

Analyses of Select Factors Affecting the Current Level of Structural Unemployment

*Daniel Aaronson, Bhashkar Mazumder, and Gadi Barlevy
Federal Reserve Bank of Chicago*

This memo reports on four recent analyses by Chicago Fed economists that are relevant to the issue of the current extent of structural unemployment.

First, we present some simple unemployment rate calculations for recent college graduates and those with usually highly sought after technical skills. The motivation is that concerns about structural unemployment are at least arguably less applicable for such workers. Nonetheless, we find that the unemployment patterns for such workers mirror those for the economy as a whole.

Then, in “A Brief Note on Homeownership and Mobility: Recent Evidence from the SIPP,” Daniel Aaronson examines the issue of *house lock*, using data from the Survey of Income and Program Participation. He finds little evidence for such a factor causing a decline in migration.

Next, in “The Effects of Unemployment Insurance Extensions on the Unemployment Rate in the 2008 to 2010 Period,” Bhashkar Mazumder refines and updates previous work on the likely effect of unemployment insurance extensions. His analysis is based on a simulation that applies to the current episode a preferred estimate from the academic literature of the effects of UI extensions on the duration of unemployment. He concludes that extensions have likely raised the equilibrium level of unemployment by about 0.8 percentage points.

Finally, in “Mismatch and Steady-State Unemployment,” Gadi Barlevy provides an upper bound on the implied steady-state unemployment rate under the assumption that all of the recent shift in the Beveridge curve is due to a decline in the match efficiency of a Cobb-Douglas matching function. Barlevy shows that even under this probably unrealistic assumption, the new steady-state unemployment rate would be less than 7.8%. He also compares the fit of this standard Beveridge curve formulation to that of Robert Shimer’s recent formulation of the concept of mismatch. He finds that Shimer’s formulation under-predicted vacancies in the earlier episodes as well as the current episode.

Unemployment Rates for Select Groups of Workers

Daniel Aaronson, Federal Reserve Bank of Chicago

This brief section explores the labor market performance of groups of workers that are likely to be less vulnerable to facing “structural” impediments to finding a job.

Figure 1 plots the unemployment rate of recent college graduates¹ (blue line) against the aggregate unemployment rate (red line). The outcomes of young, highly skilled workers are particularly interesting for three reasons. First, having just completed college, they are less likely to have obsolete skills. Second, the young and well educated are traditionally the most geographically mobile part of the population, partly because of their low homeownership rate. Finally, young workers typically lack the work experience to qualify for unemployment insurance (UI). Therefore, young college graduates may provide a representation of labor market performance that minimizes misallocation problems arising from having the wrong skills, living in the wrong location, or being discouraged to search due to UI extensions.

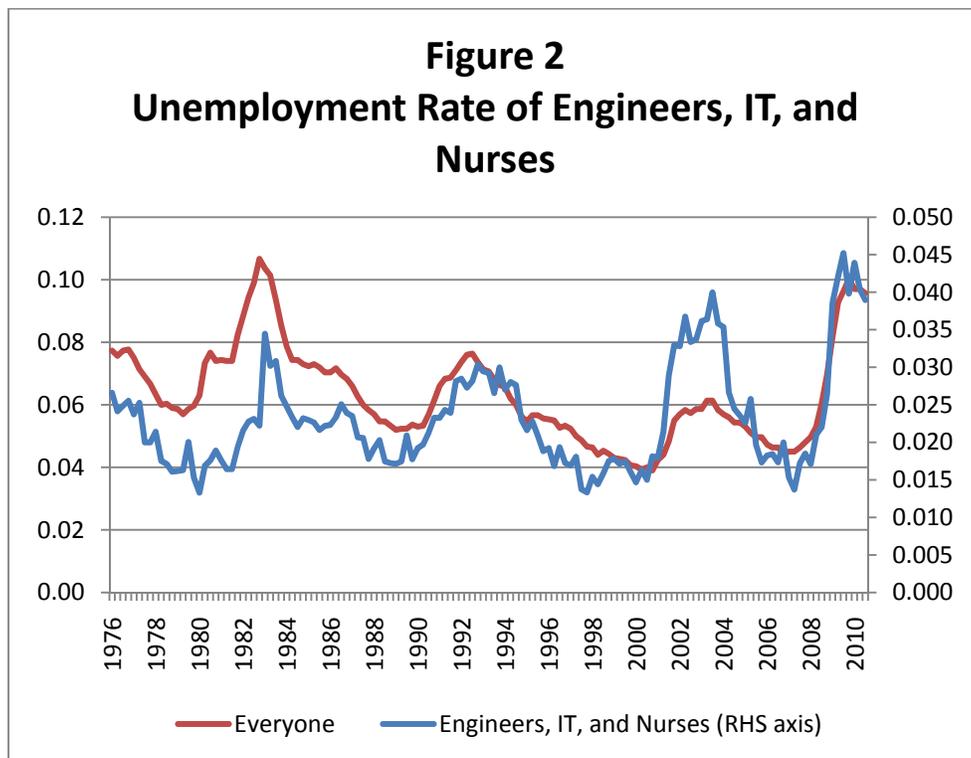
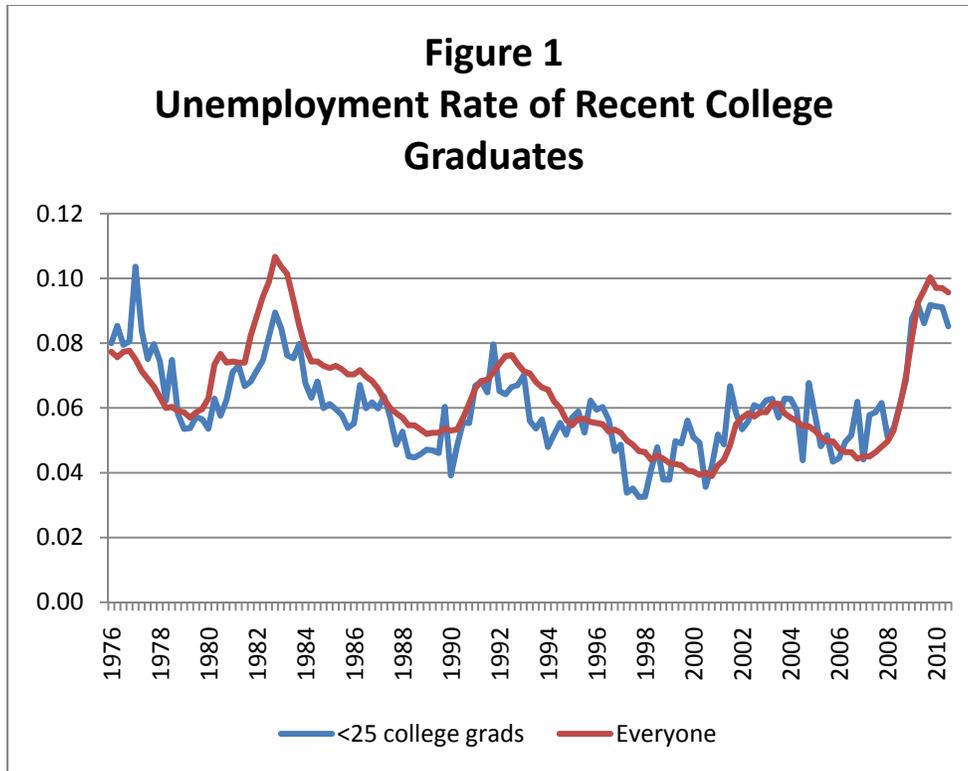
The figure shows that the unemployment rate of recent college graduates mirrors that of the entire economy, rising from roughly 5 percent in 2006 to 9 percent in the first three quarters of 2010. As of the third quarter, the rate has yet to fall appreciably and stands about 0.8 percentage points below the aggregate unemployment rate. We view these facts as inconsistent with a sizable role for mismatch and UI disincentives. Indeed, that 0.8 percentage point gap between young college graduates and the aggregate labor force is noticeably smaller than the gap in 1982-83.

Some have argued that current high unemployment reflects, in part, an unusual lack of workers with the skills that employers demand. Indeed, one reasonable concern about this analysis of young college-educated individuals is that they may lack the specialized expertise accumulated with work experience that employers are seeking. To get an idea of skill effects, in figure 2 we plot the unemployment rate for all individuals that last worked or are currently working in one of three high-skilled occupations: engineers, IT personnel, and nurses.² Again, we find that the labor market performance of these workers mirrors that of the aggregate economy, rising proportionally during 2008 and 2009 and holding steady in 2010.

Figures 1 and 2 are stylized examples, but they are consistent with the broad-based weakness exhibited in more general statistical exercises attempting to quantify reallocation, like Rissman (2009, updated in 2010) and Sahin et al. (2010). They are also consistent with the findings in Aaronson et al. (2009) and Foote and Ryan (2010), who show that the rise in unemployment spell length has been broad-based.

¹ Specifically, for each quarter, we use the CPS micro data to compute the number of college graduates under the age of 25 who are in the labor force and the number who are unemployed, seasonally adjust each series using X12, and calculate the ratio as the unemployment rate.

² Again, these are computed from the CPS micro data.



References

Aaronson, Daniel, Bhashkar Mazumder, and Shani Schechter, 2010, "What is Behind the Rise in Long-Term Unemployment?" Federal Reserve Bank of Chicago, *Economic Perspectives*, Second Quarter, pp. 28–51.

Foote, Chris, and Richard Ryan, 2010, "Understanding Recent Increases in Unemployment Duration," Federal Reserve Bank of Boston, mimeo.

Rissman, Ellen, 2009, "Employment Growth: Cyclical Movements or Structural Change?" *Economic Perspectives*, Fourth Quarter, pp. 40–57.

Şahin Ayşegül, Joseph Song, Giorgio Topa, and Gianluca Violante, 2010, "Measuring Mismatch in the U.S. Labor Market," mimeo.

The Effects of Unemployment Insurance Extensions on the Unemployment Rate in the 2008 to 2010 Period

Bhashkar Mazumder, Federal Reserve Bank of Chicago

Background

The passage and creation of the Emergency Unemployment Compensation (EUC) federal program in July 2008 and subsequent extensions substantially increased the maximum number of weeks of eligibility for unemployment insurance (UI). As of October 2010, unemployed workers in 26 states were eligible for a maximum of 99 weeks of UI benefits. The national average was about 95 weeks.¹ By contrast, during the deep recession in 1983, the maximum potential duration of UI coverage in any state was 55 weeks. This memo addresses the extent to which the UI extensions may have contributed to the rise in unemployment rate during the recent recession and early recovery period. I review some estimates made by economists in the Federal Reserve System and provide some new estimates using a simulation that applies to the current episode an estimate from the academic literature of the effects of UI extensions on the duration of unemployment.

Recent work within the Federal Reserve System

Valletta and Kuang (2010) identify the effects of UI extensions by comparing the unemployment durations of those who are likely to be eligible for UI with the durations of those who are not. They use the CPS question on the reason for unemployment to identify “involuntary job losers,” who comprise most of those who would be eligible for UI. The control group consists of “voluntary job leavers” and labor force entrants, most of whom would be ineligible for UI. They estimate that UI extensions may account for 0.8 percentage points of the increase in the unemployment rate from its pre-recession levels in 2006 and 2007 through June of 2010.

Fujita (2010b) identifies the effects of the UI extensions by comparing the unemployment durations of individuals in the CPS who were unemployed in 2004-07 with the durations of individuals who were unemployed in 2009-10. He finds that extensions can account for between 0.9 and 1.7 percentage points of the change in the steady-state unemployment rate for men. In another paper, Fujita (2010a) uses a simple steady-state model of the unemployment rate and estimates from Moffitt (1982) on the effects of UI extensions on the duration of unemployment. He estimates that the extensions may account for about 1.5 percentage points of the increase. Finally, Aaronson et al. (2010) use a range of estimates from the literature on the effect of UI extensions on unemployment duration and apply them to argue that the recent UI extensions account for between 10 and 25 percent of the increase in observed unemployment durations.

Each of these approaches has advantages and drawbacks. Valletta and Kuang (2010) utilize a relatively transparent approach, in which the data can be used to identify treatment and control groups and the

¹ This is a weighted average across all states where the weights correspond to the state share of the total number of unemployed individuals.

differences in the actual unemployment duration between the groups in the recent period can be summarized visually. They also construct a measure of expected total duration rather than relying on CPS measures, which include not-yet-completed spells of unemployment. However, one must make some arguably strong assumptions in order for their treatment and control groups to be truly comparable.²

Fujita (2010b) provides a nice complement to Valletta and Kuang's cross-sectional analysis, in that it uses the "pre" and "post" UI program differences. Fujita also deals with measurement error in the recording of spells in the CPS. Fujita's analysis, however, is much more technical in nature and it is unclear the extent to which particular assumptions may drive the results. Also, it is focused only on men, so it is unclear whether the results generalize to the entire labor force.

Fujita (2010a) provides a useful theoretical framework that uses the flows into and out of unemployment as the determinants of the steady-state unemployment rate. However, the approach implicitly assumes that *all* unemployed workers are covered by UI and it only considers extensions from 25 to 52 weeks, even though many workers can potentially collect UI for up to 99 weeks. It also relies exclusively on estimates from Moffitt (1982). These estimates are from earlier periods in the 1970s and 1980s when UI extensions were endogenous responses to high unemployment and firms may have taken into account the duration of UI benefits to strategically time the duration of temporary layoffs (Katz, 2010).

Aaronson et al. (2010) incorporate data on the fraction of the unemployed receiving UI each month to calculate a time series of the effects of extensions and consider how the effects accumulate over time. However, they implicitly assume that all increases in UI coverage are due to the extensions. They also lack a theoretical framework in which to interpret the results.

Methodology

Base case estimates

The approach considered here starts with the framework of the steady-state unemployment rate used by Fujita (2010a), but tries to incorporate many of the empirical details used by Aaronson et al. (2010). My estimates are based on previously estimated parameters from the literature on how UI extensions influence the duration of unemployment.

The steady-state unemployment rate is given by:

$$U = \frac{s}{s + f}$$

² For example, suppose those who quit their job or enter the labor force (and thus are ineligible for UI) are, on average, higher quality workers than those who are laid off (and thus eligible for UI). Then the quitter/new entrants would have shorter spells of unemployment, regardless of the effects of UI on the spells of the eligible workers.

where s is the rate at which employed workers “separate” into unemployment, f is the rate at which unemployed workers “find” employment, and U is the unemployment rate.³ I start by taking advantage of two key facts that characterized the labor market in the six months prior to the increase in UI extensions (January to June 2008). First, the unemployment rate averaged 5.1 percent, which I assume was its steady-state value. Second, the mean duration of an ongoing spell of unemployment was 17 weeks; this implies that f is about 0.285.⁴ Plugging these numbers in the equation above implies that s was about 0.0136.⁵

I then estimate how f changes as a result of the UI extensions. I use Card and Levine’s (2000) estimate that a 1 week increase in the maximum potential duration of receipt of UI leads to a 0.1 week increase in the duration of unemployment. To define the increase in the maximum potential period of UI receipt, I follow Aaronson et al. (2010) and use the triggers relevant for each state for each month to calculate a time series of the nationally weighted average of the maximum potential duration of UI receipt. As of October 2010, this stood at about 95 weeks. Therefore, the change in potential duration of UI receipt is 69 weeks.⁶ I assume that the UI coverage rate among the unemployed is 0.4, which was the coverage rate prior to the UI extensions in July 2008. I also assume that this coverage rate is the same for the UI extensions.

In combination, these assumptions imply that unemployment durations would increase from 17 weeks to 19.8 weeks with the extension of UI benefits.⁷ This implies that f would fall to 0.22. Using the equation above, the steady-state rate of unemployment would be 5.9%, or a 0.8 percentage point increase in steady-state unemployment due to UI extensions.

Alternative Assumptions

Card and Levine’s (2000) estimate of the effect of UI on unemployment duration is based on an extension of UI benefits in New Jersey in 1996. The extension was due to a political tradeoff during a period of stable economic conditions rather than a response to rising unemployment. As a result, unlike most other estimates, it is unlikely to reflect reverse causality.⁸ There may be some concern, however,

³ This formula can be derived by starting with a simple model of the evolution of the unemployment rate. The unemployment rate in the current period is equal to the unemployment rate in the previous period plus the fraction of the employed who separate, minus the fraction of the unemployed who find work: $u_t = u_{t-1} + sE - fu_{t-1}$, which implies that, in the steady state, $u_{ss} = u_{ss} + s(1 - u_{ss}) - fu_{ss}$. Solving for u_{ss} yields $u_{ss} = s/(s+f)$.

⁴ An estimate of the job-finding probability in a week is given by $1/17$. I convert this to a monthly probability by multiplying it by 4.3 (the average number of weeks in a month).

⁵ The value for f is very close to empirical estimates of the U to E (unemployment to employment) transition rates for the January to June 2008 period in Aaronson et al. (2010) using CPS microdata. The empirical E to U estimate of s is a bit lower at 0.0096. This may be due to classification error in the CPS among those who report unemployment versus non-employment. I have conducted the same experiment using the empirical estimates of f and s from the CPS and have found similar results. Although the *level* of U is much lower with the lower estimate of s , the proportional effects of UI extensions on the unemployment rate are very similar to what is obtained here.

⁶ This is the 95 weeks minus the 26 weeks available to workers prior to the UI extensions.

⁷ This is the result of multiplying 69 weeks by the 0.1 elasticity and scaling this down by the 0.4 coverage rate. This 2.8 effect on weeks is added back to the 17 week duration prevailing at the time of the UI extension.

⁸ Other estimates, such as those by Moffitt (1985) and Katz and Meyer (1990), were based on previous episodes during the 1970s and 1980s when much of the responsiveness may have been due to firms that used temporary

that the Card and Levine estimates are downwardly biased for the current situation, because workers may have a more muted behavioral response to UI extensions during expansions than recessions. Therefore, I consider as an alternative the estimate from Katz and Meyer (1990) that a week of extended benefits leads workers to remain unemployed an additional 0.16 weeks. This alternative assumption changes my estimate of the new steady-state unemployment rate to 6.3 percent, or 1.2 percentage points, 0.4 percentage points higher than my base case figure.

My baseline estimates assume a UI take-up rate of 40 percent. This was the take-up rate at the time that the extensions were implemented; since then, however, the take-up rate has risen to close to 70 percent. Some recent work by my colleague Luojia Hu finds that this increase likely reflects the worsening in economic conditions, and thus should not be used in calculating independent effects of UI extensions on unemployment.⁹ Nonetheless, as a robustness check, I reran my calculations under the alternative assumption that the UI extensions alone caused the UI take-up rate to increase from 0.4 to 0.55. This change boosts the increase in the steady-state unemployment rate from 0.8 to 1.1 percentage points.¹⁰

Finally, as Valletta and Kuang (2010) emphasize, the BLS's aggregate data on mean durations of unemployment refer to spells in progress and not to completed spells. My base case estimate of f is based on spells in progress. Valletta and Kuang have developed a method for estimating completed durations and have graciously provided me with their estimates. These suggest that completed durations prior to the extensions of UI benefits averaged 15.4 weeks, rather than the 17 weeks I had assumed. Feeding this alternative value into the steady-state model produces an estimated increase in steady-state unemployment due to UI extensions of 0.9 percentage points.

In summary, the base case and alternative estimates using the approach outlined above suggest that the extension of unemployment insurance benefits during the recent economic downturn can account for somewhere in the neighborhood of 1 percentage point of the increase in the unemployment rate, with the preferred estimate at 0.8 percentage points.

layoffs and where the recall dates were sensitive to the time limits on UI receipt. The use of temporary layoffs is much less prevalent today (Katz, 2010).

⁹ This is based on work reported in an internal Chicago Fed briefing memo prepared in August 2010.

¹⁰ If we use the Katz and Meyer estimate and also assume that the UI extensions increased the take-up rate to 55 percent, this would imply that the UI extensions led to a 1.7 percentage point increase.

References

Aaronson, Daniel, Mazumder, Bhashkar and Shani Schechter, 2010, "What is Behind the Rise in Long-term Employment?" Federal Reserve Bank of Chicago, *Economic Perspectives*, Second Quarter, pp. 28–51.

Card, David, and Phillip B. Levine, 2000, "Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefit Program," *Journal of Public Economics*, Vol. 78, Nos. 1–2, October, pp. 107–138.

Fujita, Shigeru, 2010a, "Economic Effects of the Unemployment Insurance Benefit," Federal Reserve Bank of Philadelphia, *Business Review*, Fourth Quarter.

Fujita, Shigeru, 2010b, "Effects of the UI Benefits Extensions: Evidence from the Monthly CPS," Federal Reserve Bank of Philadelphia, *Working Paper*, No. 10-35.

Katz, Lawrence F., 2010, "Long-Term Unemployment in the Great Recession," Testimony for the Joint Economic Committee U.S. Congress Hearing on "Long-Term Unemployment: Causes, Consequences and Solutions," April 29.

Katz, Lawrence F., and Bruce D. Meyer, 1990, "The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment," *Journal of Public Economics*, Vol. 41, No. 1, February, pp. 45–72.

Moffitt, Robert, 1985, "Unemployment Insurance and the Distribution of Unemployment Spells," *Journal of Econometrics*, Vol. 28, pp. 85–101.

Valletta, R., and K. Kuang, 2010, "Extended Unemployment and UI Benefits," *FRBSF Economic Letter*, No. 2010-12.

A Brief Note on Homeownership and Mobility: Recent Evidence from the SIPP

Daniel Aaronson, Federal Reserve Bank of Chicago

Introduction

This section explores new evidence on the extent to which house lock—the reluctance of households to sell their homes in a declining house price environment—has been an empirical explanation for the elevated unemployment rate over the past three years. It has been speculated that house lock may create a geographic mismatch between the location of available workers and jobs vacancies, potentially leading to persistently higher unemployment. I test for the possibility of house lock by comparing state-to-state migration rates of households that might be affected by declining home prices (i.e., homeowners, particularly those in states with large price declines) with migration rates of households that are not affected (renters). I use a dataset, the Survey of Income and Program Participation (SIPP), that has particular strengths for studying this question.

I find that state migration patterns are inconsistent with a large role for house lock through the first quarter of 2010. This result supports previous work done in the System, including Valletta (2010), Schulhofer-Wohl (2010), Foote and Ryan (2010), and Sahin et al. (2010) that use different data and methods. In particular, I find that state migration rates among both homeowners and renters barely budged during the recent recession and early recovery, in contrast to a larger relative decline between homeowner and renter migration rates during the past two recessions. There is also no evidence that migration rates varied by the magnitude of a state's recent house price decline. An important caveat is that I find a significant decline in within-state moves among homeowners relative to renters; when better data become available, I plan to test if this pattern is consistent with higher labor mismatch.

Data

The Survey of Income and Program Participation (SIPP) is a large representative sample of households interviewed every four months (called a wave) for two to four years. The first SIPP panel begins in 1984, with new cohorts added roughly when the previous cohort's survey cycle is completed.¹ The latest group entered the survey in 2008 and I use data for this group through March 2010. The sample is based on non-military households with a head between the ages of 25 and 59.² This leaves me with approximately 60,000 households per year or over 5 million household-wave observations between 1984 and 2010.

The SIPP is useful for looking at migration behavior for several reasons. First, relative to other large longitudinal datasets, it follows households for a fairly long time and includes a wealth of demographic covariates, labor market outcomes, and housing information.³ Second, unlike the Current Population

¹ I do not use the 1989 SIPP because that survey has fairly significant differences in design.

² I excluded households where the head's age at the beginning of the panel was more than five years different than the head's age at the end of the panel.

³ This includes special topical modules on migration and the household balance sheet. I do not use this information in this memo but plan to in future work. Schulhofer-Wohl (2010) uses the 2004 SIPP to show that homeowners with underwater mortgages are more likely to move than other homeowners, in contrast to the house lock story.

Survey, SIPP tracks households as they leave a residence. Therefore, we know explicitly whether the household moved, as opposed to attrition from the sample for some other reason,⁴ and also whether they left for a new labor market.

Currently, I define a labor market as a state and, therefore, a move as a change in state of residence between waves.⁵ This is not to say that I believe state borders are an ideal characterization of a labor market. Rather, this is the only geography that is currently available.⁶ It should also be noted that the overlapping nature of the survey (some households begin the 2008 SIPP in January 2008, others in February 2008, etc.) allows me to compute migration rates by month. That said, the monthly migration rate represents the share of households that moved from four months prior (between waves). Finally, there are several gaps in the time-series, because there are times when one cohort completed the full set of waves before the next cohort had begun. Regrettably, one of these gaps occurred during 2008.

Results

Figure 1 plots seasonally adjusted state migration rates for homeowners (red line) and renters (green line). The figure highlights the infrequency of moves across state lines. In a given year, just under 2 percent of all SIPP households cross a state border.⁷ State migration is particularly uncommon for homeowners; renters are about twice as likely to switch states, a pattern that holds throughout the sample period. Consequently, over the last 25 years, a significant portion of geographic reallocation has been due to households unencumbered by selling a home.

Table 1 looks specifically at recent changes, comparing the average four-month migration rate of homeowners and renters during the last expansion (columns 1 and 2) to the most recent data available, December 2008 to March 2010 (column 3). Homeowner migration rates fell from 0.0046 during the latter half of the 2000s expansion to 0.0038 during 2009 and early 2010, a 0.0008 overall decline (annualized, roughly $0.0008 \times 3 = 0.0024$ or two-tenths of a percentage point). However, relative to the entire 2000s expansion (2002-07), homeowner state migration was, on average, stable in 2009 and early 2010. Renter migration rates were also relatively flat. The last row explicitly compares the patterns among the two groups, showing that the difference over time in mobility between owners and renters is

However, the housing assets and liabilities needed to make such a calculation are not available yet for the 2008 cohort.

⁴ I do not know whether a sample attrition is associated with a move, however. Follow-up work will look at this issue more closely.

⁵ Some of the earlier SIPP panels grouped smaller states together. For example, in 1984, Idaho, New Mexico, South Dakota and Wyoming are categorized as one state. Obviously, a move between these states could not be picked up.

⁶ County or MSA codes are not publicly available. A variable that flags county moves is inaccurate according to communication between SIPP staff and Sam Schulhofer-Wohl of the Minneapolis Fed. I hope to change my definition of a labor market to county once this error is fixed by the SIPP.

⁷ Many papers have noted the declining trend in state migration (e.g., Molloy, Smith, and Wozniak, 2010), a pattern evident in the SIPP as well. Among homeowners, the four-month migration rate fell from an average of 0.0055 during 1984 to 1989 to 0.0044 from 2000 to 2009. Likewise, the renter four-month migration rate fell from 0.0105 to 0.0083 during this span. Moreover, the aggregate migration rate has fallen for compositional reasons as well—until very recently, homeownership rates had been on the rise since the 1980s.

economically small (annualized, under one-tenth of a percentage point relative to the 2002-07 period and about three-tenths of a percentage point relative to 2005-07) and statistically indistinguishable from zero.

This is fairly striking in light of the patterns in migration rates in previous business cycles, summarized in table 2. Four-month homeowner state migration rates were about 0.0008 percentage points lower during the 1991 and 2001 recessions than during the 1980s and 1990s expansions. By contrast, renter migration rates were higher during the recession months than during the expansion periods. This implies a procyclicality in the gap between homeowner and renter state migration that is not only statistically significant but economically important,⁸ in contrast to the current period. That is, migration patterns were more consistent with a “house lock” story in earlier recessions than in the most recent one.

It is possible that the renter–owner comparison still hides heterogeneity in mobility patterns based on the magnitude of the local housing bust. In particular, if house lock is important, it should bite those households residing in states with the largest house price declines. Figure 2 and table 3 stratify homeowner state migration rates by whether the household’s state house price decline between 2007:Q2 and 2010:Q2 is greater than or less than the national median price decline.⁹ Again, I find virtually no evidence of atypical migration patterns among homeowners hit with large price declines. Indeed, during 2009 and early 2010, homeowner state mobility rates declined more for households residing in states that experienced better home price performance.

Finally, figure 3 looks specifically at the five states that experienced the largest housing price decline (California, Florida, Nevada, Arizona, and Rhode Island). Although sample sizes make this series fairly volatile, it is again clear that state migration rates did not decline in an unusual fashion in these states. Four-month homeowner state migration rates decreased—from 0.0050 during 2005-07 to 0.0047 during the most recent period—at the same time that state migration out of the remaining 45 states fell even more, from 0.0045 to 0.0036 percentage points. Therefore, I see little evidence that house lock likely drove homeowner state-to-state migration rates down during the last recession and the beginning of the recent recovery.

A caveat

As I mentioned earlier, with the current data, I am restricted to using state as the definition of a local labor market. Figure 4 and table 4 suggest that a more refined measure might matter. They show that homeowner in-state migration fell during 2009 and early 2010, while renter in-state migration fell less and even increased relative to a longer baseline of 2002-07. The difference-in-difference estimate of –0.0112 is statistically significant and economically large, suggesting a potentially notable decline in residential movement among homeowners during the recession and early recovery. The precise origin and destination of these moves are important. If these disappearing moves would have been within

⁸ The point estimate in row 3 of -0.0020 is consistent with about a six-tenths of a percentage point decline in state migration of homeowners relative to renters, a substantial move in this series.

⁹ House prices are from the Federal Housing Finance Agency Index.

commuting distance of a former residence, they should have little bearing on geographic mismatch and the unemployment rate. However, if they have arisen in states with distant labor markets (i.e., San Francisco to San Diego), this might suggest some role for house lock after all. Future work will hopefully be able to disentangle this issue with new data from the SIPP.

Conclusion

Unemployment may be high due to a mismatch between the skills that employers are demanding and the skills that the available labor force can provide. Part of this mismatch may be geographic in nature; available workers may not reside where jobs vacancies are. It has been speculated that one mechanism that might lead to this situation is that potential job searchers are reluctant to sell homes that have experienced significant declines in value, referred to as house lock.

I find that state-to-state migration rates among homeowners barely budged during the latest recession and early recovery, either in absolute terms or relative to renters. Moreover, the gap between homeowner and renter state migration rates fell more during the previous two recessions than during the recent one. Finally, there is little evidence that migration varied based on the magnitude of a state's recent house price decline. I do, however, find some preliminary evidence that in-state migration fell for homeowners but not renters. Therefore, once the data become available, future work should explore whether more refined measures of labor markets, such as MSAs or counties, might alter the conclusions based on state migration patterns.

Table 1
State migration rates, recent evidence
 (four-month migration rate and sample size)

	<u>2002-07</u>	<u>2005-07</u>	<u>Dec 08 to Mar 10</u>	<u>Change since 02-07</u>	<u>Change since 05-07</u>
	(1)	(2)	(3)	(3)-(1)	(3)-(2)
Owners	0.0038 210,779	0.0046 106,697	0.0038 60,630	0.0001	-0.0008
Renters	0.0083 96,780	0.0083 47,893	0.0086 31,117	0.0003	0.0003
Difference				-0.0002 (0.0021)	-0.0011 (0.0008)

Table 2
State migration rates, 1984–2001
 (four-month migration rate and sample size)

	<u>Expansions</u>	<u>Recessions</u>	<u>Difference</u>
	(1)	(2)	(2)-(1)
Owners	0.0049 458,683	0.0041 37,510	-0.0008
Renters	0.0095 218,360	0.0107 18,818	0.0012
Difference			-0.0020 (0.0012)

Table 3
State migration rates, by size of state house price decline
 (four-month migration rate and sample size)

	<u>2002-07</u>	<u>2005-07</u>	<u>Dec 08 to</u> <u>Mar 10</u>	<u>Change</u> <u>since 02-07</u>	<u>Change</u> <u>since 05-07</u>
	(1)	(2)	(3)	(3)-(1)	(3)-(2)
Smaller house price decline states (less than median decline)					
Owners	0.0046 90,162	0.0055 46,801	0.0040 24,586	-0.0006	-0.0015
Renters	0.0077 36,301	0.0106 18,331	0.0109 11,141	0.0031	0.0002
Difference				-0.0037 [0.0031]	-0.0018 [0.0016]
Larger house price decline states (greater than median decline)					
Owners	0.0032 120,617	0.0040 59,896	0.0037 36,044	0.0005	-0.0003
Renters	0.0086 60,479	0.0071 29,562	0.0074 19,976	-0.0012	0.0003
Difference				0.0017 [0.0024]	-0.0006 [0.0010]

Table 4
In-state migration rates, recent evidence
 (four-month migration rate and sample size)

	<u>2002-07</u>	<u>2005-07</u>	<u>Dec 08 to</u> <u>Mar 10</u>	<u>Change</u> <u>since 02-07</u>	<u>Change</u> <u>since 05-07</u>
	(1)	(2)	(3)	(3)-(1)	(3)-(2)
Owners	0.0248 210,779	0.0262 106,697	0.0207 60,630	-0.0042	-0.0055
Renters	0.0594 96,780	0.0684 47,893	0.0665 31,117	0.0071	-0.0019
Difference				-0.0112 [0.0042]	-0.0036 [0.0026]

Figure 1
State Migration, seasonally adjusted

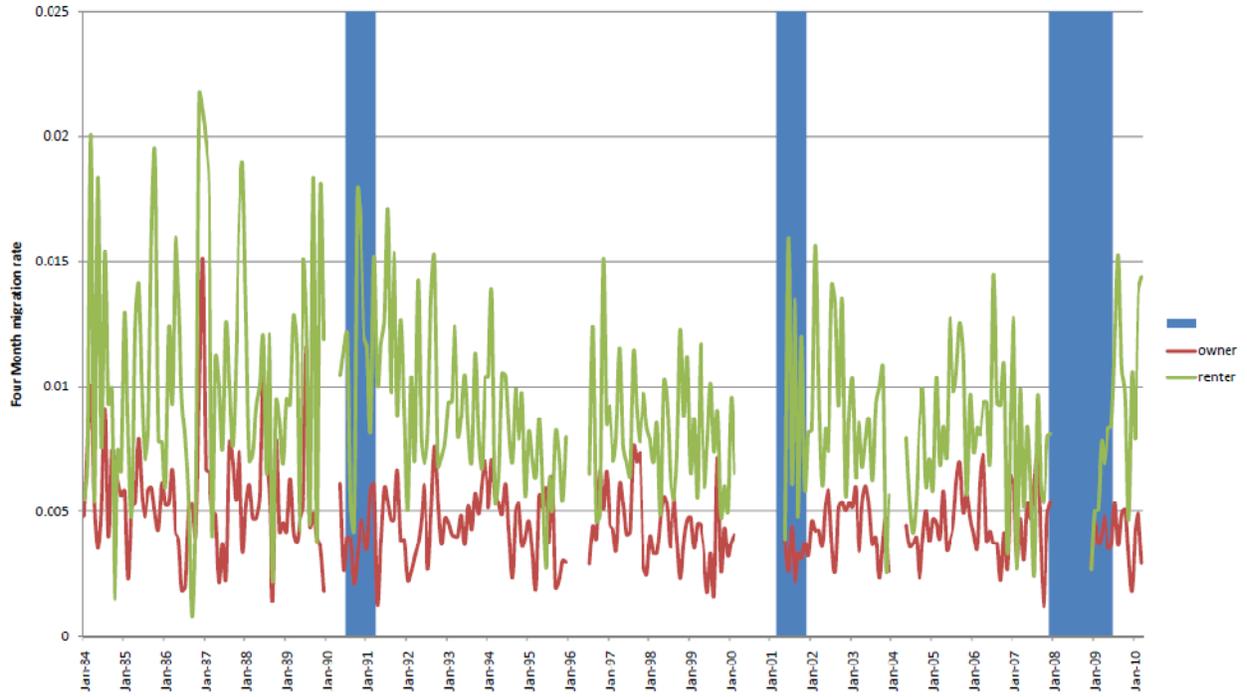


Figure 2
Homeowner State Migration, by state house price decline, SA

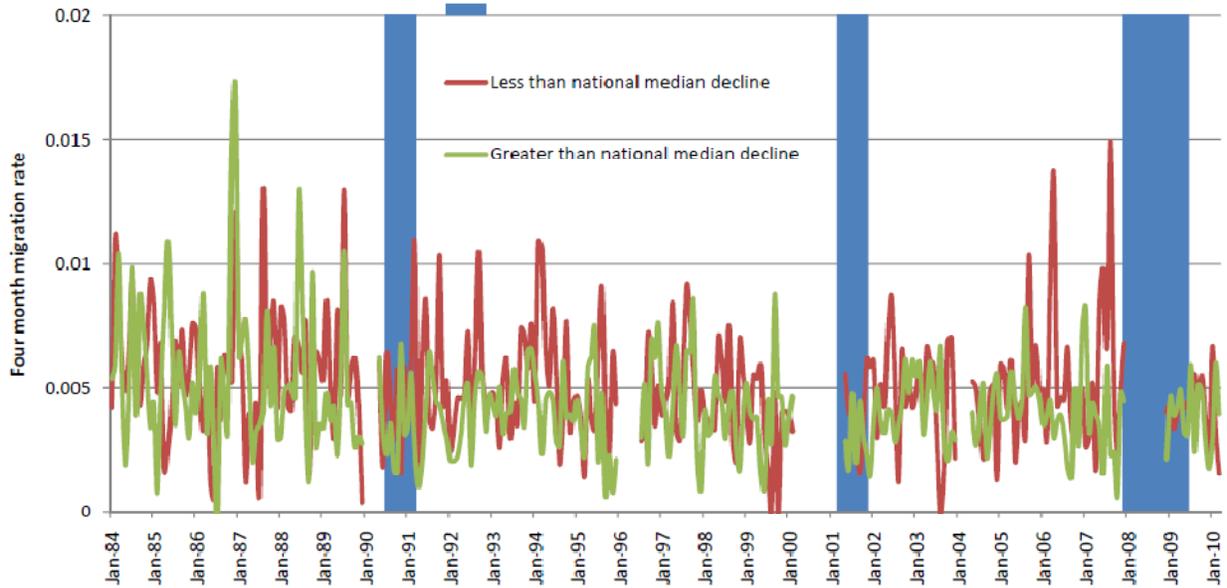


Figure 3
Homeowner State Migration, 5 largest home price decline states, SA

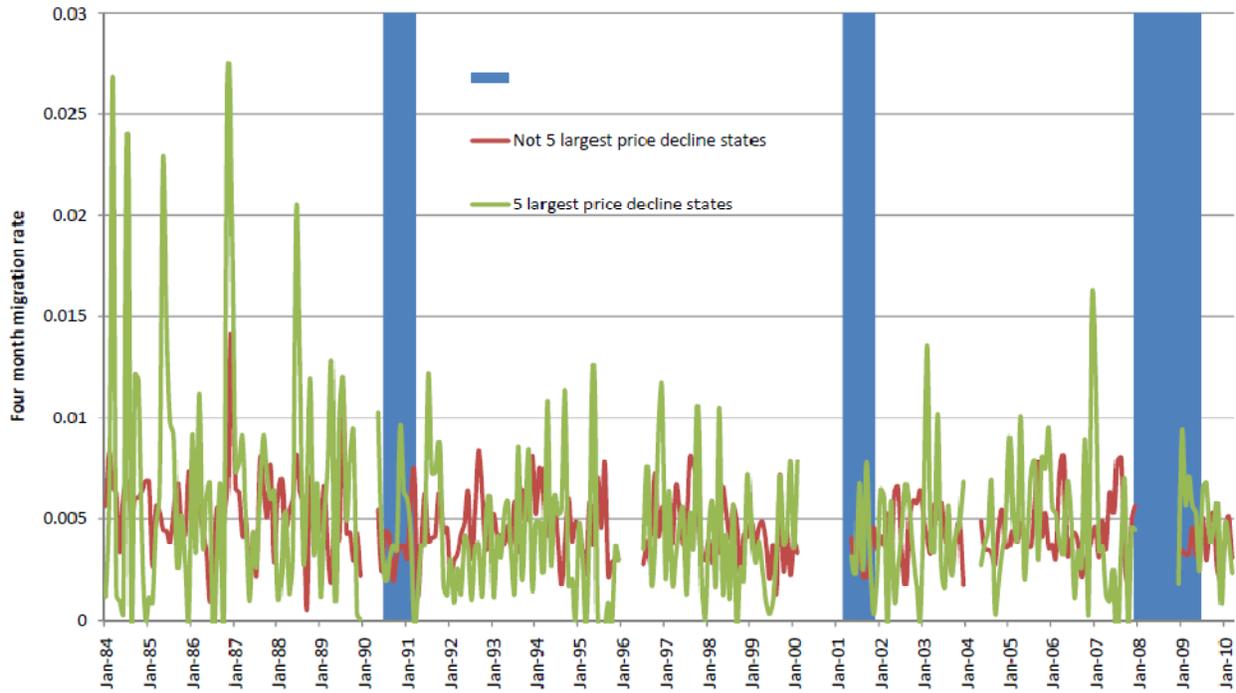
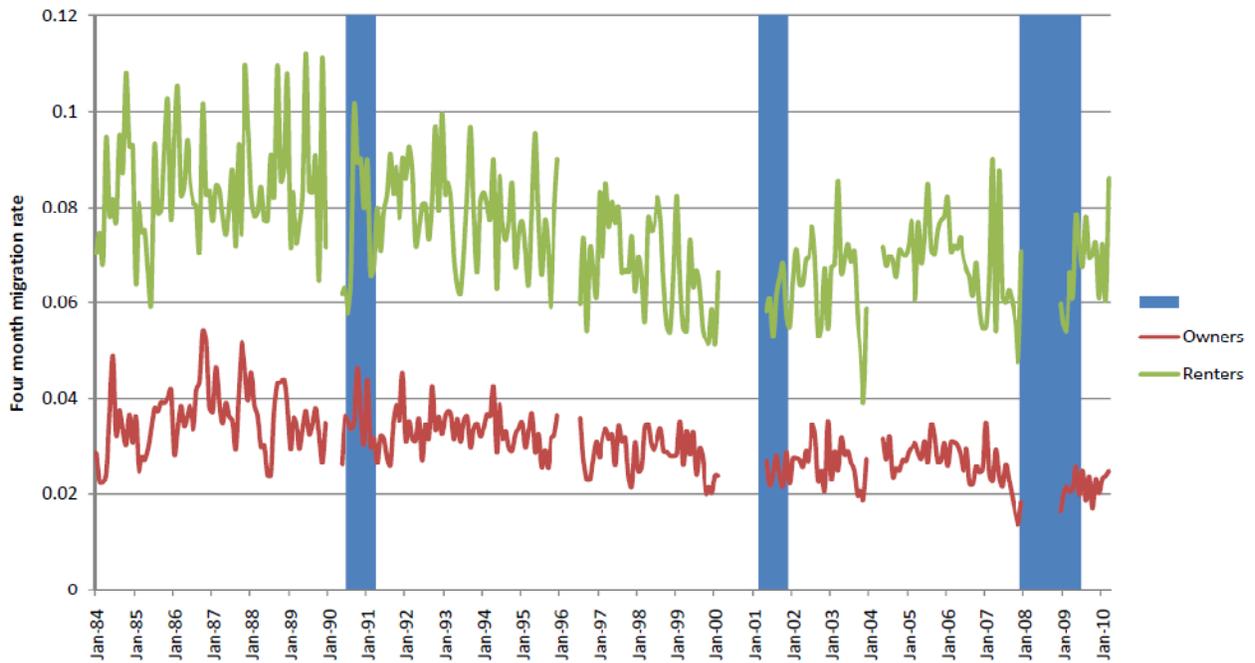


Figure 4
In-State Migration, SA



References

Foote, Chris, and Richard Ryan, 2010, "Understanding Recent Increases in Unemployment Duration," Federal Reserve Bank of Boston, mimeo.

Molloy, Raven, Christopher Smith, and Abigail Wozniak, 2010, "Internal Migration in the US: Updated Facts and Recent Trends," Federal Reserve Board of Governors, mimeo.

Schulhofer-Wohl, Sam, 2010, "Negative Equity Does Not Reduce Homeowners' Mobility," Federal Reserve Bank of Minneapolis, mimeo.

Valletta, Robert, 2010, "House Lock and Structural Unemployment," Federal Reserve Bank of San Francisco, mimeo.

Mismatch and Steady-State Unemployment

Gadi Barlevy, Federal Reserve Bank of Chicago

Introduction

Over the past year, we have witnessed a rise in the rate of vacancies posted by firms at the same time as the unemployment rate remained essentially unchanged. Some have interpreted this pattern to mean that, evidently, the problem in the labor market is not that firms don't want to hire, but that they are trying to hire and can't. By this logic, using monetary policy to encourage further hiring is unlikely to successfully drive down unemployment. In this memo, I offer a calculation that shows that even under the most charitable view that attributes all of the rise in vacancies with no change in unemployment to a decline in the ability of firms to hire, the current unemployment rate is too high to be generated solely by difficulties in hiring. That is, it still appears as though some forces are serving to discourage firms from hiring, even after taking into account the difficulties in finding suitable workers to staff positions. This suggests there may be a role for monetary policy to offset these forces and promote more hiring, although this memo is silent on whether monetary policy can or should play such a role.

More precisely, I demonstrate the following results:

1. Starting with a long-run unemployment rate of 5.0%, a shock that reduces the effectiveness of hiring would by itself result in an unemployment rate of no more than 7.0%, given the implied decline in the effectiveness of hiring using data through June 2010, and no more than 7.8%, using data through November 2010. Both estimates are considerably lower than the current unemployment rate of 9.4%.
2. A continued rise in the rate of vacancies posted without a significant decline in unemployment will make it harder to dismiss the claim that high unemployment can be attributed to difficulty in hiring. For example, at the current unemployment rate, were the vacancy rate to rise to 3.0% without any change in unemployment, the calculation I describe would not be able to rule out that the unemployment rate could be as high as 9.5% due to mismatch considerations alone.

The Matching Function Approach

The framework I use for my calculation relies on Pissarides (1985) and Mortensen and Pissarides (1994). This framework has two key elements. First, it assumes that the number of new hires h can be expressed as an increasing function of the number of unemployed workers u and the number of vacancies firms post v :

$$(1) \quad h = m(u, v).$$

The function m is referred to as a matching function and is meant to capture the various frictions in the process of filling new jobs from the ranks of the unemployed.

Second, firms are assumed to keep posting vacancies as long as it remains profitable to do so. This gives rise to what is known as the free-entry condition, which holds that the expected profits from posting a

vacancy should be zero in equilibrium, or else firms would have posted still more vacancies. Let J denote the value of a filled job to the employer who creates it, and let k denote the cost of posting the vacancy. Then the free-entry condition states that the value of a job times the probability of filling it should equal the cost of posting the vacancy. If there are $m(u, v)$ new hires, and each firm is equally likely to be the one that is matched, the probability of filling a job is given by

$\frac{m(u, v)}{v}$, and the free-entry condition is given by

$$(2) \quad \frac{m(u, v)}{v} J = k$$

In what follows, I assume the matching function has a Cobb-Douglas form, i.e.,

$$(3) \quad m(u, v) = Au^\alpha v^{1-\alpha}.$$

where A captures the productivity of matching, and α reflects the sensitivity of new hires to the number of unemployed workers. Under this specification, the notion that firms find it more difficult to hire workers corresponds to a decline in the productivity of matching, A . Petrongolo and Pissarides (2001) argue that this functional form captures the data on unemployment, vacancies, and hires quite well.¹

Since the Cobb-Douglas specification exhibits constant returns to scale, the free-entry condition can be expressed solely in terms of the ratio v/u . This ratio is often referred to as market tightness, for it reflects how many vacant positions are competing for each unemployed worker. In particular, the free-entry condition implies

$$(4) \quad A \left(\frac{u}{v} \right)^\alpha J = k$$

If the matching function remains stable over time, i.e., if the parameters α and A are time invariant, the free-entry condition (4) tells us that if the value of a job relative to the cost of a vacancy J/k varies over time for any reason, such as a change in aggregate demand or changes in worker productivity, the equilibrium ratio v/u will change as well, in a way that can be predicted using only the free-entry condition.

The Beveridge Curve

By introducing one additional assumption, I can use the framework above to predict not just how the ratio v/u changes with J/k , but also how u and v change individually. In particular, I need to assume that the rate at which employed workers move into unemployment is constant over time. This assumption may seem implausible at first, especially given the incidence of mass layoffs in recessions. However, the

¹It should be noted that a recent literature, starting with Shimer (2005a), argues that the matching function approach suffers from serious shortcomings in its ability to match various labor market facts over the business cycle. However, this critique concerns whether in these models the value of a job J varies enough over the cycle, not whether the matching function can explain how hires vary with unemployment and vacancies. My calculation does not depend on how J varies with aggregate conditions, and so is not subject to this critique.

separation rate I need to assume is constant does not involve one-off spikes of job destruction that reflect immediate adjustment to changes in economic conditions. Rather, the relevant rate is the one that corresponds to what happens in a recession once all bursts of job destruction are done. Shimer (2005b) and Hall (2005) argue that fluctuations in this separation rate contribute little to overall changes in unemployment and can be ignored. Subsequent work has found a more important but still relatively small contribution of the separation rate compared with the hiring rate, e.g., Elsby, Hobijn, and Sahin (2010). Below, I argue that assuming this separation rate remains constant as we enter a recession is consistent with data on unemployment and vacancies for three recessionary periods. Moreover, if the relevant separation rate was in fact higher during the recession, my calculation would only exaggerate the role of mismatch, and the bound I report would be too high.

To see the implications of the model for u and v individually, consider what happens if we let the relative value of a job J/k vary over time. Conditional on a given value of J/k , the free-entry condition (4) tells us that the unemployment to vacancy ratio must remain constant for as long as J/k does. However, u and v could themselves change even while J/k remains fixed, as long as they change in the right proportions. Still, one can show that as long as J/k remains fixed, u and v will converge to some steady-state values that depend on J/k , and moreover that this convergence will be rapid. This rapid convergence is not just a theoretical result, but has been confirmed empirically.² Thus, we can assume that at any point in time, whatever the value of J/k happens to be, the unemployment and vacancy rates we would observe will coincide with the steady-state levels of these variables conditional on that J/k .

To compute the conditional steady-state unemployment for a given relative value of a job J/k , note that the flow into unemployment is equal to $s(1-u)$, where s denotes the separation rate, while the flow out of unemployment is equal to $Au^\alpha v^{1-\alpha}$. Since inflows into and out of unemployment must be equal in steady state, we can use this equality to arrive at an implicit formula for the conditional steady-state unemployment rate associated with a particular J/k and, hence, with a particular v/u ratio:

$$(5) \quad u = \frac{s}{s + A(v/u)^{1-\alpha}}$$

Rearranging that expression yields an expression for the predicted vacancy rate v for a given unemployment rate u :

$$(6) \quad v = \left[\frac{s}{A} (u^{-\alpha} - u^{1-\alpha}) \right]^{\frac{1}{1-\alpha}}$$

²In particular, Shimer (2005b) shows that the steady-state unemployment level to which the economy should be converging at any point in time can be readily computed from flows into and out of unemployment at that time. He then shows that this steady state is remarkably close to actual unemployment. Those who recall unemployment as highly persistent may find it surprising that it converges quickly to its conditional steady state. However, rapid convergence to the steady state for a given level of J/k does not imply that u will not appear persistent. If the relative value of a job J/k is persistent, unemployment will appear slow to change, even though it converges to its conditional steady state at any given J/k quite rapidly.

As long as s is constant, (6) implies a negative relationship between u and v . This relationship is often called a Beveridge curve after the economist who first documented a negative empirical relationship between the two. As we vary J/k , we will trace out a negatively sloped relationship between u and v . Conversely, movements along the Beveridge curve must stem from some underlying change in the value of J/k . A high vacancy rate and low unemployment rate suggest J/k must be high to encourage firms to post a lot of vacancies, leaving fewer workers unemployed. Similarly, a low vacancy rate and high unemployment suggest J/k must be low. The value J/k , in turn, can be viewed as providing information on the state of aggregate demand or productivity.

As long as we assume s is fixed, the only way for the Beveridge curve to shift is if the matching function changes, i.e., if either A or α in (5) changes. That is, a shift in the Beveridge curve in this framework corresponds to a change in the nature of how workers and employers come together to form new hires. For example, a rise in mismatch, whereby employers begin to seek skills that unemployed workers lack, would imply that the same number of unemployed workers and vacant positions should result in fewer matches. That would appear as a decline in the productivity of matching A . The model thus delivers a clean dichotomy: shifts of the Beveridge curve correspond to changes in the efficiency with which hiring is done, i.e., changes in A , while movements along a fixed Beveridge curve correspond to changes in the incentives for firms to hire, i.e., changes in J/k .

We can now cast the debate about the relevance of monetary policy for the labor market in more precise terms. The fact that vacancies have recently risen as unemployment remained unchanged implies the Beveridge curve has shifted, i.e., A must have changed. There is arguably little monetary policy can do to affect the process by which firms and unemployed workers form new hires. Therefore, if the high rate of unemployment we see today is due to a shock to A , there is little monetary policy can do. To put it another way, monetary policy cannot shift the Beveridge curve back to its original level. But monetary policy can potentially respond to changes in economic conditions that affect J/k . Thus, the key question for determining the potential relevance of monetary policy is how much of the high unemployment rate today can be attributed to forces beyond the scope of monetary policy, i.e., a shock to A , and how much to forces that it could potentially affect, e.g., J/k . The point of the calculation below is that the implied decline in A over the past year cannot by itself account for the current level of unemployment, and thus there may be some scope for monetary policy, despite the fact that there is little monetary policy can do to address whatever forces have made the matching process less effective than it has been in the past.

Empirical Beveridge Curves and Estimating the Matching Function

To relate the Beveridge curve implied by the matching function to data, consider the Beveridge curve relationship in (6). It tells us that for each unemployment rate, the matching function can be used to predict the vacancy rate that should accompany it. If we had time series on unemployment and vacancies, we could use them to infer the two parameters of the matching function α and A that come closest to matching the empirical relationship between u and v .

As a check of whether it is reasonable to assume that s remains constant even as we transition through a recession, I opted to estimate the parameters α and A using data only for the period preceding the recession and then look at whether this relationship continues to predict the vacancy rate for observed unemployment rates into the recession. In particular, I use data between December 2000 and August 2008, just before the big run-up in unemployment. To recover α and A , I first set $s = 0.03$, per Shimer (2005a), who estimates this rate at a monthly frequency. However, the choice of s is really just a normalization of the frequency at which we are looking at the data.³ To infer A and α , I set out to match two specific moments. First, the average predicted v over all months between December 2000 and September 2008 had to match the average monthly vacancy rate in the data over this same period in the data, or 0.29. Second, the difference in v between the two end dates, which is 0.013 in the data, had to correspond to the difference in predicted v for these two dates. Matching moments this way yielded values $A = 0.752$ and $\alpha = 0.45$. The implied Beveridge curve is plotted as the blue line in the bottom panel of figure 1, together with the actual data on unemployment and vacancies (the curve and data are also presented in figure 2). As confirmed in the figure, the curve does a reasonable job of predicting the vacancy rate associated with a given unemployment rate beyond the period it was designed to match, up until the end of 2009; it only breaks down in 2010. Since the implied Beveridge curve assumes a constant separation rate, the fact that the curve fits well throughout the official recession suggests that assuming a constant separation rate s is indeed reasonable.

As further corroboration of the fit of the model, I also went back to two previous recessions, in 1975 and in 1980-81. Since JOLTS does not go back before December 2000, I use the Conference Board Help Wanted Index, normalized to coincide with JOLTS for the period in which they overlap. For each of these recessions, I follow a similar approach to estimating the matching function as above—namely taking data from the period prior to the recession to estimate the function, then seeing how the implied Beveridge curve does during the recession. However, since the Help Wanted Index is considered unreliable over long periods because of trends in the rate of newspaper advertising, I restrict attention to shorter periods in these cases. For the 1975 recession, I look at the 18-month period before the peak NBER date, i.e., May 1972 to October 1973. For the 1980 and 1981 recessions, I look at the 18-month period before the NBER peak date for the 1980 recession, i.e., July 1978 through December 1979. In both cases, I estimated A and α in the same manner, i.e., I chose these parameters so that the average predicted vacancy over the period is equal to the actual average, and the difference in the two end dates is the same for the predicted and actual series. The implied coefficients are reported in table 1, and the implied Beveridge curves are illustrated as blue lines in the remaining panels of figure 1. For the 1975 recession, the predicted Beveridge curve does well through December 1974, but then under-predicts the vacancy rate. For the 1981 recession, the predicted Beveridge curve does remarkably well through both the recession and subsequent recovery. As evident in table 1, the coefficient α is estimated to be virtually the same across all three periods. This motivates me to model the shift of the Beveridge curve in 2010 as a change in A for a fixed α .

³From (6), the curve only depends on the ratio of s/A ; hence, the frequency we use to measure flows is irrelevant.

As a comparison, I also considered an alternative model of the Beveridge curve proposed by Shimer (2007) that was used by Kocherlakota (2010) to analyze the same unemployment and vacancy data I consider. Once again, I focus on the same period prior to each recession to estimate the two parameters that pin down the predicted relationship between unemployment and vacancies in that model (these parameters are denoted m and n in Shimer's paper). Following Shimer and Kocherlakota, I choose these parameters to match the average values of u and v in the pre-recession period. The estimates for the two parameters are summarized in table 2, and the implied Beveridge curves for each period are shown in figure 1 in gray. Shimer's model produces a Beveridge curve that is too steep at high unemployment rates in all three episodes, i.e., it predicts a much sharper fall in the rate of vacancies posted as unemployment rose than we ever observe in the data. It is possible that all three episodes were associated with a rise in mismatch during both the recession and subsequent recovery. But, given that many of the explanations proposed for the most recent rise in mismatch involve elements unique to the present period, it also seems plausible that Shimer's model simply provides a poor benchmark level for vacancies at high unemployment rates.

Inferring the Extent of Mismatch

Since a rise in the vacancy rate without a change in unemployment implies a shift of the Beveridge curve, explaining the current state of the labor market requires changing the underlying matching function. For reasons cited above, I model this change as a shift in the matching productivity A , while holding α fixed. Thus, suppose that over the course of 2010, the matching function that was initially governed by one productivity parameter, A_0 , declined to some new value A_1 . From (6), we can deduce the exact value of A_1 needed to match a given unemployment and vacancy pair. For example, to match the data in June 2010 when $u = 0.095$ and $v = 0.021$, the last data point we had when we initially performed this calculation, A_1 must solve

$$(7) \quad 0.021 = \left[\frac{.03}{A_1} (0.095^{-.45} - 0.095^{.55}) \right]^{\frac{1}{.55}}$$

This implies $A_1 = 0.64$, i.e., the productivity of the matching function appears to have declined by 16% to its original level before the recession when $A = 0.752$. Figure 2 shows the Beveridge curves for both values of A_0 (the blue line) and A_1 (the light gray line).

An alternative way to infer the change in match productivity A is to introduce data on new hires. The idea is as follows. Since $m(u, v)$ corresponds to the number of new hires, which is measured in JOLTS, I can take the number of new hires and divide by the expression $u^\alpha v^{1-\alpha}$ for $\alpha = 0.45$, as I previously estimated to recover A . The time series for this ratio (shown in figure 3) suggests a decline in the productivity of matching throughout 2008-10 as opposed to only in 2010. This figure suggests a more modest decline in the productivity of matching between 2009 and 2010—12% as opposed to 16%. However, over the entire period, matching efficiency fell more substantially, on the order of 20%. In what follows, I will only use the estimate of the change in A that is based on the Beveridge curve.

The Effects of Mismatch on Unemployment

So far, I have argued that the apparent shift of the Beveridge curve over the course of 2010 implies a less productive matching process in the labor market that can be quantified. But I have yet to talk about what the model implies should happen to u and v as a result of the decline in match productivity. Should less efficient matching lead to higher equilibrium unemployment? The answer is yes, to an extent that one can infer using the free-entry condition. But it turns out that the implied rise in unemployment is smaller than the actual increase in unemployment over this same period.

Suppose we start at a steady-state unemployment rate of 5%, which is roughly the historical average of unemployment from 2000 until unemployment began to take off in August 2008. From the Beveridge curve relationship implied by (6), we know the implied vacancy rate would have to be

$$v = \left[\frac{.03}{0.752} (0.05^{-.45} - 0.05^{-.55}) \right]^{\frac{1}{.55}} = 0.03$$

and the implied ratio of u/v must equal $\frac{.05}{.03} = 1.67$. From the free-entry condition, we know that u/v will be constant, as long as we hold J/k fixed. Now, suppose we keep J/k fixed at its original level, but change A from A_0 to A_1 , so that only the productivity of the matching process is allowed to change. Not only will the Beveridge curve shift, but the implied ratio u/v that preserves the free-entry condition must change. Once we know this new ratio, we can determine where on the new Beveridge curve the economy will lie.

Before I compute this unemployment rate, note that, in a richer model, J and k would be determined endogenously, so it is not obvious that holding J and k fixed while changing A makes much sense. However, for reasons I discuss below, changes in A are likely to affect J/k in a particular direction, implying that the unemployment rate holding J/k fixed will correspond to an upper bound on unemployment.

I begin by totally differentiating the free-entry condition (4) to get

$$\frac{dA}{A} + \alpha \frac{d(u/v)}{u/v} = 0$$

This implies an elasticity of u/v with respect to match productivity of $-\frac{1}{\alpha}$. Since we estimated $\alpha = 0.45$ above, this implies that the estimated decline in the productivity of matching, holding J and k fixed, should lead the unemployment to vacancy ratio to rise by a factor of $\left(\frac{0.752}{0.640}\right)^{1/.45} = 1.43$. Since the original ratio of u/v that was consistent with a 5% unemployment rate under the original Beveridge curve was equal to 1.67, we can deduce that the new equilibrium ratio of u/v should increase by a factor of 1.43, so the new unemployment to vacancy ratio will equal

$$1.43 \times 1.67 = 2.38.$$

Using the Beveridge curve evaluated using a productivity parameter $A = 0.64$ gives us the implied unemployment rate associated with the ratio of 2.38 that must prevail in the new equilibrium:

$$u = \frac{0.03}{0.03 + 0.64 (2.38)^{-0.55}} = 0.070$$

Figure 4 provides a graphical intuition for the calculation. With each Beveridge curve is associated a given v/u ratio determined by the free-entry condition, which appears as a ray emanating from the origin. The original Beveridge curve and free-entry condition are in blue, and the ones associated with the lower A are in gray. Once we reduce A , we not only shift the Beveridge curve, but we also rotate the free-entry condition clockwise at a rate related to α . Intuitively, if hiring is rendered less effective, then firms will have an incentive to post fewer vacancies; and, as a result, more workers will remain unemployed per vacancy.

Since the unemployment rate currently stands at 9.4%, we can conclude that not only did the Beveridge curve shift during the most recent period, but there must have been some force acting to depress J/k . Hence, while it may be true that there is nothing that monetary policy can do to shift the Beveridge curve back in, one can still argue that the incentives for hiring are lower now than they would have been under normal circumstances prior to the recession (i.e., at 5% unemployment). It follows that there may be scope for monetary policy to move further along the new Beveridge curve toward the unemployment rate that corresponds to where the economy would be if economic conditions had been maintained at their average levels.

Finally, I return to the issue of the legitimacy of assuming that J and k at present should assume the same values as in normal times before the recession. In practice, both k and j are endogenous. For example, the process of creating a vacancy requires productive inputs such as labor, so this cost will depend on wages that are determined endogenously. But given that wages are procyclical both empirically and in the model, we would expect the cost of posting a vacancy k to fall as the unemployment rate rises. As for the value of a filled job J , there are various reasons to suspect it will be higher when there is more unemployment. Mortensen and Pissarides (1994) argue that the value of a job should be proportional to the surplus formed from a match. But the surplus from matching is higher when v/u is low. Intuitively, when it is easy to match, the value of matching now rather than searching for another match will be small. More generally, various features absent from the model, such as curvature in the utility function and diminishing returns to labor, would tend to make a marginal job more valuable when fewer workers are employed. Thus, at a higher unemployment rate, J is likely to be higher. But if k falls and J rises, the effect on u/v would only be smaller. As such, 7% should be viewed as an upper bound rather than a point estimate.⁴

⁴In informal communication, Rob Shimer computed the effects of a 16% drop in the productivity of the matching function in a fully worked out equilibrium model with concave utility and declining marginal product of labor. He found that unemployment would only rise from 5% to 5.8%. This suggests my bound may be a substantial overestimate of the true effect.

Updating the Calculation with More Recent Data

The calculation above was based on data through June 2010. But we now have additional data on vacancies and unemployment covering the period between July and November of 2010. During this period, the unemployment rate remained at about 9.5%. The vacancy rate, on the other hand, continued to climb, from 2.1% in June to as high as 2.5% in October. This would suggest an even greater degree of mismatch than is implied by data from the first half of 2010. Indeed, repeating the calculations above to match the average vacancy rates between July and November implied a value for A_1 of 0.61, or a 20% decline in the efficiency of matching. The implied Beveridge curve associated with this decline corresponds to the dark gray line in figure 1. The implied bound on unemployment rises from 7.0% to 7.8%, still well below the current unemployment rate.

Comparing the two bounds reveals an important caveat: the argument that a significant part of the current unemployment rate must reflect depressed economic conditions rather than mismatch depends on the extent of mismatch implied by the data. If the vacancy rate continued to rise without a significant decline in the unemployment rate, it would eventually no longer be possible to dismiss the claim that the high unemployment rate is due entirely to the difficulty firms face in hiring. For example, if we held unemployment fixed at 9.5%, the posted vacancy rate would have to rise to 3%, well within its historical range, before the bound would exceed current levels of unemployment. Thus, any conclusions reported here are conditional on the extent of mismatch implied by the data observed so far.

Conclusion

The observation that in 2010 firms kept posting more vacant positions without a commensurate decline in unemployment has led to concerns that unemployment is high because firms are unable to hire the workers they need, not because firms are reluctant to take on more workers. The matching function approach pioneered by Pissarides (1985) and Mortensen and Pissarides (1994) provides a framework for analyzing this issue. This approach indeed interprets recent developments in the labor market as evidence of less efficient matching now than in the past, as critics of more accommodative monetary policy have argued. It also confirms that greater difficulty in hiring would result in higher unemployment, as in fact occurred. But it also suggests that the apparent decline in the efficiency of matching we have witnessed so far cannot on its own account for the high unemployment rate today. The relevant issue for policy, then, is arguably not whether the shift of the Beveridge curve reflects mismatch as opposed to other phenomena. Rather, since a sizable share of current unemployment appears to be due to altogether different forces that dampen the incentives for firms to engage in hiring, the relevant question is whether monetary policy can and should offset these forces.

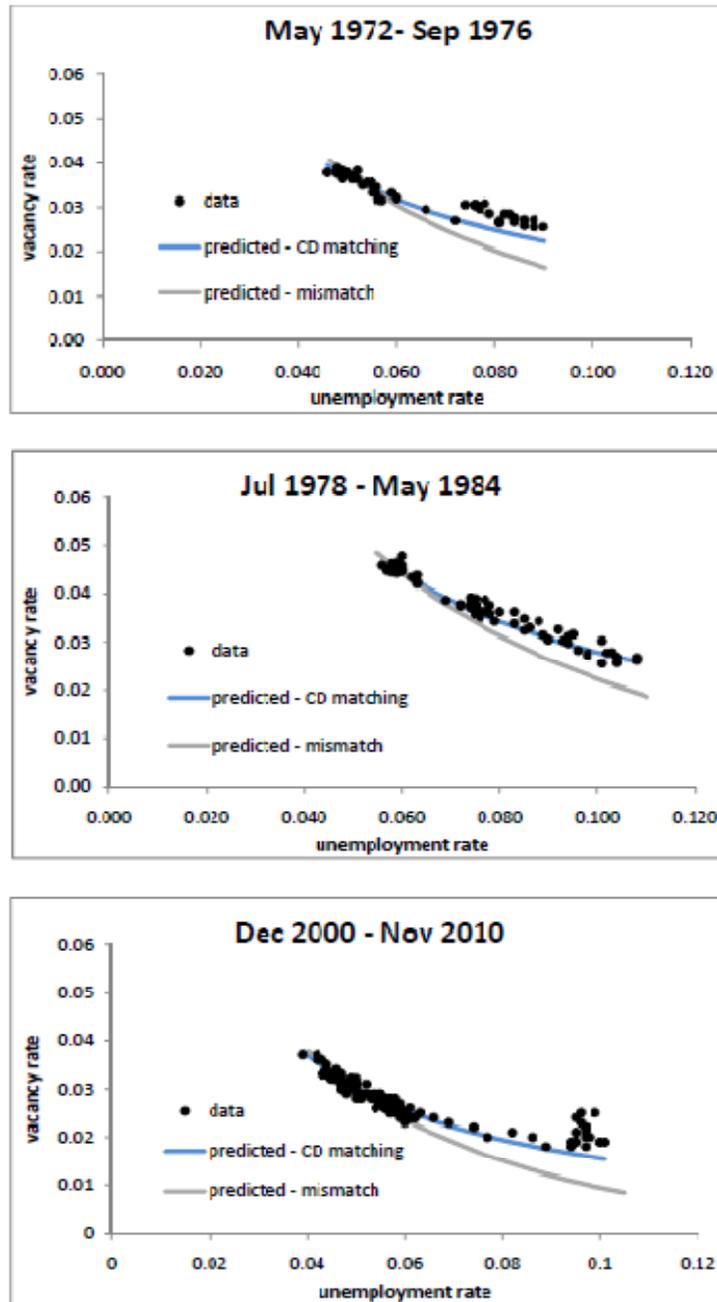
Table 1: Estimated Matching Function Parameters

	A	α
1972:05 – 1973:10	0.68	0.42
1978:07 – 1979:12	0.56	0.43
2000:12 – 2008:08	0.75	0.45

Table 2: Estimated Parameters for the Shimer (2007) Mismatch Model

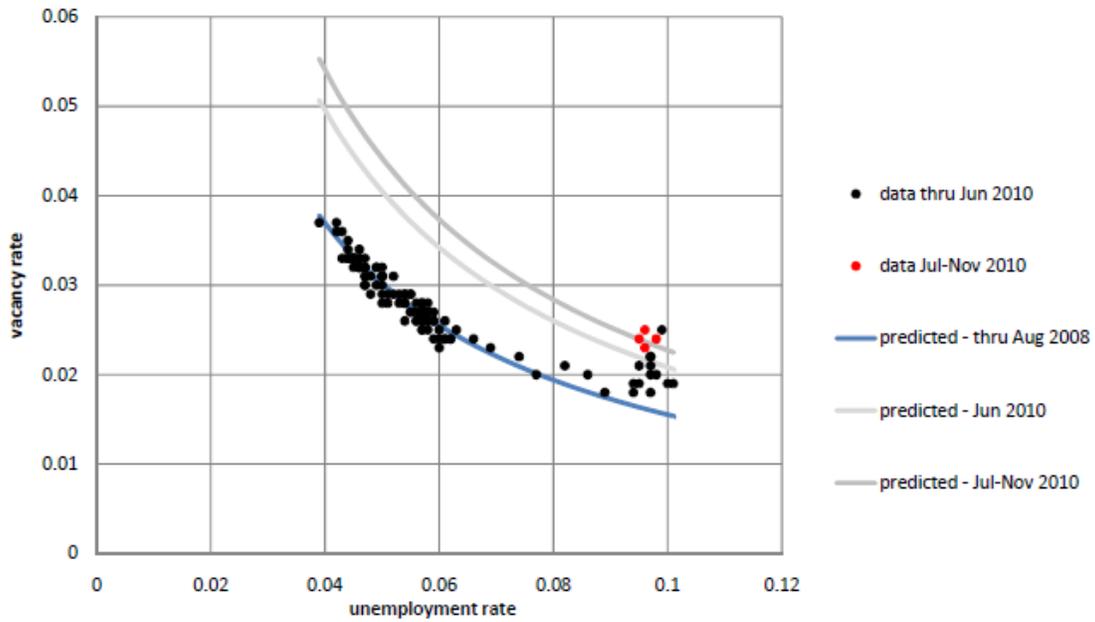
	m	n
1972:05 – 1973:10	168.7	165.8
1978:07 – 1979:12	119.7	118.0
2000:12 – 2008:08	210.5	205.5

Figure 1



Data: vacancy rate from JOLTS and Conference Board HWI normalized to match JOLTS vacancy rate, unemployment rate from BLS.

Figure 2: Beveridge Curve
Dec 2000 – Nov 2010



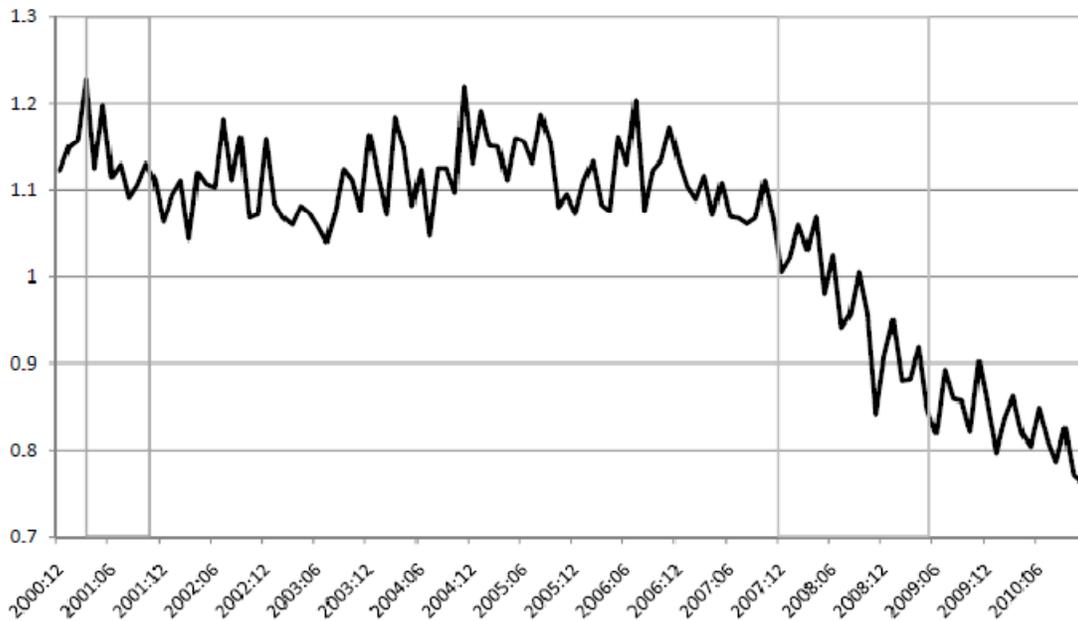
Curve through Aug 2008 assumes matching function with $A = 0.752$ and $\alpha = 0.45$ ($s = 0.03$)

Curve matched to June 2010 has same α but $A = 0.64$.

Curve matched to July–Nov 2010 has same α but $A = 0.61$.

Data source: vacancy rate from JOLTS, unemployment rate from BLS.

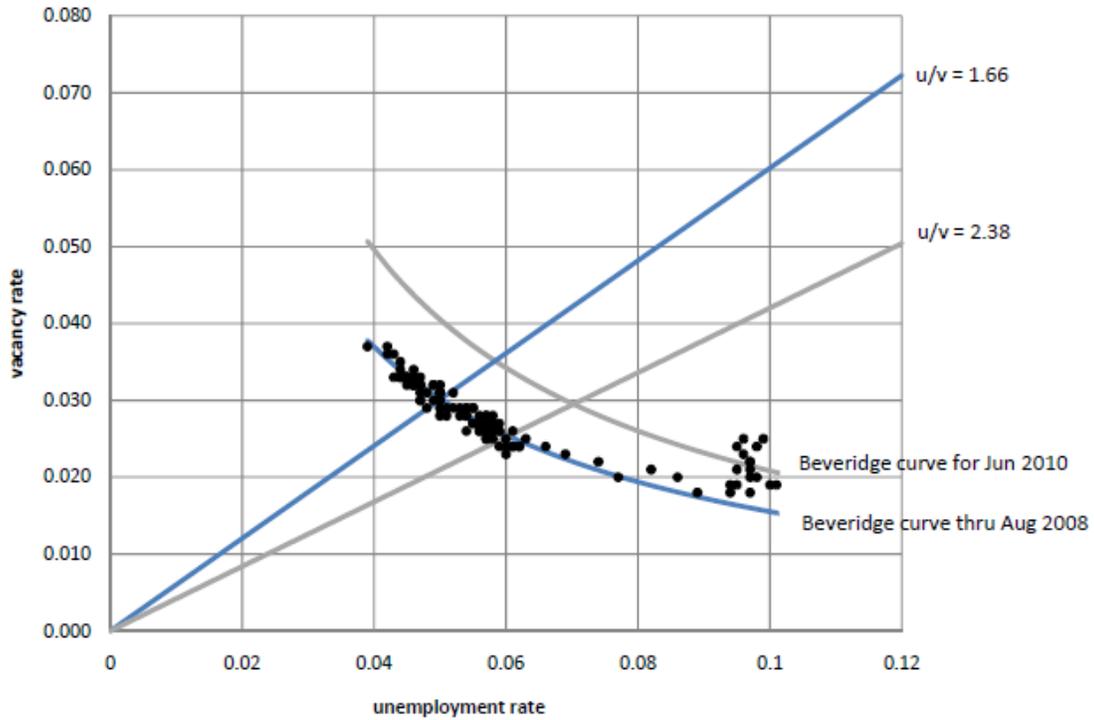
Figure 3: Productivity of Matching using Hiring Rate



Line corresponds to the reported hiring rate divided by $u^{.45}v^{.55}$.

Data source: hiring rate and vacancy rate from JOLTS, unemployment rate from BLS

Figure 4: Determination of Steady-State Equilibria



References

Elsby, Michael, Bart Hobijn and Ayşegül Sahin, 2010, "Disciplined Estimates of Gross Flows among Labor Market States," University of Michigan, Working Paper.

Hall, Robert, 2005, "Job Loss, Job-finding, and Unemployment in the U.S. Economy over the Past Fifty Years," NBER Macroeconomics Annual, pp. 101-137.

Kocherlakota, Narayana, 2010, "Inside the FOMC," speech given in Marquette, Michigan, August 17.

Mortensen, Dale, and Christopher Pissarides, 1994, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, Vol. 61, No. 3, pp. 397-415.

Petrongolo, Barbara and Christopher Pissarides, 2001, "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, Vol. 39, No. 2, pp. 390-431.

Pissarides, Christopher, 1985. "Short-run Equilibrium Dynamics of Unemployment Vacancies, and Real Wages," *American Economic Review*, Vol. 75, No. 4, pp. 676-90.

Shimer, Robert, 2005a, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, Vol. 95, No. 1, pp. 25-49.

Shimer, Robert, 2005b, "The Cyclical Behavior of Hires, Separations, and Job-to-Job Transitions," Federal Reserve Bank of St. Louis, *Review*, Vol. 87, No. 4, pp. 493-507.

Shimer, Robert, 2007, "Mismatch," 2007, *American Economic Review*, Vol. 97, No. 4, pp. 1074-1101.