Structural Unemployment in the 2008 Recession

Anton A. Cheremukhin

Federal Reserve Bank of Dallas

January 12, 2011

Summary

A number of questions have been raised recently regarding the sharp increase in unemployment and slow recovery of jobs throughout the 2008 recession.

1) How much of that unemployment surge is structural, and how much is frictional?

2) Is there mismatch between the demand for and supply of workers across industries?

The goal of this memo is to try to answer these questions exploring data on aggregate and sectoral job and worker stocks and flows.

To have a clear understanding of the questions, it is useful to define the difference between structural and frictional unemployment. Structural unemployment usually refers to unemployment that results from a mismatch between the *characteristics* of jobs supplied and demanded, while frictional unemployment is thought to be a consequence of mismatch in their *quantities*[1]. Frictional unemployment is manifested by high labor supply coexisting with slack demand for work. An indication of structural unemployment would be unusually high unmatched demand for workers coexisting with high labor supply.

A standard way to analyze changes in structural unemployment is to look at the Beveridge curve - the relationship between unemployment and vacancy rates (Figure 1). When unemployment is frictional, high labor supply coexists with slack labor demand: times of higher unemployment should be times with lower numbers of vacant jobs. This corresponds to a downward sloping relationship

[1]Here I merge the notions of frictional and cyclical unemployment. The distinction between them is not very clear to me. For a broader discussion see Rogerson (1997).

like that depicted in Figure 1 in green. The red line depicts recent developments: the economy initially follows the same downward-sloping curve it was on previously. However, starting January 2010, a substantial increase in the number of available jobs did not result in a corresponding decrease in the unemployment rate. The natural question is whether the new pattern is a manifestation of an increased degree of mismatch between unemployed workers and available jobs.

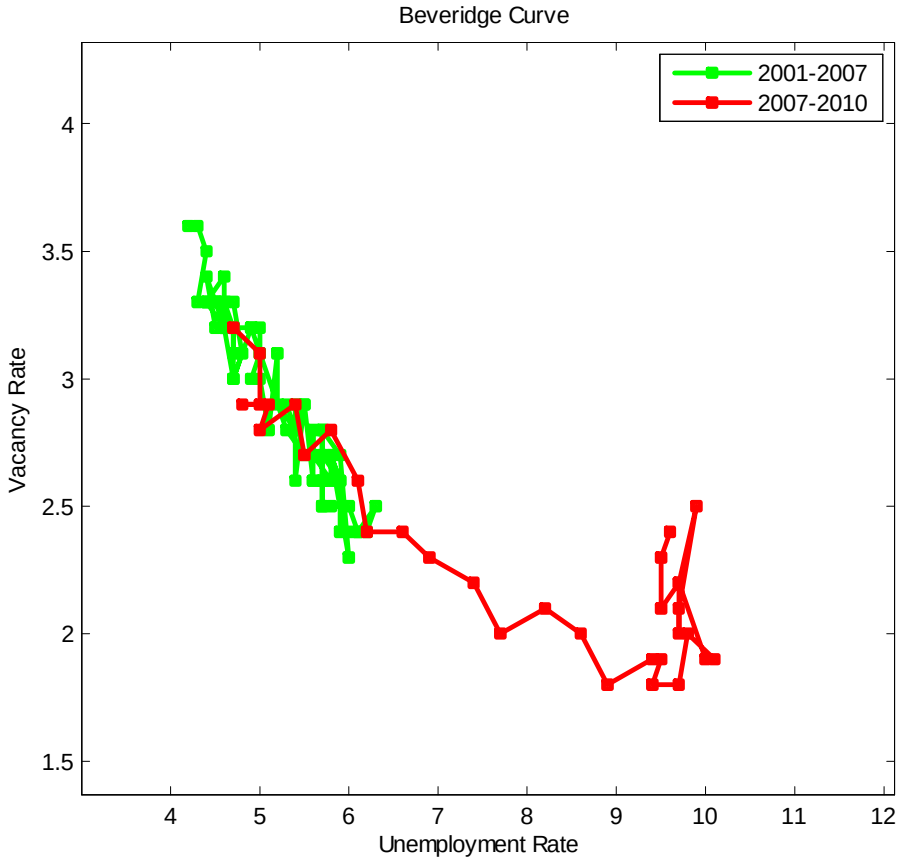

Figure 1. Unemployment and Vacancy Rates

Source: BLS, JOLTS

In this memo, I first describe the available data and compare the current recession with previous recession episodes. I show that the recent counter-clockwise loop around a downward-sloping Beveridge curve is a feature of most previous recessions. Second, I construct Beveridge curves for major subsectors of the US economy and document the same pattern in most sectors. Third, I explore

data on job and worker flows at both the aggregate and sectoral levels. I find that, as in previous recessions, the main channel of adjustment is the job destruction margin. In worsened economic conditions, firms on average prefer to fire more people rather than hire less.

Finally, I look at these patterns in the data through the lens of a search and matching model based on ideas developed by recent Nobel laureates Mortensen and Pissarides. This model incorporates decisions of firms to create and destroy jobs and can explain the path of the economy up to 2009, but not after that. The model captures the spike in job destruction in 2008, the consequent sharp increase in unemployment as well as the decline in vacancies. However, it cannot explain why the increase in vacancies since mid-2009 has not translated into lower unemployment.

Within this framework I ask what would happen if the degree of mismatch between unemployed workers and firms increased. The model predicts that firms would perceive already hired workers as much more valuable and good new matches as much more costly. Firms would then choose to both destroy and create *fewer* jobs and post *fewer* vacancies, without any change in unemployment.

The implication of the theory is that a greater mismatch ought to lead to a downward shift in the Beveridge curve and a substantial decrease in turnover. Neither of these predicted behaviors can be observed in the current recession. I reach the conclusion that there is little or no evidence of an increase in structural unemployment during the current downturn at either the aggregate or sectoral level.

In addition, I propose an alternative explanation to the observed counter-clockwise movements around the Beveridge curve. I show that these movements can be explained by the extra time it takes to set a matched worker in place and by the temporary discouraging effect unemployment insurance has on search efforts of unemployed workers.

Historical Evidence

I begin by describing available data on unemployment and vacancies. Unemployment for persons 16 years or older is measured by BLS on a monthly basis since 1948. For the period since 2001 depicted in Figure 1, a good measure of vacancies at both the aggregate and sectoral level is provided by JOLTS. For the period prior to 2001, the only available measure is the index of help wanted advertisements in newspapers constructed by the Conference Board. I merge these two datasets by normalizing the help-wanted index with the JOLTS January 2001 value. Figure 2 plots the paths of unemployment and vacancy rates during recessions (red) and recoveries (green) for all NBER recessions since 1953. The black line represents the 2008 recession.

Figure 2 puts the Beveridge relationship into a historical perspective. It shows that in *each and every* previous recession there was a counter-clockwise loop around the Beveridge curve. These loops resulted in average differences between contraction and expansion paths of 1-1.5% in terms of the unemployment rate. Most of the time, this looping did not lead to large medium-run shifts, while the apparent long-run shifts in the Beveridge curve can be attributed to demographic changes resulting from variations in participation rates across different demographic groups. This simple calculation leaves only around 1 percentage point of unemployment unexplained in the current recession, which could potentially be attributed to some sort of mismatch between the quality of labor supply and labor demand.

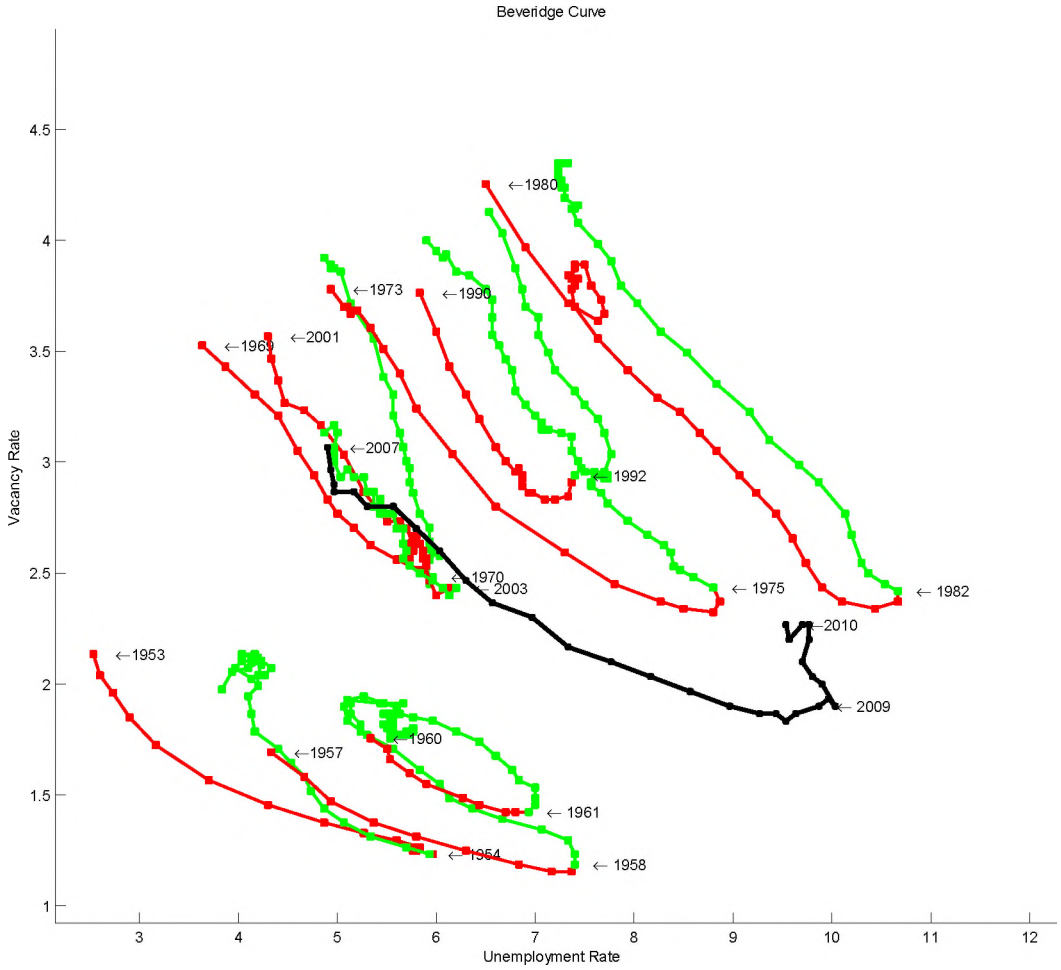

Figure 2. Beveridge Curve in Historical Perspective

Source: BLS, Conference Board, JOLTS, author's calculations

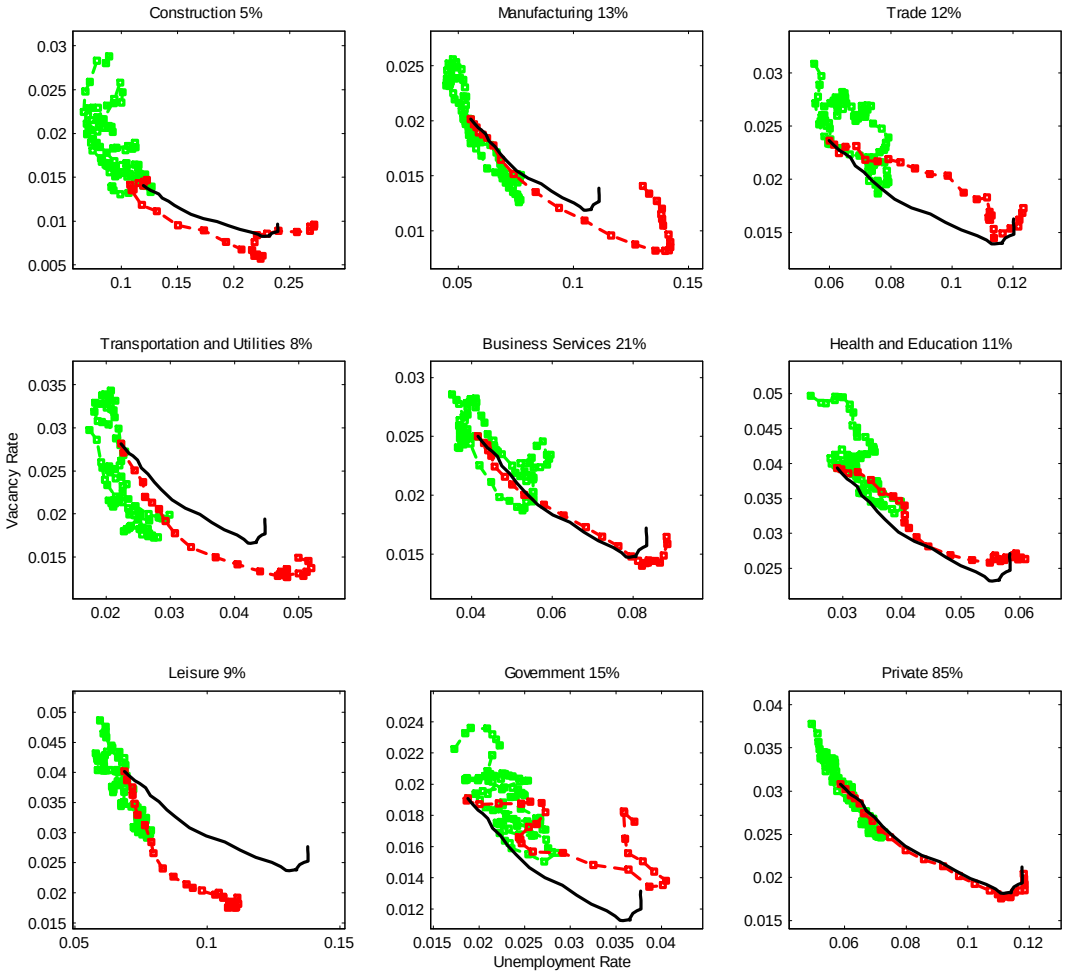

Figure 3. Sectoral Beveridge Curves

Source: Household Survey, CPS, JOLTS, author's calculations

Sectoral Evidence

In this section, I plot Beveridge curves for major sectors of the US economy. I take sectoral data on monthly stocks of unemployed workers and job openings available from the BLS and JOLTS since 2001. I convert these into rates with respect to sectoral labor forces, calculated as sums of employed and unemployed in each sector. To reduce the amount of noise in the graphs, I compute 3-month moving averages. Figure 3 plots sectoral Beveridge curves for major sub-sectors of the US economy for the available period from 2001 to 2010. It compares the sectoral paths prior to 2007 (green) and since 2007 (red) with the normalized path for the aggregate economy (black). These graphs demonstrate that even though the slopes of the Beveridge curve differ across sectors, the downward sloping relationship between unemployment and vacancies and the upward loop observed at the aggregate level are mimicked very closely by each sector of the economy. This evidence makes it much harder to make a case for mismatch across sectors.

Job and Worker Flows

In addition to worker and job stocks, for further analysis it is useful to look at worker and job flows. I plot four measures of job destruction and three measures of job creation in Figure 4. The first measure of job destruction is the difference between separations and quits measured by JOLTS. The implicit assumption behind this measure is that people quit their jobs for a purpose: because they are leaving the labor force or because they have a better job waiting for them. Separations through quits should not result in unemployment. Another measure of job destruction also from JOLTS is the number of layoffs. When a person has no choice on whether to leave the job, he will most probably end up unemployed. The third measure of job destruction is the worker flow measure

from employment to unemployment tabulated from the CPS. Finally, a measure of job losses in contracting establishments has been constructed from the BED (Business Employment Dynamics). The upper panel of Figure 4 compares these four measures of job destruction aggregated to a quarterly frequency. In addition, it plots the number of initial claims for unemployment insurance for the same time period.

The message from the first panel is that even though there is a wide discrepancy in measures of job destruction, they all show a large spike in job destruction starting in the second half of 2008 and fading off by early 2010. This spike of job destruction translated into a more than proportional increase in initial claims for unemployment insurance.

The second panel plots three measures of job creation for the same time period. The first measure is the number of hires that did not come from quits. The implicit assumption is that people who quit mostly do so because they found a better job, while the rest of the hires come from among the unemployed. There is little change in this measure during the 2008 recession and later on. The second measure is the flow of workers from unemployment to employment measured from the CPS. This measure shows a gradual increase in job creation starting in the first half of 2009. Finally, the sum of all job gains in expanding establishments constructed by the BED shows a small downward shift in job creation. Thus, the second panel shows a large discrepancy in directions of movements across different measures of job creation. The message I take from the second panel is that variations in creation are much milder than in job destruction. In worsened economic conditions, firms fire more people quickly and then slowly hire them back as conditions improve.

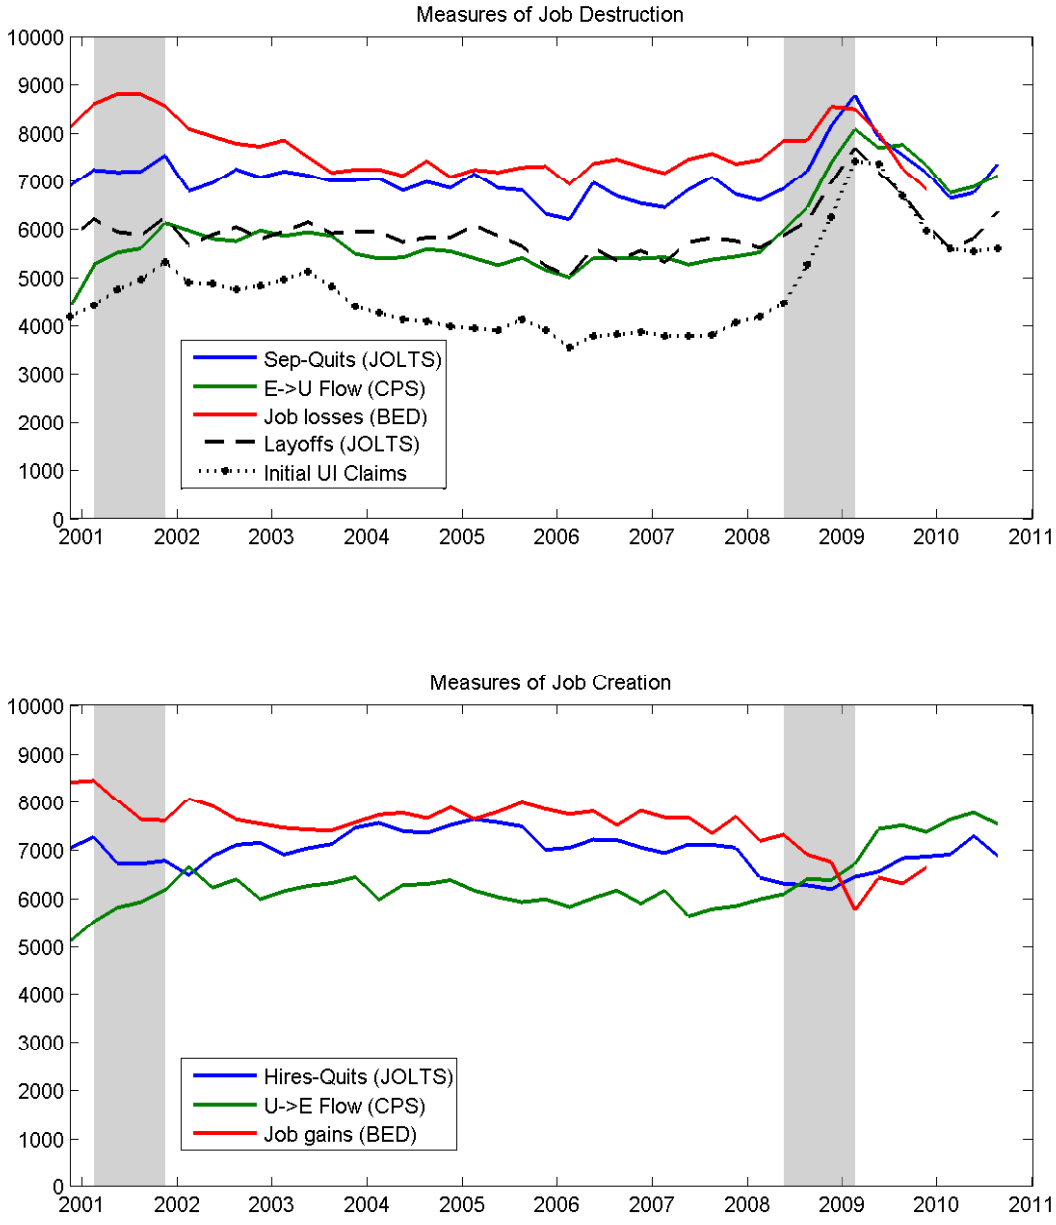

Figure 4. Worker and Job Flows

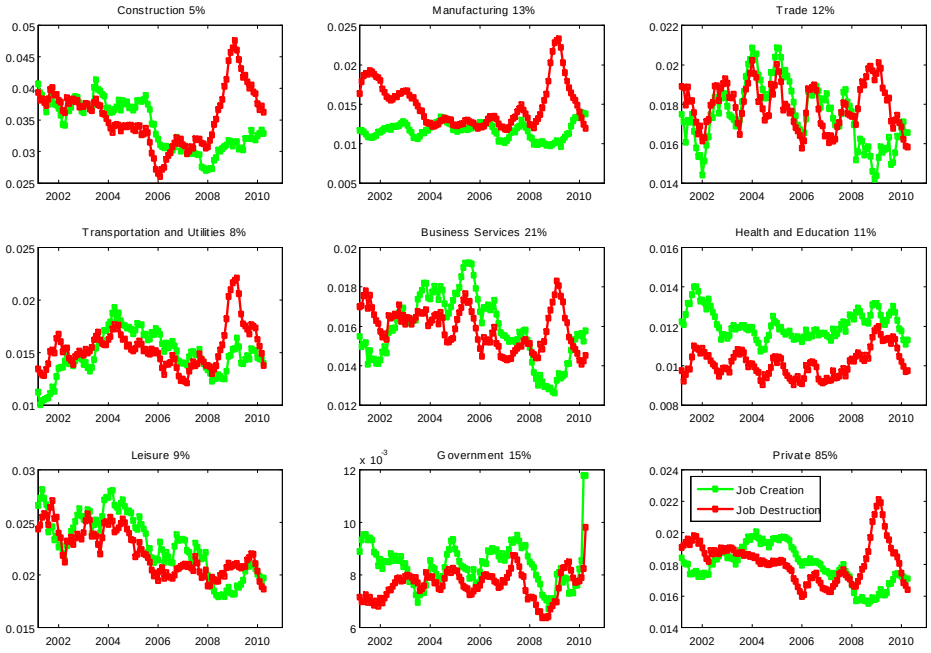

Figure 5. Job Creation and Destruction across Major Sectors

Source: JOLTS, subtracting quits from hires and separations

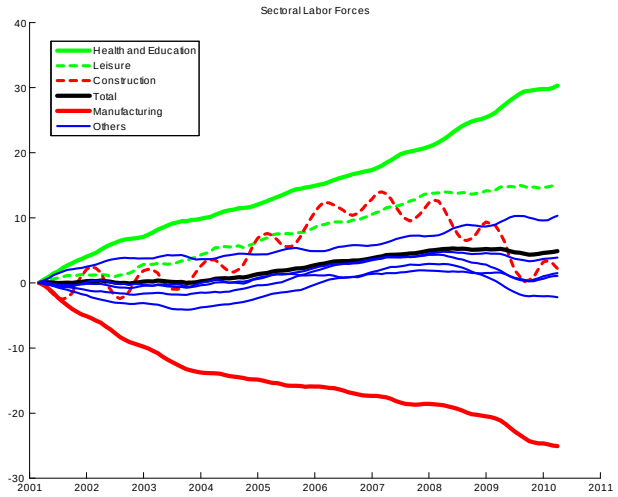

Figure 6. Sectoral Labor Forces, percentage changes

Source: CPS, Household Survey, author's calculations

Similar patterns can be easily demonstrated using two of these measures (JOLTS and BED) for most subsectors of the economy. Using JOLTS data, Figure 5 plots sectoral measures of creation and destruction computed as percentages of sectoral labor forces. For a broad range of sectors of the US economy, a large initial spike in job destruction in the second half of 2008 is followed by a pickup in job creation in the second half of 2009. Notable exceptions are government, health and education, and leisure. The government sector has a spike in job creation followed by a spike in job destruction corresponding to the temporary hires of census collectors in 2010. Health and education and leisure show remarkable stability in both creation and destruction rates. Unlike others, these sectors have creation rates on average higher than destruction rates. This might be a consequence of the large shifts in the structures of production and consequent shifts in the composition of the labor force. Figure 6 documents the decade-long gradual expansion of the health and education and leisure sectors at the expense of the manufacturing sector.

MP Model

In this section I interpret the patterns observed in the data through the lens of a search and matching model based on ideas developed in the seminal paper by Mortensen and Pissarides (1994). I construct a model in which idle firms and unemployed workers spend time searching for each other in order to form productive relationships. Matched firms and workers form employment relationships that may differ in their productivity. As time goes on, depending on their match-specific productivity, firms and workers can choose to separate and destroy their relationship. Similarly to the model constructed in Cheremukhin (2010), I incorporate the idea that most of the costs of creating a new job come from setting up the environment, buying capital, and training workers, while the cost of recruiting is relatively low. I calibrate the model using historical data for unemployment, vacancies

and job flows in order to match the path of the economy prior to the 2008 recession using a single aggregate shock.

Figure 7 depicts the response of the model economy to a persistent negative aggregate demand shock. The model predicts a spike in job destruction early in the recession, a consequent sharp increase in unemployment as well as the decline in vacancies. Figure 8 shows that the model can reproduce the negatively sloped Beveridge curve observed in the data. However, it cannot explain why the increase in vacancies since mid-2009 has not translated into lower unemployment.

The second step is to use this framework to study the consequences of mismatch. I ask the model what would happen if the degree of mismatch between unemployed workers and firms increased. That would imply that a given number of unemployed and job openings would result in a smaller number of matches. Figure 9 depicts the response of the model economy to such an exogenous shock. The model predicts that firms would perceive already hired workers as much more valuable and good matches as much more costly. As a result, firms would choose to destroy and create *fewer* jobs and post *fewer* vacancies, without any change in unemployment.

According to the model, a higher degree of mismatch leads to a downward shift in the Beveridge curve (see Figure 10) simultaneously with a substantial decrease in turnover. What we would expect to see in the data is a decrease, rather than increase, in vacancies and a simultaneous decrease in both job creation and job destruction rates. None of these predictions can be observed in the current recession. This leads me to conclude that there is little or no evidence of an increase in structural unemployment during the current downturn at either the aggregate or sectoral level.

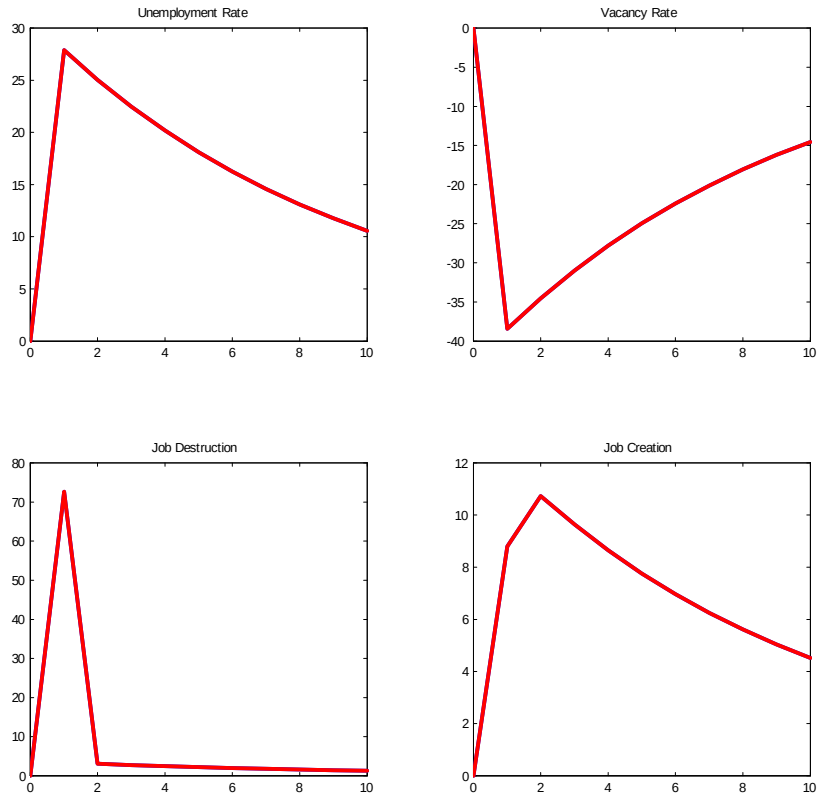

Figure 7. Response of Labor Market Variables to a Negative Demand Shock

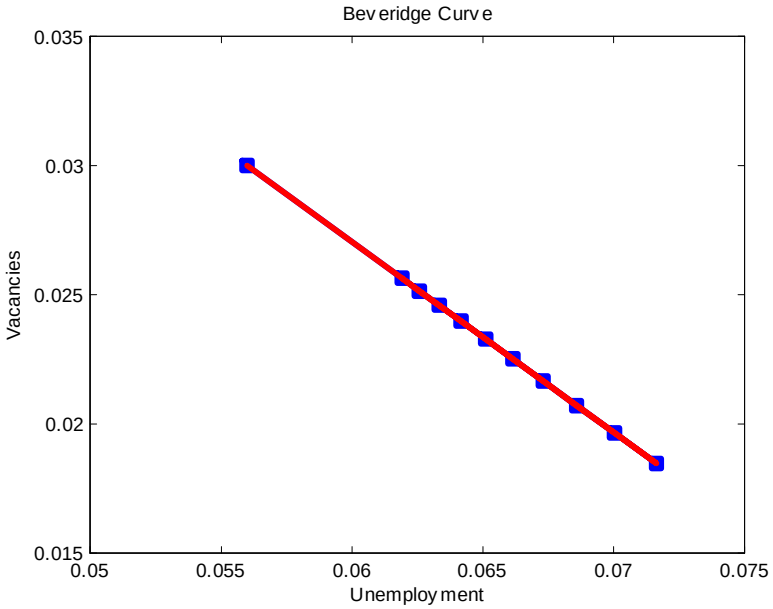

Figure 8. Beveridge Curve Produced by the Benchmark Model

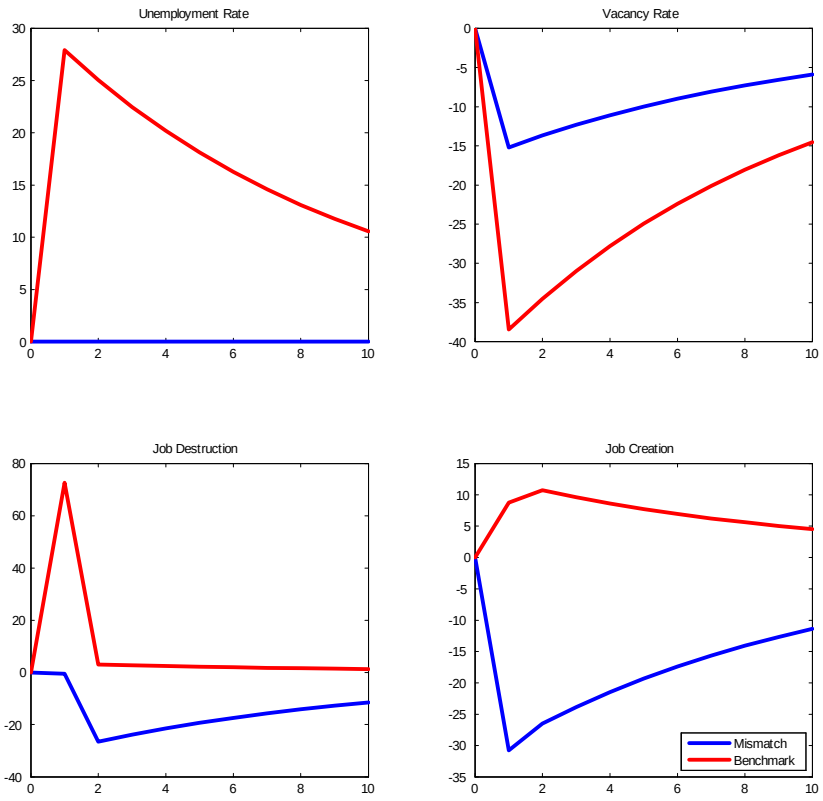

Figure 9. Response of the Model Economy to a Mismatch Shock

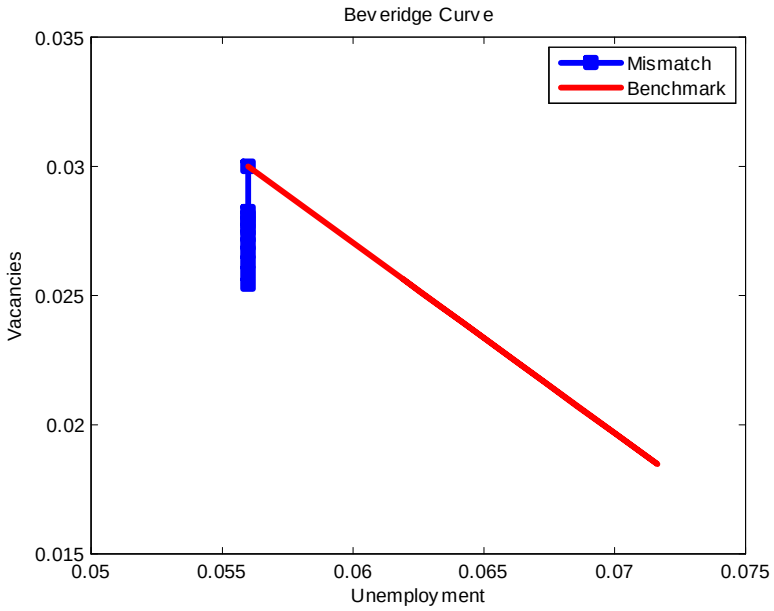

Figure 10. Response of the Beverage Curve to a Mismatch Shock

If mismatch is not the cause of the shift in the Beveridge curve, then why do job openings not result in a decrease in unemployment? I propose two potential explanations for this pattern. First, part of currently and previously observed counter-clockwise movements around the Beveridge curve can be attributed to the lag between initially posting a vacancy, then finding a good match and making all the preparations for the employee to start working. I capture this idea by introducing a 1-quarter delay between the moment the worker is matched and the moment the worker becomes productive. Figure 11 depicts the effect this delay has on the movements around the Beveridge curve. The effect is in the right direction, and might be enough to explain some previous episodes, but it cannot explain all of the current shift.

The second mechanism that can help explain the recent movements is the effect unemployment benefits can have on the effort workers put into searching for new jobs. Meyer (1990) shows that coming close to expiration of unemployment benefits increases the probability of the worker finding a job by a factor of 2 to 3. I incorporate this idea into the model economy by assuming that a worker who has just lost his job puts in three times less search effort than normal. The joint effect of the delay and temporary discouraging effects of unemployment benefits on the Beveridge curve is depicted in the second panel of Figure 11. It clearly goes in the right direction, explaining a much bigger portion of the deviation of the economy from a straight Beveridge curve.

Rather than blaming the slow recovery on the ambiguous notion of structural unemployment and mismatch between workers and job skills, one should think more carefully of the effects of existing unemployment insurance and employment protection policies.

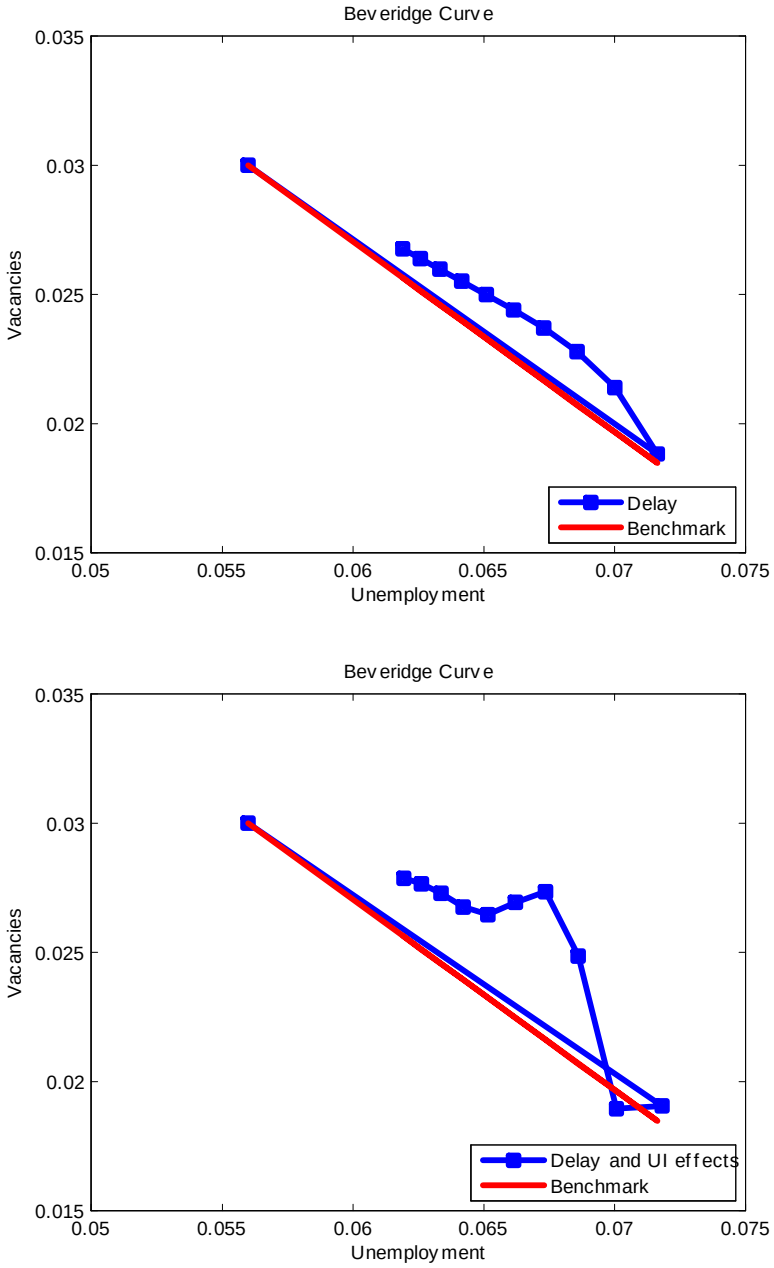

Figure 11. Effects of Delay and UI on the Beveridge Curve

Conclusion

- Counter–clockwise departures from the Beveridge curve are common in the early stages of recoveries.
- These departures are unlikely to be due to a greater mismatch between the skills of laid off workers and the requirements of available jobs.
- Instead, they may reflect effects of lags between new matches and productive employment and search disincentives due to unemployment insurance.

References

1. Cheremukhin, Anton A, 2010. "Labor Matching Model: Putting the Pieces Together," unpublished manuscript.
2. Meyer, Bruce D, 1990. "Unemployment Insurance and Unemployment Spells," Econometrica, Econometric Society, vol. 58(4), pages 757-82, July.
3. Mortensen, Dale T & Pissarides, Christopher A, 1994. "Job Creation and Job Destruction in the Theory of Unemployment," Review of Economic Studies, Blackwell Publishing, vol. 61(3), pages 397-415, July.
4. Rogerson, Richard, 1997. "Theory Ahead of Language in the Economics of Unemployment," The Journal of Economic Perspectives, vol. 11(1), pp. 73-92, Winter.