Evaluating the role of labor market mismatch in rising unemployment

Gadi Barlevy

Introduction and summary

From the second half of 2009 through the end of 2010, the U.S. labor market witnessed a systematic increase in the rate of job openings while the unemployment rate remained essentially unchanged. Some have argued that, evidently, the problem in the labor market during this period was not that firms were reluctant to hire additional workers, but that, for whatever reason, firms seemed unable to find suitable workers to staff the positions they were trying to fill. By this logic, using monetary policy to encourage further hiring by firms would have been unlikely to drive down unemployment: If firms were already trying to hire but could not, why should policy actions that mainly serve to encourage even more hiring have any impact on unemployment? The unemployment rate did finally register a decline in late 2010 and early 2011—a development that may eventually render less acute the debate about the need for monetary policy to address the problem of high unemployment. Still, constructing a framework for interpreting such labor market patterns and their policy implications remains an important goal. This is especially true given that there have been other periods in which job vacancy rates seemed to rise without a commensurately large fall in unemployment, although those episodes were not as dramatic nor as long as the most recent one.

In this article, I show how the labor market matching function approach developed by Pissarides (1985) and Mortensen and Pissarides (1994) can be used to assess the validity of the proposition that recent trends in vacancies and unemployment necessarily point to a diminished role for monetary policy. More specifically, I show that this framework indeed suggests that an increase in unemployment without a commensurate decline in vacancies can be indicative of a labor market shock that monetary policy cannot offset. However, this framework can also be used to derive a bound on how much a shock of this type can affect unemployment. Applying these insights to the period of the Great Recession reveals that this type of shock by itself would lead to an unemployment rate of 7.1 percent, considerably lower than the unemployment rate during most of this period. The higher actual unemployment rate suggests that other types of shocks, which monetary policy may be able to address, must also be operating. Hence, the recent patterns in unemployment and vacancy data do not necessarily rule out an important role for monetary policy. Whether more expansionary monetary policy would have been beneficial is a question that is beyond the scope of this article. Nevertheless, the matching function approach frames this question in a potentially useful way—that is, as a question of why the value of taking on additional workers appears to be so much lower now than in normal economic times.

My article is organized as follows. I begin by describing the matching function approach, and then I show that the shocks that affect unemployment in this framework can be decomposed into two groups—those that affect the ability of firms to find and hire qualified workers and those that affect the value to a firm of taking on an additional worker. I next explain how the model can be used to predict how a shock to the ability of firms to hire, calibrated to match the facts on unemployment and vacancies during the Great Recession, affects unemployment. Using this result, I argue that the increase in unemployment due to this shock is much smaller than the actual increase in unemployment during this period, so a shock to the ability of firms to hire cannot by itself account for the

Gadi Barlevy is a senior economist and research advisor in the Economic Research Department at the Federal Reserve Bank of Chicago. *The author thanks Dan Aaronson, Dale Mortensen, Ayşegül Şahin, Dan Sullivan, and Marcelo Veracierto for their helpful comments.*

rise in unemployment during this time. I conclude with a discussion about how measurement issues are likely to affect these conclusions.

The matching function approach

In this section, I lay out the key features of the labor market matching framework developed in Pissarides (1985) and Mortensen and Pissarides (1994). This framework rests on two key assumptions.

The first key assumption is that the total number of new hires h in any given period can be expressed as a function of the number of workers who are unemployed during that period, u, and the number of vacancies firms post over that same period, v:

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

This assumption is similar to the assumption invoked by macroeconomists that one can use an aggregate production function to express total output produced in a given period as a function of the total number of hours and the aggregate capital stock for that period. That is, the process by which unemployed workers looking for jobs and employers with vacancies looking for workers form new hires is assumed to operate with such regularity that one can reliably predict the number of new hires per period by using only data on the number of unemployed workers and the number of vacancies firms post. Empirical analysis supports the idea that the number of new hires can be related to the number of unemployed and vacant positions in a fairly predictable way. Much of this evidence is summarized in the survey by Petrongolo and Pissarides (2001).[1]

The function m is referred to as a matching function. A common restriction on the matching function is that the number of new hires h falls short of both the number of unemployed u and the number of vacant positions v. That is, some unemployed workers and some positions will remain unmatched by the end of the period. This is meant to capture various frictions in the process of filling new jobs from the ranks of the unemployed—such as a lack of coordination that leads multiple workers to apply to the same vacancies while other vacancies remain unfilled, or the fact that workers and firms do not initially know how well suited they are for each other and figuring this out can be time-consuming. Studies that explore these frictions in detail reveal that they do not always give rise to empirically plausible matching functions, and in some cases they suggest different interpretations for why hiring, vacancies, and unemployment are related. Moreover, these frictions sometimes imply that the number of new hires should depend on other variables besides just the number of unemployed workers and the number of vacant positions.[2] However, Petrongolo and Pissarides (2001) argue that the matching function approach performs quite well empirically and is suitable for analyzing certain questions concerning the labor market, just as assuming an aggregate production function is often useful for analyzing macroeconomic questions. That is, many macroeconomists are willing to posit an aggregate production function that is invariant to various shocks that affect the economy, even though the conditions under which one can ignore the decisions of individual firms in different sectors and express aggregate output as a function of aggregate inputs are quite stringent.[3] In defense of this assumption, these macroeconomists would argue that the aggregate production function performs well empirically, so it is likely to be useful in predicting how the economy would respond to shocks that only affect the aggregate capital stock and labor hours—for example, a change in income taxes that affects how much labor is supplied but does not affect the technology available for producing goods and services. Likewise, advocates of the matching function approach view the close empirical relationship between aggregate hiring and aggregate unemployment and vacancies as justification for ignoring the decisions of workers and employers that underlie the process of job creation. These advocates proceed as if hiring can be summarized with a mapping of aggregate vacancies and unemployment to aggregate new hires that is "structural," meaning that the mapping is invariant to shocks that affect unemployment and vacancies but not the frictions inherent in the matching process.

Petrongolo and Pissarides (2001) argue that the matching function is particularly well approximated by a Cobb–Douglas specification, that is,

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

where A is a scale parameter that determines the productivity of the matching process and α reflects the sensitivity of the number of new hires to the number of unemployed workers. That is, this specification will produce reasonably good predictions for the actual hiring rate given the unemployment rate u and vacancy rate v for coefficients α and A that remain stable over relatively long periods and that can be estimated from historical data.

The second key assumption of the labor market matching framework is that firms post vacancies as long as doing so remains profitable, implying that the expected discounted profits to a firm from posting a vacancy should be zero in equilibrium. Let J denote the value of a filled job to the employer who creates it, and let k denote the cost of posting (and maintaining) the vacancy, including screening and interviewing potential candidates. Then the assumption that employers are free to enter the labor market and attempt to hire workers as long as it remains profitable to do so implies that the value of a filled job times the probability of filling it should equal the cost of posting the vacancy, ensuring expected profits are equal to zero. Pissarides (1985) and Mortensen and Pissarides (1994) assume that each posted vacancy is equally likely to be matched, so with $m(u,v)$ new hires, the probability of filling a job is equal to $m(u,v)/v$. In this case, the implications of free entry can be summarized as follows:

$$3)\quad \frac{m(u,v)}{v} J = k.$$

Not all models that give rise to a matching function representation as in equation 1 imply that the probability that any given firm expects its vacancy to be filled within the relevant period corresponds to $m(u,v)/v$. For example, even when there are underlying differences in firms, such as differences in the costs of processing applicants, we might still observe a stable relationship between aggregate hiring, unemployment, and vacancies. But different firms will assign different probabilities to filling their positions within the relevant period. Still, as long as $m(u,v)/v$ reasonably captures the probability of filling a position for the marginal firm at any point in time, proceeding with this assumption will be appropriate.

Substituting the Cobb–Douglas specification for $m(u,v)$ from equation 2 into the free-entry condition as given by equation 3 reveals that the free-entry condition can be expressed solely in terms of the ratio v/u.[4] This ratio is known in the literature as market tightness, since it reflects how many vacant positions are competing for each unemployed worker. In particular, the free-entry condition as given by equation 3 can be written as

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

As long as the parameters α and A remain constant over time, the free-entry condition as given by equation 4 tells us that if the value of a filled job relative to the cost of posting a vacancy, J/k, varies over time for any reason (such as a change in aggregate demand or changes in aggregate productivity), the market tightness ratio, v/u, would have to change as well. The fact that v/u changes with J/k ensures that firms continue to expect zero cumulative discounted profits from posting additional vacancies. Hence, if we knew how a particular macroeconomic event affected the ratio J/k, we could use the free-entry condition as given by equation 4 to deduce how this event should change the ratio of the vacancy rate to the unemployment rate we observe in the labor market. The two assumptions—the existence of a matching function and free entry into the labor market—thus impose a lot of structure on how various shocks affect labor market tightness as reflected in the ratio v/u.[5]

The Beveridge curve

With the introduction of one additional assumption, the labor market matching framework can be used to predict not only how v/u changes with J/k, but also how u and v change individually. In particular, suppose that the rate at which employed workers separate from jobs into unemployment is constant over time. This assumption may seem implausible at first, especially given the incidence of mass layoffs during recessions. However, the job separation rate that I need to assume is constant does not involve one-off spikes of job destruction that reflect immediate adjustment by firms 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.[6] Shimer (2005a) and Hall (2005) argue that fluctuations in this separation rate contribute little to overall changes in unemployment and can be ignored. In subsequent work, others argue that the separation rate appears to be quite cyclically sensitive, and find the separation rate makes an important but still relatively small contribution to overall fluctuations in unemployment.[7] However, their papers all look at the role of flows of workers from employment and unemployment without accounting for spikes of job destruction. Flows into unemployment that include bursts of job destruction may account for fluctuations in total unemployment, even if the separation rate that is relevant for my analysis is fairly stable. Later, I argue that data on unemployment and vacancies from three distinct episodes of high unemployment support the claim that the relevant separation rate, s, does not rise much during recessions. Moreover, if the separation rate were in fact higher during recessions, my calculation would only exaggerate the role of labor market mismatch, and the bound I derive for the effect of a shock to the ability of firms to hire would be too high.

To see what the model predicts for the behavior of v and u as opposed to their ratio v/u, consider what happens if J/k varies over time. Conditional on a given value of J/k, the free-entry condition as given by equation 4 tells us that the vacancy-to-unemployment ratio, v/u, must remain constant as long as J/k is constant. However, u and v could themselves change even while

J/k remains fixed, as long as they change in the right proportion. 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 quick pace of convergence is not just a theoretical result; rather, it has been confirmed empirically.[8] This finding may seem odd at first, since time-series data suggest unemployment is fairly persistent over time. However, it is important to note that I am referring to *conditional* (as opposed to *unconditional*) convergence in u and v. In other words, given a value of J/k, both u and v converge quickly to the steady-state values associated with this particular value of J/k. But if J/k follows a persistent process, unemployment will still appear to change slowly over time. Rapid *conditional* convergence is thus fully consistent with unemployment appearing to be a slow-moving process. Given that convergence to a steady state for a given J/k is quick, it follows that whatever the value of J/k happens to be at any point in time, the values of u and v we would observe should roughly coincide with the steady-state levels of these variables for that J/k.

To compute the conditional steady-state unemployment for a given J/k, note that the flow into unemployment is equal to $s(1 - u)$, where s denotes the separation rate into unemployment, while the flow out of unemployment is equal to the number of new hires, $Au^{\alpha} v^{1-\alpha}$.[9] Since flows into and out of unemployment are equal in steady state, I can use this equality to arrive at an implicit formula for the conditional steady-state unemployment rate associated with a particular v/u ratio, which is associated with a particular value of J/k:

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

Rearranging equation 5 allows me to express the vacancy rate v implied by the model for a given unemployment rate u as follows:

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

As long as the separation rate into unemployment s is constant, equation 6 implies a negative relationship between u and v. This relationship, when displayed graphically as a plot of the vacancy rate against the unemployment rate, is known as a Beveridge curve after the British economist William Beveridge, who first documented the negative relationship between the two series.

The negatively sloped Beveridge curve can be seen by plotting out the u and v implied by the model for different values of J/k. In particular, according to equation 4, changes in J/k will force changes in the ratio v/u. Intuitively, as jobs become more valuable, the probability of filling a job must fall to ensure firms still expect to earn zero profits. One can then deduce, from equation 5, that higher values of the ratio v/u imply lower values of u and, from equation 6, that lower values of u imply higher values of v.

Indeed, the only thing that induces a movement along a Beveridge curve as defined by equation 5 is a change in J/k. This result holds because the Beveridge curve in equation 5 is defined as the relationship between u and v for fixed values of A, α, and s. When these values are fixed, it is apparent from equation 5 that the unemployment rate u only changes if the ratio v/u changes. But when A and α are fixed, the free-entry condition as given by equation 4 tells us that the ratio v/u is entirely determined by J/k. Thus, a movement along the Beveridge curve occurs if and only if the value of taking on an additional worker relative to the cost of posting a vacancy changes. Various events can shift this value, including a change in worker productivity, a change in the bargaining power of workers, a change in aggregate demand, and a change in the employer's operating cost (such as a change to the cost of borrowing). But, for our purposes, all of these events can be grouped into a catchall category of shocks that affect the net value of a filled job or, alternatively, shocks that move the economy along a stable Beveridge curve.

The natural counterpart to shocks that induce a movement along a Beveridge curve are shocks that shift the Beveridge curve itself. As evident from equation 5, which defines the Beveridge curve, as long as s is fixed, the only way for the Beveridge curve to shift is if the matching function $m(u,v)$ itself somehow changes. A shift in the Beveridge curve thus corresponds to a shock that changes the way in which workers and employers come together to form new hires. One example of such a shock is a disruption that gives rise to greater mismatch between the skills employers require to fill their positions and the skills that unemployed workers currently possess—such as a shift in demand away from products the labor force is already skilled at making. Such a shock would presumably result in fewer positions being filled given the same number of unemployed workers and vacant positions, and thus, the productivity term A in the matching function would decline. The model thus delivers a clean dichotomy: Shifts of the Beveridge curve correspond to shocks to the ability of firms to hire (that is, changes in A), while movements along a fixed Beveridge curve correspond to changes in the incentives for firms to hire (that is, changes in J/k).

We can now recast the debate on the role of monetary policy in the face of high unemployment, using the terminology of the matching function approach. The observation that vacancies rose while unemployment was virtually unchanged implies the Beveridge curve must have shifted, that is, the hiring process became less efficient. There is arguably little monetary policy can do to affect the process by which firms and unemployed workers match up to generate new hires. However, whether there is any role for monetary policy depends on whether a shock to match productivity, A, is the only shock responsible for high unemployment. If the increase in unemployment is also due to a change in the relative value of a filled job, J/k, there may be some scope for monetary policy after all. So, for example, if the lower J/k reflects weak aggregate demand due to some underlying frictions, then monetary policy would have a role in addressing this. In the remainder of this article, I infer the decline in A from the shift in the Beveridge curve, and then I use the matching function approach to gauge how much a shock to A of this magnitude should have raised unemployment if the parameters that govern the value of a filled job remain equal to their levels during normal economic times (that is, to pre-recession levels). Since the implied unemployment rate falls far short of the actual unemployment rate that prevailed during this time, the high unemployment rate suggests that the relative value of a filled job, J/k, must have been lower during this period than during normal economic times. Whether this finding admits a role for monetary policy depends on why the value of a filled job is lower. Still, the calculation suggests that data on unemployment and vacancies do not rule out a role for policy per se and that high unemployment is due not only to an inability to hire among employers but also to a reduced willingness to hire.

Empirical Beveridge curves and estimating the matching function

The first step in my analysis involves inferring the reduction in match productivity A over this period from shifts in the Beveridge curve. For this, I must begin with a benchmark value for A in normal times. I can do this by fitting the Beveridge curve relationship in equation 6 to data from before the recent crisis. That is, I estimate the parameters α and A of the matching function, using data only for the period before unemployment began to take off, and then I look at how this relationship holds up in predicting vacancy rates for observed unemployment rates once the unemployment rate begins to climb. To do this, I use data from the U.S. Bureau of Labor Statistics' *Job Openings and Labor Turnover Survey* (JOLTS), which begins in December 2000. To estimate the Beveridge curve, I use data through August 2008, just before the big run-up in unemployment that started a few months after the official start date of the recession according to the National Bureau of Economic Research (NBER). To estimate α and A, I follow Shimer (2005b)—who estimates the job separation rate at a monthly frequency—and set $s = 0.03$. However, the choice of s is essentially a normalization.[10] To infer A and α, I set out to match two specific aspects of the data. First, for each month, I use equation 6 to predict a vacancy rate v given an unemployment rate u in that month. The parameters A and α were chosen to ensure that the average predicted vacancy rate over all these months was equal to the actual average vacancy rate over the same period, namely, 0.029. Second, I chose the parameters to ensure that the difference in v between the start date and end date of my series, which is 0.013 in the data, matched the difference in the predicted v at these two dates. Matching the model and the data this way yielded values $A = 0.75$ and $\alpha = 0.46$. The implied (fitted) Beveridge curve corresponds to the dark gray line in panel A of figure 1, which is shown together with data on unemployment and vacancies. The points in black correspond to the data through August 2008 that were used to estimate the curve, while the points in red correspond to observations from September 2008 onward that were *not* used in estimating the curve. To help illustrate how u and v evolved from September 2008 onward, consecutive months are connected.

As evident in panel A of figure 1, the Beveridge curve implied by the model does a reasonable job initially of predicting the vacancy rate at each unemployment rate beyond the period it was estimated to match—in fact, beyond the historical range of both the unemployment and vacancy series used to estimate the curve. The forecast only starts to break down around August 2009, suggesting a change in the matching function. The fact that the curve fits well throughout the official recession as determined by the NBER—that is, from December 2007 through June 2009—and only breaks down afterward provides some reassurance that the separation rate, s, did not appear to rise significantly while the economy was contracting.

As a further check on how well the matching function approach fits the data, I went back and repeated the same exercise for two other periods with similarly high unemployment—namely, for November 1973–March 1975 and for January–July 1980 and July 1981–November 1982. Since JOLTS only begins in December 2000, I use the Conference Board's Help-Wanted Advertising Index for my measure of

FIGURE 1

Actual and predicted job vacancy rates for given unemployment rates during selected episodes

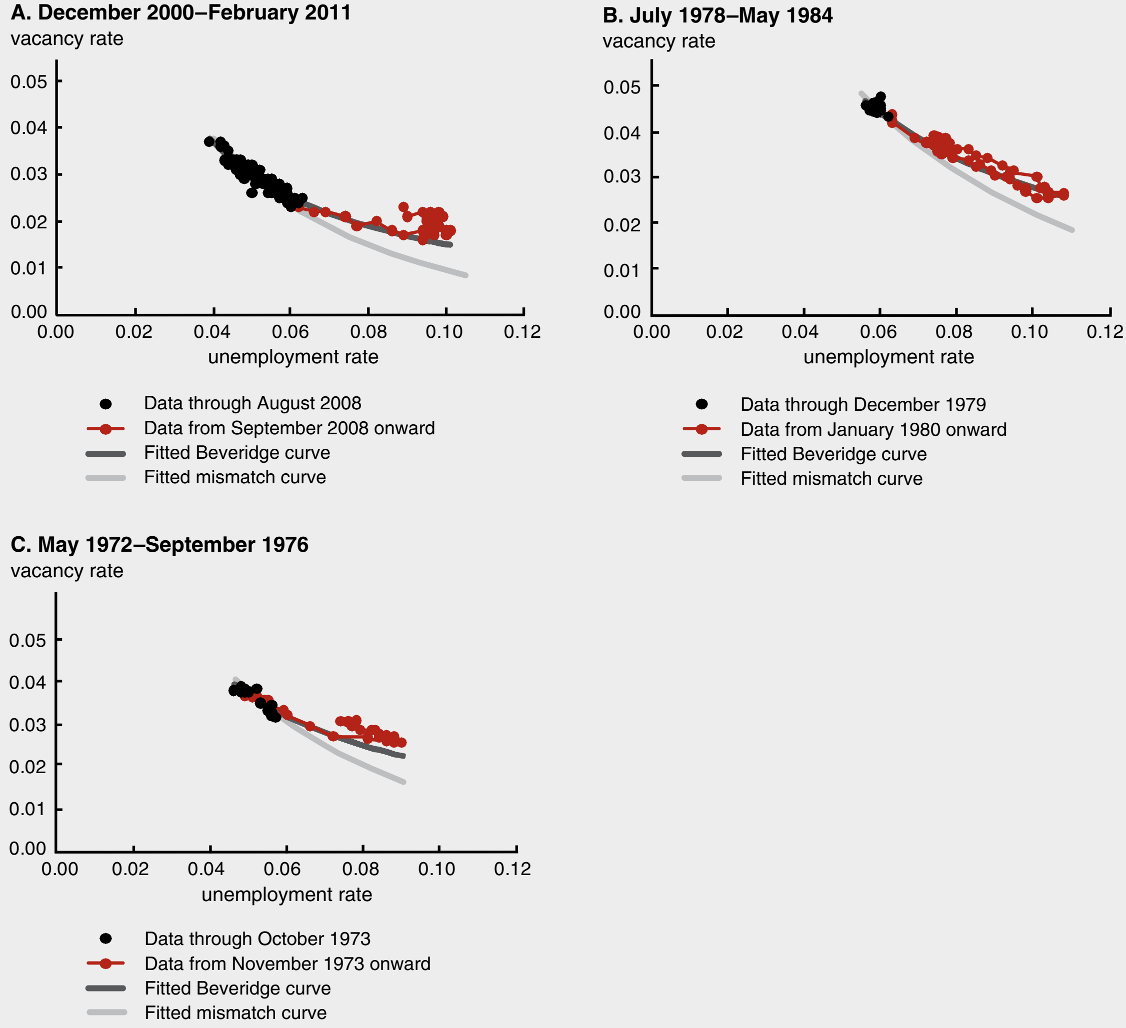

Notes: All curves are fitted only on data indicated by the black points. The fitted mismatch curves are based on Shimer (2007). See the text for further details.

Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series; and Conference Board, Help-Wanted Advertising Index, from Haver Analytics.

vacancies. This index is constructed using the number of newspaper advertisements for vacant positions. To transform this index into a vacancy rate, I normalized the series to coincide with the JOLTS vacancy rate for the period in which they overlap. For each recession, I followed a similar approach to estimating the matching function—that is, by taking data from a period prior to the recession to estimate the function and then seeing how the implied Beveridge curve does during the recession. However, since the Conference Board's Help-Wanted Advertising Index may be unreliable over long periods (given various gradual changes in the tendency of employers to rely on newspaper advertising for recruiting), I restrict attention to shorter periods for my estimation. For the 1973–75 recession, I look at the 18-month period before the NBER peak date, that is, May 1972 through October 1973. For the 1980 and 1981–82 recessions, I look at the 18-month period before the NBER peak date for the 1980 recession, that is, July 1978 through December 1979. In both cases, I estimate A and α in the same manner as for the data between December 2000 and August 2008—that is, I choose these parameters so that the average predicted vacancy rate over the period is equal to the actual average and the difference in vacancy rates between the start and end dates is the same in the predicted series as in the actual series. The estimated coefficients are reported in table 1, and the implied (fitted) Beveridge curves are illustrated as dark gray lines in the panels of figure 1 (the light gray lines are explained in the next paragraph). For the 1973–75 recession (panel C) and the 1980 and 1981–82 recessions (panel B), the data points used to estimate the curves are depicted in black, while the remaining data points are depicted in red. Note that in both panels B and C of figure 1, the original Beveridge curves are estimated from a period with little variation in the data, especially in the case of the 1980 and 1981–82 recessions. Still, the approach to using data before the recession(s) to estimate a curve performs well. For all of the recessionary periods I consider, the vacancy rate predicted for a given unemployment rate remains close to the actual vacancy rate once unemployment begins to rise. In both panels B and C of figure 1, the Beveridge curves do eventually appear to shift, although the shifts are much smaller than in the most recent episode, shown in panel A of figure 1. As evident in table 1, the coefficient α is estimated to be essentially the same in all three periods. This is consistent with my maintained approach of assuming the parameter α is fixed and that any changes in the matching function must therefore be attributed to A, the match productivity parameter.

For comparison, I also considered an alternative explanation for the matching function based on the notion of mismatch advanced by Shimer (2007), which was used by Kocherlakota (2010) to analyze the same labor market trends I consider here. The Shimer (2007) mismatch model offers a different interpretation for the relationship between new hires and unemployment and vacancies, and leads to a different zero-profit condition from equation 3.[11] Shimer's (2007) model also involves two parameters, which he denotes m and n. As I did earlier, I use the period before the recent run-up in unemployment (and the 18-month period before the start of the NBER recession for the two earlier episodes) to estimate these parameters and then consider how the model performs when unemployment rises. Following Kocherlakota (2010), in each case I choose these parameters to match the average values of u and v in the earlier period that is meant to reflect normal economic times. The estimates for the two parameters in each of the three episodes are summarized in table 2, and the implied curves relating unemployment and vacancies are shown in the respective panels in figure 1 in light gray. In all three episodes, Shimer's (2007)

TABLE 1

Estimated parameters for a Cobb–Douglas matching function

	A	α
May 1972–October 1973	0.68	0.42
July 1978–December 1979	0.56	0.43
December 2000–August 2008	0.75	0.46

Note: See the text for further details.
Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series; and Conference Board, Help-Wanted Advertising Index, from Haver Analytics.

TABLE 2

Estimated parameters for the Shimer (2007) mismatch model

	m	n
May 1972–October 1973	168.7	165.8
July 1978–December 1979	119.7	118.0
December 2000–August 2008	210.5	205.5

Note: See the text for further details.
Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series; and Conference Board, Help-Wanted Advertising Index, from Haver Analytics.

mismatch model predicts that the vacancy rate should decline more rapidly with unemployment than either what my estimated Beveridge curves predict or what we actually observe in the data. Since a shift in the curve in Shimer's model can be thought of as a shock to the ability of unemployed workers and job vacancies to match up with one another, this would suggest that all three recessionary periods and their subsequent recoveries were associated with significant rises in mismatch. While this reading of the data is certainly possible, it is striking that much of the discussion of the role of labor market mismatch during the Great Recession has tended to treat this phenomenon as exceptional; many of the explanations for the rise in mismatch in the labor market over the course of the Great Recession have emphasized features that are unique to this episode, such as the unprecedented collapse in house prices. Such views seem at odds with a specification that implies all three recessionary periods were associated with similarly large increases in labor market mismatch. The matching function approach is therefore more consistent with the view that the most recent episode is exceptional.

FIGURE 2

Fitting Beveridge curves during the Great Recession

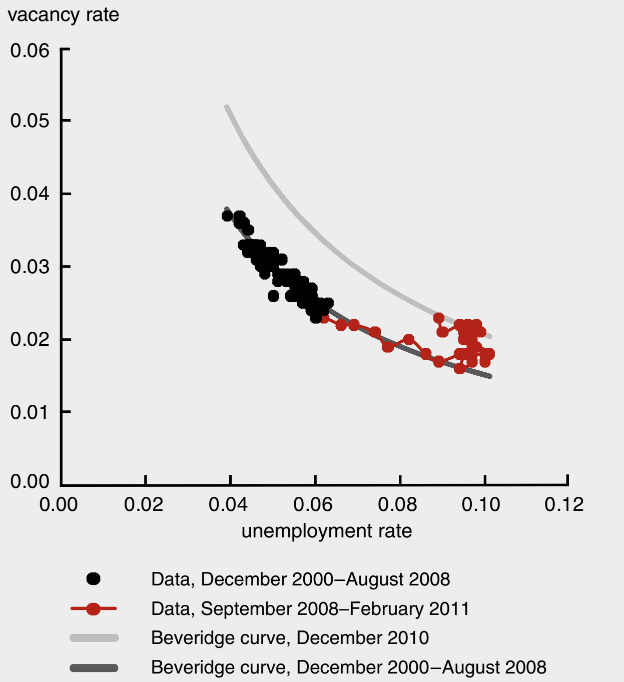

Note: See the text for further details.
Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series, from Haver Analytics.

Inferring the extent of mismatch

After estimating the parameters associated with the Beveridge curve for normal economic times, I next turn to how the decline in match productivity A can be inferred from the apparent shift in the Beveridge curve following the most recent recession. Using data through August 2008, I know that the initial productivity of the matching function is given by $A_0 = 0.75$. I can deduce the value of A_1 needed to match a given unemployment and vacancy pair at any other point in time by using equation 6. For example, to match the data for December 2010, when $u = 0.094$ and $v = 0.022$, match productivity A_1 must solve

$$7)\quad 0.022 = \left[\frac{0.03}{A_1}\left(0.094^{-0.46} - 0.094^{0.54}\right)\right]^{\frac{1}{0.54}}.$$

Solving for A_1 yields $A_1 = 0.633$, that is, by December 2010 the productivity of the matching function declined 16 percent from its original level before the recession. Figure 2 shows the Beveridge curves for both values of A_0 (the dark gray line) and A_1 (the light gray line). In principle, I can fit a new Beveridge curve through any data point. The most recent observation at the time of this writing, for February 2011, lies on the same Beveridge curve implied by the data from December 2010, as evident in figure 2. Moreover, this curve is close to the highest curve one could fit through any of the data points between September 2008 and the end of the JOLTS sample. This leads me to focus on the curve that runs though the data point corresponding to December 2010 in measuring the decline in match productivity A.

An alternative way to infer the change in A over the course of the Great Recession would be to bring in additional data on new hires rather than only rely on the data for unemployment and vacancies. 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 $u^{\alpha}v^{1-\alpha}$, using my previous estimate of $\alpha = 0.46$. In principle, this should give me a time series for match productivity A. This implied time series is depicted in figure 3. If I consider the period between August 2008 and December 2010, the implied match productivity declined by about 20 percent—a little larger than what I get without using hiring data and looking only at the implied shift in the Beveridge curve. However, as evident from figure 3, match productivity using data on new hires starts to fall around December 2007, considerably before any indications of a shift in the Beveridge curve relating unemployment and vacancies. The decline in match productivity between December 2007 and December 2010 is thus much larger, on the order of 25–30 percent. However, this decline is sensitive to the value of α, and it corresponds to 20 percent if I set $\alpha = 0.40$ instead of 0.46. But regardless of the precise value for α, data on hires suggest the decline in match productivity begins much earlier than the shift in the Beveridge curve. That said, both the magnitude and timing of the decline of matching efficiency depend on the measure of new hires used. Barnichon and Figura (2010) use the flows from unemployment to employment rather than all new hires, and find that the decline in A in 2009 is sharp, and its magnitude is comparable to what I estimate using the Beveridge curve.

Veracierto (2011) reviews several different approaches to estimating the productivity of the matching function, based on shifts of the Beveridge curve. These include accounting for flows into and out of nonemployment (see note 1, p. 94), measuring new hires based on flows into employment from either just unemployment or both unemployment and nonemployment, and using either shifts of the Beveridge curve or a comparison of changes in new hires to changes in unemployment and vacancies to deduce a time series for A. His preferred estimate suggests A had declined 15 percent since December 2007, in line with the estimate I infer from the shift in the Beveridge curve.

Since my calculations rely on the Beveridge curve specification in equation 5, I will use the estimate for the change in A based on how much the Beveridge curve shifted during the Great Recession in what follows.

FIGURE 3

Implied match productivity using data on new hires, 2001–11

Notes: The red vertical line corresponds to August 2008—the date at which the sample used to estimate the Beveridge curve in figure 1, panel A, ends. See the text for further details.
Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series, from Haver Analytics.

The effects of mismatch on unemployment

Once I determine that match productivity A declined by 16 percent between the level I estimate for normal economic times and the end of 2010, I can determine the effect of a shock of this size on the unemployment rate. To do this, I start at a steady-state unemployment rate of 5 percent, which roughly corresponds to the historical average of unemployment for the period covered by JOLTS through August 2008. From the Beveridge curve relationship implied by equation 6, I know the implied vacancy rate would have to be

$$v = \left[\frac{0.03}{0.75}\left(0.05^{-0.46} - 0.05^{0.54}\right)\right]^{\frac{1}{0.54}} = 0.03.$$

The implied ratio of u/v during these normal times will therefore equal $\dfrac{0.05}{0.03} = 1.67.$

Next, I use the free-entry condition as given by equation 4 to deduce how much a shock to A will affect the ratio u/v. To do this, suppose the shock to A had no effect on the ratio of the value of a filled job to the cost of posting a vacancy, J/k. In fact, J and k are determined endogenously, and changes in A can, and in many cases will, affect these values. However, for reasons I explain in more detail later, changes in A are likely to move J/k in a particular direction, implying that the unemployment rate holding J/k fixed will correspond to an upper bound on unemployment. Assuming J/k is constant thus offers a useful benchmark case.

Rearranging equation 4, I get $\dfrac{u}{v} = \left(\dfrac{k}{AJ}\right)^{1/\alpha}$.

Hence, given the estimated decline in the productivity of matching, holding J and k fixed, a decrease in A from 0.75 to 0.633 should lead the unemployment-to-vacancy ratio to rise by a factor of $\left(\dfrac{0.75}{0.633}\right)^{1/0.46} = 1.45.$

Given I needed the ratio u/v to equal 1.67 to support a 5 percent unemployment rate under the original Beveridge curve, I can deduce that the new equilibrium ratio of u/v will equal

$$1.45 \times 1.67 = 2.42.$$

Plugging in $u/v = 2.42$ into the Beveridge curve relationship in equation 5 when $A_1 = 0.633$ gives us the implied unemployment rate that must prevail in the new equilibrium:

FIGURE 4

Computing unemployment response from shock to match productivity A

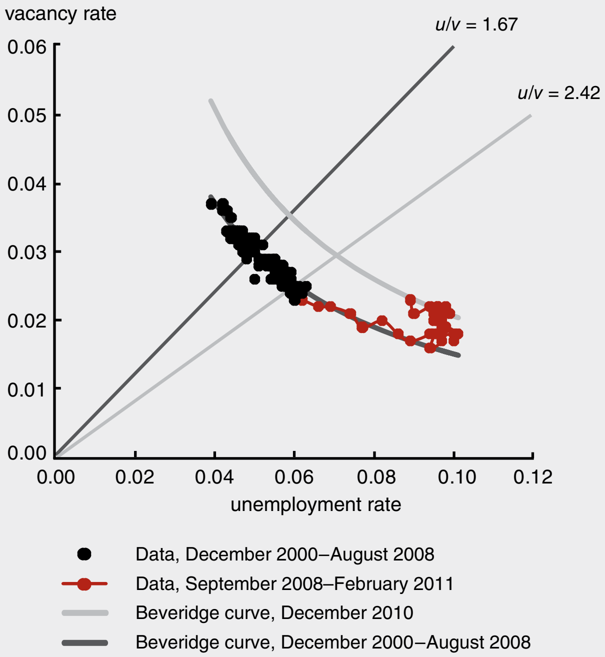

Note: See the text for further details.
Sources: Author's calculations based on data from the U.S. Bureau of Labor Statistics, *Job Openings and Labor Turnover Survey* and civilian unemployment rate series, from Haver Analytics.

$$u = \frac{0.03}{0.03 + 0.633(2.42)^{-0.54}} = 0.071.$$

Thus, a shock to A, calibrated to the magnitude implied by the patterns observed in data on unemployment and vacancies alone, will raise the unemployment rate to 7.1 percent as long as it leaves the value of a filled job unchanged. Since 7.1 percent is much lower than the actual unemployment rate, this value suggests that shocks to the productivity of matching alone cannot account for the high unemployment rate.

Figure 4 illustrates the same calculation graphically. Each level of match productivity A is associated with a distinct Beveridge curve and a distinct ratio u/v determined by the free-entry condition as given by equation 4, which in the figure corresponds to the line emanating from the origin. The original Beveridge curve and free-entry condition associated with $A = A_0$ are shown in dark gray, while the new Beveridge curve and free-entry condition associated with $A = A_1$ are shown in light gray. A decline in A not only shifts the Beveridge curve but also rotates the free-entry condition clockwise to a degree that depends on the size of α. Intuitively, if hiring becomes less effective, firms will have an incentive to post fewer vacancies per unemployed worker, ultimately leaving more workers unemployed.

As I noted earlier, both k and J are in fact determined endogenously and will likely change when A does. For example, the process of creating a vacancy requires productive inputs such as labor, so the cost k will depend on wages that are determined endogenously. Since wages tend to rise and fall with economic activity both in the data and in the original Mortensen and Pissarides (1994) model, I would expect the cost of posting a vacancy k to fall as the unemployment rate rises. As for the value of a filled job to an employer J, there are various reasons to suspect it will be higher when there is more unemployment. Mortensen and Pissarides (1994) posit that the value of a filled job is determined as the result of Nash bargaining between workers and firms over the surplus from a match.[12] But the surplus from matching is higher when v/u is low, so a fall in A will lead to a higher value for J.[13] Intuitively, when it is easy to find a match, matching immediately is only slightly more valuable than separating and letting the two parties search for new matches, which they will likely find quickly. More generally, various realistic features that are absent in the benchmark 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. Essentially, diminishing marginal utility or diminishing marginal returns make another employed worker more valuable when fewer workers are employed. If both k falls and J rises at higher unemployment rates, the effect of a shock to A on u/v would only be smaller. As such, 7.1 percent should be viewed as an *upper bound* rather than a point estimate.[14] This result only reinforces the point that the high unemployment rate that was observed during this period should not be blamed solely on a decline in the ability of firms to fill their positions, but also on greater reluctance among firms to hire as reflected in a lower J/k.

Measurement issues

The calculations presented in the preceding section are based on the assumption that the inputs that go into creating new matches—namely, unemployment and vacancies—are measured accurately. However, there are reasons to suspect both series may systematically misrepresent the nature of inputs that enter the hiring process while the empirical Beveridge curve shifted. I now discuss each of these series in turn, as well as the implications of mismeasurement for my analysis. I will argue that in both cases, measurement issues only strengthen the conclusion that the decline in the ability of firms to hire cannot by itself account for the bulk of the increase in unemployment during this period.

I first consider the unemployment series. One distinguishing feature of the current episode of high unemployment is the exceptionally long duration of unemployment insurance (UI) benefits; in some U.S. states, the unemployed can receive UI benefits for up to 99 weeks. Indeed, several research papers have sought to estimate the effect of these extensions on both the unemployment rate and unemployment durations.[15] The extension of UI benefits can matter for my analysis in several ways, including the method by which unemployment is measured. First, though, it will be useful to review the various ways in which explicitly incorporating UI benefits into the model can matter for unemployment.

One reason UI benefits can matter is that they lower the cost of remaining unemployed, allowing workers to be more selective about which job they take. As a result, relative to the case in which UI benefits remained unchanged, unemployed workers will prefer to continue searching more often, and a smaller fraction of the contacts between unemployed workers and vacant positions will result in a match, that is, a new hire. Indeed, this provides one potential explanation for the apparent decline in match productivity. Note that this effect is already taken into account in the calculation I sketched out before; that calculation tells us how much a decline in the ability of firms to hire—for *whatever* reason—ought to affect unemployment. Indeed, all the papers estimating how much the extension of UI benefits contributed to unemployment find effects that are smaller than the bound I estimate.

Second, when a worker and an employer agree to form a match, the extension of UI benefits may require an employer to offer a worker higher wages given that more generous UI benefits improve the bargaining position of workers. This effect is emphasized in Kocherlakota (2011). Unlike the first effect that appeared as a lower value for A, this effect would show up directly as a lower value for J, the value of filling a job to an employer. Indeed, this may be one reason for why the value of a filled job to an employer appears to be lower now than it is in normal times.

Neither of these two effects poses a problem for determining whether the rise in unemployment can be attributed solely to a decline in the ability of firms to hire. Rather, they merely suggest potential interpretations for what might be driving shocks to A or J. However, there is a third potential implication of extending UI benefits that may act to distort measured unemployment and could pose a problem for my calculation. In particular, extended UI benefits may encourage disaffected workers who prefer to leave the labor force to present themselves as nominally unemployed in order to qualify for UI benefits. This will be the case even if such workers are not actively looking for a job beyond whatever token steps are needed to maintain their status. Such a phenomenon would make the measured unemployment rate seem higher than its true value. Formally, let u^* denote the fraction of the labor force that is actively looking for jobs, and let u_0 denote the fraction of the labor force that is not really looking for a job but reports itself as being unemployed. If the latter fraction literally takes no steps to search for a job, the matching process will only partner up the true unemployment and vacancy positions, and the number of hires will be given by

8) $h = m\,(u^*,v)$.

At the same time, the official unemployment series will correspond to $u = u^* + u_0$, leading us to expect $m(u^* + u_0, v)$ hires. Since the matching function is increasing in both arguments, this will make the matching process appear less efficient than it truly is: We would observe surprisingly few hires given the seemingly large number of unemployed. Hence, the decline in match productivity A inferred from the shift in the Beveridge curve would exceed the true decline in A that enters the free-entry condition as given by equation 4. Since my approach provides an upper bound on the effect of a decline in the ability to hire on unemployment, though, overstating the decline in match productivity A will not overturn my results. If anything, it suggests the unemployment rate that should be expected from the decline in the ability of firms to hire is actually smaller than 7.1 percent.

Next, I turn to the time series for vacancies. Recent work by Davis, Faberman, and Haltiwanger (2010) has called into question whether vacancies provide a consistent measure of recruiting effort over time. In particular, they show that the vacancy yield, or the ratio of hires per vacancy, varies systematically across employers. For example, growing firms seem to be better at hiring, in the sense of being able to hire more workers per each vacancy posted. Davis, Faberman, and Haltiwanger (2010) argue that this pattern arises because the process of hiring requires firms to invest some effort into recruiting beyond posting the number of vacant positions they are seeking to fill.[16] They further reason that the same pattern should also occur over the business cycle: In recession times, when overall hiring is low, firms are likely to put in less effort into recruiting than in boom times. Thus, employers' hiring efforts would decline by more than would be reflected in the time series for the number of vacancies posted.[17]

Davis, Faberman, and Haltiwanger (2010) formalize these concerns as follows. Suppose that the effort that firms invest in recruiting can be summarized by the product of q and v, where q denotes recruiting intensity and v denotes the vacancy rate. The total number of hires is then given by

9) $h = m\,(u, qv)$.

That is, matching depends not on the number of vacancies, but vacancies together with how much firms invest in filling these vacancies. When recruiting intensity q falls below its historical average, the time series for vacancies v will fail to register this and will therefore overstate the overall recruiting effort. Using vacancies to proxy for recruiting efforts will then make matching efficiency appear to fall more than it in fact does. That is, we may wrongly conclude that firms find it more difficult to hire when in fact they are voluntarily choosing to search in a way that reduces the odds of hiring. Once again, this will cause us to overstate the decline in match productivity A from apparent shifts in the Beveridge curve and, therefore, to overstate the increase in the unemployment rate that can be attributed to less efficient matching now than in the past.

To provide a more quantitative illustration of this result, I can use the suggestion in Davis, Faberman, and Haltiwanger (2010) of proxying for recruiting intensity q by using the way in which the vacancy yield (hires per vacancy) varies across firms with different hiring rates. In particular, using variation in the vacancy yield across firms, they conclude that the elasticity of q with respect to overall firm hiring is given by 0.72. This implies setting $q = h^{0.72}$. Davis, Faberman, and Haltiwanger (2010) provide some evidence that this modification improves the time-series fit of the matching function. If this proxy is accurate, I can simply repeat the calculation for how much the apparent decline in match productivity should have increased unemployment, but replace the vacancy rate v in equations 1 through 6 with $h^{0.72}v$ as the second argument in the matching function.[18] Fitting a Beveridge curve to data on unemployment and this adjusted vacancy series through August 2008 yields $A_0 = 0.7$ and $\alpha = 0.54$. To match the data for December 2010 requires $A_1 = 0.605$, which is a smaller decline of only about 14 percent. For this decline, the implied unemployment rate due to just this shock to match productivity would be at most 6.3 percent. Correcting for measurement problems in vacancies can thus have a significant impact on how much unemployment is attributed to reduced effectiveness in hiring.

Conclusion

Recent labor market trends have raised concerns that the unemployment rate is high not because employers are reluctant to hire but because they are unable to hire—that is, for whatever reason, firms are unable to find suitable workers to staff the positions they are trying to fill. These concerns, if true, would cast doubt on using monetary policy to stimulate the labor market, since it works by encouraging firms to hire more. The matching function approach pioneered by Pissarides (1985) and Mortensen and Pissarides (1994) offers a framework for analyzing these issues. In particular, that framework can be used to separate the shocks that drive unemployment into two groups: shocks that affect the probability of finding a suitable worker and shocks to the value a worker generates once hired. The same framework allows us to estimate how much the probability of finding a worker declined and to compute a bound on how much this effect by itself would raise the unemployment rate. This bound as I have calculated it suggests that a decline in the ability to hire accounts for less than half of the total rise in unemployment during the Great Recession and that part of this rise in unemployment must be because firms find hiring less profitable.

While there is little monetary policy can do if firms find it more difficult to find suitable workers, there may be scope for monetary policy when firms find it less profitable to hire workers than during normal times. Whether such a role for monetary policy is warranted depends on why the value of a filled job to an employer is lower than in normal times. For example, if filled jobs are less valuable because of a shock that makes workers less productive, there is arguably little that monetary policy should do in response. But if jobs are less valuable because of insufficient aggregate demand on account of some market friction, there may be a role for monetary policy to stimulate demand. The key question for policy, then, is not what unemployment and vacancy data tell us about the possibility of mismatch, but why firms seem to find hiring workers less attractive than usual. Unfortunately, while the matching function approach is useful in pointing out the value of a filled job to an employer as an important variable, it offers little direct guidance as to why this value is so much lower now relative to normal times.

NOTES

[1]Of course, newly hired workers do not come only from the ranks of the unemployed; some were employed elsewhere, while others were not employed but did not report actively looking for a job either (that is, they were classified as "not in the labor force" by the U.S. Bureau of Labor Statistics, per the definition available at www.bls.gov/cps/cps_htgm.htm#nilf). In practice, the hiring rate can be accounted for quite well using data on unemployment, perhaps because the number of hires from out of the labor force and the number of hires of already employed workers move in opposite directions over the business cycle and tend to offset one another. One way to avoid the logical inconsistency of using data on unemployment to explain all new hires regardless of whether the worker was previously unemployed is to replace the number of new hires in equation 1 with the flow of workers from unemployment to employment, as in Barnichon and Figura (2010) and Veracierto (2011). While this approach restricts attention only to new hires who were previously unemployed, it suffers from the problem that the total number of vacancies is an imperfect measure of firm inputs into hiring the unemployed, since firms' efforts to fill these vacancies are aimed at hiring all workers and not just workers who are already unemployed.

[2]Petrongolo and Pissarides (2001) survey the microfoundations of the matching function, although several important papers in this area were published after their survey. The traditional model of coordination frictions, due to Butters (1977), assumes firms post vacancies, workers submit a single application each to some vacancy chosen at random, and each firm hires at random among the applications it receives. Burdett, Shi, and Wright (2001) emphasize that this model does not give rise to empirically plausible matching functions and that the number of hires per period will depend on additional variables, such as the size distribution of firms. Albrecht, Gautier, and Vroman (2003) assume workers can apply to multiple vacancies, but this does not give rise to empirically plausible specifications either. Lagos (2000) and Shimer (2007) model coordination frictions by letting firms and workers end up at different locations; firms choose locations at random and workers choose locations optimally to maximize their expected earnings (per Lagos) or at random (per Shimer). There are no frictions at any given location, so whichever side (firms or workers) arrives in smaller numbers winds up fully matched. Thus, each location will remain with either unemployed workers or vacant positions, but not both. Unemployed workers and vacancies are thus not inputs into forming new hires as the matching function approach implicitly assumes, but consequences of poor coordination between employers and workers on where to locate. When workers choose locations optimally, the matching function is not empirically plausible. When workers instead choose locations at random, the matching function matches the data well, at least for a certain range of unemployment and vacancies rates. Stevens (2007) develops a different theory of the matching function based on the notion that workers take time to screen heterogeneous jobs, rather than on coordination problems. She finds that the implied aggregate matching function is approximately Cobb–Douglas, as in equation 2 (p. 83). Decreuse (2010) develops a model where workers apply to jobs they do not realize are already filled. He finds that the implied matching function will depend on lagged variables beyond just the contemporaneous numbers of unemployed workers and vacant positions.

[3]For a survey that criticizes the use of aggregate production functions, see Felipe and Fisher (2003).

[4]The same is true more generally for any specification $m(u,v)$ that exhibits constant returns to scale.

[5]It should be noted that a recent body of literature, starting with Shimer (2005b), 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 the value of a filled job to the employer who creates it, J, varies enough over the cycle in these models, not whether the matching function can explain how new hires vary with unemployment and vacancies. My calculation does not depend on how J varies with aggregate conditions, nor does it impose much structure on how J ought to change over the cycle; and hence, it is not subject to this critique.

[6]More precisely, consider the Mortensen and Pissarides (1994) model where the separation rate into unemployment is endogenous. That model assumes jobs are hit with idiosyncratic shocks to the profitability of any given job at a constant rate λ per unit time. The shock term ε is drawn each time from some fixed distribution F. Firms optimally choose to terminate a job and send the worker into unemployment for severe enough shocks, that is, when ε falls below some critical level ε_d. Suppose that in a recession, firms become more demanding and raise the critical level to some higher value ε_d'. When the shock associated with the recession first hits, the unemployment rate will jump and the flow into unemployment will spike as all jobs whose ε lies between ε_d and ε_d' will be terminated immediately. The spike in the separation rate will appear large even when the regular flow into unemployment $\lambda F(\varepsilon_d)$ changes only modestly. My assumption that the separation rate is constant over time only requires that $\lambda F(\varepsilon_d)$ is relatively stable, not that flow rates from employment to unemployment (which will reflect spikes) be stable.

[7]Some examples are Mazumder (2007); Fujita and Ramey (2009); and Elsby, Hobijn, and Şahin (2010).

[8]In particular, Shimer (2005a) 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 instant. He then shows that this steady state is nearly always close to actual unemployment.

[9]Barnichon and Figura (2010) and Veracierto (2011) also take into account flows between unemployment and not in the labor force in computing steady-state unemployment, which I ignore. Acknowledging that out of the labor force is a distinct labor market state does not change my ultimate conclusion that steady-state unemployment and vacancies will appear negatively related, although it may affect the shape of the curve relating the two series and how much we should conclude it may have shifted over time. I return to these issues later.

[10]Formally, as evident in equation 6, the Beveridge curve only depends on the ratio s/A. The levels of s and A depend on the frequency used to measure flows between labor market states.

[11]In particular, the probability of profitably hiring a worker in the Shimer (2007) model will not equal $m(u,v)/v$. Instead, it corresponds to the equilibrium fraction of locations with more workers than jobs. Employers in locations with more jobs than workers may still hire, but will earn zero profits. Although the probability of a profitable hire differs from $m(u,v)/v$, this probability will still be negatively related to v in equilibrium.

[12]Nash bargaining is one rule on how to divide a given amount of resources between two parties. This particular rule for how to divide resources was proposed by Nash (1950), who showed this rule had various desirable properties. Since employers and workers must divide the surplus that results from their joint production, Nash's solution has often been applied to determine the wage that workers receive.

[13]Kocherlakota (2011) shows that under Nash bargaining, J rises by nearly as much as A falls, so labor market tightness v/u is essentially the same regardless of A. In figure 4, keeping v/u unchanged but shifting the Beveridge curve up to the value associated with A_1 would

imply an unemployment rate of no more than 6 percent. But as hinted at in note 5, Nash bargaining is a somewhat problematic assumption, since for standard parameterizations it implies that productivity shocks produce fluctuations in J that are too small to explain business cycle volatility.

[14]In informal communication, Rob Shimer computed the effects of a 16 percent 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 the unemployment rate would rise from 5 percent to 5.8 percent. This suggests my bound may be a substantial overestimate of the true effect.

[15]See, for example, Aaronson, Mazumder, and Schechter (2010); Valletta and Kuang (2010); Fujita (2011); Mazumder (2011); and Hu and Schechter (2011).

[16]More precisely, lower effort should be viewed as a change in some unobserved determinant of hiring that results in lower hiring rates for the same number of vacancies while holding unemployment fixed. This change may reflect lower effort—for example, firms may spend fewer resources on advertising a position or on screening and interviewing potential candidates. But alternatively, recruiters may raise the standards they expect from workers, which would also lower vacancy yields without representing lower effort on the part of the firms.

[17]Similarly, cyclical changes in the composition of which firms are trying to hire may lead us to incorrectly infer a change in the ability of the typical firm to hire. Suppose there was a rise in the share of hiring by firms that tend to rely more heavily on posting vacancies. In this case, measured vacancies would appear to rise more than overall hiring. Barnichon et al. (2010) provide evidence that the shift in the Beveridge curve coincided with a change in composition across industries toward industries that rely more heavily on posting vacancies to hire workers.

[18]Davis, Faberman, and Haltiwanger (2010) note that it may be difficult to ensure that both recruiting intensity varies over time and market tightness is determined by a free-entry condition such as equation 4. For example, they cite a model from Pissarides (2000, chapter 4) in which there is free entry into the labor market but recruiting intensity is constant over time. However, it is possible to get time-varying recruiting intensity in a model with free entry if we impose that both effective vacancies and the cost of creating effective vacancies to be homogeneous functions of q and v of the same degree. For example, since the product qv in equation 9 is homogeneous of degree 2 in q and v, the cost function for recruiting effort and posting vacancies must also be homogeneous of degree 2 in q and v.

REFERENCES

Aaronson, Daniel, Bhashkar Mazumder, and Shani Schechter, 2010, "What is behind the rise in long-term unemployment?," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 34, Second Quarter, pp. 28–51.

Albrecht, James, Pieter Gautier, and Susan Vroman, 2003, "Matching with multiple applications," *Economics Letters*, Vol. 78, No. 1, January, pp. 67–70.

Barnichon, Regis, Michael Elsby, Bart Hobijn, and Ayşegül Şahin, 2010, "Which industries are shifting the Beveridge curve?," Federal Reserve Bank of San Francisco, working paper, No. 2010-32, December.

Barnichon, Regis, and Andrew Figura, 2010, "What drives movements in the unemployment rate? A decomposition of the Beveridge curve," Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System, working paper, No. 2010-48.

Burdett, Kenneth, Shouyong Shi, and Randall Wright, 2001, "Pricing and matching with frictions," *Journal of Political Economy*, Vol. 109, No. 5, October, pp. 1060–1085.

Butters, Gerard R., 1977, "Equilibrium distribution of sales and advertising prices," *Review of Economic Studies*, Vol. 44, No. 3, October, pp. 465–491.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, 2010, "The establishment-level behavior of vacancies and hiring," National Bureau of Economic Research, working paper, No. 16265, August.

Decreuse, Bruno, 2010, "Matching with phantoms," GREQAM (Groupement de Recherche en Economie Quantitative d'Aix-Marseille), working paper, No. 2010-07, March.

Elsby, Michael, Bart Hobijn, and Ayşegül Şahin, 2010, "Disciplined estimates of gross flows among labor market states," presentation at Federal Reserve Bank of Chicago, September 10, available at www.econ.wisc.edu/workshop/Michigan.pdf.

Felipe, Jesus, and Franklin M. Fisher, 2003, "Aggregation in production functions: What applied economists should know," *Metroeconomica*, Vol. 54, No. 2–3, pp. 208–262.

Fujita, Shigeru, 2011, "Effects of extended unemployment insurance benefits: Evidence from the monthly CPS," Federal Reserve Bank of Philadelphia, working paper, No. 10-35/R, January.

Fujita, Shigeru, and Garey Ramey, 2009, "The cyclicality of separation and job finding rates," *International Economic Review*, Vol. 50, No. 2, May, pp. 415–430.

Hall, Robert, 2005, "Job loss, job finding, and unemployment in the U.S. economy over the past fifty years," *NBER Macroeconomics Annual*, Vol. 20, pp. 101–137.

Hu, Luojia, and Shani Schechter, 2011, "How much of the decline in unemployment is due to the exhaustion of unemployment benefits?," *Chicago Fed Letter*, Federal Reserve Bank of Chicago, No. 288, July.

Kocherlakota, Narayana, 2011, "Notes on 'Labor markets and monetary policy,'" supplementary material for speech at St. Cloud State University, March 3, available at www.minneapolisfed.org/news_events/pres/kocherlakota_notes_March3_2011.pdf.

__________, 2010, "Inside the FOMC," speech in Marquette, MI, August 17.

Lagos, Ricardo, 2000, "An alternative approach to search frictions," *Journal of Political Economy*, Vol. 108, No. 5, October, pp. 851–873.

Mazumder, Bhashkar, 2011, "How did unemployment insurance extensions affect the unemployment rate in 2008–10?," *Chicago Fed Letter*, Federal Reserve Bank of Chicago, No. 285, April.

__________, 2007, "New evidence on labor market dynamics over the business cycle," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 31, First Quarter, pp. 36–46.

Mortensen, Dale T., and Christopher A. Pissarides, 1994, "Job creation and job destruction in the theory of unemployment," *Review of Economic Studies*, Vol. 61, No. 3, July, pp. 397–415.

Nash, John F., Jr., 1950, "The bargaining problem," *Econometrica*, Vol. 18, No. 2, April, pp. 155–162.

Petrongolo, Barbara, and Christopher A. Pissarides, 2001, "Looking into the black box: A survey of the matching function," *Journal of Economic Literature*, Vol. 39, No. 2, June, pp. 390–431.

Pissarides, Christopher A., 2000, *Equilibrium Unemployment Theory*, 2nd ed., Cambridge, MA: MIT Press.

__________, 1985, "Short-run equilibrium dynamics of unemployment vacancies, and real wages," *American Economic Review*, Vol. 75, No. 4, September, pp. 676–690.

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

__________, 2005a, "The cyclicality of hires, separations, and job-to-job transitions," *Review*, Federal Reserve Bank of St. Louis, Vol. 87, No. 4, July/August, pp. 493–507.

__________, 2005b, "The cyclical behavior of equilibrium unemployment and vacancies," *American Economic Review*, Vol. 95, No. 1, March, pp. 25–49.

Stevens, Margaret, 2007, "New microfoundations for the aggregate matching function," *International Economic Review*, Vol. 48, No. 3, pp. 847–868.

Valletta, Rob, and Katherine Kuang, 2010, "Extended unemployment and UI benefits," *FRBSF Economic Letter*, Federal Reserve Bank of San Francisco, No. 2010-12, April 19.

Veracierto, Marcelo, 2011, "Worker flows and matching efficiency," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 35, Fourth Quarter, pp. 147–169.