estimates (and associated coverage regions) of the average cross-spectral gain and coherence between the two variables. The gain of a variable x_t onto another variable y_t at the frequency ω is defined as the absolute value of the OLS-coefficient in the regression of y_t on x_t at that frequency, whereas the coherence is the R^2 in that regression. Consequently, the gain has a natural interpretation in terms of the slope of the Beveridge curve, while the coherence measures the fraction of the vacancy-rate's variance at given frequencies that is accounted for by the variation in the unemployment rate. We find it convenient to express time variation in the Beveridge curve in terms of the frequency domain since it allows us to isolate the fluctuations of interest, namely policy-relevant business cycles, and therefore abstract from secular movements.

Overall, evidence of time variation is significantly stronger for the gain than for the coherence. The coherence between the two series appears to have remained broadly unchanged since the second half of the 1960s, except for a brief run-up during the Great Inflation of the 1970s, culminating in the tight posterior distribution during the Volcker disinflation of the early 1980s. Moreover, average coherence is always above 0.8, with 0.9 contained in the 68% coverage region. The high explanatory power of one variable for the other at the business cycle frequencies thus suggests that unemployment and vacancies are driven by a set of common shocks over the sample period.

The gain is large during the same periods in which the relative innovation variance of reduced-form shocks to the unemployment rate is large, namely during the first oil shock, the Volcker recession, and the Great Recession; that is, during these recessionary episodes movements in the unemployment rate are relatively larger than those in the vacancy rate. This points towards a flattening of the Beveridge curve in downturns, when small movements in vacancies are accompanied by large movements in unemployment. Time variation in the gain thus captures the shifts and tilts in the individual Beveridge curves highlighted in Figure 1 in one simple statistic.

As a side note, our evidence does not indicate fundamental differences between the Volcker disinflation and the Great Recession, that is, between the two deepest recessions in the post-war era. This is especially apparent from the estimated gain in Figure 4, which shows a similar time profile during both episodes. The relationship between vacancies and unemployment, although clearly different from the years leading up to the financial crisis, is broadly in line with that of the early 1980s.

We can now summarize our findings from the reduced-form evidence as follows. The correlation pattern between unemployment and vacancies shows a significant degree of time variation. It strengthens during downturns and weakens in upswings. This is consistent with the idea that over the course of a business cycle, as the economy shifts from a peak to a trough, labor market equilibrium moves downward along the Beveridge curve. This movement creates a tight negative relationship between unemployment and vacancies. As the economy recovers, however, vacancies start rising without much movement in unemployment. Hence, the correlation weakens. The economy thus goes off the existing Beveridge curve, in the manner of a counter-clockwise loop, as identified by Blanchard and Diamond (1989), or it moves to a new Beveridge curve, as suggested by the recent literature on mismatch, e.g. Furlanetto and Groshenny (2012), Lubik (2012), and Sahin et al. (2012). Evidence from the frequency domain suggests that the same shocks underlie movements in the labor market, but that over the course of the business cycle shocks change in their importance. During recessions movements in unemployment dominate, while in upswings vacancies play a more important role. We now try to identify the structural factors determining this reduced-form behavior.

5 Identification

A key focus of our analysis is to identify the underlying sources of the movements in the Beveridge curve. In order to do so we need to identify the structural shocks behind the behavior of unemployment and vacancies. Our data set contains a nonstationary variable, GDP, and two stationary variables, namely the unemployment and vacancy rates. This allows us to identify one permanent and two transitory shocks from the reduced-form innovation covariance matrix. While the permanent shock has no effect on the two labor market variables in the long run, it can still lead to persistent movements in these variables, and thus in the Beveridge curve, in the short to medium run.[3] More specifically, we are interested in which shocks can be tied to the changing slope and the shifts in the Beveridge curve. We let our identification strategy be guided by the implications of the simple search and matching model, which offers predictions for the effects of permanent and transitory productivity shocks as well as for other transitory labor market disturbances.

5.1 A Simple Theoretical Framework

We organize the interpretation of our empirical findings around the predictions of the standard search and matching model of the model labor market as described in Shimer (2005). The model is a data-generating process for unemployment and vacancies that is driven by a variety of fundamental shocks. The specific model is taken from Lubik (2012). The specification and derivation is described in more detail in the appendix.

The model can be reduced to three key equations that will guide our thinking about the empirics. The first equation describes the law of motion for employment:

$$N_{t+1} = (1 - \rho_{t+1}) \left[N_t + m_t U_t^{\xi} V_t^{1-\xi} \right]. \tag{13}$$

The stock of existing workers N_t is augmented by new hires $m_t U_t^{\xi} V_t^{1-\xi}$, which are the result of a matching process between open positions V_t and job seekers U_t via the matching function. The matching process is subject to exogenous variation in the match efficiency m_t, which affects the size of the workforce with a one-period lag. Similarly, employment is subject to exogenous variations in the separation rate ρ_t, which affects employment contemporaneously. It is this timing convention that gives rise to an identifying restriction.

The second equation is the job-creation condition, which describes the optimal vacancy posting decision by a firm:

$$\frac{\kappa_t}{m_t} \theta_t^{\xi} = \beta E_t (1 - \rho_{t+1}) \left[(1 - \eta) (A_{t+1} - b_{t+1}) - \eta \kappa_{t+1} \theta_{t+1} + \frac{\kappa_{t+1}}{m_{t+1}} \theta_{t+1}^{\xi} \right]. \tag{14}$$

[3] While this rules out strict hysteresis effects, in the sense that temporary shocks can have permanent effects, it can still lead to behavior that looks over typical sample periods as hysteresis-induced. Moreover, the empircial evidence concerning hysteresis is decidedly mixed.

$\theta_t = V_t/U_t$ is labor market tightness and a crucial statistic in the search and matching model. Effective vacancy creation costs $\frac{\kappa_t}{m_t}\theta_t^{\xi}$ are increasing in tightness since firms have to compete with other firm's hiring efforts given the size of the applicant pool. Hiring costs are subject to exogenous variations in the component κ_t and have to be balanced against the expected benefits, namely the right-hand side of the above equation. This consists of surplus of a worker's marginal product A_t over his outside options (unemployment benefits b_t), net of a hold-up term, $\eta\kappa_t\theta_t$, accruing to workers and extracted from the firm on account of the latter's costly participation in the labor market, and the firm's implicit future cost savings $\frac{\kappa_t}{m_t}\theta_t^{\xi}$ when having already hired someone. η is a parameter indicating the strength of worker's bargaining power. Finally, production is assumed to be linear in employment, but subject to both permanent and temporary productivity shocks, A_t^P and A_t^T, respectively:

$$Y_t = A_t N_t = \left(A_t^P A_t^T\right) N_t. \tag{15}$$

The permanent shock is a pure random walk, while the temporary shock is serially correlated but stationary.

Figure 5 depicts the theoretical impulse response functions of the unemployment and vacancy rate to each of the shocks. We can categorize the shocks in two groups, namely into shocks that move unemployment and vacancies in the same direction, and those that imply opposite movements of these variables. This classification underlies the identification by sign restrictions that we use later on. Both productivity shocks increase vacancies on impact and lower unemployment over the course of the adjustment period. The effect of the temporary shock is much more pronounced since it is calibrated at a much higher level of persistence than the productivity growth rate shock. Persistent productivity shocks increase vacancy posting because they raise the expected value of a filled position. As more vacancies get posted, new employment relationships are established and the unemployment rate falls. We note that permanent shocks have a temporary effect on the labor market because they tilt the expected profit profile in a manner similar to temporary shocks. However, they are identified by their long-run effect on output, which by definition no other shock can muster.

Shocks to match efficiency, vacancy posting costs and unemployment benefits lead to negative comovement between unemployment and benefits. Increases in match efficiency and decreases in the vacancy costs both lower effective vacancy creation cost $\frac{\kappa_t}{m_t}\theta_t^{\xi}$ *ceteris paribus* and thereby stimulate initial vacancy creation. These vacancies then lead to lower unemployment over time. In the case of κ_t there is additional feedback from wage setting since the hold-up term $\kappa_t\theta_t$ can rise or fall. Similarly, increases in match efficiency have an additional effect via the matching function as the higher level of vacancies is now turned into even more new hires, so that employment rises. Movements in benefits also produce negative comovements between the key labor market variables, but the channel is via wage setting. Higher benefits increase the outside option of the worker in bargaining which leads to higher wages. This reduces the expected profit stream to the firm and fewer vacancy postings and higher unemployment.

On the other hand, a persistent increase in the separation rate drives both unemployment and vacancy postings higher. There is an immediate effect on unemployment, which *ceteris paribus* lowers labor market tightness, thereby reducing effective vacancy posting cost. In isolation, this effect stimulates vacancy creation. At the same time, persistent increases in separations reduce expected profit streams from filled positions which has a dampening effect on desired vacancies. This is balanced, however, by persistent declines in tightness because of increased separations. The resulting overall effect is that firms take advantage of the larger pool of potential hires and increase vacancy postings to return to the previous long-run level over time.

5.2 Disentangling Permanent and Transitory Shocks

We now describe how we implement identification of a single permanent shock and two transitory shocks in our time-varying parameter VAR model, based on the theoretical insights derived in the previous section. The permanent shock is identified from a long-run restriction as originally proposed by Blanchard and Quah (1989). We label a shock as permanent if it affects only GDP in the long run, but not the labor market variables. The short- and medium-run effects on all variables is left unrestricted. In terms of the simple model, the identified permanent shock is consistent with the permanent productivity shock A_t^P which underlies the stochastic trend in output. We follow the procedure proposed by Galí and Gambetti (2009) for imposing long-run restrictions within a time-varying parameter VAR model.

Let $\Omega_t = P_t D_t P_t'$ be the eigenvalue-eigenvector decomposition of the VAR's time-varying covariance matrix Ω_t in each time period and for each draw from the ergodic distribution. We compute a local approximation to the matrix of the cumulative impulse-response functions (IRFs) to the VAR's structural shocks as:

$$\bar{C}_{t,\infty} = \underbrace{[I_N - B_{1,t} - ... - B_{p,t}]^{-1}}_{C_0} \bar{A}_{0,t}, \tag{16}$$

where I_N is the $N \times N$ identity matrix. The matrix of the cumulative impulse-response functions is then rotated via an appropriate Householder matrix H in order to introduce zeros in the first row of $\bar{C}_{t,\infty}$, which corresponds to GDP, except for the (1,1) entry. Consequently, the first row of the cumulative impulse-response functions,

$$C^P_{t,\infty} = \bar{C}_{t,\infty} H = C_0 \bar{A}_{0,t} H = C_0 A^P_{0,t}, \tag{17}$$

is given by $[x\ 0\ 0]$, with x being a non-zero entry. By definition, the first shock identified by $A^P_{0,t}$ is the only one exerting a long-run impact on the level of GDP. We therefore label it the permanent output shock.

5.3 Identifying the Transitory Shocks Based on Sign Restrictions

We identify the two transitory shocks by assuming that they induce a different impact pattern on vacancies and the unemployment rate. Our theoretical discussion of the search and matching model has shown that a host of shocks, e.g., temporary productivity, vacancy cost, or match efficiency shocks, imply negative comovement for the two variables, while separation rate shocks increase vacancies and unemployment on impact. We transfer these insights to the structural VAR identification scheme.

Let $u_t \equiv [u^P_t, u^{T_1}_t, u^{T_2}_t]'$ be the vector of the structural shocks in the VAR: u^P_t is the permanent output shock, $u^{T_1}_t$ and $u^{T_2}_t$ are the two transitory shocks; let $u_t = A^{-1}_{0,t}\epsilon_t$, with $A_{0,t}$ being the VAR's structural impact matrix. Our sign restriction approach postulates that $u^{T_1}_t$ induces the opposite sign on vacancies and the unemployment rate contemporaneously, while $u^{T_2}_t$ induces an impact response of the same sign. We compute the time-varying structural impact matrix $A_{0,t}$ by combining the methodology proposed by Rubio-Ramirez et al. (2005) for imposing sign restrictions, and the procedure proposed by Galí and Gambetti (2009) for imposing long-run restrictions in time-varying parameter VARs.

Let $\Omega_t = P_t D_t P_t'$ be the eigenvalue-eigenvector decomposition of the VAR's time-varying covariance matrix Ω_t, and let $\tilde{A}_{0,t} \equiv P_t D_t^{\frac{1}{2}}$. We draw an $N \times N$ matrix K from a standard-normal distribution and compute the QR decomposition of K; that is, we find matrices Q and R such that $K = Q{\cdot}R$. A proposal estimate of the time-varying structural impact matrix can then be computed as $\bar{A}_{0,t} = \tilde{A}_{0,t} \cdot Q'$. We then compute the local approximation to the matrix of the cumulative IRFs to the VAR's structural shocks, $\bar{C}_{t,\infty}$, from (16). In order to introduce zeros in the first row of $\bar{C}_{t,\infty}$, we rotate the matrix of the cumulative IRFs via an appropriate Householder matrix H. The first row of the matrix of the cumulative IRFs, $C^P_{t,\infty}$ as in (17), is given by $[x\ 0\ 0]$. If the resulting structural impact matrix $A_{0,t} = \bar{A}_{0,t}H$ satisfies the sign restrictions, we store it; otherwise it is discarded. We then repeat the procedure until we obtain an impact matrix that satisfies both the sign restrictions and the long-run restriction at the same time.

6 Structural Evidence

Our identification strategy discussed in section 5 allows us distinguish between one permanent and two transitory shocks. The permanent shock is identified as having a long-run effect on GDP, while the transitory shocks are identified from sign restrictions derived from a simple search and matching model. A side product of our strategy is that we can identify the natural rate of output as its permanent component. Figure 6 shows real GDP in logs together with the median of the posterior distribution of the estimated permanent component and the 68% coverage region. We also report the corresponding transitory component together with the output gap estimate from the CBO.

Our estimate of the transitory component is most of the time quite close to the CBO output gap, which is produced from a production function approach to potential output, whereas our estimate is largely atheoretical. The main discrepancy between the two estimates is in the wake of the Great Recession, particularly the quarters following the collapse of Lehman Brothers. Whereas the CBO estimate implies a dramatic output shortfall of around 7.5% of potential output in the first half of 2009, our estimated gap is much less at between 3-4% with little change since then. The reason behind our smaller estimate of the current gap is a comparatively large role played by permanent output shocks in the Great Recession. As the first panel shows, the time profile of the permanent component of log real GDP is estimated to have been negatively affected in a significant way by the Great Recession, with a downward shift in the trend path; that is, natural output is now permanently lower. The question we now investigate is whether and to what extent these trend shifts due to permanent output shocks seep into the Beveridge curve.

6.1 Impulse Response Functions

As a first pass, we report IRFs to unemployment and vacancies for each of the three shocks in Figures 7-9. Because of the nature of the time-varying parameter VAR, there is not a single IRF for each shock-variable combination. We therefore represent the IRFs by collecting the time-varying coefficients on impact, two quarters ahead, one year ahead, and five years ahead in individual graphs in order to track how the dynamic behavior of the labor market variables changes over time. An IRF for a specific period can then be extracted by following the impulse response coefficient over the four panels. The IRFs are normalized such that the long-run effect is attained at a value of one, while transitory shocks eventually return the responses to zero.

In Figure 7, an innovation to the permanent component of output raises GDP on impact by one-half of the long-run effect, which is obtained fairly quickly after around one year in most periods. A permanent shock tends to raise the vacancy rate on impact, after which it rises for a few quarters before falling to its long-run level. The unemployment rate rises on impact, but then quickly settles around zero. The initial, seemingly counterfactual response is reminiscent of the finding by Galí (1999) that positive productivity shocks have negative employment consequences, which in our model translates into an initial rise in the unemployment rate. Furthermore, the behavior of the estimated impulse responses is broadly consistent with the results from the calibrated theoretical model, both in terms of direction and size of the responses. As we will see below, compared to the transitory shocks the permanent productivity shock, which in the theoretical model takes the form of a growth rate variation, exerts only a small effect on unemployment and vacancy rates. Notably, the coverage regions for both variables include zero at all horizons. Overall, the extent of time variation in the IRFs appears small. It is more pronounced at shorter horizons than in the long run.

We report the IRFs to the first transitory shock in Figure 8. This shock is identified as inducing an opposite response of the vacancy and the unemployment rate on impact. In the theoretical model, this identified empirical shock is associated with a transitory productivity shock, variations in match efficiency, hiring costs, or benefit movements. The IRFs of all three variables in the VAR are hump-shaped, with a peak response after one year. Moreover, the amplitudes of the responses are much more pronounced than in the previous case. The vacancy rate is back at its long-run level after 5 years, while there is much more persistence in the unemployment rate and GDP. We also note that our simple theoretical framework cannot replicate this degree of persistence.

The vacancy rate exhibits the highest degree of time variation. What stands out is that its response is asymmetric over the business cycle, but only in the pre-1984 period. During the recessions of the early and mid-1970s, and the deep recession of the early 1980s culminating in the Volcker disinflation, the initial vacancy response declines (in absolute value) over the course of the downturn before increasing in the recovery phase. That is, the vacancy rate responds less elastically to the first transitory shock during downturns than in expansions - which is not the case for the unemployment rate. This pattern is visible at all horizons. Between the Volcker disinflation and shortly before the onset of the Great Recession the impact response of the vacancy rate declines gradually from -1% to almost -2% before rising again sharply during a recession.

The second transitory shock is identified by imposing the same sign response on unemployment and vacancies. In the context of the theoretical model, such a pattern is due to movements in the separation rate. The IRFs in Figure 9 show that the vacancy rate rises on impact, then reaches a peak four quarters out before returning gradually over the long run. The unemployment rate follows the same pattern, while the shock induces a large negative response of GDP. None of the responses exhibits much time variation, at best there are slow-moving changes in the IRF-coefficients towards less elastic responses. Interestingly, the impact behavior of the vacancy rate declines over the course of the Great Recession. We note, however, that the coverage regions are very wide and include zero for the unemployment rate and GDP at all horizons.

6.2 Variance Decompositions

Figure 10 provides evidence on the relative importance of permanent and transitory shocks for fluctuations in vacancies and the unemployment rate. We report the median of the posterior distributions of the respective fractions of innovation variance due to the permanent shock and the associated coverage regions. For the vacancy rate, permanent shocks appear to play a minor role, with a median estimate of between 10% and 20%. The median estimate for the unemployment rate exhibits a greater degree of variation, oscillating between 10% and 40%. Despite this large extent of time variation, it is difficult to relate fluctuations in the relative importance of permanent shocks to key macroeconomic events. Possible candidates are the period after the first oil shock, when the contribution of permanent shocks increased temporarily, and the long expansion of the 1980s until the late 1990s, which was temporarily punctured by the recession in 1991. Moreover, there is no consistent behavior of the permanent shock contribution over the business cycle. Their importance rises both in downturns and in upswings. On the other hand, this observation gives rise to the idea that all business cycles, at least in the labor market, are different along this dimension.

We now turn to the relative contribution of the two transitory shocks identified by sign restrictions. The evidence is fairly clear-cut. Given the strongly negative unconditional relationship between vacancies and the unemployment rate, we would expect the contribution of $u_t^{T_2}$, that is, the shock that induces positive contemporaneous comovement between the two variables, to be small. This is, in fact, borne out by the second column of the graph in Figure 11. The median estimate of the fraction of innovation variance of the two series due to $u_t^{T_2}$ is well below 20%. Correspondingly, the first transitory shock appears clearly to be dominant for both variables. Based on the theoretical model, we can associate this shock with either temporary productivity disturbances or with stochastic movements in hiring costs, match efficiency, or unemployment benefits. Given the parsimonious nature of both the theoretical and empirical model, however, we cannot further disentangle this.

6.3 Structural Shocks and Beveridge Curve Shifts

We now turn to one of the main results of the paper, namely the structural sources of time variation in the Beveridge curve. We first discuss the relationship between the business cycle, as identified by the transitory component in GDP and measures of the Beveridge curve. We then decompose the estimated gain and coherence of unemployment and vacancies into their structural components based on the identification scheme discussed above.

Figure 12 reports two key pieces of evidence on the cyclical behavior of the slope of the Beveridge curve. The left panel shows the fraction of draws from the posterior distribution for which the transitory component of output is positive. This is plotted against the fraction of draws for which the average cross-spectral gain between vacancies and the unemployment rate at the business-cycle frequencies is greater than one. The graph thus gives an indication of how the slope of the Beveridge curve moves with aggregate activity over the business cycle.

We can differentiate two separate time periods. During the 1970s and the early 1980s, that is, during the Great Inflation, the slope of the Beveridge curve systematically comoves contemporaneously with the state of the business cycle. It is comparatively larger (in absolute value) during business-cycle upswings, and comparatively smaller during periods of weak economic activity. Similarly, the Great Recession is characterized by very strong comovement between the slope of the Beveridge curve and the transitory component of output, but this time the slope slightly leads the business cycle.

On the other hand, in the long expansion period from 1982 to 2008, labelled the Great Moderation that was only marred by two minor recessions, the slope of the Beveridge curve comoves less clearly with the business cycle, a pattern that is especially apparent during the 1990s. In the early and late part of this period the Beveridge curve appears to lag the cycle. This is consistent with the notion of jobless recoveries after the two mild recessions. Despite upticks in economic activity, the labor market did not recover quickly after 1992 and, especially, after 2001. In the data, this manifests itself in a large gain between unemployment and vacancies (see Figure 13). Moreover, this is also consistent with the changing impulse response patterns to structural shocks discussed above. In a sense, the outlier is the Great Recession, which resembles more the recessions of the Great Inflation rather than those of the Great Moderation.

The second panel reports additional evidence on the extent of cyclicality of the slope of the Beveridge curve. It shows the distribution of the slope coefficient in the LAD (Least Absolute Deviations) regression of the cross-spectral gain on a constant and the transitory component of output. Overall, the LAD coefficient is greater than zero for 82.5% of the draws from the posterior distribution, which points towards the pro-cyclicality of the slope of the Beveridge curve.

Figure 13 shows how the two types of shocks shape the evolution of the Beveridge curve. We plot the average gain and coherence between vacancies and the unemployment rate at business-cycle frequencies over time together with the fraction of draws for which the average gain is greater than one. The upper row of the panel reports the statistics conditional on the permanent shock, the lower panel contains those conditional on the two transitory shocks. Whereas the coherence conditional on the permanent shock does not show much time variation, conditioning on transitory shocks reveals a pattern that is broadly similar to the reduced-form representation. This suggests that the comparatively greater coherence between the two series around the time of the Great Inflation and of the Volcker disinflation is mostly due to transitory shocks.

Time variation in the gain, on the other hand, appears to be due to both types of shocks. Although the middle column suggests that the extent of statistical significance of the fluctuations in the gain is similar, the first column shows a different magnitude. In particular, fluctuations in the gain conditional on permanent output shocks, which accounted for a comparatively minor fraction of the innovation variance of the two series, is significantly wider than the corresponding fluctuations conditional on transitory shocks. Moreover, and unsurprisingly in the light of the previously discussed evidence on the relative importance of the two types of shocks, both the magnitude and the time-profile of the fluctuations of the gain conditional on transitory shocks are very close to the reduced-form evidence.

6.4 Interpreting Changes in the Beveridge Curve Based on an Estimated DSGE Model

One of our key contributions is to have demonstrated the presence and the extent of non-linearity in the Beveridge-curve relationship over time in U.S. data. We did so in a structural VAR, where the identification restrictions were derived from the behavior of a simple search and matching model for the labor market. Nevertheless, VARs are by their very nature largely atheoretical, in the sense that they represent the reduced form of a potentially much richer underlying dynamic stochastic general equilibrium (DSGE) model. The best that we can do is to identify structural innovations, but they do not necessarily reveal much about the structure of the theoretical model. However, researchers may want to go further than this. One particular point of interest is the source of the non-linearity in the data. We now make some forays in this direction.

Following Fernández-Villaverde and Rubio-Ramirez (2007) we assume that all parameters in the DSGE model in Section 4 are first-order autoregressive processes. The set of parameters thus includes the original primitive parameters, their first-order autocorrelation coefficients, and their innovation variances. We then estimate the model using Bayesian methods.[4] We draw from the posterior distribution for each individual parameter and compute the associated gain of the unemployment rate onto vacancies at business cycle frequencies. Figure 14 contains a graph that shows, for parameter intervals around the modal estimates generated by the Random Walk Metropolis algorithm, the average gain of the unemployment rate onto vacancies at the business-cycle frequencies, as a function of each individual parameter.

The parameters with the largest impact on the gain, as a measure of the slope of the Beveridge curve, are the separation rate ρ, the match efficiency m, and the match elasticity ξ. The estimated gain is independent of the vacancy creation cost κ and the bargaining share η. It is also largely inelastic to variations in the parameters' autocorrelation coefficients, with the exception of the match efficiency and separation rate parameters and the serial correlation of the permanent productivity shock. The innovation variances do not affect the gain, the exception being the separation rate. These results are in line with the observation of Lubik (2012) that the key driver of shifts in the Beveridge curve are variations in the matching function parameters. While productivity shocks can generate movements along the Beveridge curve, movements *of* the Beveridge curve have to come through changes in the law of motion for employment. In contrast to his findings, our exercise puts additional weight on variations in the separation rate. As Figure 14 suggests, variations in the level, the persistence and the volatility of the separation rate are the main factor underlying the non-linearity and the time variation in the Beveridge curve relationship. We regard this as a crucial starting point for future research.

[4] The specification of the prior follows Lubik (2012). Posterior estimates and additional results are available from the authors upon request.

7 Conclusion

We have used a Bayesian time-varying parameter structural VAR with stochastic volatility to investigate changes in both the reduced-form relationship between vacancies and the unemployment rate, and in their relationship conditional on permanent and transitory output shocks, for the post-WWII United States. Evidence points towards both similarities and differences between the Great Recession and the Volcker disinflation, and a widespread time-variation along two key dimensions. First, the slope of the Beveridge curve, as captured by the average cross-spectral gain between vacancies and the unemployment rate at business-cycle frequencies, exhibits a large extent of variation since the second half of the 1960s. Moreover, it is broadly pro-cyclical, with the gain being positively correlated with the transitory component of output. The evolution of the slope of the Beveridge curve during the Great Recession appears to be very similar to its evolution during the Volcker recession in terms of its magnitude and its time profile. Second, both the Great Inflation episode, and the subsequent Volcker disinflation, are characterized by a significantly larger (in absolute value) negative correlation between the reduced-form innovations to vacancies and the unemployment rate than the rest of the sample period. Those years also exhibit a greater cross-spectral coherence between the two series at the business-cycle frequencies, thus pointing towards them being driven, to a larger extent than the rest of the sample, by common shocks.

Appendix

A The Data

The series for real GDP ('GDPC96, Real Gross Domestic Product, 3 Decimal, Seasonally Adjusted Annual Rate, Quarterly, Billions of Chained 2005 Dollars') is from the U.S. Department of Commerce: Bureau of Economic Analysis. It is collected at quarterly frequency and seasonally adjusted. A quarterly seasonally adjusted series for the unemployment rate has been computed by converting the series UNRATE ('Civilian Unemployment Rate, Seasonally Adjusted, Monthly, Percent, Persons 16 years of age and older') from the U.S. Department of Labor: Bureau of Labor Statistics to quarterly frequency (by taking averages within the quarter). A monthly seasonally adjusted series for the vacancy rate has been computed as the ratio between the 'Help Wanted Index' (HWI) and the civilian labor force. The HWI index is from the Conference Board up until 1994Q4, and from Barnichon (2010) after that. The labor force series is from the U.S. Department of Labor: Bureau of Labor Statistics ('CLF16OV, Civilian Labor Force, Persons 16 years of age and older, Seasonally Adjusted, Monthly, Thousands of Persons'). The monthly seasonally adjusted series for the vacancy rate has been converted to the quarterly frequency by taking averages within the quarter.

B Deconvoluting the Probability Density Function of $\hat{\lambda}$

This appendix describes the procedure we use in section 2 to deconvolute the probability density function of $\hat{\lambda}$. We consider the construction of a $(1-\alpha)\%$ confidence interval for $\hat{\lambda}$, $[\hat{\lambda}^L_{(1-\alpha)}, \hat{\lambda}^U_{(1-\alpha)}]$. We assume for simplicity that λ_j and $\hat{\lambda}$ can take any value over $[0;\infty)$. Given the duality between hypothesis testing and the construction of confidence intervals, the $(1-\alpha)\%$ confidence set for $\hat{\lambda}$ comprises all the values of λ_j that cannot be rejected based on a two-sided test at the $\alpha\%$ level. Given that an increase in λ_j automatically shifts the probability density function (pdf) of $\hat{L}_j$ conditional on λ_j upwards, $\hat{\lambda}^L_{(1-\alpha)}$ and $\hat{\lambda}^U_{(1-\alpha)}$ are therefore such that:

$$P\left(\hat{L}_j > \hat{L} \mid \lambda_j = \hat{\lambda}^L_{(1-\alpha)}\right) = \alpha/2, \tag{B1}$$

and

$$P\left(\hat{L}_j < \hat{L} \mid \lambda_j = \hat{\lambda}^U_{(1-\alpha)}\right) = \alpha/2. \tag{B2}$$

Let $\phi_{\hat{\lambda}}(\lambda_j)$ and $\Phi_{\hat{\lambda}}(\lambda_j)$ be the pdf and, respectively, the cumulative pdf of $\hat{\lambda}$, defined over the domain of λ_j. The fact that $[\hat{\lambda}^L_{(1-\alpha)}, \hat{\lambda}^U_{(1-\alpha)}]$ is a $(1-\alpha)\%$ confidence interval automatically implies that $(1-\alpha)\%$ of the probability mass of $\phi_{\hat{\lambda}}(\lambda_j)$ lies between $\hat{\lambda}^L_{(1-\alpha)}$ and $\hat{\lambda}^U_{(1-\alpha)}$. This, in turn, implies that $\Phi_{\hat{\lambda}}(\hat{\lambda}^L_{(1-\alpha)}) = \alpha/2$ and $\Phi_{\hat{\lambda}}(\hat{\lambda}^U_{(1-\alpha)}) = 1-\alpha/2$. Given that this holds for any $0 < \alpha < 1$, we therefore have that:

$$\Phi_{\hat{\lambda}}(\lambda_j) = P\left(\hat{L}_j > \hat{L} \mid \lambda_j\right). \tag{B3}$$

Based on the *exp*-Wald test statistic, $\hat{L}$, and on the simulated distributions of the $\hat{L}_j$'s conditional on the λ_j's in Λ, we thus obtain an estimate of the cumulative pdf of $\hat{\lambda}$ over the grid Λ, $\hat{\Phi}_{\hat{\lambda}}(\lambda_j)$. Finally, we fit a logistic function to $\hat{\Phi}_{\hat{\lambda}}(\lambda_j)$ via nonlinear least squares and we compute the implied estimate of $\phi_{\hat{\lambda}}(\lambda_j)$, $\hat{\phi}_{\hat{\lambda}}(\lambda_j)$, whereby we scale its elements so that they sum to one.

C Details of the Markov-Chain Monte Carlo Procedure

We estimate (4)-(12) using Bayesian methods. The next two subsections describe our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data. The third section discusses how we check for convergence of the Markov chain to the ergodic distribution.

C.1 Priors

The prior distributions for the initial values of the states, θ_0 and h_0, which we postulate to be normally distributed, are assumed to be independent both from each other and from the distribution of the hyperparameters. In order to calibrate the prior distributions for θ_0 and h_0 we estimate a time-invariant version of (4) based on the first 15 years of data. We set:

$$\theta_0 \sim \mathcal{N}\left[\hat{\theta}_{OLS}, 4 \cdot \hat{V}(\hat{\theta}_{OLS})\right], \tag{C1}$$

where $\hat{V}(\hat{\theta}_{OLS})$ is the estimated asymptotic variance of $\hat{\theta}_{OLS}$. As for h_0, we proceed as follows. Let $\hat{\Sigma}_{OLS}$ be the estimated covariance matrix of ϵ_t from the time-invariant VAR, and let C be its lower-triangular Cholesky factor, $CC' = \hat{\Sigma}_{OLS}$. We set:

$$\ln h_0 \sim \mathcal{N}(\ln \mu_0, 10 \times I_N), \tag{C2}$$

where μ_0 is a vector collecting the logarithms of the squared elements on the diagonal of C. As stressed by Cogley and Sargent (2002), this prior is weakly informative for h_0.

Turning to the hyperparameters, we postulate independence between the parameters corresponding to the two matrices Q and A for convenience. Further, we make the following standard assumptions. The matrix Q is postulated to follow an inverted Wishart distribution:

$$Q \sim \mathcal{IW}\left(\bar{Q}^{-1}, T_0\right), \tag{C3}$$

with prior degrees of freedom T_0 and scale matrix $T_0\bar{Q}$. In order to minimize the impact of the prior, we set T_0 equal to the minimum value allowed, the length of θ_t plus one. As for $\bar{Q}$, we calibrate it as $\bar{Q} = \gamma \times \hat{\Sigma}_{OLS}$, setting $\gamma = 1.0 \times 10^{-4}$, as in Cogley and Sargent (2002). This is a comparatively conservative prior in the sense of allowing little random-walk drift. We note, however, that it is smaller than the median-unbiased estimates of the extent of random-walk drift discussed in section 2, ranging between 0.0235 and 0.0327 for the equation for the vacancy rate, and between 0.0122 and 0.0153 for the equation for the unemployment rate. As for α, we postulate it to be normally distributed with a large variance:

$$f\left(\alpha\right) = \mathcal{N}(0, 10000 \cdot I_{N(N-1)/2}). \tag{C4}$$

Finally, we follow Cogley and Sargent (2002, 2005) and postulate an inverse-Gamma distribution for $\sigma_i^2 \equiv Var(\nu_{i,t})$ for the variances of the stochastic volatility innovations:

$$\sigma_i^2 \sim \mathcal{IG}\left(\frac{10^{-4}}{2}, \frac{1}{2}\right). \tag{C5}$$

C.2 Simulating the Posterior Distribution

We simulate the posterior distribution of the hyperparameters and the states conditional on the data using the following MCMC algorithm (see Cogley and Sargent, 2002). x^t denotes the entire history of the vector x up to time t, that is, $x^t \equiv [x_1', x_2', \quad , x_t']'$, while T is the sample length.

1. *Drawing the elements of* θ_t: Conditional on Y^T, α, and H^T, the observation equation (4) is linear with Gaussian innovations and a known covariance matrix. Following Carter and Kohn (1994), the density $p(\theta^T|Y^T, \alpha, H^T)$ can be factored as:

$$p(\theta^T|Y^T, \alpha, H^T) = p(\theta_T|Y^T, \alpha, H^T) \prod_{t=1}^{T-1} p(\theta_t|\theta_{t+1}, Y^T, \alpha, H^T). \tag{C6}$$

   Conditional on α and H^T, the standard Kalman filter recursions determine the first element on the right hand side of (C6), $p(\theta_T|Y^T, \alpha, H^T) = N(\theta_T, P_T)$, with P_T being the precision matrix of θ_T produced by the Kalman filter. The remaining elements in the factorization can then be computed via the backward recursion algorithm found in Cogley and Sargent (2005). Given the conditional normality of θ_t, we have:

$$\theta_{t|t+1} = \theta_{t|t} + P_{t|t}P_{t+1|t}^{-1}\left(\theta_{t+1} - \theta_t\right), \tag{C7}$$

and

$$P_{t|t+1} = P_{t|t} - P_{t|t}P_{t+1|t}^{-1}P_{t|t}, \tag{C8}$$

which provides, for each t from $T-1$ to 1, the remaining elements in (4), $p(\theta_t|\theta_{t+1}, Y^T, \alpha, H^T) = N(\theta_{t|t+1}, P_{t|t+1})$. Specifically, the backward recursion starts with a draw from $\mathcal{N}(\theta_T, P_T)$, call it $\tilde{\theta}_T$. Conditional on $\tilde{\theta}_T$, (C7)-(C8) give us $\theta_{T-1|T}$ and $P_{T-1|T}$, thus allowing us to draw $\tilde{\theta}_{T-1}$ from $N(\theta_{T-1|T}, P_{T-1|T})$, and so on until $t = 1$.

2. *Drawing the elements of* H_t: Conditional on Y^T, θ^T, and α, the orthogonalized innovations $u_t \equiv A(Y_t - X_t'\theta_t)$, with $Var(u_t) = H_t$, are observable. Following Cogley and Sargent (2002), we then sample the $h_{i,t}$'s by applying the univariate algorithm of Jacquier et al. (1994) element by element.

3. *Drawing the hyperparameters*: Conditional on Y^T, θ^T, H^T, and α, the innovations to θ_t and to the $h_{i,t}$'s are observable, which allows us to draw the hyperparameters, namely the elements of Q and the σ_i^2, from their respective distributions.

4. *Drawing the elements of* α: Finally, conditional on Y^T and θ^T, the ϵ_t's are observable. They satisfy:

$$A\epsilon_t = u_t, \tag{C9}$$

with the u_t being a vector of orthogonalized residuals with known time-variying variance H_t. Following Primiceri (2005), we interpret (C9) as a system of unrelated regressions. The first equation in the system is given by $\epsilon_{1,t} \equiv u_{1,t}$, while the following equations can be expressed as transformed regressions:

$$\begin{aligned} \left(h_{2,t}^{-\frac{1}{2}}\epsilon_{2,t}\right) &= -\alpha_{2,1}\left(h_{2,t}^{-\frac{1}{2}}\epsilon_{1,t}\right) + \left(h_{2,t}^{-\frac{1}{2}}u_{2,t}\right) \\ \left(h_{3,t}^{-\frac{1}{2}}\epsilon_{3,t}\right) &= -\alpha_{3,1}\left(h_{3,t}^{-\frac{1}{2}}\epsilon_{1,t}\right) - \alpha_{3,2}\left(h_{3,t}^{-\frac{1}{2}}\epsilon_{2,t}\right) + \left(h_{3,t}^{-\frac{1}{2}}u_{3,t}\right), \end{aligned} \tag{C10}$$

where the residuals are independently standard normally distributed. Assuming normal priors for each equation's regression coefficients the posterior is also normal and can be computed as in Cogley and Sargent (2005).

Summing up, the MCMC algorithm simulates the posterior distribution of the states and the hyperparameters, conditional on the data, by iterating on (1)-(4). In what follows, we use a burn-in period of 50,000 iterations to converge to the ergodic distribution. After that, we run 10,000 more iterations sampling every 10th draw in order to reduce the autocorrelation across draws.

D A Simple Search and Matching Model of the Labor Market

The model specification follows Lubik (2012). Time is discrete and the time period is a quarter. The model economy is populated by a continuum of identical firms that employ workers, each of whom inelastically supplies one unit of labor. Output Y_t of a typical firm is linear in employment N_t:

$$Y_t = A_t N_t. \tag{D1}$$

A_t is a stochastic aggregate productivity process. It is composed of a permanent productivity shock, A_t^P, which follows a random walk, and a transitory productivity shock, A_t^T, which is an AR(1)-process. Specifically, we assume that $A_t = A_t^P A_t^T$.

The labor market matching process combines unemployed job seekers U_t with job openings (vacancies) V_t. This can be represented by a constant returns matching function, $M_t = m_t U_t^{\xi} V_t^{1-\xi}$, where m_t is stochastic match efficiency, and $0 < \xi < 1$ is the match elasticity. Unemployment is defined as those workers who are not currently employed:

$$U_t = 1 - N_t, \tag{D2}$$

where the labor force is normalized to one. Inflows to unemployment arise from job destruction at rate $0 < \rho_t < 1$, which can vary over time. The dynamics of employment are thus governed by the following relationship:

$$N_t = (1 - \rho_t) \left[N_{t-1} + m_{t-1} U_{t-1}^{\xi} V_{t-1}^{1-\xi} \right]. \tag{D3}$$

This is a stock-flow identity that relates the stock of employed workers N_t to the flow of new hires $M_t = m_t U_t^{\xi} V_t^{1-\xi}$ into employment. The timing assumption is such that once a worker is matched with a firm, the labor market closes. This implies that if a newly hired worker and a firm separate, the worker cannot re-enter the pool of searchers immediately and has to wait one period before searching again.

The matching function can be used to define the job finding rate, i.e., the probability that a worker will be matched with a firm:

$$p(\theta_t) = \frac{M_t}{U_t} = m_t \theta_t^{1-\xi}, \tag{D4}$$

and the job matching rate, i.e., the probability that a firm is matched with a worker:

$$q(\theta_t) = \frac{M_t}{V_t} = m_t \theta_t^{-\xi}, \tag{D5}$$

where $\theta_t = V_t/U_t$ is labor market tightness. From the perspective of an individual firm, the aggregate match probability $q(\theta_t)$ is exogenous and unaffected by individual decisions. Hence, for individual firms new hires are linear in the number of vacancies posted: $M_t = q(\theta_t)V_t$.

A firm chooses the optimal number of vacancies V_t to be posted and its employment level N_t by maximizing the intertemporal profit function:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[A_t N_t - W_t N_t - \kappa_t V_t \right], \tag{D6}$$

subject to the employment accumulation equation (D3). Profits are discounted at rate $0 < \beta < 1$. Wages paid to the workers are W_t, while $\kappa_t > 0$ is a firm's time-varying cost of opening a vacancy. The first-order conditions are:

$$N_t \quad : \qquad \mu_t = A_t - W_t + \beta E_t \left[(1-\rho_{t+1})\mu_{t+1} \right], \tag{D7}$$

$$V_t \quad : \qquad \kappa_t = \beta q(\theta_t) E_t \left[(1-\rho_{t+1})\mu_{t+1} \right], \tag{D8}$$

where μ_t is the multiplier on the employment equation.

Combining these two first-order conditions results in the *job creation condition* (JCC):

$$\frac{\kappa_t}{q(\theta_t)} = \beta E_t \left[(1-\rho_{t+1}) \left(A_{t+1} - W_{t+1} + \frac{\kappa_{t+1}}{q(\theta_{t+1})} \right) \right]. \tag{D9}$$

This captures the trade-off faced by the firm: the marginal effective cost of posting a vacancy, $\frac{\kappa_t}{q(\theta_t)}$, that is, the per-vacancy cost κ adjusted for the probability that the position is filled, is weighed against the discounted benefit from the match. The latter consists of the surplus generated by the production process net of wage payments to the workers, plus the benefit of not having to post a vacancy again in the next period.

In order to close the model, we assume in line with the existing literature that wages are determined based on the Nash bargaining solution: surpluses accruing to the matched parties are split according to a rule that maximizes their weighted average. Denoting the workers' weight in the bargaining process as $\eta \in [0,1]$, this implies the sharing rule:

$$\mathcal{W}_t - \mathcal{U}_t = \frac{\eta}{1-\eta}\left(\mathcal{J}_t - \mathcal{V}_t\right), \tag{D10}$$

where $\mathcal{W}_t$ is the asset value of employment, $\mathcal{U}_t$ is the value of being unemployed, $\mathcal{J}_t$ is the value of the marginal worker to the firm, and $\mathcal{V}_t$ is the value of a vacant job. By free entry, $\mathcal{V}_t$ is assumed to be driven to zero.

The value of employment to a worker is described by the following Bellman equation:

$$\mathcal{W}_t = W_t + E_t\beta[(1-\rho_{t+1})\mathcal{W}_{t+1} + \rho_{t+1}\mathcal{U}_{t+1}]. \tag{D11}$$

Workers receive the wage W_t, and transition into unemployment next period with probability ρ_{t+1}. The value of searching for a job, when the worker is currently unemployed, is:

$$\mathcal{U}_t = b_t + E_t\beta[p_t(1-\rho_{t+1})\mathcal{W}_{t+1} + (1-p_t(1-\rho_{t+1}))\mathcal{U}_{t+1}]. \tag{D12}$$

An unemployed searcher receives stochastic benefits b_t and transitions into employment with probability $p_t(1-\rho_{t+1})$. Recall that the job finding rate p_t is defined as $p(\theta_t) = M(V_t, U_t)/U_t$ which is decreasing in tightness θ_t. It is adjusted for the probability that a completed match gets dissolved before production begins next period. The marginal value of a worker $\mathcal{J}_t$ is equivalent to the multiplier on the employment equation, $\mathcal{J}_t = \mu_t$, so that the respective first-order condition defines the Bellman-equation for the value of a job. Substituting the asset equations into the sharing rule (D10) results in the wage equation:

$$W_t = \eta\left(A_t + \kappa_t\theta_t\right) + (1-\eta)b_t. \tag{D13}$$

Wage payments are a weighted average of the worker's marginal product A_t, which the worker can appropriate at a fraction η, and the outside option b_t, of which the firm obtains the portion $(1-\eta)$. Moreover, the presence of fixed vacancy posting costs leads to a hold-up problem where the worker extracts an additional $\eta\kappa_t\theta_t$ from the firm.

Finally, we can substitute the wage equation (D13) into (D9) to derive an alternative representation of the job creation condition:

$$\frac{\kappa_t}{m_t}\theta_t^{\xi} = \beta E_t(1-\rho_{t+1})\left[(1-\eta)\left(A_{t+1} - b_t\right) - \eta\kappa_t\theta_{t+1} + \frac{\kappa_t}{m_{t+1}}\theta_{t+1}^{\xi}\right]. \tag{D14}$$

References

[1] Andrews, Donald K. (1991): "Heteroskedasticity and Autocorrelation-Consistent Covariance Matrix Estimation". *Econometrica*, 59, 817-858.

[2] Barnichon, Regis (2010): "Building a Composite Help-Wanted Index". *Economics Letters*, 109, 175-178.

[3] Benati, Luca (2007): "Drifts and Breaks in Labor Productivity". *Journal of Economic Dynamics and Control*, 31(8), 2847-2877.

[4] Blanchard, Olivier J., and Peter Diamond (1989): "The Beveridge Curve". *Brookings Papers on Economic Activity*, 1, 1-60.

[5] Blanchard, Olivier J., and Danny Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances". *American Economic Review*, 79(4), 655-673.

[6] Carter, Charles K., and Robert Kohn (1994): "On Gibbs Sampling for State Space Models". *Biometrika*, 81(3), 541-553.

[7] Cogley, Timothy, and Thomas J. Sargent (2002): "Evolving Post-World War II U.S. Inflation Dynamics". *NBER Macroeconomics Annual*, 16, 331-388.

[8] Cogley, Timothy, and Thomas J. Sargent (2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US". *Review of Economic Dynamics*, 8, 262-302.

[9] Fernández-Villaverde, Jesús and Juan Rubio-Ramirez (2007): "How Structural are Structural Parameter Values?". *NBER Macroeconomics Annual*, 22, 83-132.

[10] Furlanetto, Francesco and Nicolas Groshenny (2012): "Mismatch Shocks and Unemployment During the Great Recession". Manuscript, Norges Bank.

[11] Galí, Jordi (1999): "Technology, Employment, and the Business Cycle: Do Technology Shock Explain Aggregate Fluctuations?" *American Economic Review*, 89(1), 249-271.

[12] Galí, Jordi and Luca Gambetti (2009): "On the Sources of the Great Moderation". *American Economic Journal: Macroeconomics*, 1(1), 26-57.

[13] Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi (1994): "Bayesian Analysis of Stochastic Volatility Models". *Journal of Business and Economic Statistics*, 12(4), 371-89 .

[14] Lubik, Thomas A. (2012): "The Shifting and Twisting Beveridge Curve". Manuscript, Federal Reserve Bank of Richmond.

[15] Newey, Whitney, and Kenneth West (1987): "A Simple Positive-Semi-Definite Heteroskedasticity and Autocorrelation-Consistent Covariance Matrix". *Econometrica*, 55, 703-708.

[16] Primiceri, Giorgio (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy". *Review of Economic Studies*, 72, 821-852.

[17] Rubio-Ramirez, Juan, Dan Waggoner, and Tao Zha (2005): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference". *Review of Economic Studies*, 77(2), 665-696.

[18] Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante (2012): "Mismatch Unemployment". *Federal Reserve Bank of New York Staff Reports*, No. 566.

[19] Shimer, Robert (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies". *American Economic Review*, 95, 25-49.

[20] Stock, James, and Mark M. Watson (1996): "Evidence of Structural Instability in Macroeconomic Time Series Relations". *Journal of Business and Economic Statistics*, 14(1), 11-30.

[21] Stock, James, and Mark M. Watson (1998): "Median-Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model". *Journal of the American Statistical Association*, 93(441), 349-358.

Table 1 Results based on the Stock-Watson TVP-MUB methodology: *exp*- and *sup*-Wald test statistics, simulated *p*-values, and median-unbiased estimates of λ

Equation for:	*exp*-Wald (p-value)	$\hat{\lambda}$	*sup*-Wald (p-value)	$\hat{\lambda}$
	Newey and West (1987) correction			
vacancy rate	9.40 (0.0053)	0.0286	28.91 (0.0028)	0.0327
unemployment rate	4.97 (0.1661)	0.0153	16.17 (0.1770)	0.0153
	Andrews (1991) correction			
vacancy rate	7.65 (0.0195)	0.0235	25.40 (0.0086)	0.0286
unemployment rate	4.68 (0.1987)	0.0133	14.61 (0.2594)	0.0122

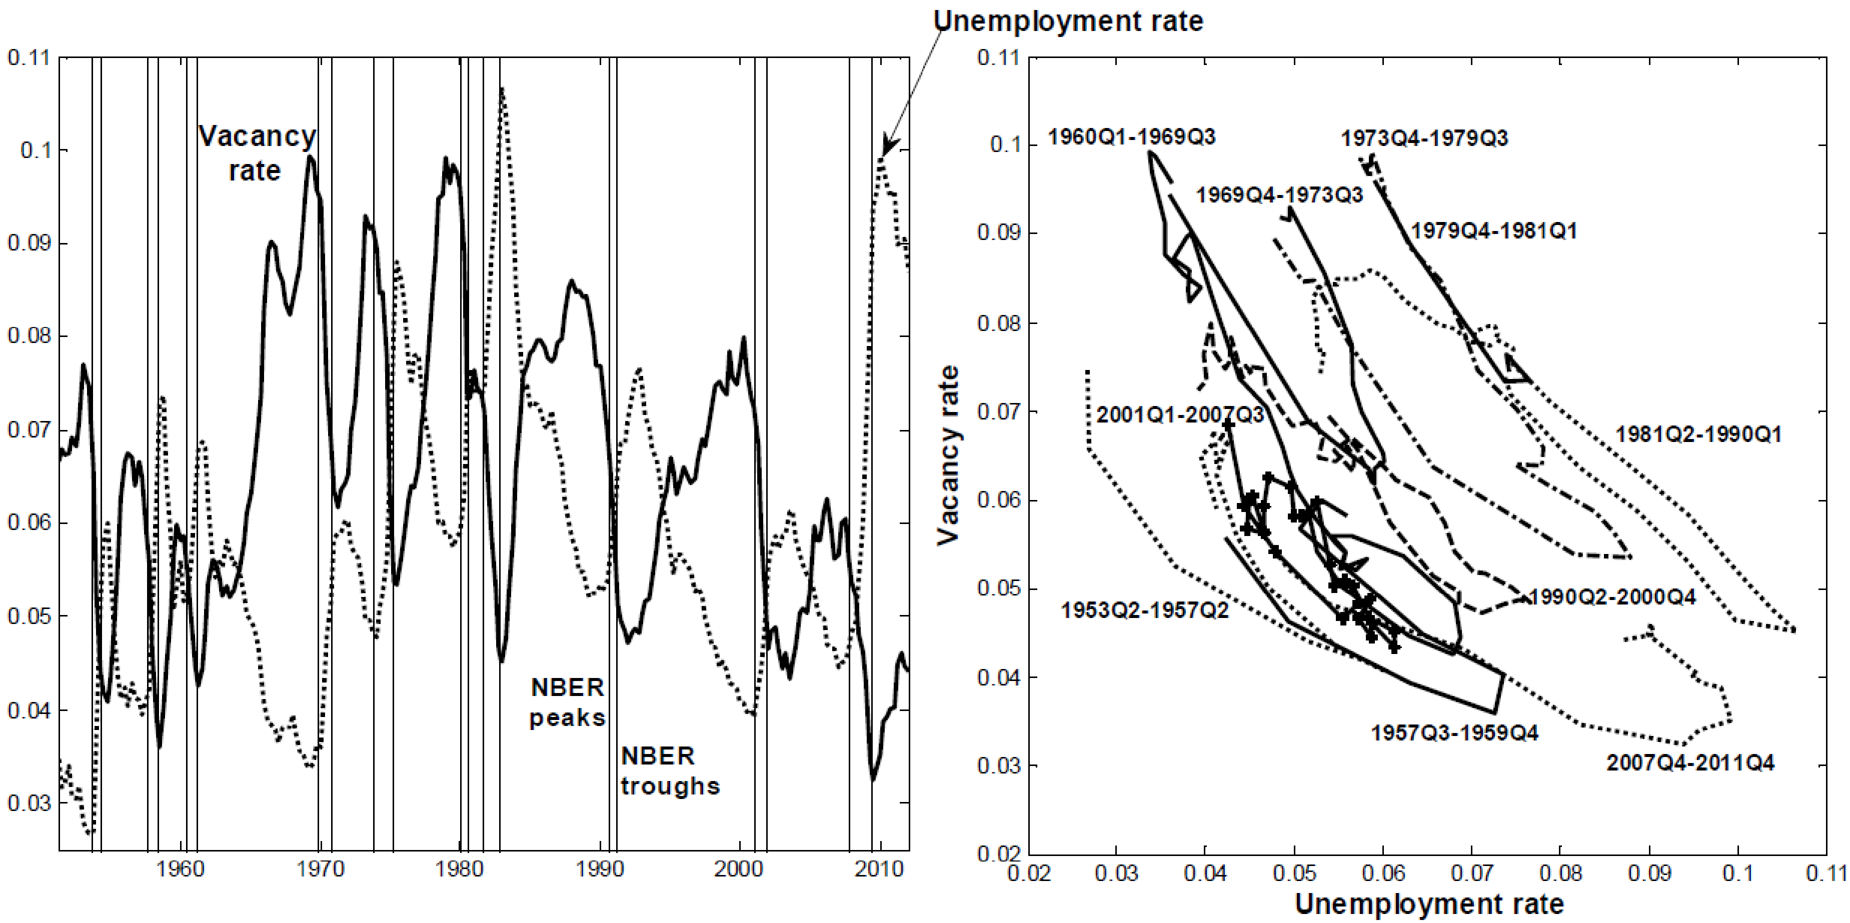

Figure 1 The unemployment rate and the vacancy rate

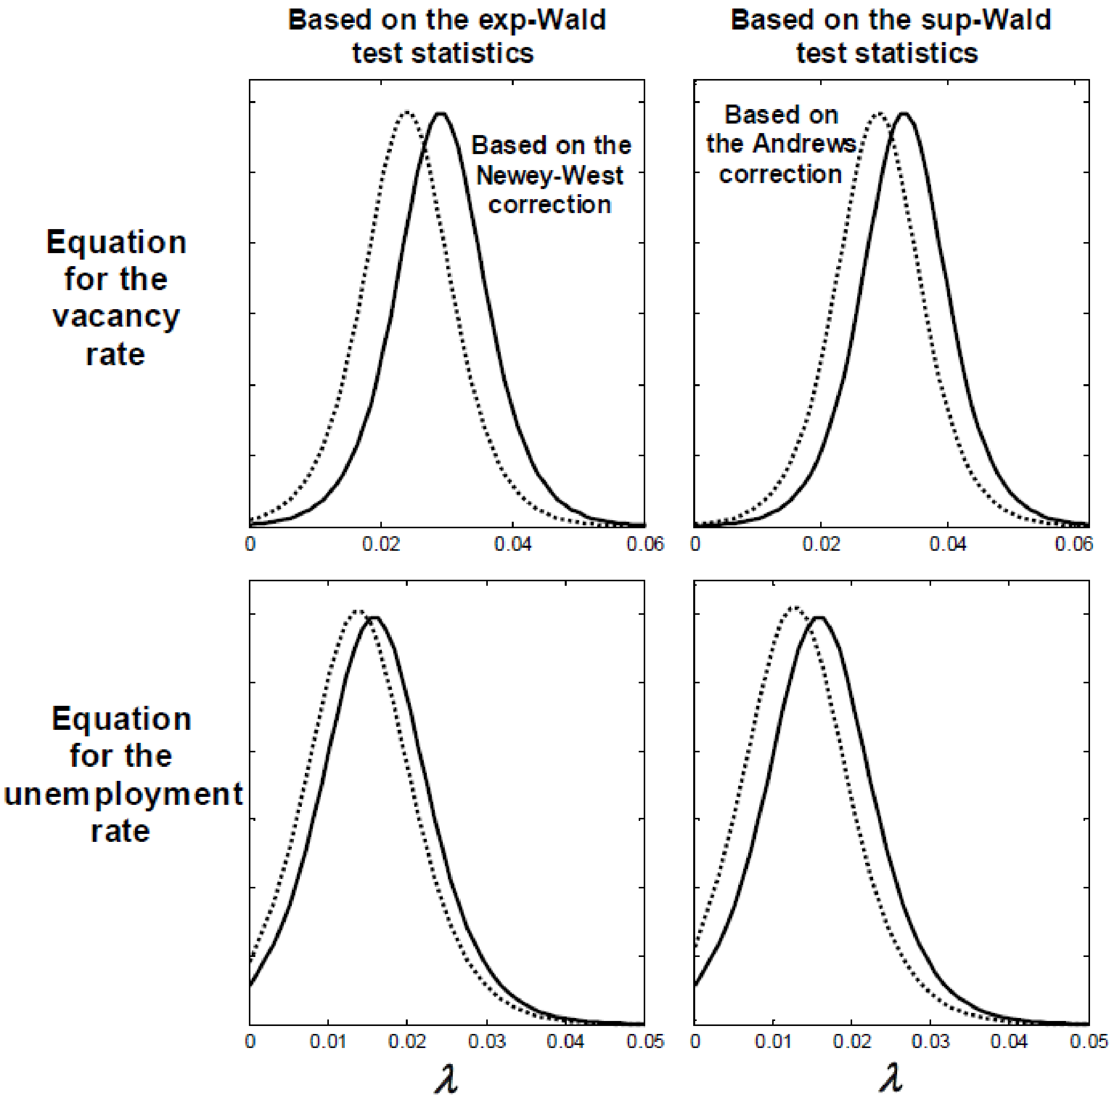

Figure 2 Deconvoluted PDFs of λ

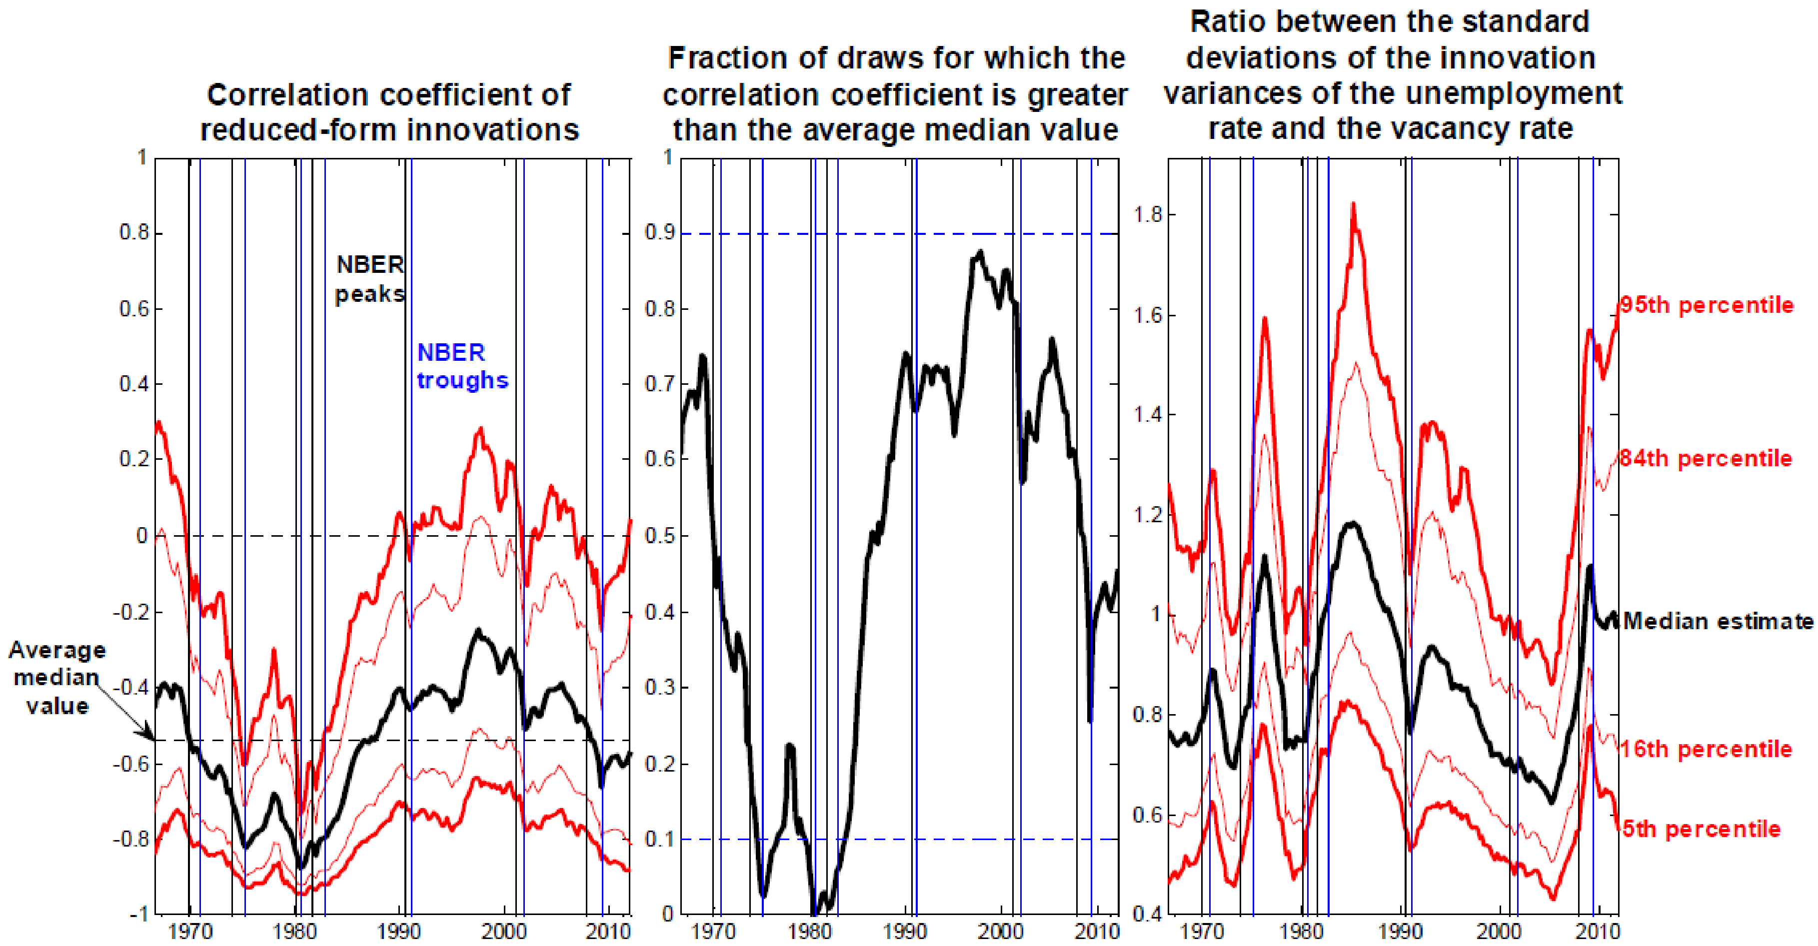

Figure 3 Correlation coefficient of reduced-form innovations to vacancies and the unemployment rate, and ratio between the standard deviations of reduced-form innovations to the two variables

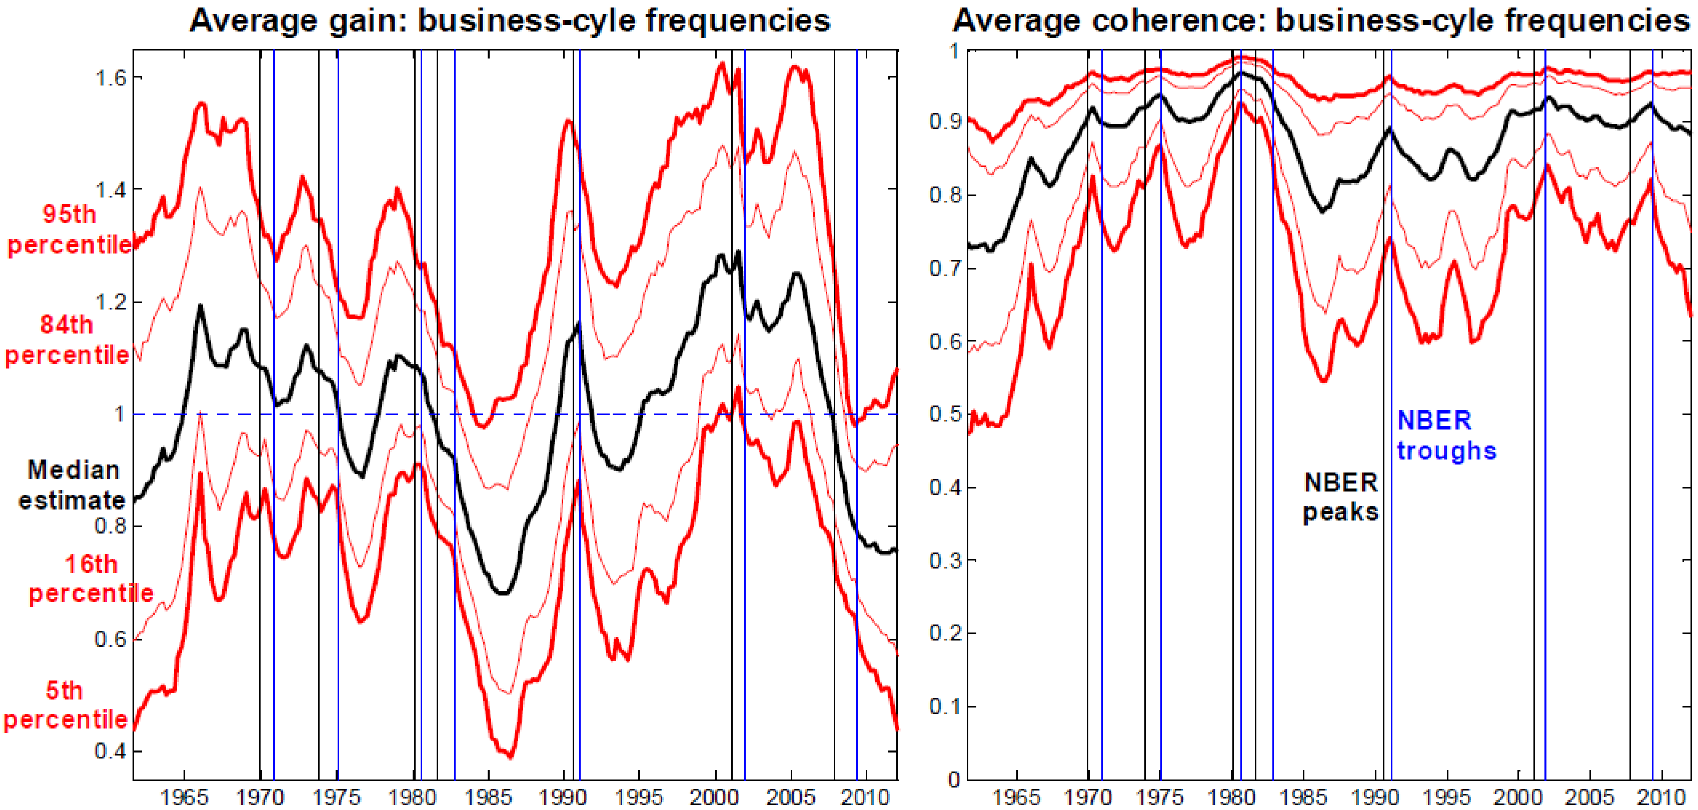

Figure 4 Average gain of the unemployment rate onto vacancies, and average coherence between vacancies and the unemployment rate, at the business-cycle frequencies

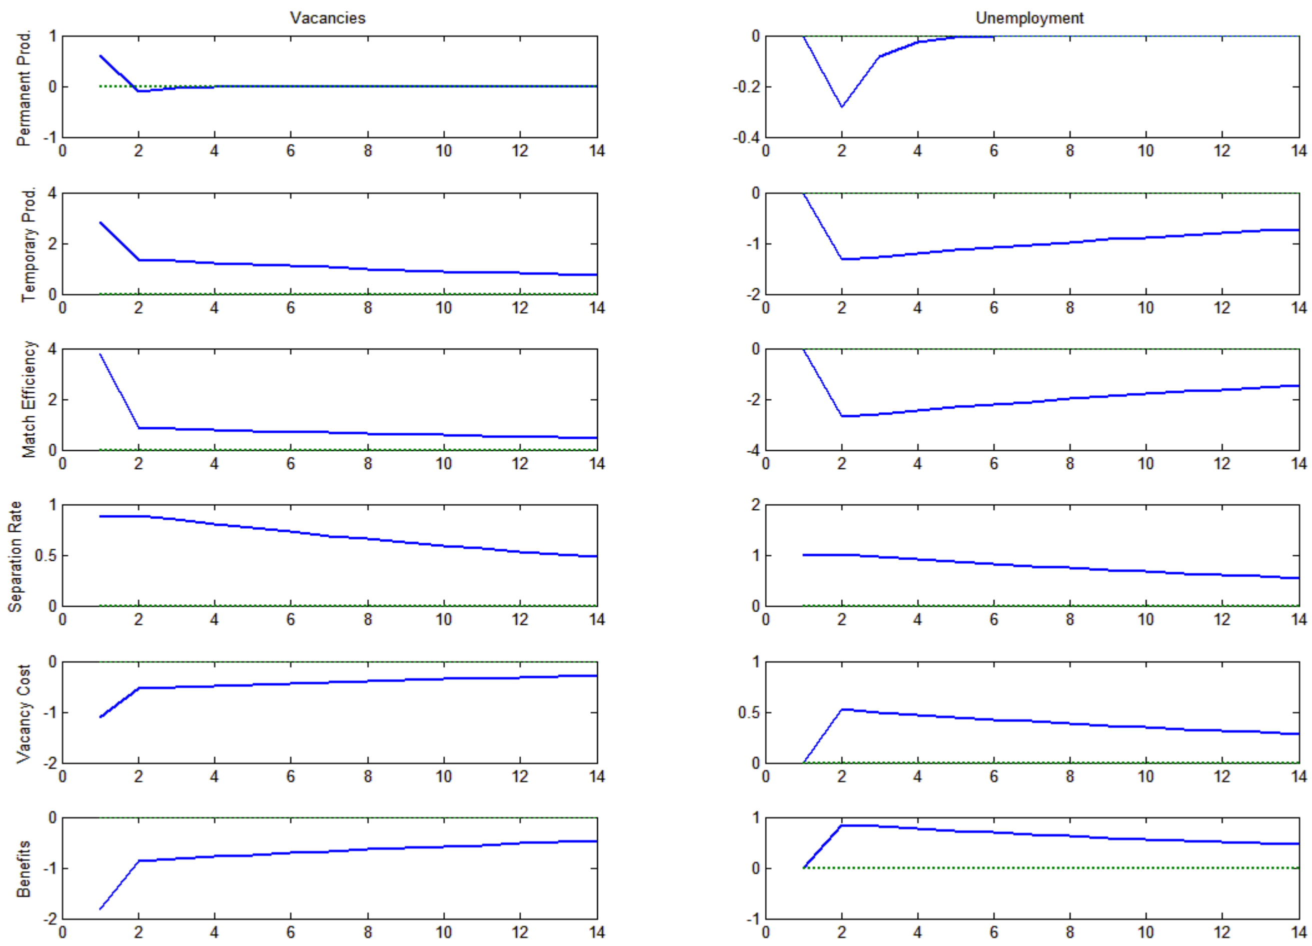

Figure 5 Impulse response functions of the calibrated search and matching model

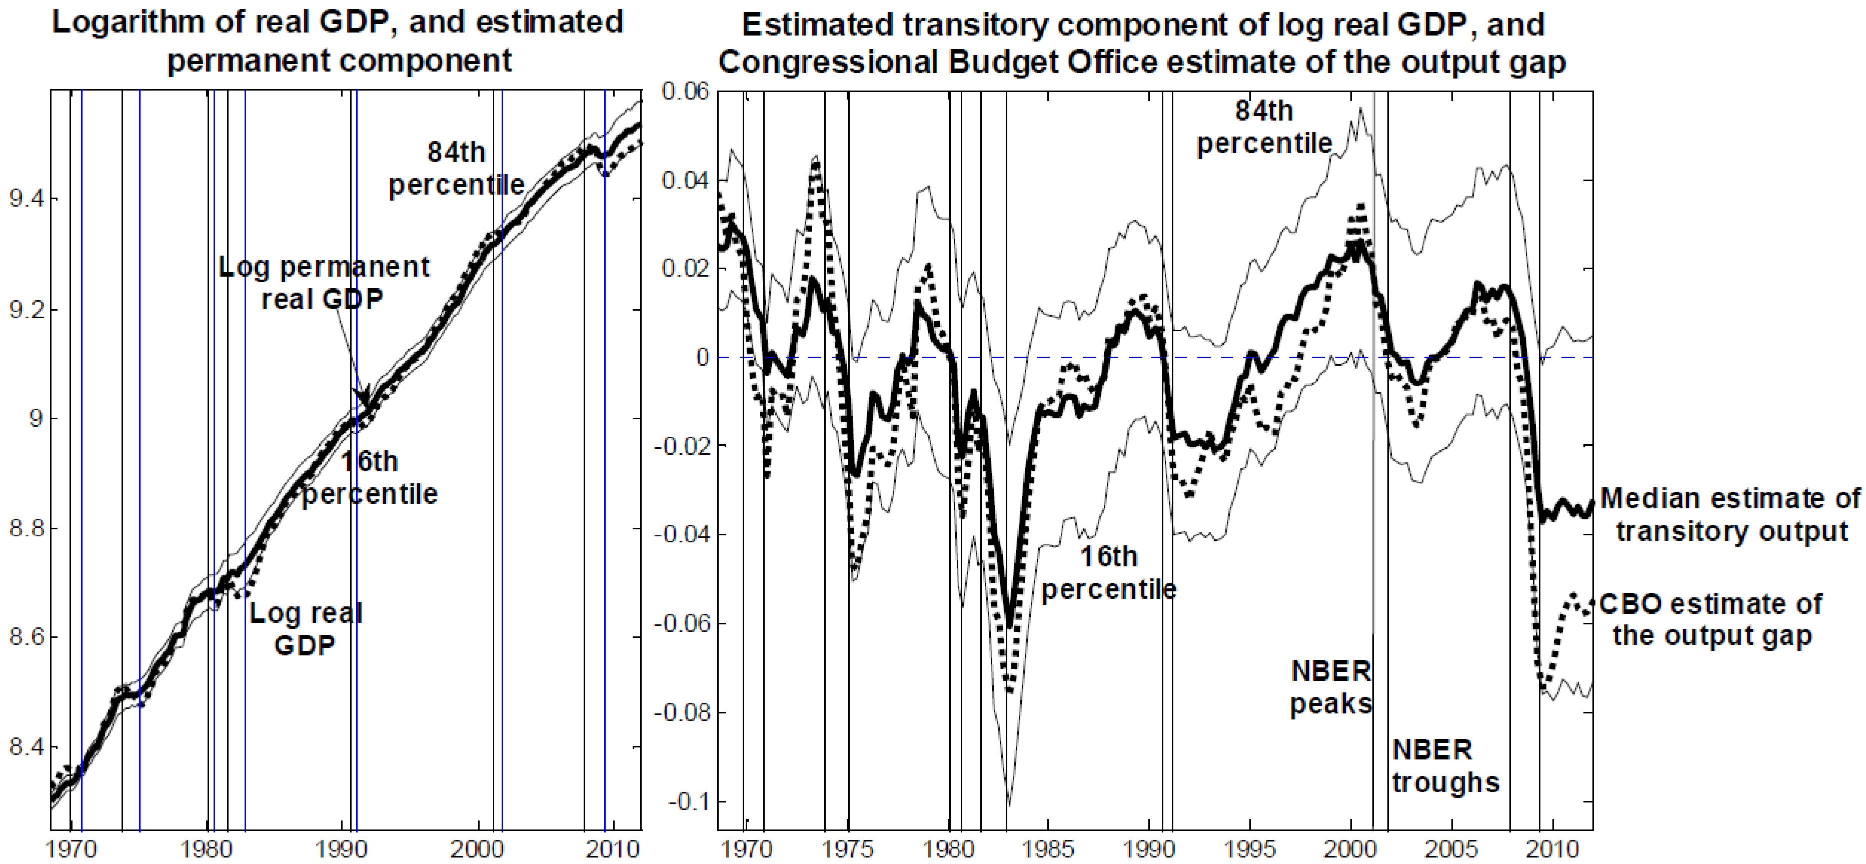

Figure 6 Logarithm of real GDP and estimated permanent component, estimated transitory component of log real GDP, and Congressional Budget Office estimate of the output gap

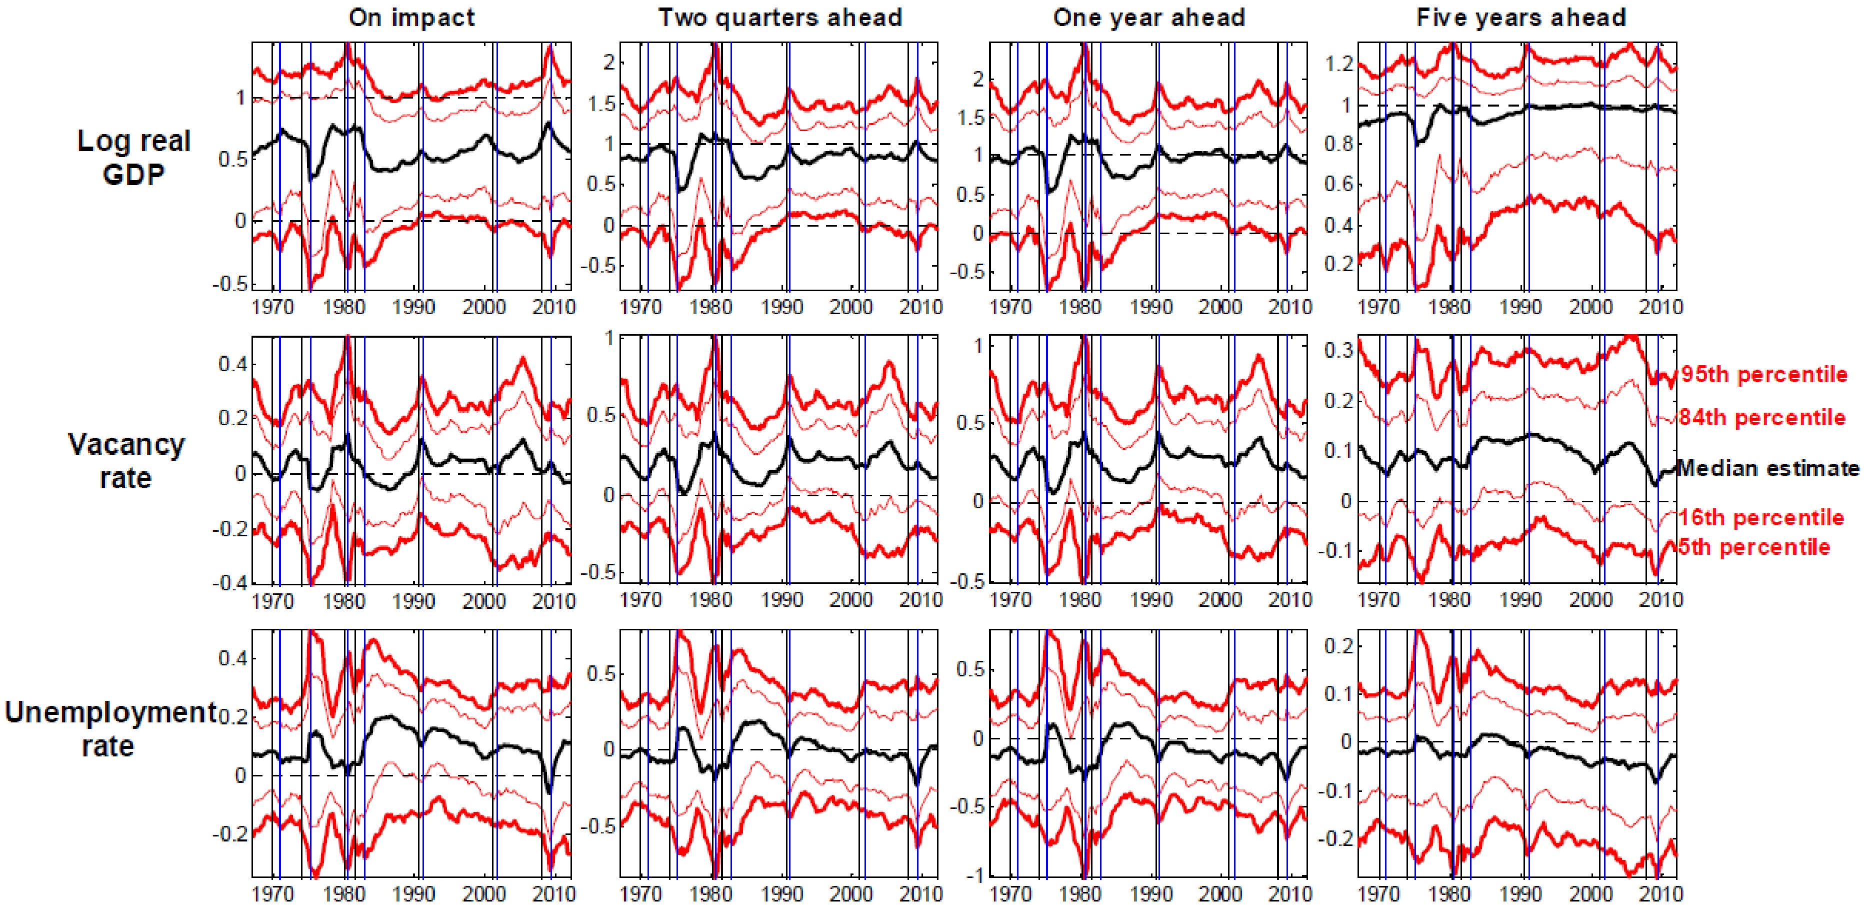

Figure 7 Impulse response functions to a permanent output shock

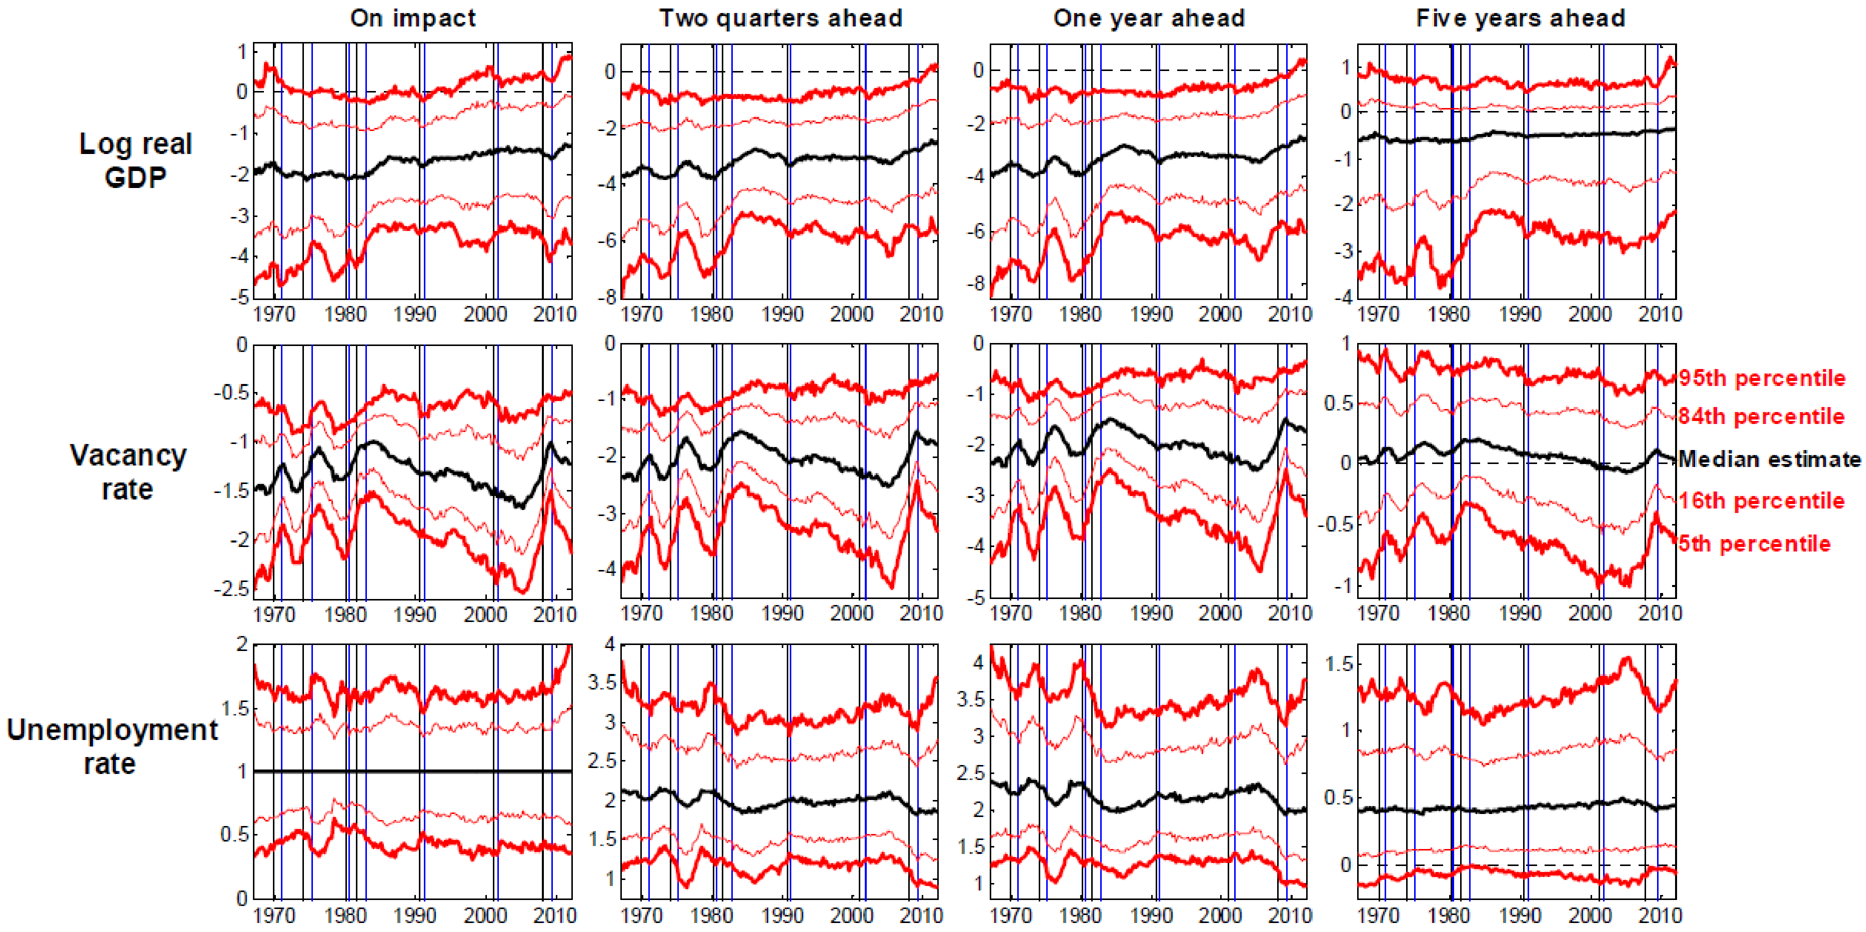

Figure 8 Impulse response functions to the first transitory shock

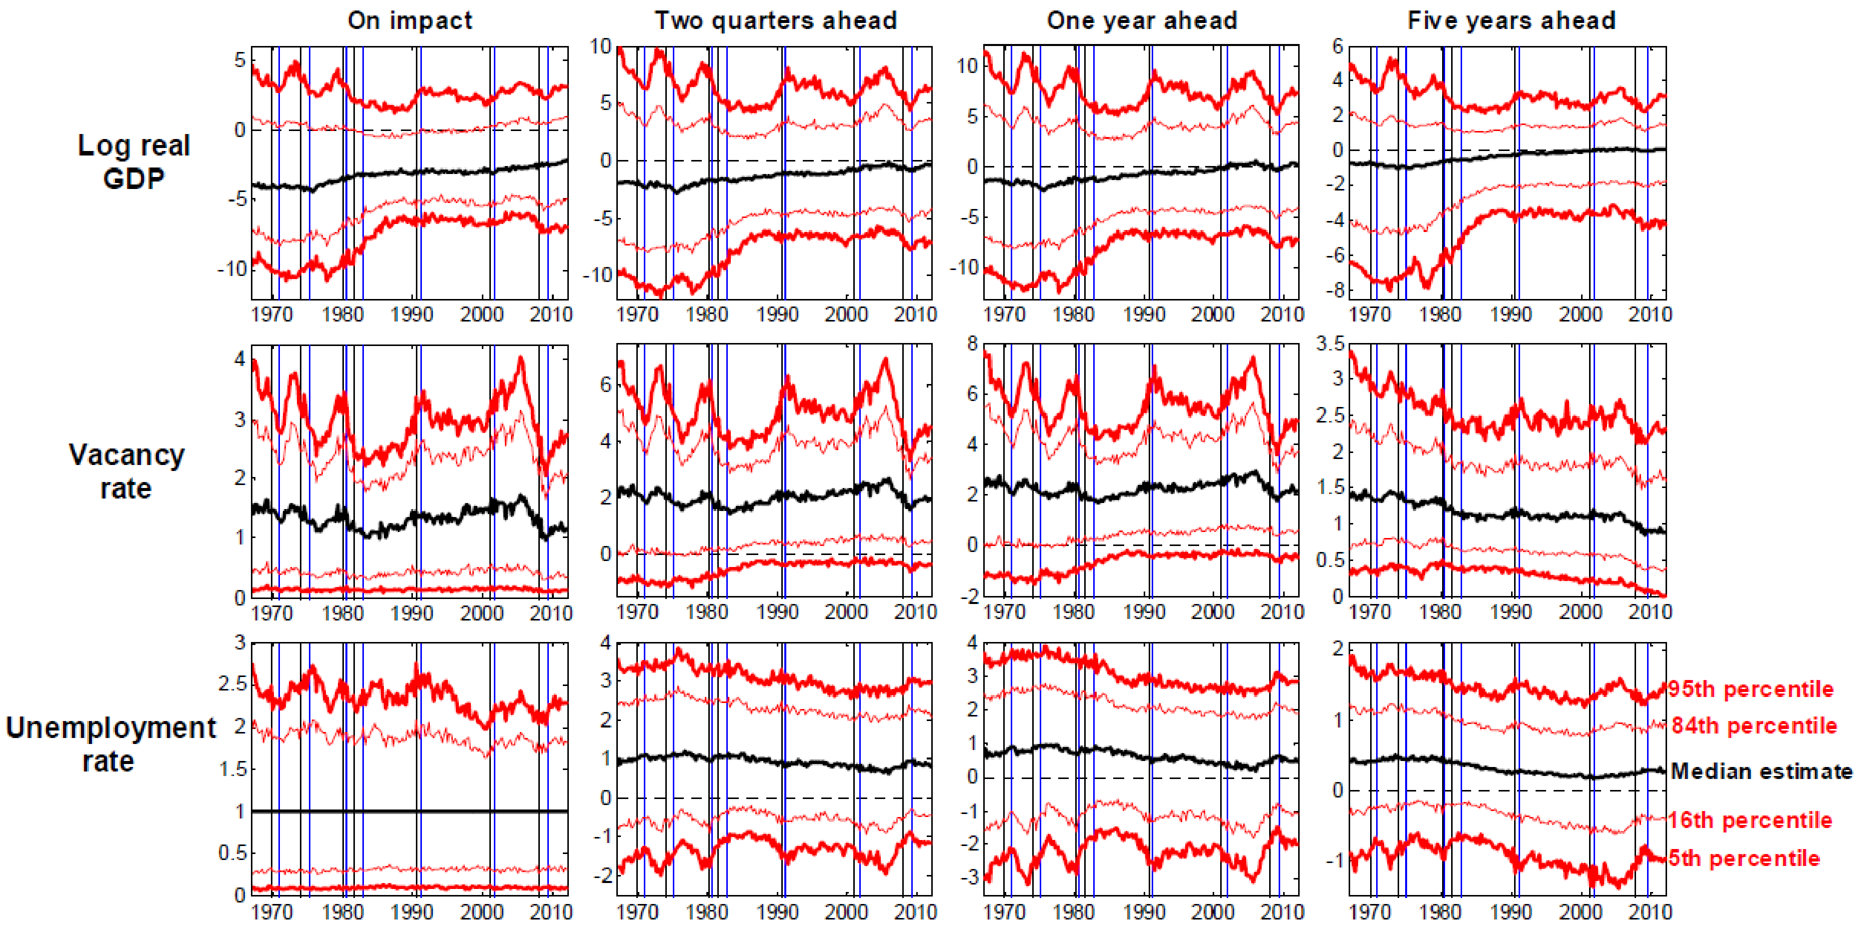

Figure 9 Impulse response functions to the second transitory shock

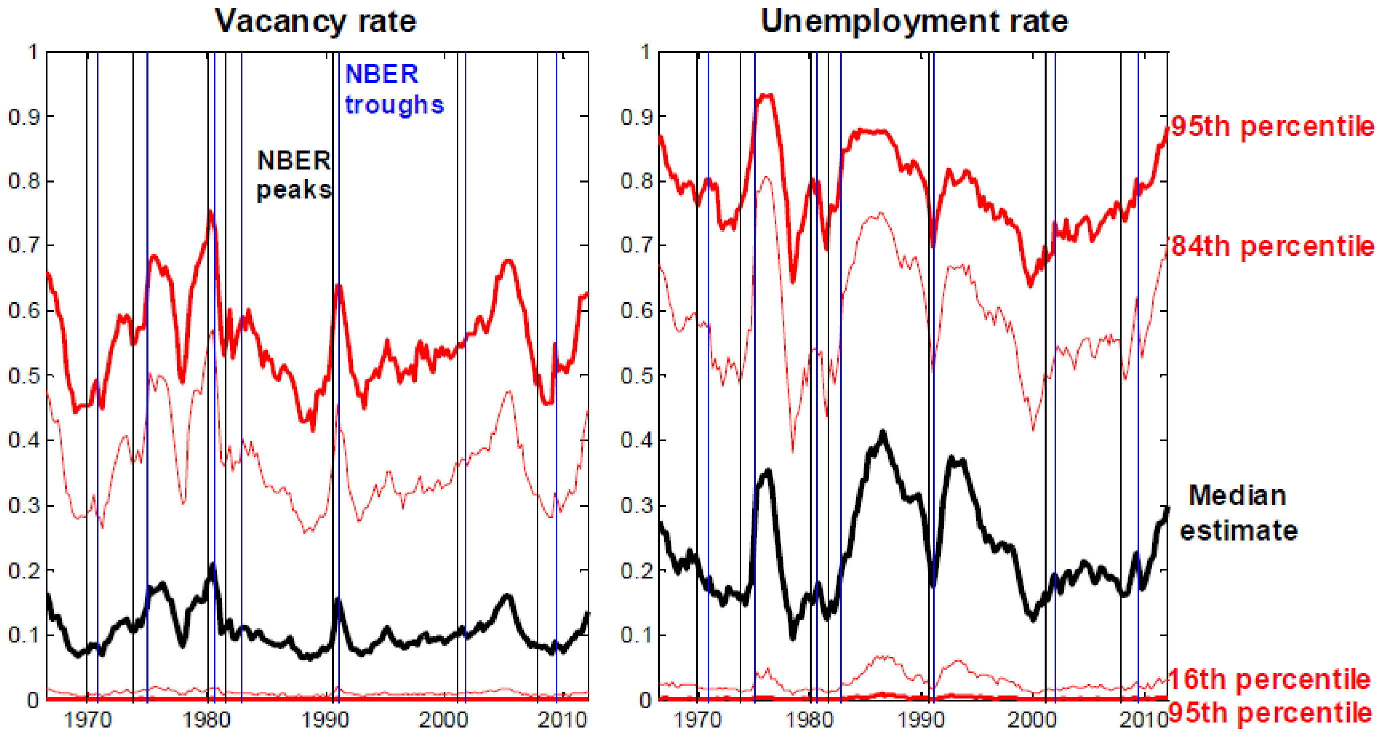

Figure 10 Fractions of innovation variance due to the permanent output shock

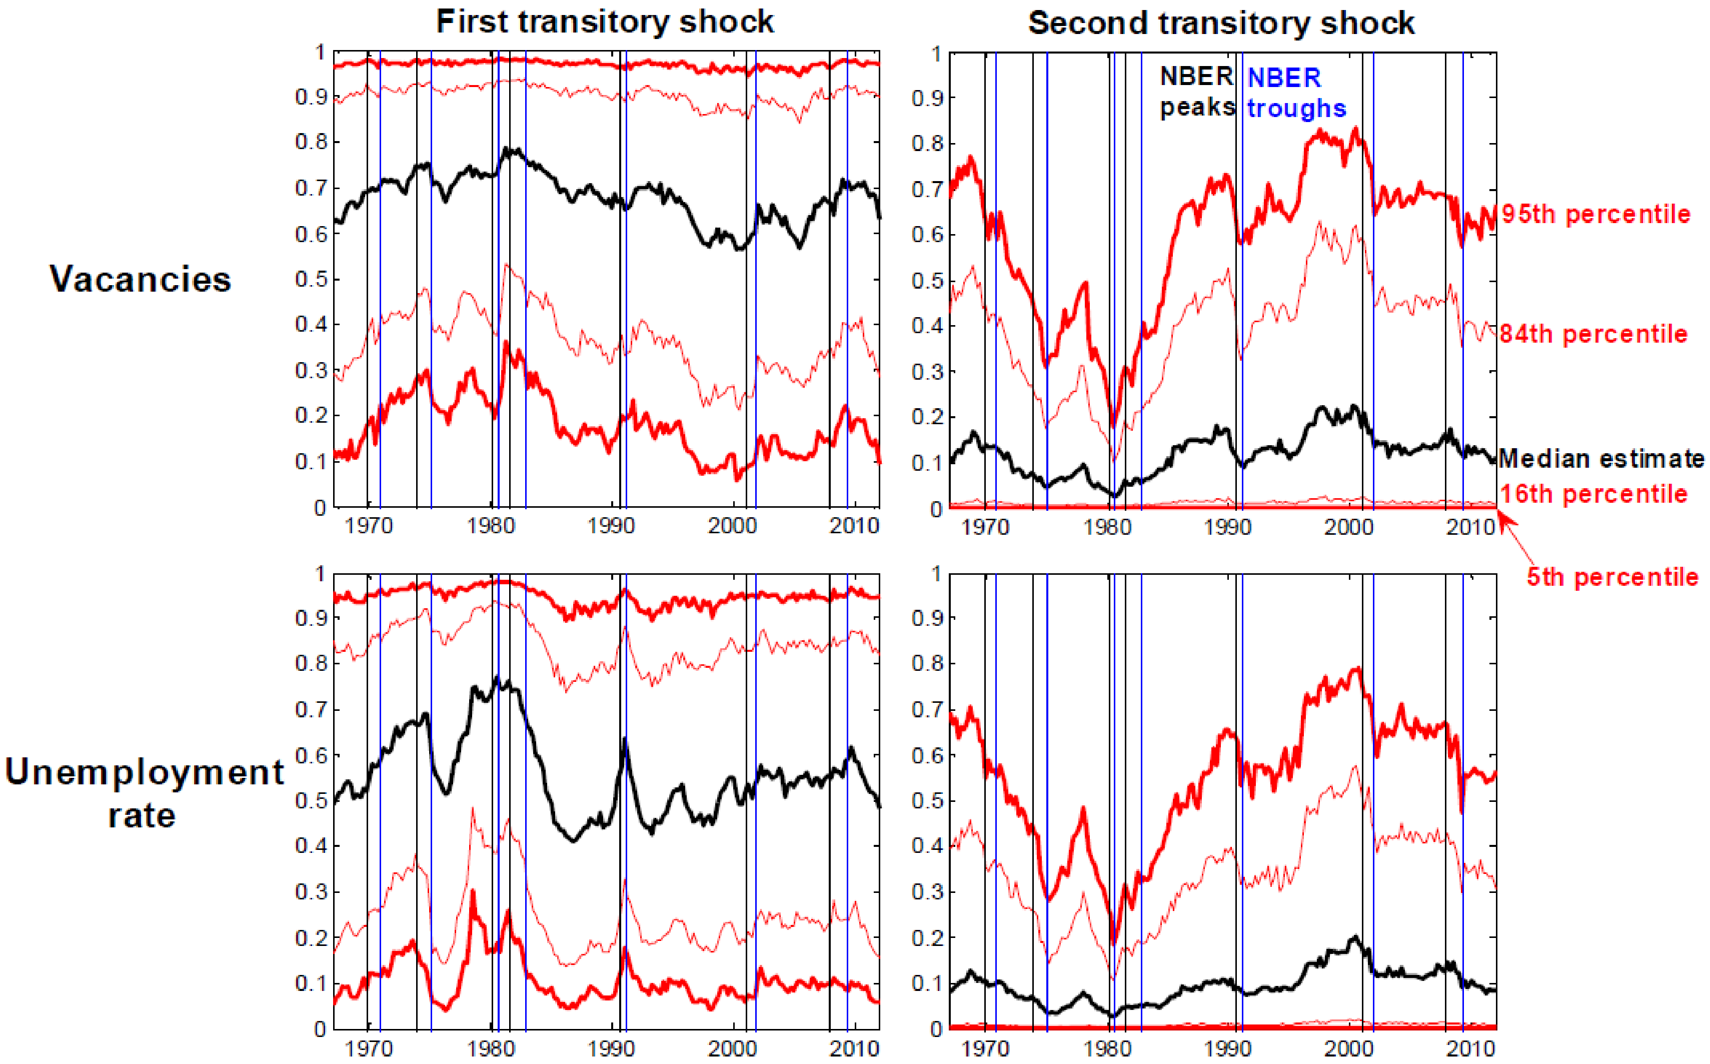

Figure 11 Fractions of innovation variance due to the two transitory shocks identified from sign restrictions

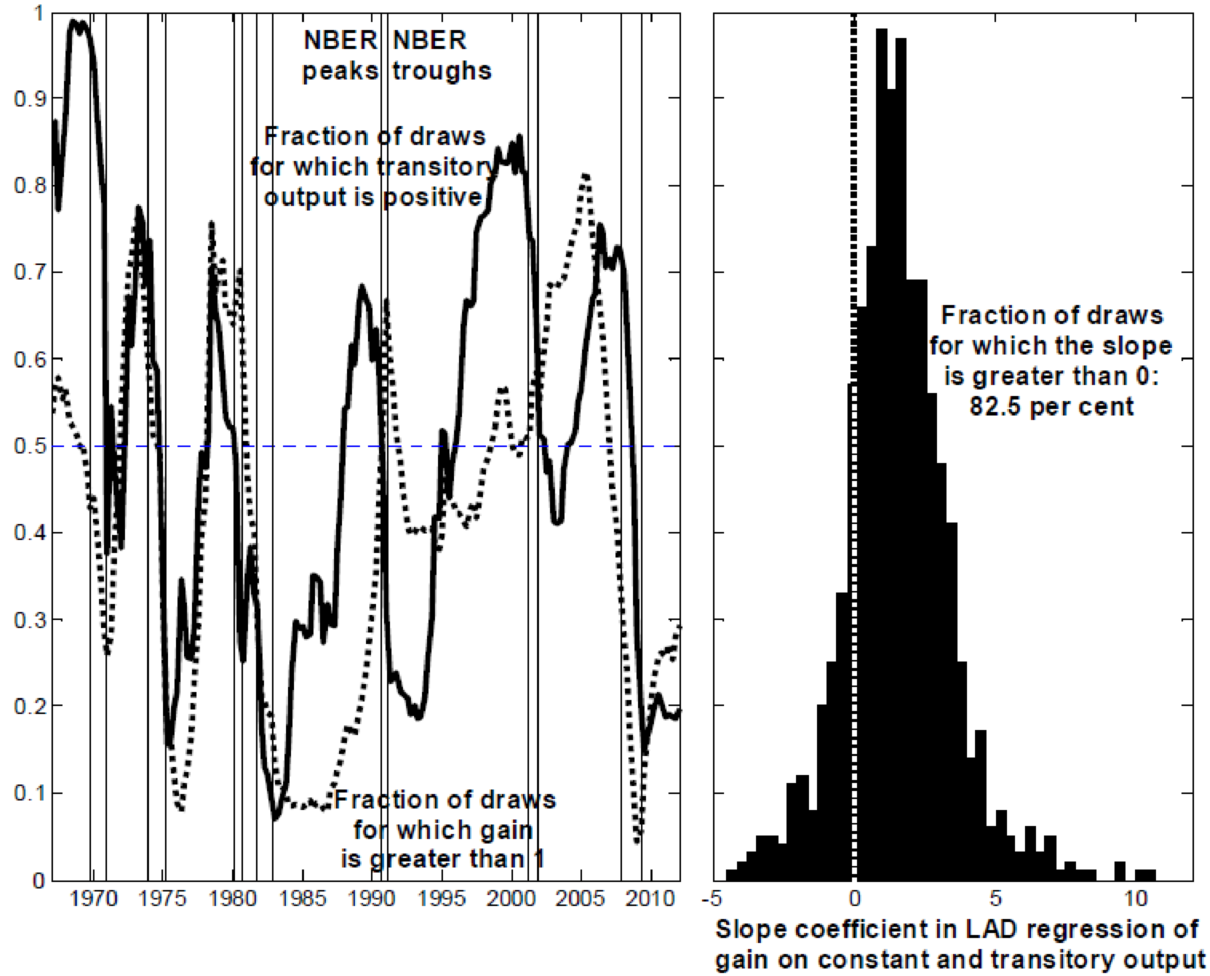

Figure 12 Evidence on the pro-cyclicality of the Beveridge curve

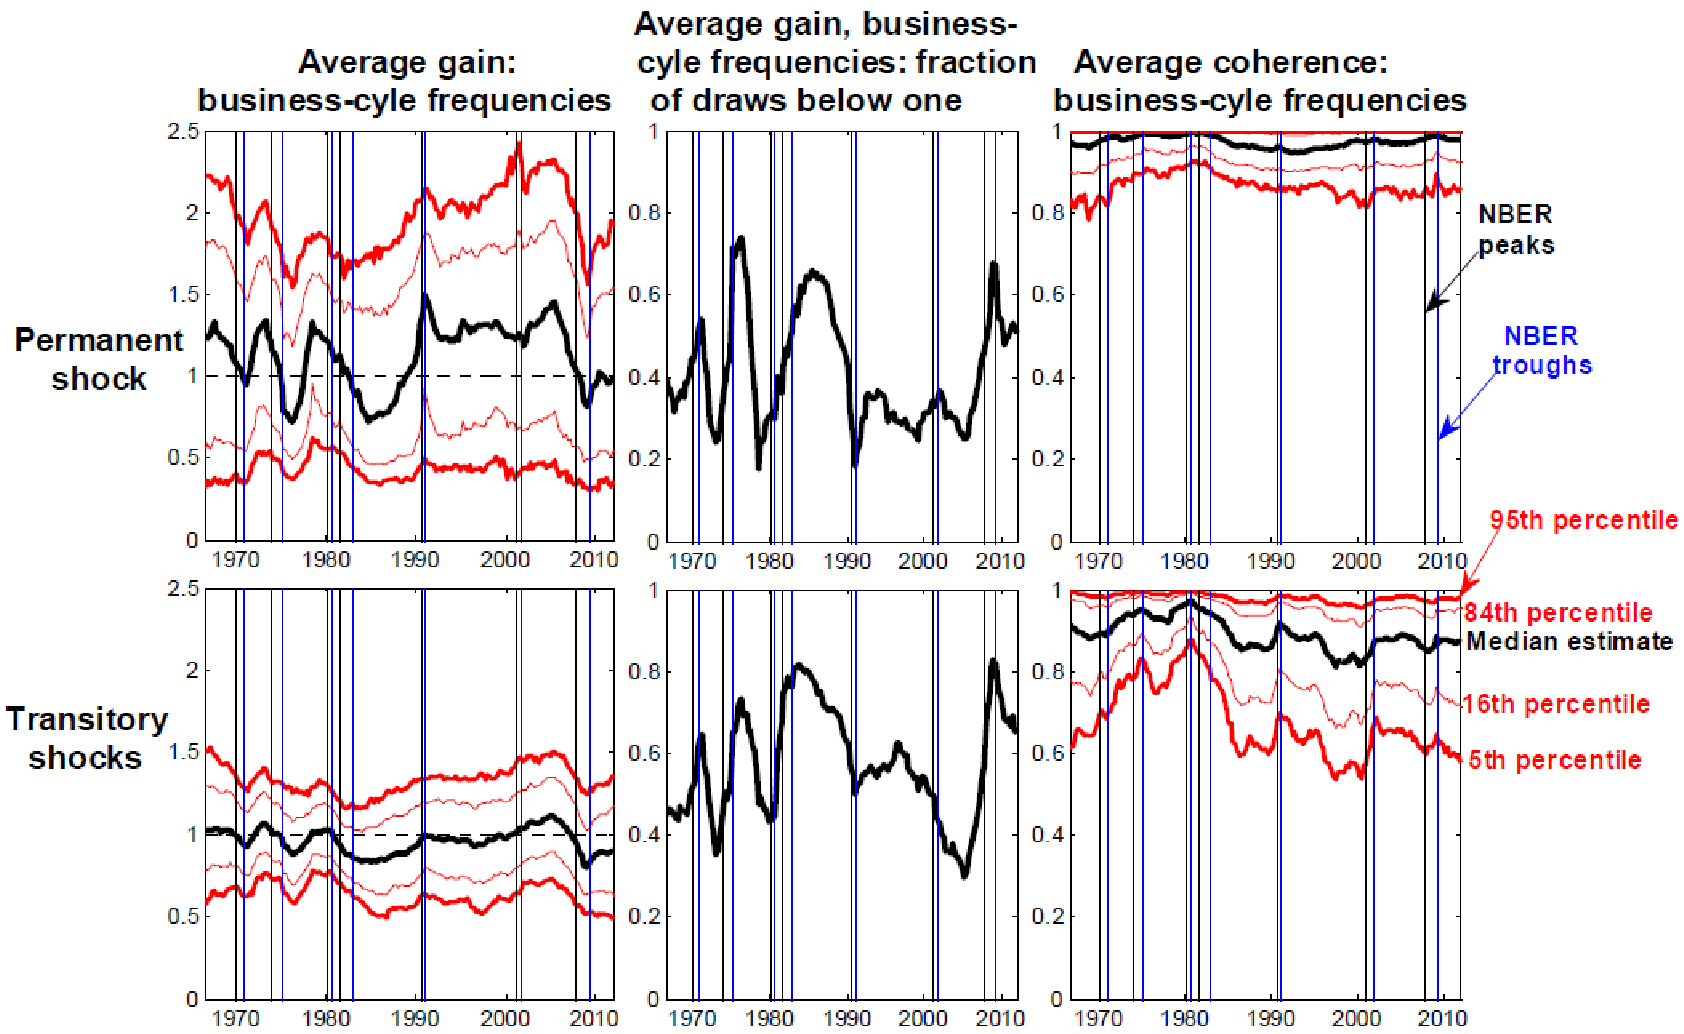

Figure 13 Business-cycle frequencies: average gain and coherence between vacancies and the unemployment rate conditional on the permanent and the transitory output shocks

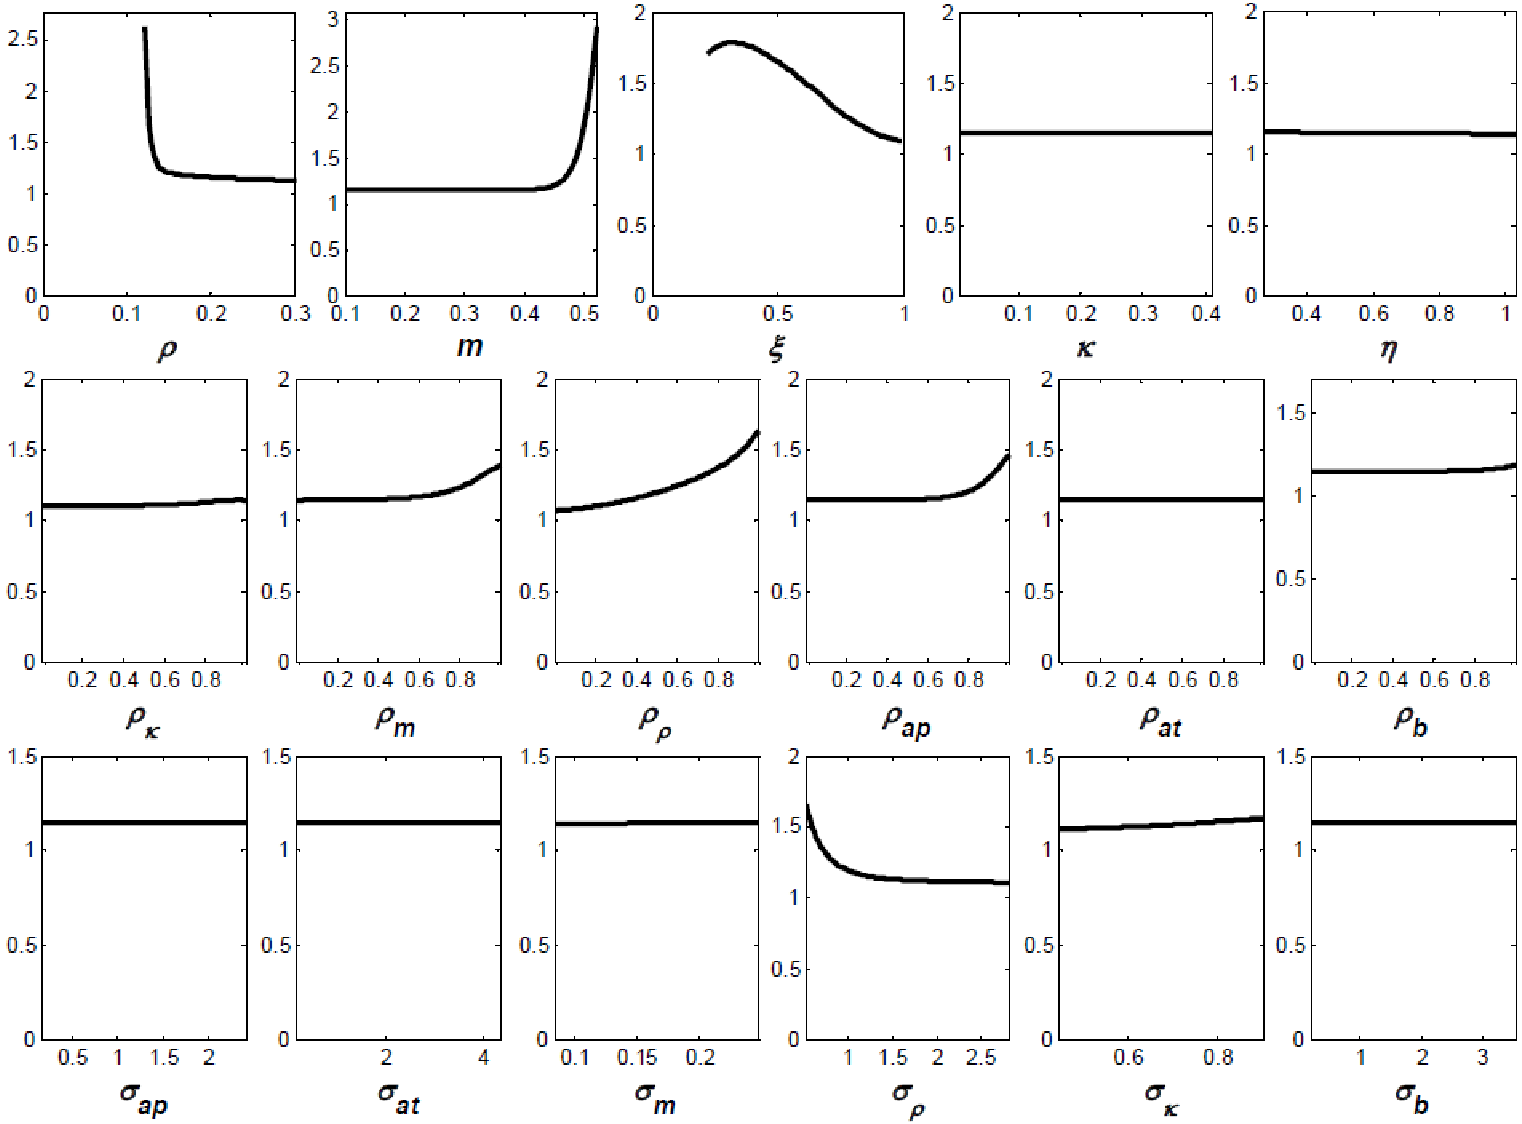

Figure 14 Average gain of the unemployment rate onto vacancies at the business-cycle frequencies, as a function of individual parameters of the DSGE model (for parameter intervals around the modal estimates generated by the Random Walk Metropolis)