

The Shifting and Twisting Beveridge Curve: An Aggregate Perspective*

Thomas A. Lubik
Federal Reserve Bank of Richmond[†]

October 2013
Working Paper No. 13-16

Abstract

One of the most striking aspect of the Great Recession in the United States is the persistently high level of unemployment despite an uptick in economic activity and an increased willingness by firms to hire. This has stimulated a debate on mismatch in the labor market. The argument is that despite the high unemployment rate the effective pool of job seekers is considerably smaller due to adverse effects of long-term unemployment, high unemployment benefits or structural change. Despite high vacancy postings firms are therefore unable to hire desired workers. I study this issue from an aggregate perspective by deriving the Beveridge curve from a discrete-time search and matching model of the labor market driven by a variety of shocks. I first establish that the observed pattern in the data can only be described in the context of the model by the interaction of a cyclical decline in productivity and a decline in match efficiency. I then estimate the model using Bayesian methods on unemployment and vacancy data before the onset of the Great Recession. The posterior estimates indicate that the recent behavior of the Beveridge curve is most likely explained by a structural decline in match efficiency.

JEL CLASSIFICATION: C11, C32, E20, E24,
KEYWORDS: Mismatch; unemployment; vacancies;
Great Recession; Bayesian estimation

*I am grateful to Nika Lazaryan and Tim Hursey for exceptional research assistance. I wish to thank seminar audiences at the Federal Reserve Bank of Kansas City, the University of Virginia, City University of Hong Kong, the Hong Kong University for Science and Technology, Elon University, and participants at the 2011 Conference of the Society for Computational Economics in San Francisco, for useful comments. The views in this paper are those of the author and not necessarily those of the Federal Reserve Bank of Richmond, or the Federal Reserve System.

[†]Research Department, 701 E. Byrd St., Richmond, VA 23219, USA. Tel: +1 804 697 8246. Email: Thomas.Lubik@rich.frb.org.

1 Introduction

The Beveridge curve captures the relationship between aggregate unemployment and vacancies. Plotting the former against the latter shows a downward-sloping relationship that appears tightly clustered around a concave curve (see Figure 1). The curve reflects the highly negative correlation, -0.91 , between unemployment and vacancies that is a hallmark of labor markets. Moreover, the Beveridge curve also encapsulates the logic of the search and matching approach to modeling labor markets. In times of economic expansions, unemployment is low and vacancies, that is, open positions offered by firms, are high. 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 worsen and demand slows down, firms post fewer vacancies and unemployment rises, commensurate 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.

This pattern is evident in much of the business cycles after World War II (see Benati and Lubik, 2013). The aftermath of the recession of 2007-2009 appears, however, strikingly different. The recession was exceptional for its length and depth, and in that the unemployment rate almost reached its previous post-war peak. One of the most striking aspects of the recovery is that the labor market picture has not improved as fast as had been expected. After staying close to its recession peak for more than year, the unemployment rate has come down slowly since then by 2.5 percentage points to its current level of 7.3%. At the same time, as the economy has picked up, vacancy postings have been rising. In terms of the Beveridge curve, we observe a cluster of data points above where we expect them to be, based on the standard relationship (see Figure 1). Moreover, the off-Beveridge curve cluster has persisted for a long while now and appears to evolve in a manner parallel to the pre-recession curve.

I therefore ask the question whether the recent behavior of unemployment and vacancies is consistent with a normal Beveridge curve relationship or whether something fundamental has changed in the labor market. I use a simple search and matching framework to study this issue. I show that the recent behavior of unemployment and vacancies can qualitatively be

explained by a combination of two influences. Negative aggregate productivity shocks reduce output, increase unemployment, lower vacancy postings, and are thus in general consistent with movements along the Beveridge curve. Persistent clustering of unemployment-vacancy combination off a Beveridge curve are consistent with shifts in the Beveridge curve caused by a decline in match efficiency.

From a quantitative viewpoint, however, it is difficult to ascertain whether the recent data signify a break in the functioning of the labor market, are the outcome of a sequence of unfortunate shocks, or simply reflect the inability of a researcher confronted with parameter uncertainty to draw conclusions. In order to address the last point, I take a Bayesian approach and treat the model parameters as random variables. I then conduct a Bayesian prior analysis to see whether the model is in principle capable of capturing the Beveridge curve. The answer is affirmative, both for the steady state version of the model as well as the fully dynamic and stochastic version. I then estimate the latter and extract the implied shock processes that are needed to replicate the data. The model is capable of explaining both the pre- and post-recession unemployment and vacancies data, but only with a seemingly structural shift in match efficiency.

The evolution of the labor market over the course of the Great Recession can therefore be told as follows. The business cycle downturn induced a typical movement along the Beveridge curve as vacancy postings dried up and unemployment rose. The length and depth of the recession induced changes in the way the labor market functions. This can be seen as either a sequence of persistent and negative shocks to match efficiency or as an induced break in a parameter. The latter interpretation, in particular, can capture the idea of labor market mismatch as laid out in Sahin et al. (2013). The depth of the recession led to the emergence of wide-spread long-term unemployment, so that despite increased willingness to hire, the afflicted workers so not find their match.

In order to develop my argument conceptually, I proceed in several steps. I first perform a steady-state analysis based on a simple accounting equation for employment and an optimality condition for vacancy postings. A steady-state analysis is convenient since it allows me to derive analytical predictions that are helpful in gaining intuition. Moreover, a steady-state analysis can be justified (e.g., as in Shimer, 2005) when adjustment dynamics are fairly rapid, which arguably is the case in monthly labor market data. The results from this analysis are clear-cut. The cluster of data points can only be explained as the outcome of two separate effects: a productivity-driven movement along the normal Beveridge curve, and a shift or twist to a new Beveridge curve. I demonstrate that this shift is consistent

with a decrease in match efficiency, an increase in the (fixed and exogenous) separation rate, or a decrease in the match elasticity. I argue that the most plausible source of the shift is a decline match efficiency, that is, the ability of the labor market to transform open positions and job seekers into productive employment relationships.

My paper is closely related to Barlevy (2011). He conducts a steady-state analysis similar to mine and derives the same qualitative conclusions. His data end in 2010 and can therefore not address the more recent evolution of the unemployment and vacancy combination in a manner parallel to the pre-recession data. I also conduct a more extensive empirical analysis based on a fully specified dynamic and stochastic general equilibrium model. My empirical approach is based on the Bayesian analysis of the simple search and matching model as developed in Lubik (2009, 2012). In that respect, my paper is similar to Furlanetto and Groshenny (2013). These authors develop a New Keynesian model with search and matching features in which match efficiency is but one shock. Bayesian estimation of their model delivers similar results to mine, namely a sharp drop in the match efficiency process during the Great Recession. Finally, Barnichon and Figura (2013) and Sahin et al (2013) approach the mismatch debate using labor market heterogeneities and micro-level data on occupations and sectors. By aggregation of individual matching functions they can derive an aggregate matching function and impute match efficiency. The resulting time series for the thus constructed series is very similar to the one derived in my paper.

The paper is organized as follows. In the next two sections, I develop intuition for the shifts and twists in the Beveridge curve and their determinants from a steady-state analysis. Section 2 derives the curve from a simply accounting equation, while Section 3 adds the optimal vacancy posting decision of a firm. In Section 4 I conduct a Bayesian prior analysis that attempts to determine whether the model in its various specifications is in principle capable of replicating the pattern seen in the data. I estimate the model using Bayesian methods in Section 5 and assess to what extent the data can be explained by shocks alone or by shifts in parameters. The final section concludes and speculates on future areas of research.

2 Beveridge Curve Shifts: Preliminary Insights

I organize my discussion around a simple search and matching model of the labor market. This model has been popularized in the macroeconomics literature by Pissarides (2000) and Shimer (2005), but it dates back to earlier work by Diamond (1982) and Mortensen and Pissarides (1994). The specific framework I use in this paper is drawn from Lubik

(2012). The model describes the evolution of vacancies and unemployment by means of two relationships that are derived from the optimizing decisions of workers and firms. The model and all derivations are presented in the Appendix.

The first equation, and the one that underlies the Beveridge curve, is a relationship that describes the dynamics of employment:

$$N_t = (1 - \rho) \left[N_{t-1} + mU_{t-1}^\xi V_{t-1}^{1-\xi} \right]. \quad (1)$$

This is a stock-flow identity that relates the stock of employed workers N to the flow of new hires $M = mU^\xi V^{1-\xi}$ into employment. ρ is the (constant) separation rate that captures employment outflows. New hires are determined as the outcome of a matching process that combines vacancies V with the unemployed U who are searching for jobs. The matching function is of the Cobb-Douglas type with match elasticity $0 < \xi < 1$ and a level parameter $m > 0$ that measures the efficiency of the matching process. The measure of the unemployed is defined as $U_t = 1 - N_t$, where the labor force is normalized to one.

The theoretical Beveridge curve is usually depicted as the steady-state counterpart of this dynamic employment equation. It is posited that shocks are short-lived enough and adjustment dynamics rapid enough so that departures from the steady state are resolved within the sampling frequency. Consequently, steady-state relationships are good enough approximations for the underlying, richer dynamics. After substituting the definition of the unemployment rate into (1), I can rearrange this expression to derive a steady-state relationship between vacancies and unemployment:

$$V = \left(\frac{1 - \rho}{m} \right)^{\frac{1}{1-\xi}} \left(\frac{1 - U}{U} \right)^{\frac{1}{1-\xi}} U. \quad (2)$$

This Beveridge curve describes an equilibrium locus of combinations of U and V such that inflows and outflows to the (un)employment pool are balanced.

I now use this relationship to interpret the empirical behavior of unemployment and vacancies over the last decade. I use monthly data for the sample period December 2000 to July 2013, which is governed by the availability of vacancy data in JOLTS. The unemployment rate series is headline unemployment from the household survey. I normalize the seasonally adjusted JOLTS vacancy series by the labor force in order to render it comparable to the unemployment rate. A time series plot of the two series can be found in Figure 1.

The Beveridge curve (2) is characterized by three parameters. The match efficiency m and the separation rate ρ determine the intercept and thus its location, while the match

elasticity ξ governs its curvature. I calibrate these parameters as follows. I set the separation rate to a value of $\rho = 0.036$. This follows the value reported in Shimer (2005) for monthly data, which is effectively just a normalization of the model for a given frequency. I estimate the remaining two parameters directly from the data using non-linear least squares. However, I split the sample in two. The first sub-sample, from which the parameters are estimated, runs from December 2000 to September 2008. In my analysis I treat this as the ‘normal’ period. From the last data point on, which, incidentally, lies on the estimated Beveridge curve, the unemployment rate rose rapidly. The second half of the sample thus covers the Great Recession and the recovery. The point estimates are $\hat{m} = 0.80$ and $\hat{\xi} = 0.49$.¹ Based on these parameters I extend the Beveridge curve outward for ranges of the unemployment rate attained during the recession.

Figure 1 shows the estimated Beveridge curve over a scatter plot of the data. The data exhibit the familiar downward-sloping pattern. By the logic of the search and matching model, firms post fewer vacancies and the unemployment rate rises as business conditions deteriorate. The impact of the Great Recession on the labor market is clearly visible in the sharp and continuous increase in unemployment starting in September 2008. Noticeably, all data points from this date on lie above the predicted Beveridge curve. During this period unemployment has been rising faster than the pre-recession Beveridge curve would suggest. At the height of the labor market downturn, the unemployment rate appears to have stalled between 9.5 and 10%, while the vacancy rate was creeping upward. From then on, the evolution of unemployment and vacancies has been consistently off the pre-recession Beveridge curve.

This observation has given rise to a bevy of arguments about ‘mismatch’ in the labor market. The argument is that in the course of the Great Recession the normal functioning of the labor market became disrupted. Firms willing to hire would find it difficult to be matched with potential employees despite the large pool of the unemployed. This notion can be captured by a shift in the Beveridge curve, which could produce the pattern seen in the data through a combination of a new steady-state equilibrium and accompanying adjustment dynamics. I now analyze the potential factors behind this argument using Eq. (2).

¹I take the natural logarithm of the Beveridge curve and collect terms. This results in the regression equation:

$$\log v = \frac{1}{1-\xi} \log \left(\frac{1}{m} \frac{\rho}{1-\rho} \right) + \frac{1}{1-\xi} \log(1-u) - \frac{\xi}{1-\xi} \log u.$$

Conditional on the calibrated value of ρ , I can then identify ξ and m from the estimated slope and intercept terms. I ignore obvious issues of collinearity and endogeneity for the purposes of this simple exercise.

An up and outward shift of the curve can be explained by a change in the two locational parameters m and ρ . A decrease in the match efficiency m would shift the Beveridge curve upward, since for a given number of unemployed the number of vacancies would have to be higher in order to generate the same number of new hires. Arguably, a shift in the match efficiency parameter captures the narrative of mismatch in the labor market well. Alternatively, an increase in the separation rate ρ would also shift the curve up. A given level of employment would now produce a higher number of inflows into unemployment, which would have to be balanced by a larger number of vacancy postings in order to generate the same steady-state level of flows out of unemployment.

I assess the size of the shifts required using the following experiment. I assume that the data point consistent with the new location of the Beveridge curve is July 2010, with an unemployment rate of 9.5% and a vacancy rate of 1.9%. Imposing these two values on Eq. (2), I can then back out the new parameter values, $m' = 0.67$ and $\rho' = 0.042$. The behavior of vacancies and the unemployment rate would thus be consistent with a decrease of the efficiency of the matching process by 16% and an increase in the (monthly) separation rate by 12%. While either parameter break appears plausible, the actual behavior of the separation rate in the data tells a different story. The total separation rate as measured in JOLTS actually *fell* over the course of the Great Recession from an average of 3.6% to slightly above 3%, where it has remained stable since.

An alternative source of the shift in the Beveridge curve is a change in the match elasticity ξ , which has a level effect but also affects the curvature. Following the same procedure as before, I find that a fall of ξ by 22% to $\xi' = 0.38$ is needed to be consistent with the data point of July 2010. Figure 2 depicts the shifts of the curve for the two parameter constellations.² The Beveridge curve is notably flatter in the case of a decline in the match elasticity. At a given unemployment rate, vacancy postings are now less reactive. Using the matching function, one can define the hiring rate as $q(\frac{V}{U}) = \frac{mU^\xi V^{1-\xi}}{V} = m(\frac{V}{U})^{-\xi}$. For a given vacancy-unemployment ratio, the hiring rate becomes less elastic and thus less responsive to aggregate labor market conditions with a fall in match elasticity. Furthermore, the JOLTS hiring rate shows a substantial peak-to-trough decline from 3.8% to 2.9% which would be consistent with a fall in the match elasticity.

I can draw a few preliminary conclusions at this point. The recent behavior of the unemployment and vacancy rates is consistent with an upward shift in the Beveridge curve that originated either from a decline in match efficiency or a decrease in the elasticity of

²Since the separation rate and the match efficiency have the same effect on the intercept term of the curve, the two curves exactly overlay each other.

the matching function. This lends support to the notion of mismatch in the labor market, which finds its theoretical counterpart in breaks in the parameters of the matching function. The two parameter shifts, however, imply different time paths going forward. A decline in match efficiency imparts a ‘new normal’ and movements in unemployment and vacancies that are no longer consistent with pre-recession labor market equilibrium. A fall in match elasticity, on the other hand, eventually returns the two labor market variables close to the original Beveridge curve along a counter-clockwise adjustment path as in Blanchard and Diamond (1989). While an increase in the separation rate is a theoretical possibility, this explanation is ruled out by the actual behavior of separations in the data.

However, the preceding analysis should be taken with a degree of caution. First, the empirical analysis is very rudimentary. In particular, the suggested values for the match elasticity are far below typical numbers in the literature. Second, and more importantly, the analysis does not distinguish between adjustment dynamics and business cycle dynamics. That is, a comparative statics analysis only provides limited insight into the behavior over time. Third, the analysis does not distinguish between exogenous shocks and breaks in parameters as sources of the perceived shifts in the comovement pattern of unemployment and vacancies. Finally, the Beveridge curve describes an equilibrium outcome, but it is an open question whether the data are consistent with the behavior of a general equilibrium model. I will now address these concerns using the full model.

3 Beveridge Curve Shifts: A Broader Perspective

The Beveridge curve describes the downward-sloping relationship between unemployment and vacancies. The contemporaneous correlation between the two time series over the JOLTS sample period is -0.85 . This relationship can be described using a matching function and the stock-flow accounting of the employment equation (1). At the same time, the levels of unemployment and open positions are the endogenous outcome of the decision making processes of firms and workers. I therefore investigate now whether the parameter shifts identified above can at least qualitatively reproduce the data pattern within the confines of a simple search and matching model of the labor market.

The second equation of the search and matching model is the job creation condition (JCC):

$$\frac{\kappa}{m_t} \left(\frac{V_t}{U_t} \right)^\xi = \beta(1 - \rho)E_t \left[(1 - \eta)(A_{t+1} - b) - \eta\kappa \frac{V_{t+1}}{U_{t+1}} + \frac{\kappa}{m_{t+1}} \left(\frac{V_{t+1}}{U_{t+1}} \right)^\xi \right]. \quad (3)$$

I describe in the Appendix how it can be derived from the optimal vacancy posting decision of a firm and a wage setting equation.

The steady-state representation of the JCC is:

$$\frac{1 - \beta(1 - \rho)}{\beta(1 - \rho)} \frac{\kappa}{m} \left(\frac{V}{\bar{U}}\right)^\xi + \eta \kappa \frac{V}{\bar{U}} = (1 - \eta)(A - b). \quad (4)$$

Note that $\theta = V/U$ defines labor market tightness. The JCC determines the steady-state level θ as a function of the model parameters. The job creation condition can thus be written as a linear function in V - U space: $V = \theta^*U$, where shifts in the parameters only affect the slope, but not the intercept of the JCC.³

I calibrate the JCC in a similar manner as above. The sampling period for the calibration is December 2000 to September 2008. The three parameters of the Beveridge curve also appear in the JCC and take on the same values: $\rho = 0.03$, $m = 0.80$, and $\xi = 0.49$. I set the discount factor $\beta = 0.99$, and choose the bargaining parameter by imposing the Hosios-condition for social efficiency, $\eta = \xi = 0.49$.⁴ There are three remaining parameters. I normalize the level of productivity to $A = 1$. Next, I assume that the outside option of the worker makes up 90% of the productivity level, $b/A = 0.9$. My calibration is therefore close to that of Hagedorn and Manovskii (2008), who argue that a high outside option for the worker is needed to match the cyclical properties of the data. The JCC (4) can then be used to back out the cost parameter κ for a given level of unemployment and vacancies. I compute these from the sample averages, $\bar{V} = 2.6\%$ and $\bar{U} = 5.2\%$. This implies $\kappa = 0.18$.

Figure 3 depicts the JCC and the Beveridge curve. They intersect at the steady-state equilibrium and sample mean for the unemployment and vacancy rates. The Figure also shows how the JCC shifts with changes in the productivity level, whereas the employment equation (1) is unaffected. With technology as the driving process, movements in the JCC thus trace out what appears in the data as the Beveridge curve. While the combinations of U and V have to lie *on* the theoretical curve (2), they are generated by the interaction of the employment equation with employers' vacancy posting decisions in (4). A focus on the relationship (2) alone can therefore only reveal a partial picture of the causes of mismatch.

The figure also shows shifts in the JCC as the level of productivity changes from its baseline value of $A = 1$. I consider both increases and decreases of 2.5% and 5% each. When productivity rises to $A = 1.025$ the JCC shifts and tilts upwards (dashed line), and shifts further out when $A = 1.05$ (dash-dotted line). Decreases in productivity of the

³I solve Eq. (4) numerically since analytical solutions are available only for the special case $\xi = 1/2$.

⁴A typical calibration in the literature is an agnostic $\eta = 0.5$, which is close to mine.

same percentage move the JCC in the opposite direction. Variations in productivity thus trace out a downward-sloping relationship between unemployment and vacancies along the employment curve (1), to wit, the Beveridge curve. This is the kind of experiment that underlies most discussions of the dynamic properties of the search and matching model (see Shimer, 2005). What productivity movements alone cannot accomplish, however, is to explain the recent pattern in the data as evidenced by the data cluster off the ‘normal’ Beveridge curve in Figure 3.

I therefore analyze the effects of the parameter changes considered above on both the JCC and the employment equation. The results are graphed in Figure 4. I perform two sets of experiments. In Panels A and B I only vary the match efficiency and the match elasticity, respectively. The parameter values are the same as in the previous section. Specifically, I consider a decline of m to 0.67 and of ξ to 0.38 in order to be consistent with the July 2010 data point on the Beveridge curve. As established before, declines in these two parameters shift and twist the employment equation outward. What is striking, however, is that these changes only have minor effects on the JCC, which marginally shifts outward. These parameter changes alone can therefore not generate an equilibrium outcome consistent with the July 2010 data point. To do so, an additional movement along the new employment equation is required.

Panels C and D depict the combined effects of declines in the two parameters and a decrease of productivity to $A = 0.945$, which in both cases leads to an intersection of the JCC at the target data point of July 2010. The interpretation of the patterns in the recent data now seems straightforward. At the onset of the Great Recession, declining GDP, interpreted in the model as fall in productivity, led to a down and outward shift of the JCC along the employment equation, tracing out a typical Beveridge curve. As the downturn deepened, however, the process of matching the unemployed with open vacancies appears to have been affected adversely. Consequently, the employment equation shifts upward along a JCC that has settled at a lower productivity level. What emerges from this narrative is an explanation of mismatch grounded in the behavior of the matching function. The shift appears to kick in at a percentage decline in productivity of roughly 5%. Once this threshold is crossed, the Beveridge curve shifts to a new equilibrium.

This exercise also holds a useful insight for the specification of empirical search and matching models. Movements in productivity lead to sharp movements in the JCC, but leave the employment equation unaffected. Vice versa, shifts in the parameters of the matching function move and twist the employment equation, but have only a minor effect

on the JCC. This suggests that in order to explain labor market behavior via the search and matching framework at least two independent driving forces are needed, specifically, productivity shocks and movements in the parameters of the matching function that act directly on the employment equation.⁵

A priori, the pattern in the data is consistent with shifts in either the match efficiency m or the match elasticity ξ . This identification problem will have to be resolved by an empirical analysis of the fully specified dynamic model. I can gain some insight, however, by introducing additional information. The matching function allows me to define the rate at which the unemployed are matched with employers. This job-finding rate can be expressed as $p(\frac{V}{U}) = \frac{mU^\xi V^{1-\xi}}{U} = m(\frac{V}{U})^{1-\xi}$. I gather data on the total number of hires from JOLTS and scale it by the total number of unemployed. I then regress this rate on labor market tightness $\theta = V/U$ in logs over the sample period until the onset of the downturn. The estimated values of the match efficiency $\hat{m} = 0.97$ and the match elasticity $\hat{\xi} = 0.38$. Figure 5 depicts the fitted hiring rate from December 2000 to September 2008 (data points in red) and its prediction for the recession sample. The graph clearly shows that the predicted hiring rate is substantially above the actual data (green dots). In other words, transitions from unemployment into employment are lower in the wake of this downturn than the baseline calibration can explain. This is, of course, the flip side of the mismatch pattern in the Beveridge curve. The pattern in the second sample period can be matched by a downward shift in the match efficiency, while a change in ξ cannot quite capture the level shift.

I therefore draw the following conclusions at this point. The recent behavior of the vacancy rate can be understood from the perspective of a search and matching model as the interaction of the shifts of at least two structural parameters. The *a priori* most plausible combination is a decline in aggregate productivity and a subsequent decrease in match efficiency. In what follows, I will use this as my working hypothesis. However, the preceding analysis comes with important caveats. First, the steady-state analysis is only suggestive of the dynamics involved and the previous conclusions should at best be seen in a qualitative light. I will therefore analyze the data with a fully specified dynamic and stochastic model. Second, the previous results rest on specific point calibrations. Yet, there is large uncertainty regarding the parameter values. I will therefore proceed to a Bayesian approach which treats structural parameters as random variables.

⁵This argument has also been made by Lubik (2009).

4 A Bayesian Prior Analysis

The analysis presented so far gave qualitative insight into whether the recently observed pattern of unemployment and vacancies is consistent with what we would have expected to be based on pre-recession data. My suggested answer is that it is not. More specifically, it appears that the Beveridge curve has shifted outwards, seemingly permanently. However, it is not obvious whether the recently observed pattern is quantitatively different in a statistical sense. In other words, it may be the case that the post-recession data were generated from the same ‘normal’ process as, say, the simple search and matching model, but because of sampling uncertainty, model uncertainty, and parameter uncertainty it appears to the researcher that the data clusters are different. What appears as a Beveridge curve shift may thus simply be a less likely, but nevertheless consistent, realization of the ‘normal’ curve.

In the rest of the paper, I therefore investigate along these lines. I ask the question whether the search and matching model is consistent with both pre- and post-recession data under the assumption of parameter uncertainty. As a preliminary step, I perform the following exercise. In the spirit of the simple analysis above, I compute the steady-state distribution of the Beveridge curve from the employment equation alone. This exercise is parsimonious in that I only need to deal with three structural parameters. I then add the job creation condition to compute the steady-state distribution from the general equilibrium conditions. Both computations should reveal some information on whether the stochastic specification is in principle capable of replicating the pattern in the data.

I first need to select the priors. The prior distributions are reported in Table 1. I could pursue a largely agnostic strategy by choosing diffuse priors. By doing so, the implied prior density of the variables would be determined by the model and the implicit cross-equation restrictions. I prefer, however, to restrict the prior choice by bringing to bear additional information from a variety of sources. As before, I set the separation rate to a value of $\rho = 0.03$ and keep it fixed for the duration of this exercise. Since the match efficiency m has to be strictly positive I chose a Gamma-distribution with a mean of 0.80. As explained above, this is derived as the intercept from fitting the Beveridge curve relationship (2) to the pre-October 2008 data. Similarly, I chose a Beta-distribution for the match elasticity ξ with a mean of 0.49. I set the standard deviation for both parameters at 0.1, which implies 95% coverage intervals for m and ξ of, respectively, [0.60, 0.99] and [0.29, 0.68]. Given the Beveridge curve relationship (2) it is straightforward to compute the implied joint distribution for unemployment and vacancies. These are depicted in Figure 6 for an

unemployment range from 4% to 12%.

The contour plot in Panel A is overlaid with the data from the full sample period. During normal times in the first part of the sample, the data lie along the ridge of the joint distribution. This simply reflects that the parameter distributions were chosen based on the fitted Beveridge curve in Figure 1. As the Great Recession takes its course, however, the graph shows that the data points fall off the ridge. Given the joint density, I can assign a probability to specific outcomes under this prior distribution. As before, I pick the unemployment rate in July 2010, namely 9.5%, as a reference point. The estimated Beveridge-curve relationship would predict a vacancy rate of 1.3%, whereas the actual observed value was 1.9%.

Panel B graphs the implied prior distribution of the vacancy rate conditional on an unemployment rate of 9.5%. The two vertical lines are the actual and predicted vacancy rates. The deviation of actual from predicted vacancies is thus around the one-standard deviation mark, which may not be considered statistically significant. We therefore cannot rule out with a fair degree of certainty that the recent pattern has not been generated by small stochastic variations in the parameters. I also note that the distribution is slightly right-skewed, so that there is relatively more probability mass on vacancy rates above the predicted mode. This reflects the tendency of the Beveridge-relationship to engender vacancy postings when unemployment is high. Naturally, this observation needs to be addressed in terms of the job creation condition, to which I now turn.

I compute the implied density of unemployment and vacancy rates for parameter draws from the priors using the steady-state equilibrium conditions (2) and (4). As before, I fix the discount factor $\beta = 0.99$ and I normalize productivity $A = 1$. Additionally, I set the prior mean of the bargaining share at the Hosios-condition $\eta = \xi = 0.46$, but allow for a wider dispersion on account of the greater degree of uncertainty surrounding this parameter. The prior mean of the benefit ratio remains at $b/A = 0.9$. Since the parameter is bounded above by productivity, I assume a Beta-distribution for this ratio with a standard deviation of 0.1. This implies a 95% coverage interval of $[0.62, 0.99]$. Finally, I set the prior mean of the vacancy posting cost κ at 0.18 with a reasonably tight standard deviation. Recall that this value was chosen based on capturing the sample means in the data. I will use this specification later on as the benchmark prior distribution in the structural estimation.

Figure 7 contains a scatter plot of unemployment and vacancy rates for 20,000 draws from the prior distributions. I superimpose the JOLTS-sample, while the horizontal and vertical lines correspond to the sample means from the pre-recession period. Two observa-

tions stand out. First, the prior draws trace out the familiar downward-sloping relationship in V-U space. Second, there is a large degree of dispersion, which is not readily apparent from the prior analysis based on the employment equation alone. This is to a large extent driven by the additional variation introduced by the remaining parameters of the model, but also reflects the specific prior chosen. The data points during the normal sample period lie along the previously identified ridge. At the same time, the Great Recession sample also coincides with a high density region. It is thus a priori not obvious that it is inconsistent with the standard model. The scatter plot also offers insights for the structural estimation of the model. Given the prior distribution, the data appear very informative for the estimation of the structural parameters. The posterior distribution can be expected to be much more concentrated than the prior, which results in sharp inference for the parameters.

5 An Estimated Beveridge Curve Model

The conclusion that can be drawn from the results in the previous section is that the standard search and matching model can in principle reproduce the patterns found in the data over the entire sample period. That is, the recent behavior of the Beveridge curve appears statistically indistinguishable from its earlier behavior. I am now taking a more deeply quantitative look at this issue. To this end, I estimate the full dynamic and stochastic search and matching model using Bayesian methods. Details on the econometric method and the specifics of the empirical analysis can be found in Lubik (2009), while a comparable exercise is performed in Lubik (2012). The model specification is the same as the one used in the previous simulation exercise, as is the choice of the prior distribution for the parameters to be estimated (see Table 1). I only use the unemployment rate and the vacancy rate as observables in the estimation. The model is driven by two shocks, labor productivity A_t and exogenous variations in match efficiency m_t , which are both assumed to be AR(1) processes. The full description of the model and the state-space representation of the estimated model can be found in the Appendix.

I estimate the model for the baseline sample period from December 2000 to September 2008 and then assess how well it fits the data, what the driving forces of movements in unemployment and vacancies are, and how the posterior estimates differ from the priors. Next, I estimate the model for the full sample up to July 2013, which includes the Great Recession and its aftermath. I compare the differences in the posterior estimates and try to identify which shocks, if at all, are consistent with the pattern in the data. I also conduct an experiment that estimates the model only over the latter part of the full sample.

This should give an indication whether the model's behavior is driven more by shocks or parameter shifts.

The Bayesian estimates are reported in Table 2 for each sample period. I omit the estimates for the separation rate and the vacancy creation costs as the prior and posterior means coincide in each sample period. The posterior distribution for the separation rate ρ is more concentrated, while that of the cost parameter κ overlays the prior, which suggests a lack of identification. The posterior mean of the match elasticity ξ is 0.81, with a tight 90% coverage interval, that does not include the prior mean of 0.49, nor do the coverage intervals overlap. Recall that the prior mean was chosen based on fitting a regression line to a scatter plot of unemployment and vacancy combinations. Its value was backed out based on a steady-state relationship. The posterior estimate thus suggests that the dynamic model offers additional information that requires a highly elastic response of the matching rate $q(\theta)$ to labor market tightness θ during normal times. While this estimate is somewhat higher than the range suggested in calibration studies - for instance, Petrongolo and Pissarides (2001) recommend a value in the range $[0.5, 0.7]$ - my estimate is consistent with other empirical studies.

The match efficiency m and the benefit parameter b posterior estimates are shifted away from the prior means, which in both cases lie outside the coverage intervals. For the former, the posterior mean is higher than the prior at 0.91, while the estimate of b is below the prior mean. In fact, the latter estimate is almost exactly halfway between the value in Shimer (2005) and the one espoused by Hagedorn and Manovskii (2008), upon which my choice of the prior is based. This is again consistent with previous estimates, which tend to favor a high replacement ratio. The posterior estimates clearly reject the Hosios-condition $\eta = \xi$ to the effect that the overlapping probability mass between the two distributions is essentially zero. The posterior of $\eta = 0.47$ is below the prior but not as much as previous studies would indicate. On the other hand, this parameter tends to be hard to pin down in the absence of other sources of information, such as wages. Finally, the posterior estimates of the shock parameters indicate more persistent and more volatile disturbances than assumed in the prior.

The estimates for the Great Recession sample period are in the fourth column of Table 2 (labeled Period II). Both the match elasticity at $\xi = 0.69$ and the match efficiency $m = 0.7$ are below the estimates from the first sample period. This is consistent with the argument made in Sections 2 and 3 using the steady-state model. In order to capture the data points that emerge above the implied Beveridge curve during and after the Great Recession, the

estimation algorithm sets the estimates at lower values. In fact, m is even below the prior mean. Moreover, the estimate of the benefit parameter b is now much closer to the prior mean. Since this increases ceteris paribus the volatility of unemployment and vacancies (see Hagedorn and Manovskii, 2008), it aids in matching the data. A higher b is also consistent with a narrative of more generous and extended unemployment benefits as the source of the persistently high unemployment rate. The posterior mean for η is higher than before, but there is substantial overlap of the coverage intervals. The AR-coefficients are unchanged relative to the first sample period. At the same time, the standard deviation of the productivity shock falls, while that of m_t rises. This supports the conclusion that the recent behavior of unemployment and vacancies is driven by a fall in the mean level of match efficiency, and an increase in their volatility. I will look at this again from the point of view of variance decompositions.

Finally, I estimate the model over the full sample period. The results are contained in the last column of Table 2. The two sub-samples are almost of equal length (the normal period being longer by two and a half years), so that it is not obvious to what extent the first sample dominated the estimation. Yet the substantial cluster of data points off the normal Beveridge curve has to be matched by the estimation algorithm. The posterior means are fairly unequivocal as to how this gets resolved. The match elasticity does not change from the first sub-sample, nor do the benefit and bargaining parameters, while match efficiency falls between the estimates of the two sub-samples. Surprisingly, the persistence of the two shock processes is higher than in both sub-samples, while the innovation variances lie in between.

Based on these observations, I can derive two conclusions. First, it is unlikely that changes in the match elasticity drive the recent Beveridge-curve behavior. This only leaves changes in match efficiency as the source, which confirms the previous analysis. In other words, the full sample can only be matched by averaging over the sub-sample match efficiency estimates. This implies a break in this parameter between the two sample periods. This point will become more obvious further below. Second, while changes in the benefit parameters (and to a lesser degree in the bargaining share) are contributing factors, the volatility and incidence of shocks to the match efficiency are contributing factors.

Next, I compute the long-run variance decomposition for the three samples before I look at the filtered series for the shocks. The variance decomposition is reported in Table 3. The main driving force for unemployment is productivity, while vacancies are largely explained

by shocks to match efficiency.⁶ The role of productivity shocks in explaining unemployment declines from the first to the second sample, with the full sample contribution close to that in the Great Recession. The picture for the contribution of match efficiency shocks is the mirror image of this. As for vacancies, the contribution of productivity shocks rises substantially, while that of match efficiency shocks falls. The variance decomposition is thus broadly consistent with the previous discussion of the posterior estimates. The outward shift of the Beveridge curve during the Great Recession is therefore captured over the full sample by a higher contribution of match efficiency shocks. For a given unemployment rate and productivity shocks, the data cluster off the normal Beveridge curve is the outcome of a higher incidence of shocks to m_t .

I now extract the implied series for the unobserved shock processes, productivity A_t and match efficiency m_t , by filtering the model at the posterior means. The series are reported in Figure 8. The graph at the top contains the series for the productivity shock for each of the three samples, while the bottom graph contains the filtered match efficiency series. Mean productivity is normalized to one in the model and is assumed to be unchanged over the sample period. The index numbers can therefore be interpreted as (gross) percentages.

The behavior of productivity over the full sample tells the story of the labor market in the last decade. As the economy was entering the 2001 recession, productivity fell from its initial peak of 10% above its long-run value to its mean, where it hovered for several years. This is commonly known as the jobless recovery period when unemployment stayed elevated and would come down only slowly. Starting in the second half of 2004, growth picked up, reaching its peak in early 2007. At the onset of the recession in the last quarter of 2007, productivity rapidly fell until it reached its trough in the middle of 2009 with an estimated productivity level of 12% below its long-run trend. Afterward, productivity started rising slowly, but gradually, commensurate with the economy's improved performance. This rise is consistent with the uptick in vacancy postings, and, in fact, is one factor behind the data cluster off the Beveridge curve. Yet, in the middle of 2013, four years after the trough, the economy still operates around 2% below its long-run level. The two series for the sub-samples very much follow the pattern established for the full sample period. Since the means are restricted to one, the apparent level shifts in the filtered productivity series simply reflect the fact that the estimation algorithm adjusts for variations in the location

⁶In a sense, this observation is comforting, since disturbances to parameters in the employment equation (1) tend to act as residuals in the specified equation and therefore simply reflect the inability of the model to capture the data on its own (see the argument in Lubik, 2012). In other words, the model, given this specific data set, provides enough restrictions so that the exogenous shocks do not fulfill the role of simply capturing the distance between the data and the model structure.

of the data series over the sub-sample. Absent of this, the three series would overlay each other closely.

In the bottom graph I report the actual level of the estimated match efficiency parameter. Since I allow the mean to change over the sample periods the differences in the levels reflect the different posterior mean estimates over the three sampling periods. The mean for the first sub-sample is estimated at 0.91. Interestingly, the filtered match efficiency closely follows the pattern of productivity, their correlation being 0.91. As we saw above, the posterior mean for the Great Recession sample drops to 0.7. When compared with filtered productivity, however, the pattern is different. Match efficiency depicts a noticeable downward trend, continuing even after productivity picks up again. In fact, correlation between the two series is now only 0.30. This supports the argument that unemployment is kept elevated because of the pressure from declining match efficiency despite the productivity-driven uptick in vacancies.

Finally, the filtered series also shows how estimation over the full sample accounts for the data. At the onset of the Great Recession, the labor market suffers a sequence of negative match efficiency shocks, which bring the series down to close to the level consistent with the second sub-sample. A priori, this would seem consistent with a structural break, in particular as the match efficiency has remained low for an extended period. I also note that the series for the full sample is showing a slight upward trend driven by the mean reversion in the model. There is no trend in match efficiency for the second sub-sample, which provides further evidence of a new normal Beveridge curve as is visible in Figure 1. What I cannot conclusively answer, however, is whether the off-Beveridge pattern is generated by a break in the mean or by a sequence of persistently negative innovations. The algorithm fits the data by imposing a high persistence parameter on the matching parameter (see Table 2). Of course, this is related to the difficulties that structural break tests have in distinguishing between true breaks and random walk behavior. An analysis of this issue is relegated to future research.⁷

The final question I ask of the estimated labor market model, is whether it can capture Beveridge curve dynamics. Figure 9 depicts the impulse response functions of unemployment and vacancies to the two exogenous shocks, estimated from the first sample period.⁸

⁷Progress in this direction has been made by Benati and Lubik (2013) who estimate a time-varying parameter VAR for the Beveridge curve relationship. The source of time variation in their set-up is random-walk behavior in the coefficients of a structural VAR. While they do find evidence of time variation in the Beveridge curve for post-WWII U.S. data, and especially for deep recessions such as the Great Recession, it is not statistically strong.

⁸The impulse responses are qualitatively very similar across the different sub-periods, although there

A one-standard deviation productivity shock raises vacancy postings on impact, which then decline gradually as the initial impulse runs out. Productivity does not affect employment contemporaneously since it takes time to create jobs. Unemployment falls one period after the initial shock and rises gradually back to its long-run level. Similarly, unemployment does not react to matching shocks contemporaneously because of the timing assumption embedded in the employment equation. Positive shocks lower unemployment since more potential workers are being matched for a given level of vacancies. At the same time, vacancy creation falls because firms need to post fewer open positions now that the job-matching rate rises on account of higher match efficiency.

The crux of the matter is that in response to matching shocks, unemployment and vacancies move in the same direction, thus implying a positive correlation contrary to what we see in the Beveridge curve. A negative correlation does emerge conditional on productivity shocks only. This becomes apparent when I plot the time paths for the two variables against each other in the lower panel of Figure 9. The straight lines reflect the mean unemployment and vacancy rates of the second sub-sample, namely 9.1% and 2.1%, respectively. The shocks are scaled to represent unit innovations to the respective terms. Positive productivity shocks induce an initial rapid rise in vacancies and an accompanying fall in unemployment. Beyond the turning point, which occurs three periods after the initial impulse, vacancies fall and unemployment rises slowly back to their long-run values. The latter pattern is of course consistent with a typical movement along the Beveridge curve, while the initial periods induce what Blanchard and Diamond (1989) refer to as the counter clockwise loop in unemployment and vacancy dynamics. In other words, the off-Beveridge curve pattern can be explained by recurrent positive productivity shocks, which keep vacancy postings elevated and the unemployment rate hovering at its high level.

On the other hand, matching shocks do not induce Beveridge curve-type dynamics, as unemployment and vacancies comove. This is evident from the location of the curve in the lower left quadrant in the lower panel of Figure 9. In response to positive matching shocks, adjustment to the long-run value implies positive comovement, which rules out matching shocks as the sole driver of the recent data. However, combinations of the two shocks can produce the pattern as I have demonstrated above. The estimated impulse responses do, however reveal one aspect that the qualitative discussion missed. Unemployment reacts relatively more strongly to the shocks than vacancies do in response to either shock, which is reminiscent of the argument in Shimer (2005). This would suggest the data to cluster

are small quantitative differences. The responses for the other periods are available from the author upon request.

more along the direction of unemployment in a Beveridge curve diagram rather than the vertical direction, which we see in the data immediately after the business cycle trough.

6 Conclusion

The behavior of unemployment and vacancy postings, when put in relation to a standard Beveridge curve, has been one of the most puzzling and troubling features of the U.S. economy in recent years. In this article, I attempt to explain this pattern from within the confines of a simple search and matching model of the labor market. I show that the pattern in the data can only be understood as the outcome of two distinct shocks. Productivity shocks, which are a stand-in for the overall level of economic activity, determine the level of unemployment and vacancies along the Beveridge curve, while disturbances to match efficiency determine shifts of the Beveridge curve itself. I demonstrate that the pattern in the data can be explained by temporary variations in those two shocks alone, and that the seeming unusual behavior of the data is consistent with a stochastic view of a standard search and matching framework.

Bayesian estimation of the simple search and matching model over both the full JOLTS sample period and over two sub-samples before and after the onset of the Great Recession shows that the behavior of the match efficiency process is such that it persistently needs to be below its full-sample mean to capture the recent Beveridge curve behavior. Since the model implies strong mean reversion in the match efficiency process, this has to be accomplished by sequences of large and negative innovations to this parameter. On the other hand, sub-sample estimation shows a clear level difference between the respective values without any evidence of movements toward the full sample mean. This is not the case for other parameters. The estimation results therefore suggest that there has been a break in the labor market matching process to the effect that the Beveridge curve has permanently shifted outward.

My approach in this paper is silent as to what the underlying reasons for this break are. One hypothesis is that structural breaks in match efficiency only occur during times of extreme economic duress as we have witnessed in the Great Recession. Sharp falls in output thus trigger fundamental changes in the labor market, which are represented by a permanent decline in match efficiency and a higher long-run unemployment rate. This can be modeled by a regime-switching or threshold-switching approach, which is the topic of ongoing research.

References

- [1] Barlevy (2011): “Evaluating the Role of Labor Market Mismatch in Rising Unemployment”. *Federal Reserve Bank of Chicago Economic Perspectives*, 35(3), 82-96.
- [2] Barnichon, Regis, and Andrew Figura (2013): “Labor Market Heterogeneities and the Aggregate Matching Function”. Manuscript.
- [3] Benati, Luca, and Thomas A. Lubik (2013): “The Time-Varying Beveridge Curve”. Federal Reserve Bank of Richmond Working Paper No. 13-12.
- [4] Blanchard, Olivier J., and Peter A. Diamond (1989): “The Beveridge Curve”. *Brookings Papers on Economic Activity*, 1, 1-60.
- [5] Diamond, Peter A. (1982): “Aggregate Demand Management in Search Equilibrium”. *Journal of Political Economy*, 90(5), 881-894.
- [6] Furlanetto, Francesco, and Nicolas Groshenny (2013): “Mismatch Shocks and Unemployment during the Great Recession”. Norges Bank Research Working Paper No. 2013-16.
- [7] Hagedorn, Marcus, and Iouri Manovskii (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited”. *American Economic Review*, 98, 1692-1706.
- [8] Lubik, Thomas A. (2009): “Estimating a Search and Matching Model of the Aggregate Labor Market”. *Federal Reserve Bank of Richmond Economic Quarterly*, 95(2), 101-120.
- [9] Lubik, Thomas A. (2012): “Aggregate Labor Market Dynamics in Hong Kong”. *Pacific Economic Review*, 17(2), 257-279.
- [10] Mortensen, Dale and Christopher A. Pissarides (1994): “Job Creation and Job Destruction in the Theory of Unemployment”. *Review of Economic Studies*, 61(3), 397-415.
- [11] Petrongolo, Barbara, and Christopher Pissarides (2001): “Looking Into the Black Box: A Survey of the Matching Function”. *Journal of Economic Literature*, 39, 390-431.
- [12] Pissarides, Christopher A. (2000): *Equilibrium Unemployment*. MIT Press. Cambridge, Ma.

- [13] Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante (2013): “Mismatch Unemployment”. Federal Reserve Bank of New York Staff Report No. 566.
- [14] Shimer, Robert (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies”. *American Economic Review*, 95, 25-49.
- [15] Sims, Christopher A. (2001): “Solving Linear Rational Expectations Models”. *Computational Economics*, 20, 1-20.

Appendix: Model Description

The model is a variant of the one employed in Lubik (2009, 2012). As shown in that paper, the model is capable of capturing labor market relationships reasonably well. Time is discrete and the time period is a quarter. The economy is populated by a continuum of identical firms that employ workers each of whom inelastically supplies one unit of labor. Output Y of a typical firm is linear in employment N :

$$Y_t = A_t N_t. \quad (\text{A1})$$

A is an aggregate productivity process, common to all firms, that obeys the law of motion:

$$\log A_t = (1 - \rho_A) \log \bar{A} + \rho_A \log A_{t-1} + \varepsilon_{A,t}, \quad (\text{A2})$$

where $0 < \rho_A < 1$, $\bar{A} > 0$, and $\varepsilon_{A,t} \sim \mathcal{N}(0, \sigma_A^2)$.

The labor market matching process combines unemployed job seekers U with job openings (vacancies) V . This can be represented by a constant returns matching function, $M_t = m_t U_t^\xi V_t^{1-\xi}$, where $m > 0$ is match efficiency, and $0 < \xi < 1$ is the match elasticity. The matching process is assumed to be follow:

$$\log m_t = (1 - \rho_m) \log \bar{m} + \rho_m \log m_{t-1} + \varepsilon_{m,t}, \quad (\text{A3})$$

where $0 < \rho_m < 1$, $\bar{m} > 0$, and $\varepsilon_{m,t} \sim \mathcal{N}(0, \sigma_m^2)$. Unemployment is defined as:

$$U_t = 1 - N_t, \quad (\text{A4})$$

where the labor force is normalized to one. Inflows to unemployment arise from exogenous job destruction at rate $0 < \rho < 1$.

The dynamics of employment are therefore governed by the following relationship:

$$N_t = (1 - \rho) \left[N_{t-1} + m_{t-1} U_{t-1}^\xi V_{t-1}^{1-\xi} \right]. \quad (\text{A5})$$

This is a stock-flow identity that relates the stock of employed workers N to the flow of new hires $M = m U^\xi V^{1-\xi}$ into employment. The timing assumption is such that variations in match efficiency do not affect employment contemporaneously. 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:

$$p(\theta_t) = \frac{M_t}{U_t} = m_t \theta_t^{1-\xi}, \quad (\text{A6})$$

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

$$q(\theta_t) = \frac{M_t}{V_t} = m_t \theta_t^{-\xi}, \quad (\text{A7})$$

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 hence new hires are linear in number of vacancies posted for individual firms: $M_{it} = q(\theta_t)V_{it}$.

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:

$$E_0 \sum_{t=0}^{\infty} \beta^t [A_t N_t - w_t N_t - \kappa V_t]. \quad (\text{A8})$$

subject to the employment accumulation equation (A5). Profits are discounted at rate $0 < \beta < 1$. Wages paid to the workers are W , while $\kappa > 0$ is a firm's fixed cost of opening a vacancy. The first order conditions are:

$$N_t : \quad \mu_t = \beta E_t [A_{t+1} - w_{t+1} + (1 - \rho)\mu_{t+1}], \quad (\text{A9})$$

$$V_t : \quad \kappa = \mu_t(1 - \rho)q(\theta_t), \quad (\text{A10})$$

where μ_t is the multiplier on the employment equation. Combining these two first-order conditions results in the *job creation* condition:

$$\frac{\kappa}{q(\theta_t)} = \beta(1 - \rho)E_t \left[A_{t+1} - w_{t+1} + \frac{\kappa}{q(\theta_{t+1})} \right]. \quad (\text{A11})$$

This captures the trade-off faced by the firm: the marginal, effective cost of posting a vacancy, $\frac{\kappa}{q(\theta_t)}$, 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.

Wages are determined based on the Nash bargaining solution: surpluses accruing to the matched parties are split according to a rule that maximizes the weighted average of the respective surpluses. Denoting the workers' weight in the bargaining process as $\eta \in [0, 1]$, this implies the sharing rule:

$$\mathcal{W}_t - \mathcal{U}_t = \frac{\eta}{1 - \eta} \mathcal{J}_t, \quad (\text{A12})$$

where \mathcal{W}_t is the asset value of employment, \mathcal{U}_t is the value of being unemployed, and \mathcal{J}_t is, as before, the value of the marginal worker to the firm.⁹

⁹In models with one-worker firms, the net surplus of a firm is given by $\mathcal{J}_t - \mathcal{V}_t$, with \mathcal{V}_t the value of a vacant job. By free entry, \mathcal{V}_t is then assumed to be driven to zero.

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

$$\mathcal{W}_t = w_t + \beta E_t[(1 - \rho)\mathcal{W}_{t+1} + \rho\mathcal{U}_{t+1}]. \quad (\text{A13})$$

Workers receive the wage w_t , and transition into unemployment next period with probability ρ . The value of searching for a job, when currently unemployed, is:

$$\mathcal{U}_t = b + \beta E_t[p_t(1 - \rho)\mathcal{W}_{t+1} + (1 - p_t(1 - \rho))\mathcal{U}_{t+1}]. \quad (\text{A14})$$

An unemployed searcher receives benefits b and transitions into employment with probability $p_t(1 - \rho)$. It is adjusted for the probability that a completed match gets dissolved before production begins next period. Substituting the asset equations into the sharing rule (A12), results, after some algebra, in the wage equation:

$$W_t = \eta(A_t + \kappa\theta_t) + (1 - \eta)b. \quad (\text{A15})$$

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 , 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 $\eta\kappa\theta_t$ from the firm.

I can substitute the wage equation and the job-matching rate into the job-creation condition to obtain:

$$\frac{\kappa}{m_t}\theta_t^\xi = \beta(1 - \rho)E_t \left[(1 - \eta)(A_{t+1} - b) - \eta\kappa\theta_{t+1} + \frac{\kappa}{m_{t+1}}\theta_{t+1}^\xi \right]. \quad (\text{A16})$$

Note that this expression is a first-order expectational difference equation in labor market tightness, with productivity as a driving process. Firms are more willing to post vacancies if productivity shocks increase the wedge to the outside option of the worker. In the empirical analysis, I make use of the simple structure of the model. The dynamics can be fully described by two equations, the employment accumulation equation (A5) and the job creation condition (A16), after convenient substitutions.

The full general equilibrium model consists of the following equations:

$$N_t = (1 - \rho) \left[N_{t-1} + m_{t-1} U_{t-1}^\xi V_{t-1}^{1-\xi} \right], \quad (\text{A17})$$

$$\frac{\kappa}{m_t} \theta_t^\xi = \beta E_t \left[(1 - \rho) \left(A_{t+1} - w_{t+1} + \frac{\kappa}{m_t} \theta_{t+1}^\xi \right) \right], \quad (\text{A18})$$

$$w_t = \eta (A_t + \kappa \theta_t) + (1 - \eta) b, \quad (\text{A19})$$

$$\theta_t = \frac{V_t}{U_t}, \quad (\text{A20})$$

$$N_t = 1 - U_t, \quad (\text{A21})$$

$$\log A_t = (1 - \rho_A) \log \bar{A} + \rho_A \log A_{t-1} + \varepsilon_{A,t}, \quad (\text{A22})$$

$$\log m_t = (1 - \rho_m) \log \bar{m} + \rho_m \log m_{t-1} + \varepsilon_{m,t}, \quad (\text{A23})$$

The first equation is the employment accumulation equation, followed by the job creation condition. The third equation describes the Nash-bargained wage, while the following two equations define labor market tightness and relate employment to unemployment. The description of the model is completed by the law of motions of the two exogenous stochastic processes.

The model is linearized around the steady state. I normalize aggregate productivity $A = 1$ and fix β , but estimate all remaining parameters. Denote the log-deviations of a variable from its steady state as $\tilde{x}_t = \log x_t - \log x$. I find it convenient to reduce the system to one in two endogenous variables only, viz. \tilde{N}_t and $\tilde{\theta}_t$, by substituting the wage equation into the job creation condition, the two definitional equations into the employment equations. The resulting linearized equation system is:

$$\tilde{N}_t = \left(1 - \frac{\rho}{1 - N} \right) \tilde{N}_{t-1} + \rho(1 - \xi) \tilde{\theta}_{t-1} - \frac{\rho}{1 - \rho} \tilde{\rho}_t, \quad (\text{A24})$$

$$\xi \tilde{\theta}_t = \beta(1 - \rho) \left(\xi - \eta m \theta^{1-\xi} \right) E_t \tilde{\theta}_{t+1} + \beta(1 - \rho) \frac{m}{\kappa \theta^\xi} E_t \tilde{A}_{t+1} - \frac{\rho}{1 - \rho} E_t \tilde{\rho}_{t+1}, \quad (\text{A25})$$

$$\tilde{A}_t = \rho_A \tilde{A}_{t-1} + \varepsilon_{A,t}, \quad (\text{A26})$$

$$\tilde{m}_t = \rho_m \tilde{m}_t + \varepsilon_{m,t}. \quad (\text{A27})$$

The four equations form a linear rational expectations (LRE) model in two endogenous variables and two exogenous shocks. The model has the special feature that it is block-diagonal: the job creation condition is an expectational difference equation in $\tilde{\theta}_t$ only and could thus be solved independently from the rest of system. The solution for tightness then feeds into employment dynamics as a driving process. It can easily be demonstrated that the system possesses a unique solution over the admissible parameter space. The system is solved using the method described in Sims (2001).

Table 1: Benchmark Prior

| Parameter | Mean | St.Dev. | Distr. | Source |
|----------------------------|------|---------|--------|-----------------------------------|
| Separation Rate ρ | 0.03 | 0.001 | Beta | Shimer (2005); monthly JOLTS data |
| Match Elasticity ξ | 0.49 | 0.1 | Beta | Beveridge Curve Estimation |
| Match Efficiency m | 0.80 | 0.1 | Gamma | Beveridge Curve Estimation |
| Benefit b | 0.9 | 0.1 | Beta | Hagedorn and Manovskii (2008) |
| Bargaining η | 0.49 | 0.1 | Beta | Hosios-condition, wide prior |
| Job Creation Cost κ | 0.18 | 0.001 | Gamma | Steady-State Sample Mean |
| Discount Factor β | 0.99 | -- | Fixed | Annual Real Interest Rate |
| Productivity A | 1 | -- | Fixed | Normalized |

Table 2: Posterior Means

| Parameter | Prior Mean | Period I | Period II | Full Sample |
|------------------------|------------|----------------------|----------------------|----------------------|
| Match Elasticity ξ | 0.49 | 0.81 [0.75, 0.88] | 0.69 [0.60, 0.78] | 0.83 [0.77, 0.88] |
| Match Efficiency m | 0.80 | 0.91 [0.77, 1.05] | 0.70 [0.57, 0.84] | 0.82 [0.70, 0.95] |
| Benefit b | 0.90 | 0.70 [0.55, 0.86] | 0.82 [0.72, 0.92] | 0.73 [0.59, 0.87] |
| Bargaining η | 0.49 | 0.47 [0.32, 0.62] | 0.60 [0.45, 0.75] | 0.51 [0.35, 0.66] |
| AR-coeff. A | 0.90 | 0.96 | 0.94 | 0.98 |
| AR-coeff. m | 0.90 | 0.94 | 0.93 | 0.98 |
| StD. A | 0.01 | 0.009 | 0.006 | 0.008 |
| StD. m | 0.01 | 0.013 | 0.016 | 0.014 |

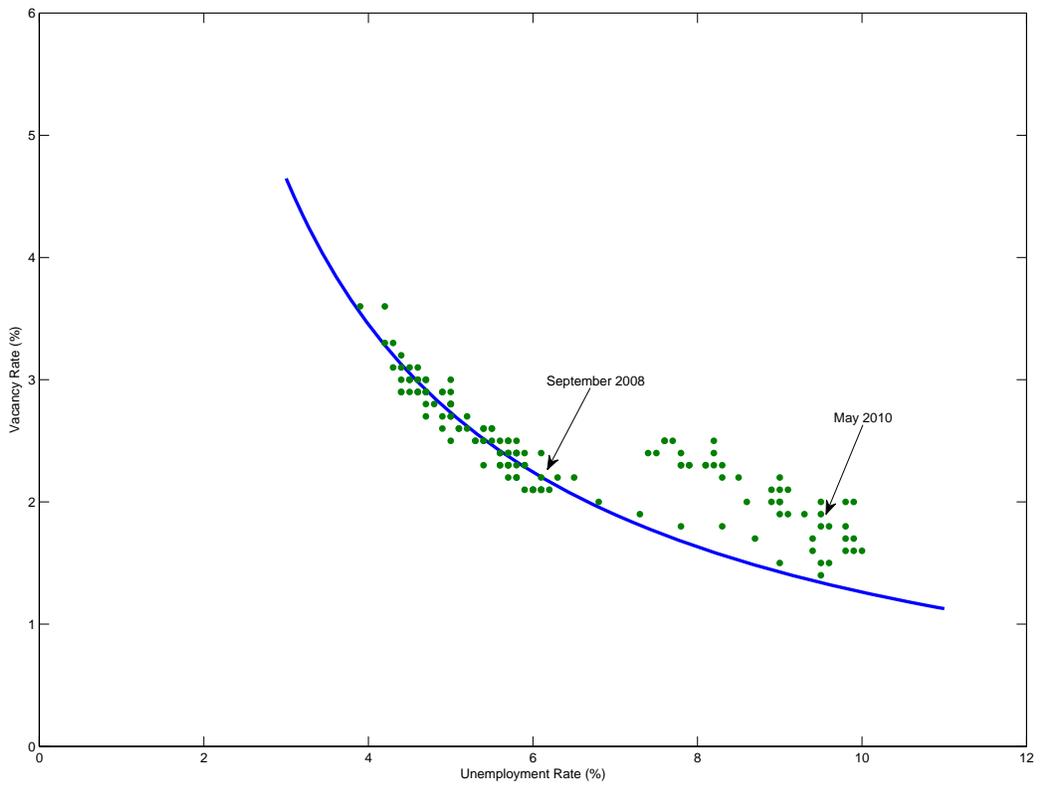


Figure 1: Estimated Beveridge Curve

Table 3: Variance Decompositions

| Variable | Shock | Period I | Period II | Full Sample |
|--------------|------------------|----------|-----------|-------------|
| Unemployment | Productivity | 0.85 | 0.81 | 0.79 |
| | Match Efficiency | 0.15 | 0.19 | 0.21 |
| Vacancies | Productivity | 0.21 | 0.31 | 0.13 |
| | Match Efficiency | 0.79 | 0.69 | 0.87 |

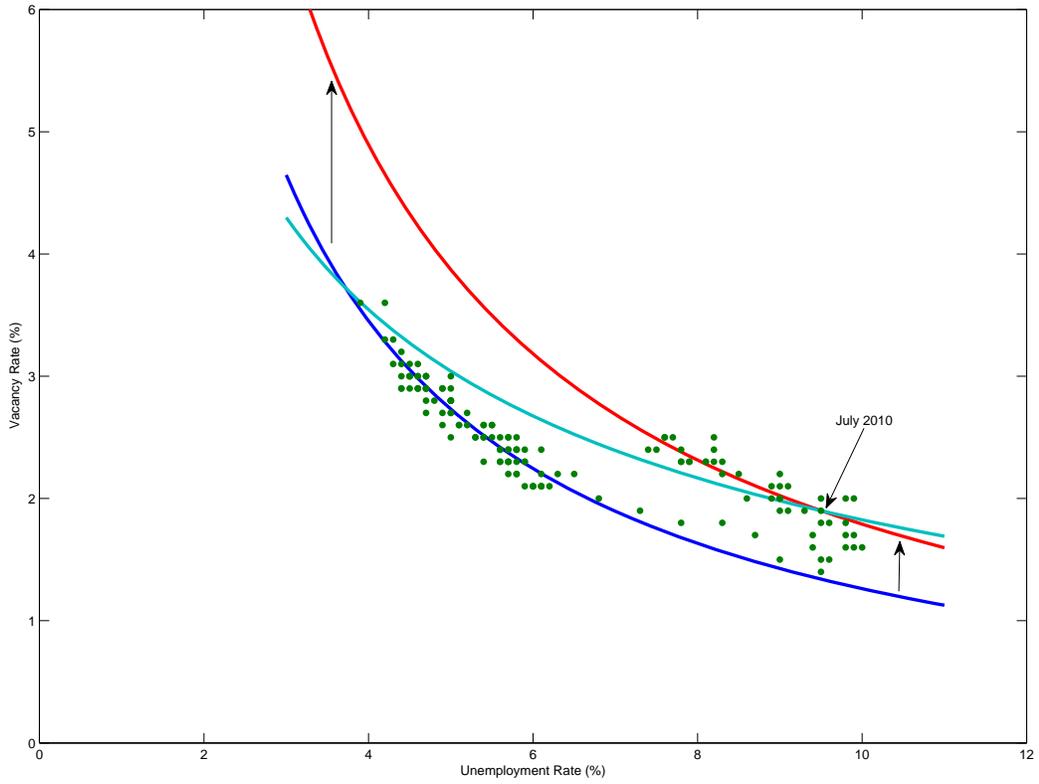


Figure 2: Beveridge Curve Shifts

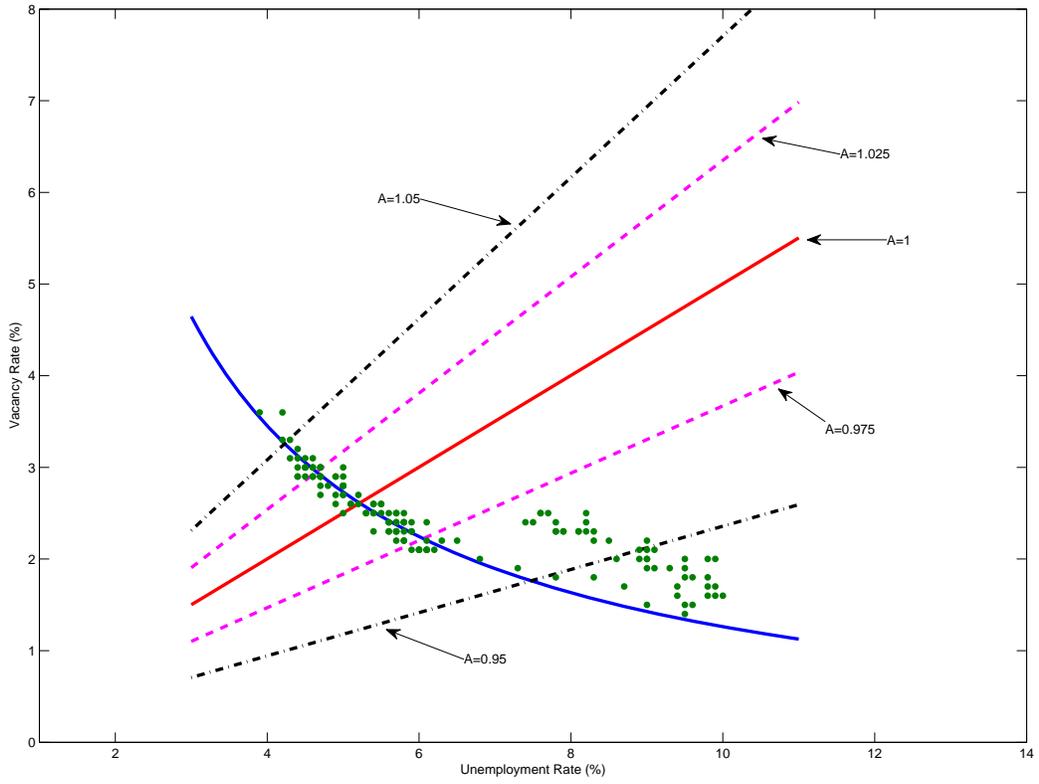


Figure 3: Beveridge Curve and Job Creation Conditions

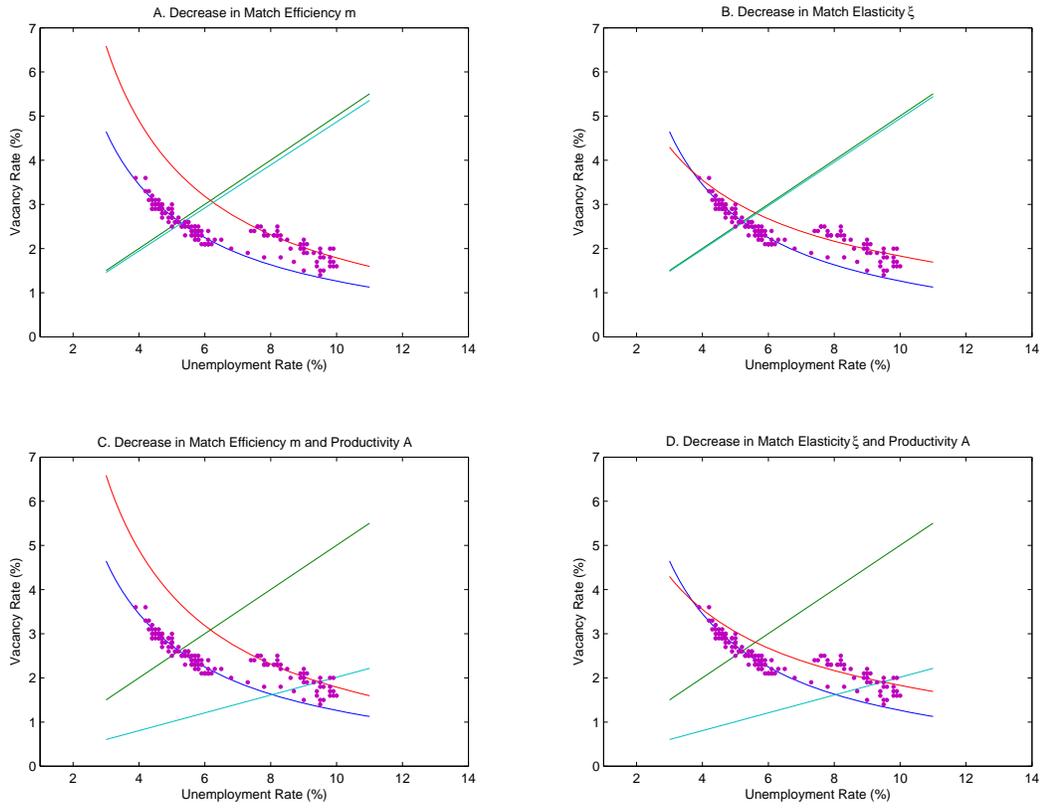


Figure 4: The Beveridge Curve and Parameter Shifts

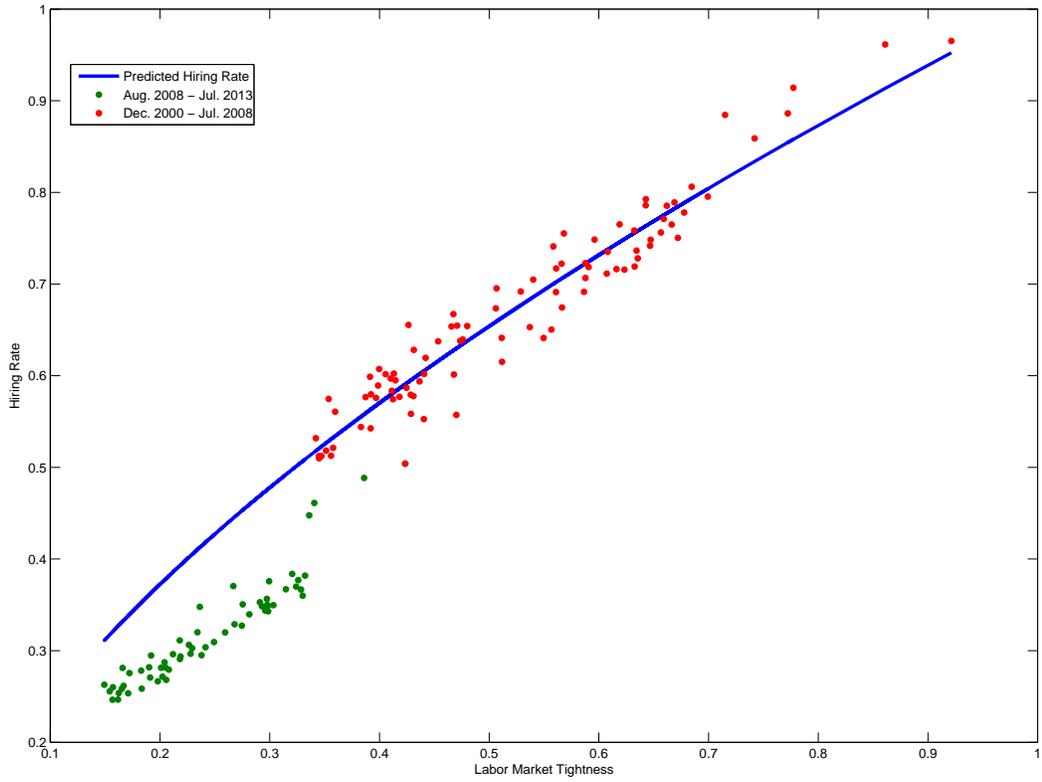


Figure 5: Labor Market Tightness and Hiring

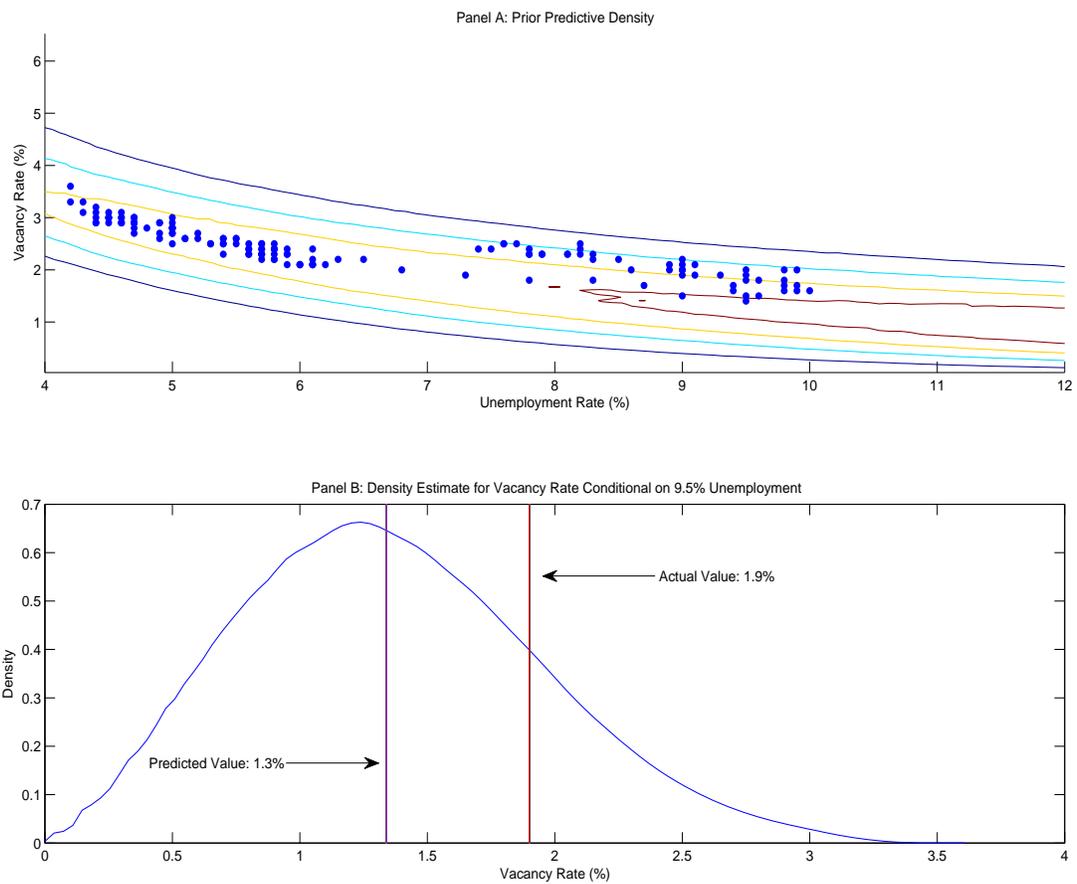


Figure 6: Prior Predictive Density

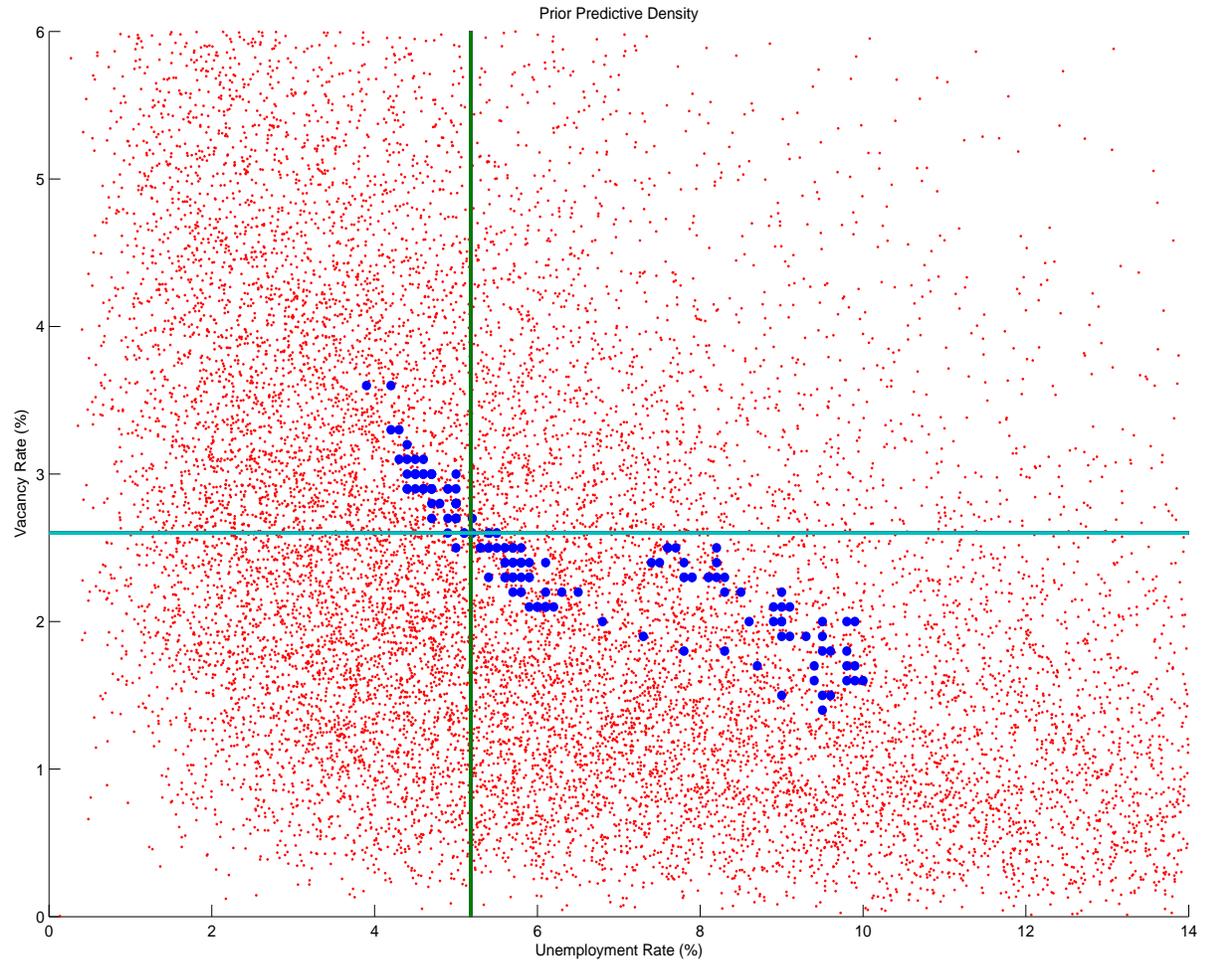


Figure 7: Prior Predictive Density: Scatter Plot

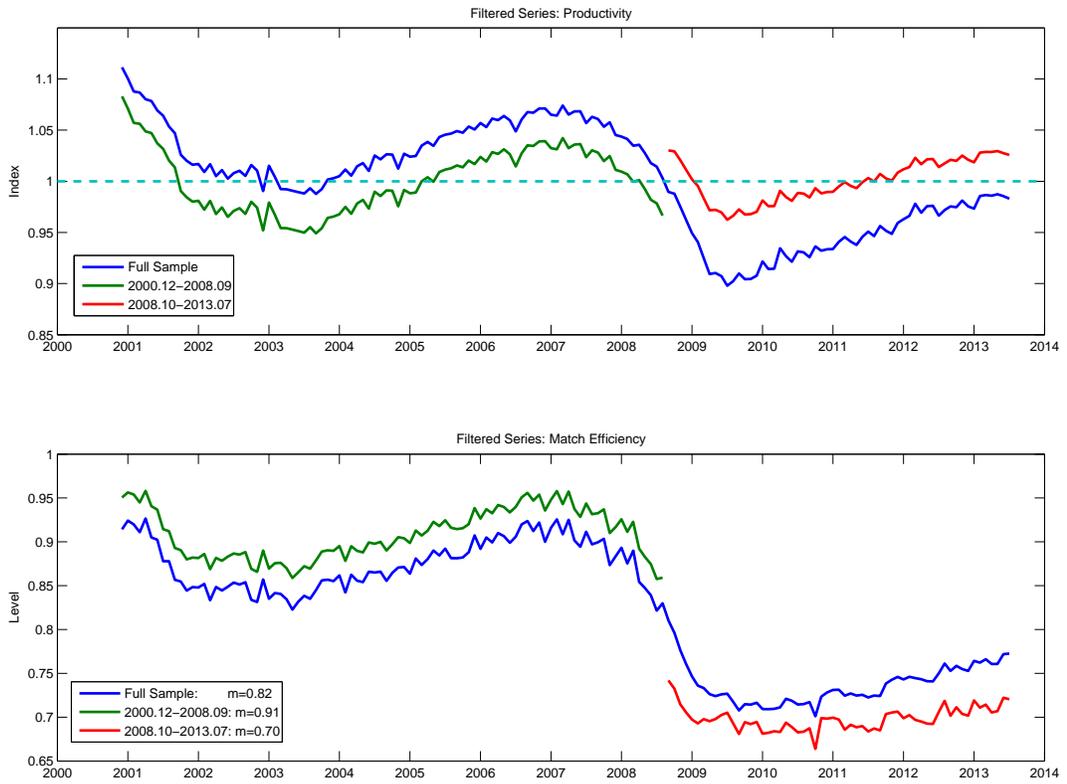


Figure 8: Filtered Shocks

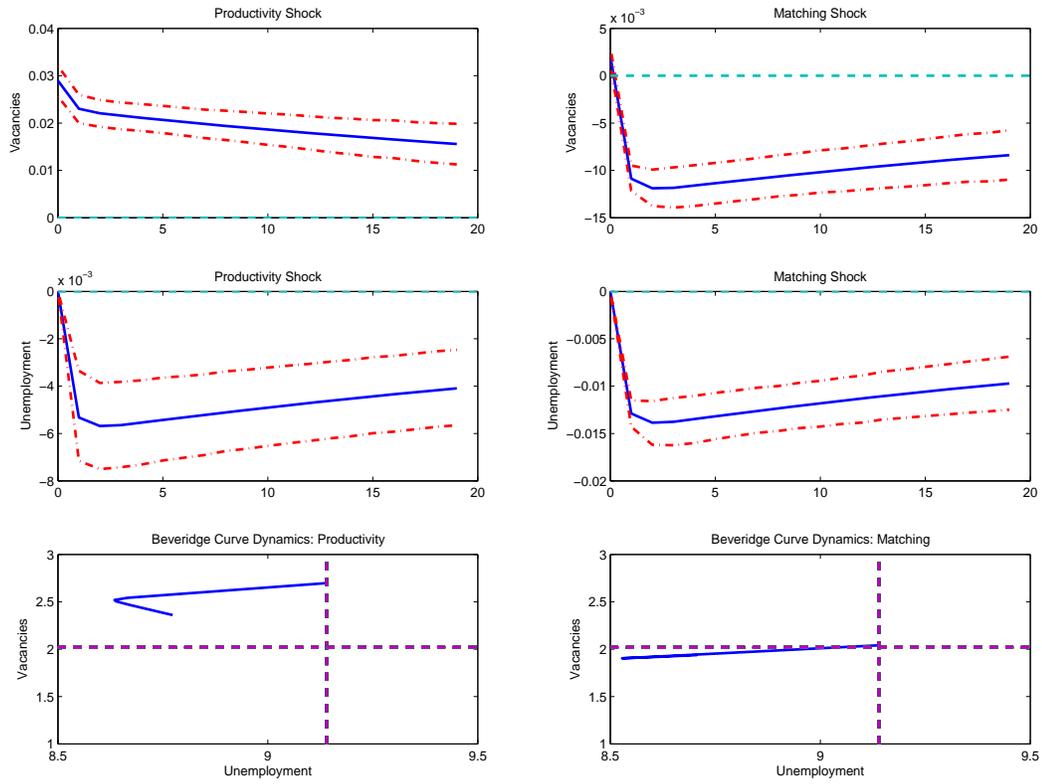


Figure 9: Impulse Response Functions and Beveridge Curve Dynamics