KATHARINE G. ABRAHAM
University of Maryland

JOHN C. HALTIWANGER
University of Maryland

LEA E. RENDELL
University of Maryland

How Tight Is the US Labor Market?

ABSTRACT We construct a generalized measure of labor market tightness based on the ratio of vacancies to effective searchers. Our generalized measure exhibits substantially less volatility than the standard measure defined as the ratio of vacancies to unemployment. Effective searchers include not only the unemployed but also those who are out of the labor force and the employed. These groups account for a substantial share of hires and their presence mutes the effects of the pronounced countercyclical movements in unemployment. The effective searcher measure also distinguishes different groups among the unemployed. During protracted contractions, the distribution of unemployment shifts toward the long-term unemployed, a group with lower relative search intensities, contributing to the smaller proportional increase in effective searchers as compared to the simple unemployment count. The Beveridge curve constructed using effective searchers is much more stable than the standard Beveridge curve. Further, the matching function for hires based on our generalized measure outperforms the matching function based on the ratio of vacancies to unemployment. Our approach thus reduces the unexplained residual variation required in the matching function to be consistent with real world data.

Conflict of Interest Disclosure: Katharine G. Abraham is professor of economics and survey methodology at the University of Maryland; John C. Haltiwanger is the Dudley and Louisa Dillard Professor of Economics and Distinguished University Professor of Economics at the University of Maryland; Lea E. Rendell is a graduate student in economics at the University of Maryland. Beyond these affiliations, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation. The views expressed in this paper are those of the authors and do not necessarily reflect those of the University of Maryland.

For a simple summary of labor market conditions, observers and analysts long have turned to the unemployment rate. Unemployment exhibits clearly cyclical behavior, rising during downturns and falling during recoveries. By this metric, the labor market was tighter at the end of 2019 than it had been for half a century. Search and matching models of the labor market (Diamond 1982; Blanchard and Diamond 1989, 1992; Mortensen and Pissarides 1994; Pissarides 2000) imply that unemployment (or more generally job searchers) should be considered together with job openings in assessing labor market tightness. In these models, a higher ratio of job openings to unemployment makes it harder to recruit workers and easier to find jobs, thus indicating that the labor market is tighter. Official US statistics on job openings are available only since 2000, but the ratio of vacancies to unemployment at the end of 2019 was substantially higher than at any point since they began to be collected.

While both the unemployment rate and the ratio of vacancies to unemployment imply that the labor market was extraordinarily tight as of the end of 2019, there is reason to question whether these standard measures in fact capture the true degree of labor market tightness. This paper will focus on the measure at the heart of conventional search and matching models—the ratio of vacancies (V) to unemployment (U). In the search and matching framework, both the job-filling rate and the job-finding rate should vary systematically with V/U. In fact, however, these rates often deviate substantially from the rates implied by calibrated matching functions with V/U as the driving variable. This was especially true during and following the Great Recession, a finding often explained by an appeal to unobserved fluctuations in matching efficiency (Elsby, Michaels, and Ratner 2015). One of this paper's central goals is to explore whether and to what extent the use of broader measures of effective vacancies and effective searchers in the matching function can improve its fit.

One reason V/U might perform poorly in tracking job-filling and job-finding rates is that the number of unemployed people may be a poor proxy for the availability of workers to fill vacant jobs. For one thing, different groups among the unemployed may be more or less attached to the labor market and more or less likely to move into employment. In a seminal paper published in *Brookings Papers on Economic Activity*, Perry (1970) noted that women's rising labor force participation and the entry of the baby boom generation into the labor force could have raised measured unemployment independently of underlying labor market conditions. In recent years, researchers have argued that the higher-than-usual share of long-term unemployment among the unemployed following the Great

Recession implied effective unemployment lower than suggested by the unemployment rate (Krueger, Cramer, and Cho 2014). Further, as noted by Blanchard and Diamond (1989) and emphasized in a number of recent empirical studies, most new hires originate from out of the labor force or from another job (job-to-job flows) rather than from unemployment. The standard model ignores these job seekers altogether. Another potentially important factor is that the search intensity of potential workers in any given category may vary over time (Krueger and Mueller 2010, 2011; Davis 2011; Hall and Schulhofer-Wohl 2018).

Simply looking at job vacancies may be problematic as well. Existing evidence implies that the intensity with which firms recruit to fill their vacancies varies over time. Along with doing a better job of measuring the number of effective searchers, accounting for this variation also could help to account for observed variation in both job-filling and job-finding rates (Davis, Faberman, and Haltiwanger 2013; Gavazza, Mongey, and Violante 2018; Mongey and Violante 2020).

Building on the search and matching literature, including several prior studies that have made use of augmented measures of either effective searchers or effective vacancies, we propose a generalized measure of labor market tightness that addresses some important limitations of the standard measures. We begin our analysis by broadening the concept of effective job searchers to account not only for job candidates drawn from among the unemployed but also from among those currently out of the labor force or already working, as well as for the heterogeneity within each of these three groups of searchers. In defining the groups used for our analysis, we disaggregate the population relatively finely by labor market status in order to capture as much of the heterogeneity in job search behavior as possible. We construct a measure of effective searchers by taking a weighted sum across twenty-two different groups within the population age sixteen and older, with the weights based on the relative base period job-finding rates for each of the different groups as a proxy for relative job search intensities.

This generalized measure of effective searchers, expressed as a share of the population, exhibits much less volatility than the unemployment rate constructed on the same basis. One reason is that searchers who are out of the labor force or employed mute the effects of the pronounced countercyclical movements in unemployment on the broader generalized measure. Second, composition effects matter. In particular, during a deep and extended contraction such as the Great Recession, the composition of unemployment shifts toward the long-term unemployed, a group with lower relative job-finding rates. This further dampens the increase in the number of effective searchers in the broader measure during economic downturns.

One of the puzzles of recent labor market history has been the pronounced and persistent outward shift in the Beveridge curve relating job vacancies and unemployment following the Great Recession. We find that over the period from 1994 through 2019 the Beveridge curve constructed using vacancies and effective searchers is much more stable than the curve constructed using vacancies and unemployment. Further, using our generalized measure of effective searchers rather than unemployment in the measure of labor market tightness reduces the unexplained residual variation in the job-filling rate (hires per vacancy) and in the job-finding rate of unemployed workers (hires from unemployment per unemployed person). Over the 1994–2019 period, the residual unexplained variation in the job-filling rate calibrated using our generalized measure of effective searchers is only about half as large as that based on the standard measure. For the job-finding rate among the unemployed, the residual unexplained variation of the calibrated series using our generalized effective searchers measure is about a third lower than that based on the standard measure.

Our baseline analysis is conservative in that it neglects some of the important factors that a generalized measure of labor market tightness ideally would take into account. First, the relative intensity of search on the part of job seekers in any given group may change over time. Second, the intensity with which employers recruit to fill their vacant jobs also may change over time. We extend our baseline analysis with an exploratory investigation of time-varying relative search and recruiting intensities. Given the measurement challenges associated with this exercise, we view the results as suggestive rather than conclusive. Even so, our efforts further reduce the unexplained residual variation in job-filling rates. The evidence regarding the job-finding rates of the unemployed is mixed. Our accounting for time-varying recruiting intensity reduces the unexplained residual variation in job-finding rates, but our accounting for time-varying search intensity actually increases it. Better measures of both search intensity and recruiting intensity could well produce better results.

I. A Broader Perspective on Labor Market Tightness

A matching function that relates hires to unemployment and vacancies is the cornerstone of modern macro models of the labor market. Especially in recent years, however, it has been necessary to posit substantial fluctuations in matching efficiency to account for observed fluctuations in job finding and job filling. Fluctuations in matching efficiency similarly have been offered as an explanation for shifts in the Beveridge curve that describes the inverse relationship between vacancies and unemployment (Elsby, Michaels, and Ratner 2015).

An alternative approach that we build on here is to consider whether using more suitable arguments in the matching function can eliminate the apparent instability of the standard model. Rather than equating effective searchers with the unemployed, this broader perspective on labor market tightness recognizes that the people available to fill open jobs also include those out of the labor force or in another job. Constructing an aggregate measure of effective searchers also requires a way to measure the search intensity of those in the effective searcher pool. Relatedly, a broader perspective should recognize that labor market tightness depends not only on the number of posted job openings but also on how hard employers are trying to fill those jobs.

I.A. The Pool of Effective Searchers

A lengthy literature has examined how changes in the composition of the unemployed may affect the interpretation of the official unemployment rate. One strand of the literature, launched by the seminal work of Perry (1970) and further developed by Shimer (2001), Aaronson and others (2015), and Barnichon and Mesters (2018), among others, focuses on the demographic composition of the unemployed. Another strand focuses on the relative numbers of long-term and short-term unemployed (Kaitz 1970; Krueger, Cramer, and Cho 2014). Whether because of lower search intensity, loss of human capital, or employer unwillingness to hire them (Abraham and others 2019), the long-term unemployed have lower job-finding rates and may contribute proportionately less than the short-term unemployed to the pool of effective searchers. How a person entered unemployment also may be important. As an example, the job-finding pattern among those laid off from a job differs considerably from the pattern for other groups among the unemployed (Katz 1986; Katz and Meyer 1990; Fujita and Moscarini 2017).

A comprehensive measure of effective job searchers also needs to account for people who are outside of the labor force. Those in this group have a much lower job-finding rate than the average unemployed person, but there are many more of them. In a typical month, the number of people who enter employment directly from out of the labor force is much larger than the number entering directly from unemployment

(Hornstein, Kudlyak, and Lange 2014). People who are out of the labor force but say they want a job are much more likely to enter employment than the rest of the out-of-the-labor-force population (Blanchard and Diamond 1989; Jones and Riddell 1999; Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017; Hall and Schulhofer-Wohl 2018).

The employed are a final group of searchers (Sedlacek 2016; Hall and Schulhofer-Wohl 2018). In the canonical search-and-matching model, vacancies include the job openings created by departing employees. Symmetrically, the measurement of effective searchers should take into account on-the-job searchers who may fill those jobs. Available survey data suggest that on-the-job search is prevalent (Black 1980; Blau and Robins 1990; Faberman and others 2017) and administrative data show that a large share of hires are people moving from one job to another (Haltiwanger, Hyatt, and McEntarfer 2018; Haltiwanger and others 2018).

I.B. Job Search Intensity

In addition to properly identifying the effective searcher population, a full accounting of effective search activity also needs to incorporate search intensity. Approaches used in the literature to do this include directly measuring search activities, making use of information on the gap between individuals' desired and actual hours, and the approach that we adopt—inferring relative search intensity from relative job-finding rates.

In an early effort to measure job search intensity directly, Shimer (2004) uses information from the Current Population Survey (CPS) on the number of different search methods reported by the unemployed as a search intensity proxy. Several studies, including DeLoach and Kurt (2013), Gomme and Lkhagvasuren (2015), and Mukoyama, Patterson, and Şahin (2018), measure search intensity among the unemployed using data from the American Time Use Survey (ATUS) on time devoted to job search. ATUS data are available only beginning in 2003, but Mukoyama, Patterson, and Şahin (2018) combine them with CPS data on job search methods to construct a longer search intensity series. They model the relationship of search time to the search methods reported by the unemployed, then use that estimated relationship to construct a longer search intensity series for the unemployed. Studies using ATUS data reach conflicting conclusions about whether search intensity among the unemployed is procyclical or countercyclical. Ahn and Shao (2017) use ATUS data to study the cyclicality of job search among the employed. Because the ATUS does not ask what respondents are doing while they are at work, ATUS measures of job search among the employed seem especially likely to miss at least some job search activity.

Faberman and others (2020) use data on the gap between desired and actual hours to proxy for job search intensity. They show that, in 2013–2015 data from the Survey of Consumer Expectations (SCE), this gap is correlated with a direct measure of search intensity. They use the SCE data to calculate the average difference between desired and actual hours for each of thirty-nine groups defined based on labor force status and demographic characteristics. Treating the gaps for defined groups as constant over time, the authors use CPS data to construct a time series of aggregate slack, which they define as the total gap between desired and actual hours divided by total desired hours.

Finally, job-finding rates have been used to proxy for search intensity. The simplest version of this approach uses group-specific job-finding rates in a base period to weight the number of people in each group to produce an aggregate measure of effective searchers. A notable example of this approach is the Federal Reserve Bank of Richmond's Hornstein-Kudlyak-Lange Non-Employment Index (NEI), which uses long-run average job-finding rates to aggregate its nine groups of effective searchers among the unemployed and those out of the labor force (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017).

Several studies, including Veracierto (2011), Hornstein and Kudlyak (2016), and Sedlacek (2016), have sought to infer the variation in within-group search intensities from changes in relative job-finding rates. Among the studies adopting this general approach, Hall and Schulhofer-Wohl (2018) offer the most comprehensive characterization of the job searcher pool, considering sixteen groups of job seekers—thirteen groups among the unemployed and two among those out of the labor force plus the employed.

The basic strategy in all of these studies is, in effect, to infer what is happening to group-specific search intensities based on how having more or fewer people in any given group affects the number of matches. If adding people to a group makes a larger than expected contribution to the number of matches realized when the labor market is tight, for example, procyclicality in search intensity is a plausible explanation.[1] A limitation

1. Alternatively, the cross-group differences that are the basis for the suggested inference about job search intensity could be attributable to differences in the pattern of the shocks experienced by different groups of searchers. This is a less parsimonious explanation, and it is not entirely apparent what the source of such shocks might be, though it cannot be ruled out.

is that cyclical variation in search intensity that is common across groups cannot be distinguished empirically from the elasticity of matching with respect to the (properly measured) ratio of vacancies to searchers in the standard matching function or common changes in matching efficiency. A modeler can hope to quantify changes in aggregate search intensity that result from changes in the relative sizes of groups with relatively procyclical or relatively countercyclical job-finding rates. The potentially more important changes in search intensity that are common across groups, however, cannot be quantified using this approach.

I.C. Time-Varying Employer Recruiting Intensity

A final factor missing from the standard search-and-matching model is employer recruiting intensity. Empirical implementations of the standard model use data on the number of job openings. The intensity with which employers recruit to fill their openings can vary considerably, however, depending both on the company's own circumstances and on aggregate labor market conditions.

The most literal interpretation of recruiting intensity is the time and effort devoted to advertising the firm's job openings, processing applications, and so on, but other aspects of firms' recruiting behavior may be even more important. As an example, in a tighter labor market, employers may choose to consider job candidates with criminal records (Casselman 2018; Smialek 2019), lower the levels of education and experience they require, offer better working conditions, or raise wages. All of these can be viewed as changes in recruiting intensity, in the sense that employers are trying harder to fill their vacant jobs.

Modestino, Shoag, and Balance (forthcoming) show that, controlling for occupation, the shares of online job advertisements stating a requirement for a college degree or four or more years of experience rose during the Great Recession. The changes were larger in states and occupations that experienced a larger increase in the supply of workers. Using establishment-level Job Openings and Labor Turnover Survey (JOLTS) data, Davis, Faberman, and Haltiwanger (2013) show that, holding aggregate conditions constant, employers with a larger gross hiring rate fill their jobs at a faster pace. They interpret this finding through the lens of recruiting intensity—that is, they infer that recruiting intensity is positively associated with the gross hiring rate. Their findings suggest that the decline in hiring rates during the Great Recession should have led to a corresponding decline in recruiting intensity. Cross-sectional heterogeneity in recruiting intensity also may imply changes over time in overall recruiting intensity attributable to composition effects (Davis, Faberman, and Haltiwanger 2012b). Building on the results of Davis, Faberman, and Haltiwanger (2013), Gavazza, Mongey, and Violante (2018) develop an equilibrium model of the recruiting intensity response to negative shocks. Their model incorporates both the gross hiring rate channel and a second channel in which, in a weaker labor market, firms of all types exert less effort to fill their jobs. Based on a calibration exercise, they conclude that the latter effect is more important than the former. Mongey and Violante (2020) argue similarly that the effect of labor market slackness on firms' chosen recruiting intensity is a key driver of the residual variation in match rates during the Great Recession.

II. An Organizing Framework

In the standard specification of the canonical search-and-matching model (Diamond 1982; Blanchard and Diamond 1989, 1992; Mortensen and Pissarides 1994; Pissarides 2000), employers create job openings they would like to fill (V), and unemployed individuals (U) search among these job openings for employment. The process of matching unemployed workers to vacant jobs is represented by a production function, often assumed to be Cobb-Douglas in form, with vacancies and unemployment as the inputs and matches (hires) as the output:

$$H_t = m(V_t, U_t) = \mu_t V_t^{1-\alpha} U_t^{\alpha}, \tag{1}$$

where H is hires, V is the number of job openings, U is the number of unemployed people, t is the time period, μ_t is a time-varying match efficiency parameter, and $\alpha(1 - \alpha)$ is the elasticity of the matching function with respect to unemployment (vacancies). In this framework, labor market tightness (θ_t) is expressed as:

$$\theta_t = \frac{V_t}{U_t}. \tag{2}$$

This relationship may be viewed through the lens of the job-finding rate, expressed as hires per unemployed worker:

$$\frac{H_t}{U_t} = \mu_t \left(\frac{V_t}{U_t}\right)^{1-\alpha} = \mu_t (\theta_t)^{1-\alpha}. \tag{3}$$

All else the same, when the labor market is tighter (when θ_t is larger), an unemployed individual is more likely to find a job. An alternative but equivalent approach is to view this relationship through the lens of the job-filling rate, expressed as hires per vacant job:

$$\text{(4)} \qquad \frac{H_t}{V_t} = \mu_t \left(\frac{U_t}{V_t} \right)^{\alpha} = \mu_t \left(\frac{1}{\theta_t} \right)^{\alpha}.$$

All else the same, when the labor market is tighter, an employer is less likely to be able to recruit an unemployed person to fill a vacant job.

Since the matching function as written in equation (1) has constant returns to scale, it can be expressed as a relationship among the hiring rate h, vacancy rate v, and unemployment rate u:

$$\text{(5)} \qquad h_t = m(v_t, u_t) = \mu_t v_t^{1-\alpha} u_t^{\alpha},$$

where $h = H/\mathrm{E}$, $v = V/\mathrm{E}$, and $u = U/\mathrm{E}$ and E is employment. An additional constraint is that, in steady state, the number of separations (inflows to unemployment) must equal the number of hires (outflows from unemployment). This steady state relationship can be expressed:

$$\text{(6)} \qquad \delta_t = h_t = m(v_t, u_t) = \mu_t v_t^{1-\alpha} u_t^{\alpha},$$

where δ is the separation rate (in this case separations from employment into unemployment expressed as a fraction of employment) and the other terms are as previously defined.[2] The downward sloping relationship between the unemployment rate and the vacancy rate implied by equation (6) commonly is termed the Beveridge curve.

In this framework, shifts in either δ or μ will shift the position of the Beveridge curve. Improvement in the matching function (an increase in μ), for example, shifts the Beveridge curve inward (lowering unemployment for given vacancies), while deterioration in the matching function shifts the Beveridge curve outward (raising unemployment for given vacancies). Shifts in μ also will affect the job-filling rate and the job-finding rate.

2. Nothing fundamental changes if this expression is modified to allow for steady state growth at rate g in desired employment, in which case the left-hand side becomes $\delta + g$. Although the standard Beveridge curve specification expresses vacancies and unemployment relative to employment, as in equation (6), for comparison with the generalized Beveridge curve examined later, our empirical implementation works with vacancies and unemployment relative to population.

In the model as just sketched out, $\theta_t = \frac{V_t}{U_t}$, but unemployment and vacancies are imperfect proxies for the measures of effective searchers and effective vacancies that we believe should be the objects of interest. Empirically, the standard framework requires substantial variation in μ_t to account for observed shifts in the Beveridge curve, as well as to explain the considerable residual variation between actual job-filling and job-finding rates and those implied by the model. One reason for this may be that, by ignoring heterogeneity among the unemployed and job search among those who are out of the labor force or employed, the standard specification misrepresents the stock of potential job candidates. If effective job seekers of each of the different types moved together with unemployment over time, it would not be important to account for them separately (Broersma and van Ours 1999; Sedlacek 2016), but this cannot be the case as increases in any one group imply decreases in others. In addition, the standard tightness measure does not account either for temporal variation in search intensity or for temporal variation in employer recruiting intensity.

We can elaborate the simple model to account for these complexities. Building on the standard hiring function, we can write:

$$H_t = m\left(\sum_j \rho_t^{v_j} V_{jt}, \sum_i \rho_t^{s_i} S_{it}\right) = \mu_t \left(\sum_j \rho_t^{v_j} V_{jt}\right)^{1-\alpha} \left(\sum_i \rho_t^{s_i} S_{it}\right)^{\alpha}, \tag{7}$$

where V_j represents the number of job openings from firms of type j, S_i represents the number of job searchers of type i, $\rho_t^{v_j}$ represents the intensity of employer recruiting effort for firms of type j at time t, and $\rho_t^{S_i}$ represents the intensity of job search on the part of searchers of type i at time t.[3] Note that the elasticity of the matching function specified in equation (7) as well as the unexplained residual variation in the actual number of matches as compared to the number implied by the model are likely to differ from those associated with the standard model specified in equation (1).

In this expanded framework, labor market tightness can be written as:

$$\tilde{\theta}_t = \frac{\sum_j \rho_t^{v_j} V_{jt}}{\sum_i \rho_t^{s_i} S_{it}}. \tag{8}$$

3. This formulation builds on specifications of generalized matching functions in Davis (2011), Davis, Faberman, and Haltiwanger (2013), and Abraham (2015).

We refer to the numerator of this expression as effective vacancies (EV_t) and the denominator as effective searchers (ES_t). The steady state equilibrium of hires equal to separations is now:

$$(9) \qquad \delta_t = h_t = m\left(\sum_j \rho_t^{v_j} v_{jt}, \sum_i \rho_t^{S_i} s_{it}\right) = \mu_t \left(\sum_j \rho_t^{v_j} v_{jt}\right)^{1-\alpha} \left(\sum_i \rho_t^{S_i} s_{it}\right)^{\alpha},$$

where separations are now all separations from employment and variables are rates expressed as fractions of the population. Over the course of a business cycle, absent changes in matching efficiency or other factors that shift the position of the generalized Beveridge curve, effective searchers and effective vacancies will move inversely as implied by equation (9).[4]

In this generalized setting, the job-filling rate is given by:

$$(10) \qquad \frac{H_t}{V_t} = \mu_t \left(\frac{1}{\tilde{\theta}_t}\right)^{\alpha} \rho_t^{v}, \text{ where } \rho_t^{v} = \frac{\sum_j \rho_t^{v_j} V_{jt}}{V_t}.$$

For the generalized model, the ratio of hires to unemployment is no longer a job-finding rate since not all hires come from among the unemployed. To characterize the job-finding rate for any subgroup or collection of subgroups, we build on the transformation of the matching function used by Hall and Schulhofer-Wohl (2018). Following Hall and Schulhofer-Wohl (2018), job-finding rates differ across groups only because of differences in their effective search intensity, meaning that they are the same on a search intensity–adjusted basis. The common search intensity–adjusted job-finding rate varies with the tightness of the labor market:

$$(11) \qquad \frac{H_t}{ES_t} = \frac{H_{it}}{ES_{it}} = \frac{H_{it}}{\rho_t^{S_i} S_{it}} = \tilde{f}_t = \tilde{f}_{it} = A_t T_t^{\eta},$$

where $T_t = V_t/H_t$ is average vacancy duration and A_t are any common time effects on job-finding rates not captured by vacancy duration. Equation (11)

4. Something we have not considered explicitly is the possibility of mismatch between vacant jobs and effective job seekers. Although commonly cited by business leaders and policy officials as an important contributor to unemployment, especially during periods when the labor market is weak (Abraham 2015), available evidence suggests that mismatch plays at most a modest role in explaining aggregate unemployment fluctuations (Şahin and others 2014; Crump and others 2019). In our framework, we will think of mismatch as captured by μ.

is a transformation of (7). To see this, define $A_t = \mu_t^{1+\eta}\ (\rho_t^v)^\eta$. Then with appropriate substitution we have:

$$H_t = \mu_t \left(\sum\nolimits_j \rho_t^{v_j} V_{jt}\right)^{\eta/(1+\eta)} \left(\sum\nolimits_i \rho_t^{S_i} S_{it}\right)^{1/(1+\eta)}, \tag{12}$$

where $\alpha = 1/(1 + \eta)$. Returning to the job-finding rate for group i, we can write:

$$\frac{H_{it}}{S_{it}} = f_{it} = \rho_t^{S_i} A_t T_t^\eta, \tag{13}$$

Equation (13) can be used to quantify the implied job-finding rate for the unemployed. Using the subgroups among the unemployed defined in the general model, we have:

$$\frac{H_{ut}}{U_t} = \frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta}\left(\rho_t^v\right)^\eta \left(V_t/H_t\right)^\eta \sum\nolimits_{i\in u}\left(\rho_t^{S_i}\, S_{it}/S_{ut}\right). \tag{14}$$

The standard model with a single group and ρ_t^v equal to 1.0 is just a special case of equation (14) given by:

$$\frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta}\left(V_t/H_t\right)^\eta. \tag{15}$$

Equations (7)–(15) lay out the aspirational general model that measures both effective vacancies and effective searchers, allowing for cross-sectional and time series variation in search intensity for searchers and recruiting intensity for vacancies. Our primary goal is to evaluate the extent to which this generalization overcomes some of the known limitations of the standard framework. We report empirical analyses that compare and contrast calibrations of the standard and general models, emphasizing the residual unexplained variation in the matching function associated with each model manifest in the Beveridge curve, job-filling rate and job-finding rate.

In our empirical analysis, we focus first on a simpler version of equations (7)–(15) that allows for cross-sectional variation in relative search intensities across groups but not for variation in those relative search intensities over time. That is, we begin with a specification in which $\rho_t^{S_i} = \gamma_i$ for all t and $\rho_t^{v_j} = 1$ for all j and t. Our goal in this portion of the analysis is to quantify how much allowing for a broader group of effective searchers can improve the performance of the matching function. Then we explore specifications that permit relative search and recruiting intensities to vary over time.

III. Creating a Measure of Labor Market Tightness Based on Effective Searchers

Our measures of effective searchers build on several earlier papers, including the research underlying the Federal Reserve Bank of Richmond's Non-Employment Index (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017) and most especially the work of Hall and Schulhofer-Wohl (2018). Similar to these other papers, we use CPS micro data to track flows across labor market states and from job to job. Our analysis makes use of job-finding rates for each of twenty-two groups—thirteen groups among the unemployed (as in Hall and Schulhofer-Wohl 2018); three among those who are out of the labor force but say they want a job and four among others who are out of the labor force (as in the Richmond Fed NEI); and two among the employed.[5] Table 1 shows the full list of twenty-two groups. Following Hall and Schulhofer-Wohl (2018), we adjust each group's job-finding rate to hold demographics constant at the group's 2005–2007 values. In addition, following Fujita, Moscarini, and Postel-Vinay (2019), we make an adjustment for a change in the procedures used to collect CPS data that otherwise would lead to an understatement in job-changing rates among employed individuals. Details of the construction of the demographically adjusted job-finding rates are provided in online appendix B. As already described, we interpret the cross-group variation in (demographically adjusted) job-finding rates as variation in search intensity.

Table 1 shows, for 2006 and 2010, the population shares of each of our twenty-two groups, together with each group's demographically adjusted job-finding rate. Those who were recently temporarily laid off have the highest rate and those not in the labor force (retired) the lowest rate. Involuntary part-time workers have a job-finding rate that is twice that of other employed people, though still relatively low compared to other identified groups. Not surprisingly, average job-finding rates fell between

5. The featured Richmond Fed index distinguishes only two groups among the unemployed (short-term and long-term) and does not include employed searchers, though there is a second version that allows for search among those working part-time for economic reasons. In their work, Hall and Schulhofer-Wohl (2018) distinguish only two groups among those out of the labor force (want a job and do not want a job) and treat the employed as a single group.

Table 1. Estimated Relative Job-Finding Rates

	Share	*JFR*	*Rel. JFR*	*Rel. JFR (raw)*
2006				
Unemployed: Recently left job	0.16	39.46	0.71	0.61
Unemployed: Recently permanently laid off	0.21	32.80	0.59	0.52
Unemployed: Recently temporarily laid off	0.23	55.22	1.00	1.00
Unemployed: Temporary job recently ended	0.12	38.63	0.70	0.60
Unemployed: Recently newly entered	0.11	21.26	0.38	0.33
Unemployed: Recently reentered	0.34	29.89	0.54	0.46
Unemployed: Left job months ago	0.15	27.86	0.50	0.43
Unemployed: Permanently laid off months ago	0.36	21.19	0.38	0.33
Unemployed: Temporarily laid off months ago	0.16	44.28	0.80	0.69
Unemployed: Temporary job ended months ago	0.13	26.03	0.47	0.40
Unemployed: Newly entered months ago	0.12	14.75	0.27	0.23
Unemployed: Reentered months ago	0.45	23.44	0.42	0.37
Unemployed: Long-term unemployed	0.43	17.41	0.32	0.27
Want Job: Discouraged	0.15	14.74	0.27	0.23
Want Job: Looked last 12 months	0.43	14.24	0.26	0.22
Want Job: Other	1.24	15.26	0.28	0.24
Not in Labor Force: In school	4.34	9.41	0.17	0.15
Not in Labor Force: Retired	15.51	1.56	0.03	0.02
Not in Labor Force: Disabled	4.67	1.96	0.04	0.03
Not in Labor Force: Other	7.44	8.87	0.16	0.14
Employed: Involuntary part-time	1.79	5.12	0.09	0.08
Employed: Not involuntary part-time	61.44	2.22	0.04	0.03
2010				
Unemployed: Recently left job	0.09	27.81	0.54	0.48
Unemployed: Recently permanently laid off	0.29	23.12	0.45	0.38
Unemployed: Recently temporarily laid off	0.28	51.80	1.00	1.00
Unemployed: Temporary job recently ended	0.13	32.88	0.63	0.56
Unemployed: Recently newly entered	0.12	12.65	0.24	0.22
Unemployed: Recently reentered	0.27	21.30	0.41	0.37
Unemployed: Left job months ago	0.16	19.29	0.37	0.32
Unemployed: Permanently laid off months ago	0.90	14.41	0.28	0.24
Unemployed: Temporarily laid off months ago	0.26	36.15	0.70	0.60
Unemployed: Temporary job ended months ago	0.24	20.06	0.39	0.33
Unemployed: Newly entered months ago	0.24	9.41	0.18	0.16
Unemployed: Reentered months ago	0.57	16.45	0.32	0.28
Unemployed: Long-term unemployed	2.14	10.92	0.21	0.18
Want Job: Discouraged	0.47	11.33	0.22	0.19
Want Job: Looked last 12 months	0.52	9.76	0.19	0.17
Want Job: Other	1.27	12.30	0.24	0.21
Not in Labor Force: In school	5.07	6.28	0.12	0.11
Not in Labor Force: Retired	15.56	1.41	0.03	0.02
Not in Labor Force: Disabled	5.17	1.42	0.03	0.02
Not in Labor Force: Other	7.26	6.76	0.13	0.12
Employed: Involuntary part-time	3.73	3.63	0.07	0.06
Employed: Not involuntary part-time	55.27	1.77	0.03	0.03

Source: Authors' calculations using CPS.

Notes: Job-finding rates (JFR) estimated using CPS survey data linking households month to month. Relative job-finding rate (Rel. JFR) calculated by dividing all job-finding rates by job-finding rate for unemployed recently laid off. Recently unemployed groups refer to those unemployed 0–4 weeks. "Unemployed months ago" refers to those unemployed 5–26 weeks. "Long-term unemployed" refers to those unemployed 27 weeks or more.

2006 and 2010 as the economy worsened following the onset of the Great Recession. Our primary interest, however, lies with the relative job-finding rates across the different groups. These are much more stable—the correlation in relative job-finding rates between 2006 and 2010 is 0.98. We construct our measure of effective searchers by weighting each of the twenty-two groups by its relative 2006 job-finding rate, shown in the third column of the top panel of table 1. All of the relative job-finding rates are defined with reference to the demographically adjusted 2006 job-finding rate of those who were recently temporarily laid off.

Table 1 also reports relative job-finding rates for 2006 and 2010 calculated directly from the CPS micro data not controlling for changing demographics, shown in the column denoted "Rel. JFR (raw)." Although our baseline analysis uses the demographically adjusted 2006 relative job-finding rate estimates, the results are very similar if we instead use the raw 2006 relative job-finding rates. Later in the paper, we present results based on demographically adjusted relative job-finding rates that are allowed to vary over time.

Figure 1 displays the standard measure of searchers (the simple unemployment count) along with several alternative measures that move toward our fully generalized effective searcher measure. The first of these considers only the unemployed as effective searchers but allows for changes in the composition of the unemployed across the thirteen groups we have specified. The second alternative measure incorporates the three groups among those out of the labor force who say they want a job, and the third adds the four groups among others who are out of the labor force. The fourth and final effective searcher measure is the fully generalized measure that adds the two groups of employed people. Each of the generalized measures weights the different groups it includes in accord with their relative 2006 job-finding rates.

The alternative series shown in figure 1 are highly correlated but distinctly different in their volatility. Allowing for heterogeneity among the unemployed yields a measure that is less cyclically volatile than the standard unemployment measure. Including, in turn, those who are out of the labor force but want a job, others who are out of the labor force, and finally the employed yields progressively less volatile measures. The first column of table 2 reports the standard deviations over the 1994–2019 period of the normalized series plotted in figure 1. Whereas unemployment as a share of the population, normalized to equal 1 in 2006, has a standard deviation of 0.34 over the 26-year period, the standard deviations of the alternative measures, also expressed as a share of the population and

Figure 1. Standard versus Generalized Measures of Searchers (Constant Relative Search Intensities)

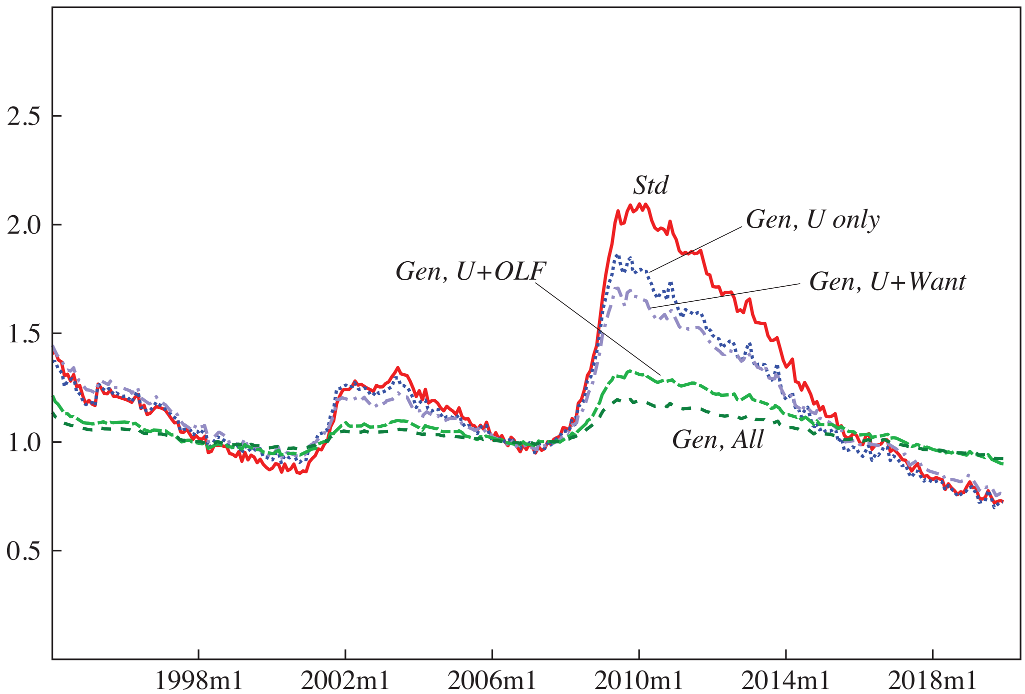

Source: Authors' calculations using Current Population Survey (CPS).

Notes: All measures ratios to population age 16+ normalized to 1 in 2006. Std = unemployed; Gen, U only = generalized measure, unemployed only; Gen, U + Want = generalized measure, unemployed plus want a job; Gen, U + OLF = generalized measure, excludes employed; Gen, All = generalized measure, all twenty-two groups.

Table 2. Cyclical Volatility of Alternative Searcher Measures: Constant Relative Job Search Intensities

	Standard deviation, 1994–2019	*March 2001*	*June 2009*	*December 2019*
Standard	0.34	0.94	2.06	0.72
General, unemployed only	0.26	1.00	1.86	0.72
General, unemployed + want job	0.22	0.99	1.71	0.76
General, unemployed + OLF	0.10	0.97	1.32	0.90
General, all	0.06	0.99	1.19	0.92
U6	0.35	0.89	2.00	0.81
Richmond Fed NEI	0.11	0.96	1.34	0.89

Source: Authors' calculations using CPS and Richmond Fed NEI.

Notes: All measures normalized to 1 in 2006. Standard = unemployed; General, unemployed only = effective searchers with unemployed only; General, unemployed + want job = effective searchers with unemployed and want a job; General, unemployed + OLF = effective searchers excluding employed; General, all = all effective searchers; U6 = U6 measure of labor underutilization.

normalized to equal 1 in 2006, are progressively lower. The standard deviation of our fully generalized measure, calculated on a basis comparable to that of the unemployment measure, is just 0.06.

Another way to look at the alternative measures is to ask how the current level of effective searchers as a share of the population compares to the level at points of time in the past. The end of the prolonged expansion that lasted through most of the 1990s and into the early 2000s offers one interesting point of comparison. The standard measure of effective searchers (unemployment) relative to the population recorded in December 2019 was about 23 percent below its March 2001 value. In contrast, the December 2019 fully generalized measure of effective searchers relative to the population was only about 6 percent below its March 2001 value. Another interesting point of comparison is with the values for June 2009, the trough of the Great Recession. Whereas the standard measure of effective searchers fell by about 65 percent between June 2009 and December 2019, the fully generalized measure fell by just 23 percent.

One important reason for the muted cyclicality of the baseline generalized measure with fixed relative job search intensities as compared to the standard measure is that the generalized measure counts more people as effective searchers. The largest proportional fluctuations in group size over a typical cycle are the fluctuations among the unemployed. In the standard measure, any proportional increase in the number of unemployed people is de facto a proportional increase in the number of effective searchers. In the baseline version of the generalized measure, in contrast, the unemployed are only a fraction of all effective searchers, and increases in unemployment thus mechanically have a smaller proportional effect on the aggregate number of effective searchers. Composition also matters. When economic conditions are weak, the share of the unemployed who are long-term unemployed rises, and the long-term unemployed have lower job-finding rates than other unemployed people. In addition, in the baseline version of the generalized measure, even though the employed have lower relative search intensities than the unemployed, their search intensities are positive. This means that, during a downturn, reductions in the number of employed effective searchers partially offset increases in the number of effective searchers associated with rising unemployment.

Figure 2 compares our generalized measure of effective searchers with two alternative measures. The first is an index based on the Bureau of Labor Statistics U6 measure of slackness. The U6 measure counts the marginally attached and involuntary part-timers along with the unemployed but like

Figure 2. Standard, U6, Richmond Fed, and Generalized Effective Searcher Measures (Constant Relative Search Intensities)

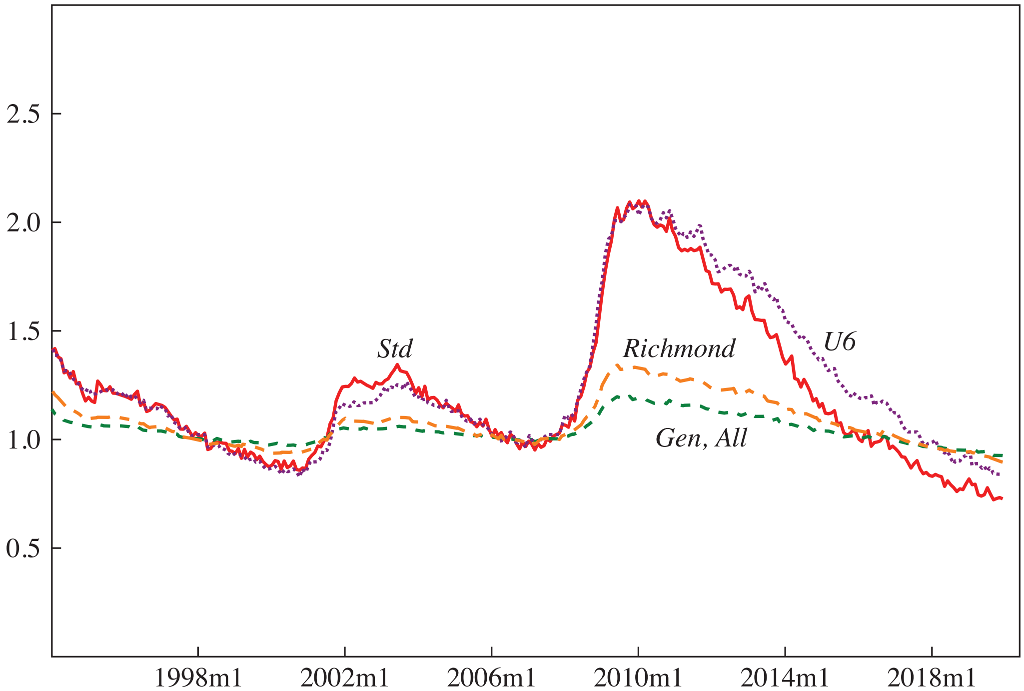

Sources: Authors' calculations using CPS and Richmond Fed NEI.
Notes: All measures ratios to population age 16+ normalized to 1 in 2006. Std = unemployed; Gen, All = generalized measure, all twenty-two groups; U6 = U6 measure of labor underutilization; Richmond = Richmond Fed NEI.

the official unemployment rate weights all of them equally. The headline Richmond Fed NEI (Hornstein, Kudlyak, and Lange 2014; Kudlyak 2017) incorporates both the unemployed and those out of the labor force but distinguishes fewer groups among the unemployed than our generalized measure (two versus thirteen) and does not incorporate employed searchers. Similar to our measure, the different groups used in the Richmond Fed index are weighted based on persistent differences in their average relative job-finding rates. For consistency, we normalize both the U6 index and the Richmond Fed index by the population age sixteen and older.

As can be seen in figure 2, the U6 index has about the same volatility as the standard measure. The cyclical variation in the Richmond Fed index is proportionally less than that of the standard measure but greater than that of our generalized index of effective searchers. Table 2 reports summary statistics for both the U6 and the Richmond Fed indexes. Consistent with the visual impression conveyed by figure 2, over the 1994–2019 period as a whole, the volatility of the U6 index is very similar to that of the unemployment index. The volatility of the Richmond Fed index lies between that of the standard measure and our generalized measure.

As already mentioned, in constructing our generalized measures of effective searchers, we have used estimated relative job-finding rates that hold demographics constant over time. We also have constructed similar measures using the simple average 2006 job-finding rates for the twenty-two groups. Comparing these measures, shown in online appendix figure A.1, to those in figure 1 suggests that controlling for demographics is relatively unimportant. As a further sensitivity analysis, we also have asked how much difference it makes that we have broken the unemployed into thirteen different groups, as opposed to distinguishing just between the short-term and the long-term unemployed. Measures constructed using the latter approach are shown in online appendix figure A.2. In addition, we have constructed an effective searcher measure that breaks the population into just five groups—short-term unemployed, long-term unemployed, out of the labor force and want a job, out of the labor force and do not want a job, and employed. The comparison between this measure and our fully generalized measure is shown in online appendix figure A.3. Over our time period, effective searcher measures constructed using these alternative approaches look broadly similar to measures constructed using our baseline approach. We note, however, that the similar behavior of these alternative searcher measures over this period does not necessarily imply they always will behave so similarly. There are marked differences in base period job-finding rates across the groups that are pooled together in the alternative measures, and, in a different period, distinguishing among them could make more of a difference. Later in the paper, we consider the performance of each of the alternative job searcher measures in explaining job-filling and job-finding rates.

We now are ready to compare labor tightness measured using effective searchers as opposed to unemployment. The numerator for all of the tightness measures shown in figure 3 is vacancies. This is the published JOLTS series from 2001:M1 to 2019:M12 and the Davis, Faberman, and Haltiwanger (2012a) backcast series for 1994:M1–2000:M12.[6] The denominators of the alternative generalized measures incorporate successively

6. We also have produced a labor market tightness measure using the Barnichon (2010) series for the pre-JOLTS period. We find results that are quite similar whether we use the Davis, Faberman, and Haltiwanger (2012a) or the Barnichon (2010) series.

Figure 3. **Standard versus Generalized Measures of Labor Market Tightness (Constant Relative Search/Intensities, Recruiting Intensity = 1)**

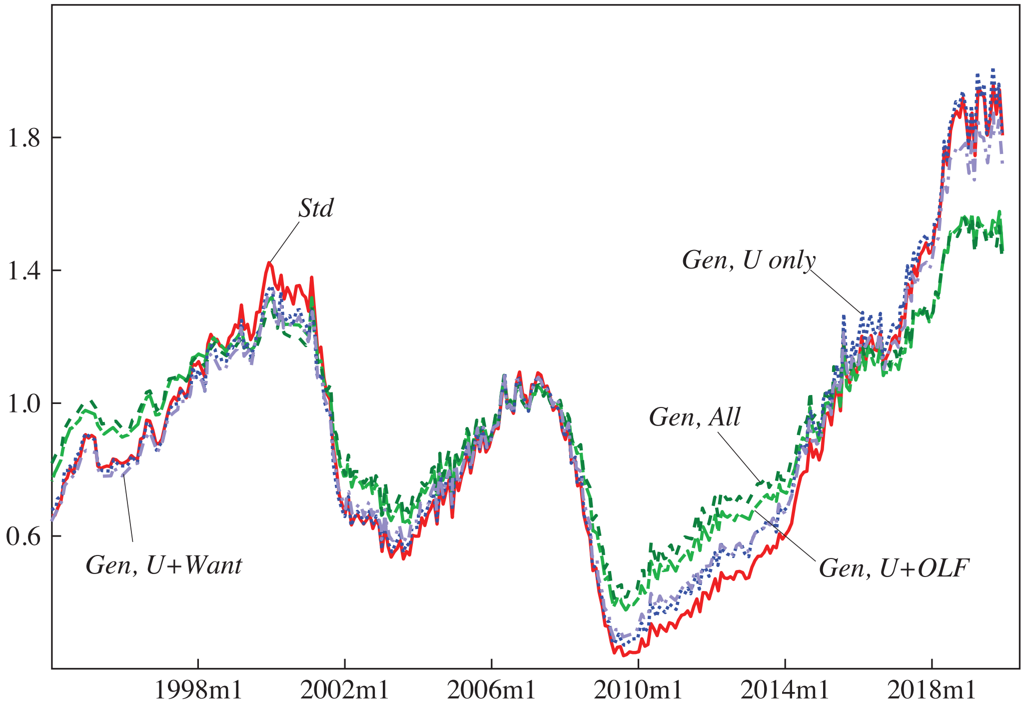

Sources: Authors' calculations using CPS, Job Openings and Labor Turnover Survey (JOLTS) for 2000:M12–2019:M12 and backcast vacancies from Davis, Faberman, and Haltiwanger (2012a) for 1994:M1–2000:M11 (DFH (2012a) vacancies hereafter).

Notes: All measures normalized to 1 in 2006. Std = V/U; Gen, U only = V/ES (unemployed only); Gen, U + Want = V/ES (unemployed plus want a job); Gen, U + OLF = V/ES (excludes employed); Gen, All = V/ES (all twenty- two groups).

more encompassing pools of effective searchers, in each case with the subgroups we have defined weighted in accord with their relative search intensities. The first generalized labor market tightness measure shown in figure 3 considers only the unemployed as effective searchers, while allowing for changes in unemployment composition. The second incorporates people who are out of the labor force but say they want a job, and the third adds the remainder of those out of the labor force. The final, fully generalized tightness measure also treats the employed as effective searchers. Once again, for ease of comparison, all of the measures in figure 3 have been normalized to equal 1 on average in 2006.

The more inclusive generalized tightness measures displayed in figure 3 are markedly less cyclical than the standard tightness measure—in particular, they fell much less steeply during the Great Recession and subsequently have risen less. The December 2019 values of the generalized measures incorporating only unemployed searchers or adding just the want a job group are not very different from the value of the standard measure, but the other measures are substantially lower, implying that the labor market was not as tight at that point as implied by the standard measure. To put this into context, the December 2019 value of the standard labor market tightness measure is about 42 percent higher than in March 2001. In contrast, the generalized measure using effective searchers is only about 17 percent higher. In short, our generalized measure using effective searchers in place of unemployment suggests a significantly different evolution of labor market tightness than the standard measure.

For comparison, online appendix figure A.4 plots labor market tightness measures constructed using the U6 index and the Richmond Fed index along with the standard and our generalized measure. The value of the tightness measure based on the U6 index is about 27 percent higher in December 2019 than it was in March 2001, and the value of the tightness measure based on the Richmond Fed index is about 18 percent higher.

IV. Beveridge Curve

A closely related but distinct way to look at the properties of the effective versus standard measures of searchers and job openings is through the lens of the Beveridge curve. Figure 4 displays the standard Beveridge curve using monthly data on vacancies and unemployment from 1994:M1 to 2019:M12. For this purpose, we use normalized unemployment and vacancy series defined relative to their 2006 average values. Plotting these series against one another makes clear their inverse relationship. In addition to the familiar downward sloping relationship between vacancies and unemployment, the figure also shows the substantial outward shift in that relationship during the long, slow recovery from the Great Recession. To illustrate, consider the period two to four years after the trough of the Great Recession (from June 2011 through May 2013), as compared to the comparable period following the trough of the 2001 recession (from November 2003 through October 2005). Job openings are only slightly lower over this portion of the recovery from the Great Recession than during the corresponding period following the 2001 recession (by about 8 percent) but unemployment is much higher (by about 48 percent). The marked increase in unemployment compared to that associated in the past with a similar level of vacancies led many to speculate that, following the Great Recession, there had been a decline in matching efficiency in the labor market.

Figure 4. Standard Beveridge Curve

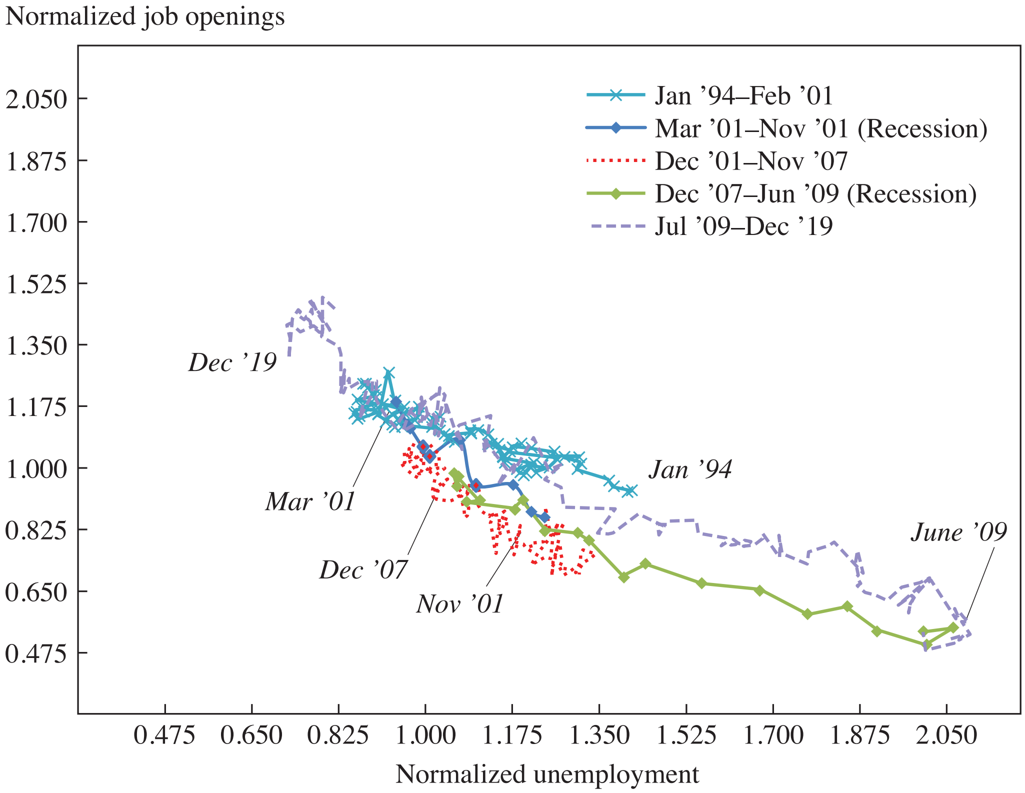

Source: Authors' calculations using job openings (vacancies) from JOLTS for 2000:M12–2019:M12 and DFH (2012a) vacancies for 1994:M1–2000:M11; unemployed from CPS.
Note: Both series as rates relative to population age 16+ normalized to 1 in 2006.

Figure 5 depicts the generalized Beveridge curve using effective searchers (using our baseline constant relative search intensities measure) with series normalized so the values plotted are all relative to their 2006 averages. The generalized Beveridge curve shown in figure 5 is much steeper than the standard version shown in figure 4. As noted previously when discussing figures 1 and 2, the proportional variation in effective searchers over time is much smaller than the proportional variation in unemployment. This translates into a normalized Beveridge curve that spans a much shorter distance along the horizontal axis than does the standard Beveridge curve and also is much more stable than the standard curve during the period following the Great Recession. Consider again the period two to four years after the trough of the Great Recession as compared to the period following the 2001 recession. As before, vacancies are slightly lower from two to four years after the Great Recession (by about

Figure 5. Beveridge Curve using Vacancies and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)

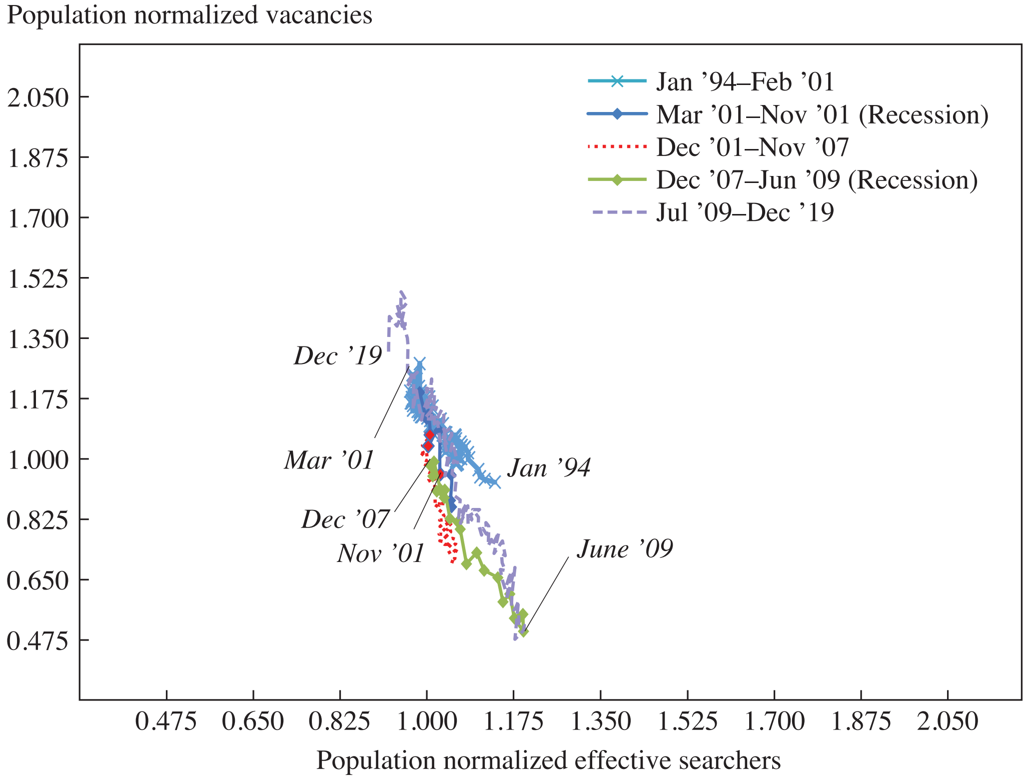

Source: Authors' calculations using job openings (vacancies) from JOLTS for 2000:M12–2019:M12 and DFH (2012a) vacancies for 1994:M1–2000:M11; effective searchers from CPS.

Notes: Both series as rates relative to population age 16+ normalized to 1 in 2006. Effective searchers as described in text.

8 percent), but whereas unemployment was *much* higher over the same period (about 48 percent), effective searchers were only slightly higher (about 10 percent).

One way to summarize the general shape of the Beveridge curve using different effective searcher measures is to fit descriptive regressions of vacancies on the various measures. Online appendix table B.3 shows that, consistent with figures 4 and 5, the standard Beveridge curve has a slope that is well below 1.0 in absolute value, whereas the Beveridge curve based on our fully generalized effective searcher measure is much steeper. The table also shows how moving in steps from the standard measure to our fully generalized measure leads to a Beveridge curve that is increasingly steep, reflecting the progressively lower volatility of the more encompassing measures. The slope of the Beveridge curve using the U6 measure is similar to that for the standard unemployment measure; the slope using the Richmond Fed index lies between that for the standard and the fully generalized measures.

V. Implications of Redefining Effective Searchers for the Matching Function

We have argued that the measure of labor market tightness using our generalized effective searcher series should be preferred conceptually to the standard measure, but we would like to have evidence that it actually does a better job of explaining the temporal variation in job-filling and job-finding rates. To evaluate the performance of the alternative measures in the matching function, we return to equation (1), the standard specification, and equation (7), the generalized specification, and ask how well each performs in tracking actual job-filling and job-finding rates. The targets we seek to match are the job-filling rate based on JOLTS data and the job-finding rate among the unemployed based on CPS gross flows data. We begin with our baseline model that defines $\rho_t^{V_i} = 1$ and $\rho_t^{S_i} = \gamma_i$. The next section of the paper will consider time-varying recruiting and relative search intensities.

V.A. *Actual versus Model-Based Patterns in the Job-Filling Rate*

Both the standard and the generalized matching function have implications for the evolution of the job-filling rate (H/V) as illustrated in equations (4) and (10). Because the left-hand sides of equations (4) and (10) are the same and are based on readily available data, we can compare the residual variation in the calibrated job-filling rates obtained using the standard and the generalized matching functions. For the present case, the specification of equation (10) is given by:

$$\frac{H_t}{V_t} = \mu_t \left(\frac{\sum_i \gamma_i S_{it}}{V_t} \right)^{\alpha}. \tag{10′}$$

In addition to our generalized measure with twenty-two different groups of effective searchers, we also consider versions of the generalized labor market tightness measure based on the U6 index and the Richmond Fed index.

In addition to the vacancy and unemployment (effective searcher) measures appearing in equations (4) and (10′), the calibrated job-filling rate associated with each of these tightness measures also depends on the elasticity of the matching function. To give each set of matching function arguments the best possible chance to fit the data well, we have estimated a separate α for each using a simplified version of the method proposed by Hall and Schulhofer-Wohl (2018). This method relates the job-finding rates for specific groups to vacancy duration from the JOLTS. As described more fully in online appendix B, this yields an estimate of the elasticity η of the job-finding rate with respect to vacancy duration for each measure of effective searchers. Given this estimate, we compute $\alpha = 1/(1 + \eta)$. For the standard model, this gives us an estimate for η of 1.04 (standard error 0.05) implying a value for α of 0.49. For the model with our effective searcher specification, the estimate is $\eta = 0.75$ (0.04) implying $\alpha = 0.57$. For the U6 measure, we obtain $\alpha = 0.48$ ($\eta = 1.1$ [0.05]) and for the Richmond Fed measure we obtain $\alpha = 0.60$ ($\eta = 0.67$ [0.05]). All of these estimates are reasonably similar and well within the middle of the range of estimates in the matching function literature (Petrongolo and Pissarides 2001). Our main results are broadly unchanged if we apply a common value of the matching function elasticity within the range of the separate estimates to calibrate the job-filling rates using the different tightness measures.[7]

Figure 6 presents the actual and calibrated job-filling rates from equations (4) and (10′) using the standard, generalized, U6, and Richmond Fed measures. Once again, all series have been normalized to average 1 in 2006.[8] The calibrated job-filling rate based on equation (10′) and our generalized tightness measure tracks the actual job-filling rate much more closely than the calibrated rate based on equation (4) and the standard tightness measure. The U6 measure performs no better than the standard measure. The Richmond Fed measure performs substantially better than the standard measure, but not as well as our generalized measure.

To quantify the improvements in performance, panel A of table 3 reports the root-mean-square error (RMSE) of the calibrated job-filling rates as compared to the actual rates based on the different tightness measures. The generalized measure produces an RMSE of 0.13, a little more than half

7. In their evaluations of job-filling and job-finding rates using alternative measures of searchers and vacancies, Davis (2011), Davis, Faberman, and Haltiwanger (2013), and Mongey and Violante (2020) use $\alpha = 0.5$.

8. The model-specific normalization for the calibrated job-filling rates implies that we are permitting the mean matching efficiency for each model to be different. More specifically, it is permitted to vary in such a way that the calibrated job-filling rate in all models is equal to the actual job-filling rate on average in 2006.

Figure 6. Actual versus Calibrated Job-Filling Rates with Standard and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)

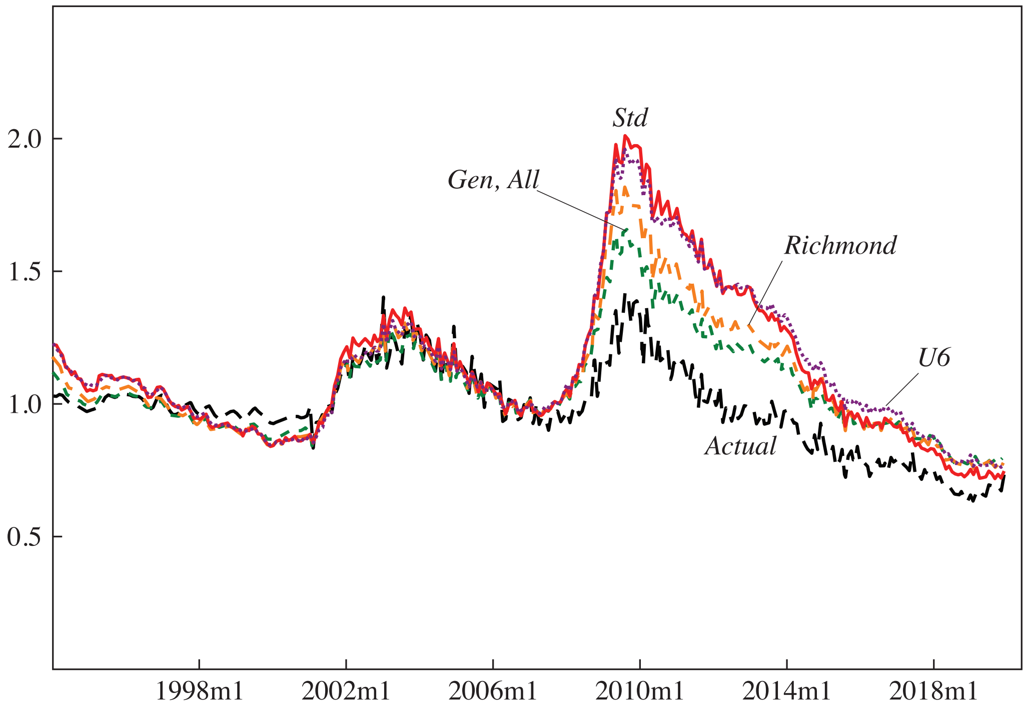

Source: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and Richmond Fed NEI.

Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups); U6 = calibrated using V/U6 index; Richmond = calibrated using V/Richmond Fed NEI.

as large as the RMSE of 0.25 produced using the standard measure. In other words, using the generalized measure with effective searchers in the calibration substantially reduces the unexplained residual variation in the job-filling rate.[9]

For comparison purposes, we also show the RMSEs in the calibrated job-filling rates for the U6 and Richmond Fed indexes. The RMSE for the U6 index is identical to that for the standard measure; the RMSE for the Richmond Fed index is intermediate between those for the standard and the generalized model. Table 3 also reports summary statistics for the alternative calibrated job-filling rates compared to the actual. The calibrated

9. Note that the calibration exercise we carry out is quite different in nature from regressing the actual job-filling rate on the calibrated values. See online appendix B for further discussion.

Table 3. Relative Performance for Job-Filling and Job-Finding Rates using Standard versus Effective Searchers: Constant Relative Search Intensities, Recruiting Intensity = 1

	Std dev	*Corr w/ actual*	*RMSE*	*Ratio RMSE to standard*
A. Job-filling rate				
Actual	0.17	1.00	NA	NA
Standard	0.29	0.75	0.25	1.00
General, all	0.19	0.81	0.13	0.54
U6	0.27	0.70	0.25	1.00
Richmond Fed NEI	0.23	0.78	0.18	0.71
B. Job-filling rate using $\alpha = 0.57$ for all measures				
Standard	0.35	0.74	0.31	1.00
Gen, U only	0.31	0.79	0.25	0.81
Gen, U + want	0.28	0.78	0.23	0.75
Gen, U + OLF	0.21	0.78	0.16	0.53
Gen, all	0.19	0.81	0.13	0.43
C. Job-finding rate for the unemployed				
Actual	0.16	1.00	NA	NA
Standard ($\alpha = 0.49$)	0.20	0.31	0.26	1.00
Standard ($\alpha = 0.57$)	0.14	0.32	0.22	0.85
General, all	0.15	0.62	0.17	0.67

Source: Authors' calculations using CPS; JOLTS; Davis, Faberman, and Haltiwanger (2012a) vacancies; Richmond Fed NEI; and Bureau of Labor Statistics Gross Flows.

Notes: Panel A shows statistics for calibrated job-filling rates using matching function elasticities specific to each. Panel B shows statistics using same matching function elasticity for all measures. Panel C shows statistics for calibrated job-finding rate for the unemployed.

job-filling rate using effective searchers has a higher correlation and a standard deviation closer to the actual job-filling rate than the alternative calibrated series.

To help with understanding the factors underlying the improvement in performance of the generalized versus the standard tightness measure in tracking the job-filling rate, figure 7 presents calibrated rates based on a series of measures that incorporate the differences between the two in stages. For these figures and the associated analysis, we use the same matching function elasticity of $\alpha = 0.57$ for all of the counterfactual measures and for the standard measure. Panel B of table 3 reports the RMSEs in the calibrated job-filling rate using each of the different versions of the generalized tightness measure shown in figure 7. On its own, allowing for heterogeneity among the unemployed reduces the RMSE of the residual unexplained variation by about 20 percent. We gain an additional 5 percent by including in the pool of effective searchers those out of the labor force who want a job, an additional 20 percent by including the remaining people who are out of the labor force, and an additional 10 percent by including

Figure 7. Actual versus Calibrated Job-Filling Rates with Standard and Effective Searchers, Using $\alpha = 0.57$ for All Measures (Constant Relative Search Intensities, Recruiting Intensity = 1)

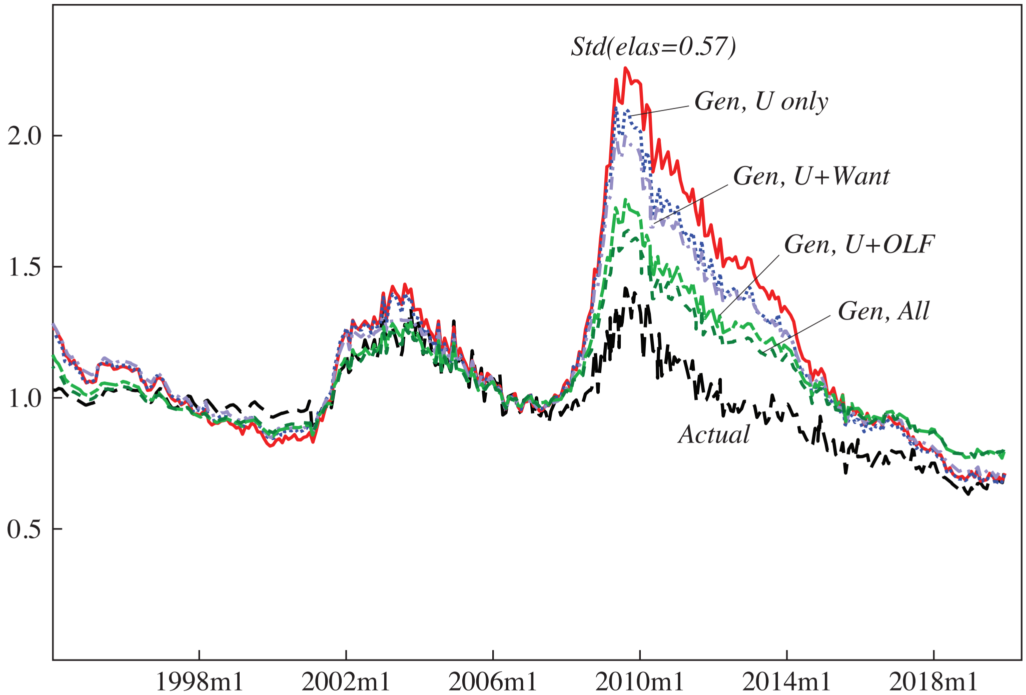

Source: Authors' calculations using CPS, JOLTS, and DFH (2012a) vacancies.

Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, U only = calibrated using V/ES (unemployed only); Gen, U + Want = calibrated using V/ES (unemployed plus want a job); Gen, U + OLF = calibrated using V/ES (excludes employed); Gen, All = calibrated V/ES (all twenty-two groups).

the employed. Appropriate caution should be used in interpreting these figures, as the numbers we have reported are not an exact decomposition, but they do provide useful guidance with respect to which features of our generalized measure account for its better performance.

We also have explored a number of sensitivity checks that we summarize briefly here; details are shown in online appendix A. First, as shown in online appendix figure A.5, we replicate the analysis of job-filling rates using simple averages of the direct (raw) relative job-finding rates from table 1 rather than the relative job-finding rates that abstract from demographics. We also replicate the findings with a generalized measure that breaks the unemployed into just two groups, the short-term and the long-term unemployed, rather than thirteen more disaggregated groups (online appendix figure A.6) and with another version that disaggregates the population into only five groups (short-term unemployed, long-term unemployed, out of the labor force and want a job, out of the labor force and do not want a job, and employed; online appendix figure A.7). In both cases, the results are broadly similar to those we have just reported. The unexplained residual variation in the job-filling rate is the same using the raw job-finding rates and slightly higher when either limiting the disaggregation of the unemployed to just two groups or using the five-group disaggregation described above (RMSE = 0.14 for both as compared to RMSE = 0.13 for the fully disaggregated specification).

In addition, we have replicated the job-filling rate analysis using the Barnichon (2010) vacancy estimates based on help-wanted advertising for the 1994:M1–2000:M12 period in place of the series based on the methodology described by Davis, Faberman, and Haltiwanger (2012a). Again, the results are broadly similar (see online appendix figure A.8), though over the 1994:M1–2000:M12 period for which we must use projected vacancies, the generalized measure using Davis, Faberman, and Haltiwanger's (2012a) methodology performs better in tracking the job-filling rate than the Barnichon (2010) series. Over that period, the RMSE using the Davis, Faberman, and Haltiwanger series is 0.04 while the RMSE using the Barnichon series is 0.09.

V.B. Actual versus Model-Based Patterns in the Job-Finding Rate

We now turn to investigating the performance of the generalized versus standard matching function for tracking job-finding rates. As noted in section II, overall hires per unemployed person is not a meaningful outcome measure for the generalized model with effective searchers. Instead, we use equations (14) and (15) to calibrate the job-finding rate of the unemployed. In the case we are currently considering, the specification of equation (14) is given by:

$$\frac{H_{ut}}{U_t} = \frac{H_{ut}}{S_{ut}} = \mu_t^{1+\eta}\left(V_t/H_t\right)^{\eta}\left(\sum_{i\in u}\gamma_i S_{it}/S_{ut}\right). \tag{14$'$}$$

Like the standard model of equation (15), the right-hand side of equation (14′) includes vacancy duration—vacancies relative to hires—but with the difference that this is now hires from all sources, not just hires from unemployment. This reflects the fact that, in the generalized model, the job-finding rate per effective searcher is assumed to be the same across all effective searcher groups. The generalized model also has an extra term that reflects the ratio of effective searchers among the unemployed to the number unemployed, with a higher ratio implying a larger number of hires per unemployed person.

Figure 8. Actual versus Calibrated Job-Finding Rates for the Unemployed with Standard and Effective Searchers (Constant Relative Search Intensities, Recruiting Intensity = 1)

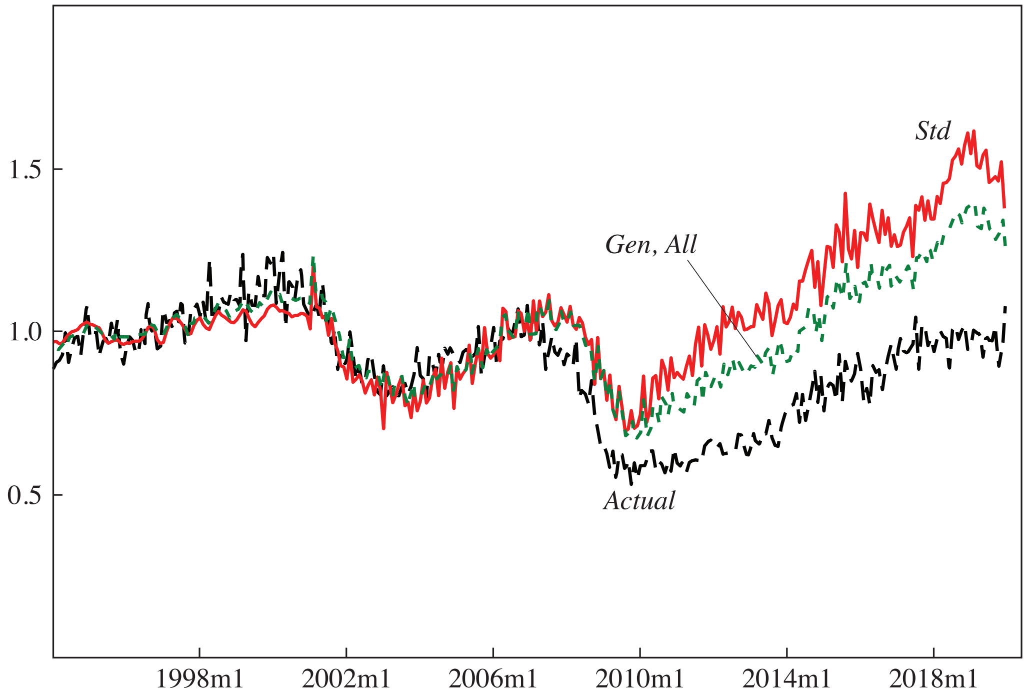

Sources: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and BLS Gross Flows.

Notes: All rates normalized to 1 in 2006. Actual = job-finding rate for unemployed from published BLS Gross Flows; Std = calibrated using V/U with $\alpha = 0.49$ ($\eta = 1.04$); Gen, All = calibrated using V/ES (all twenty-two groups).

Figure 8 shows the actual and calibrated job-finding rates for the unemployed for the generalized and standard matching functions based on using equations (14′) and (15), respectively. As with our analysis of job-filling rates, we normalize both the actual and the calibrated series to be equal to 1 in 2006.[10] The actual job-finding rate is highly procyclical and falls especially sharply in the Great Recession. Both the generalized and the standard matching function track the job-finding rate among the unemployed reasonably well during the period prior to the Great Recession, though the generalized matching function performs somewhat better over that period. The generalized and the standard matching function track the sharp decline in the job-finding rate among the unemployed during the Great Recession about equally well, but the standard model implies a faster

10. As before, this implies that we are permitting the mean matching efficiency to differ across the standard and generalized models in such a way that both yield values in 2006 with mean equal to 1.

recovery whereas the calibrated generalized model lies closer to what actually happened.

Panel C of table 3 quantifies the improvement in performance in the calibrated job-finding rate for the unemployed from using the generalized rather than the standard matching function. The RMSE of the generalized calibration of the job-finding rate from unemployment is about a third lower than the RMSE for the standard calibration. As with the calibrated job-filling rate, the calibrated job-finding rate using effective searchers has a higher correlation and a standard deviation closer to the actual job-finding rate than the standard calibrated series.

Summary statistics reported in table 3 also shed light on the difference that the choice of α can make in the results obtained. Using $\alpha = 0.57$ (the elasticity for the generalized matching function) reduces the RMSE of the standard calibration by about 15 percent compared to its RMSE when using $\alpha = 0.49$. On the other hand, comparing the results in panels A and B of table 3, increasing α from 0.49 to 0.57 worsens the relative performance of the standard model for calibrating the job-filling rate. The implication of these conflicting effects is that the relatively poor performance of the standard model cannot be rescued with an alternative estimate of the matching function elasticity.

As already noted in comparing figure 1 with online appendix figures A.1, A.2, and A.3, the time series behavior of effective searchers is quite similar whether we use raw or demographically adjusted base period job-finding rates to weight the different groups of searchers; two groups or thirteen groups among the unemployed; or five broader groups of searchers rather than twenty-two groups as in our baseline analysis. In online appendix figures A.9, A.10, and A.11, we show results for the job-finding rate from unemployment for all three of these variants. The RMSE for the measure based on raw base period job-finding rates is nearly identical to that for our baseline specification. Both the RMSE for the generalized model with only two unemployment groups and the RMSE when we use just five groups rather than twenty-two groups are about 70 percent of that for the standard model, as compared to 67 percent for the fully general model.

VI. Exploring Time Variation in Recruiting Intensity and Search Intensity

The generalized model we outlined in section II includes both cross-sectional and time series variation in both recruiting intensities and search intensities, but thus far we have considered only the cross-sectional variation.

Research on alternative approaches to measuring the variation in recruiting intensity and relative search intensities is still in its early stages. In this section, we present results from an exploratory analysis of the variation in these intensities over our 1994:M1–2019:M12 sample period. For recruiting intensity, we use the Davis, Faberman, and Haltiwanger (2013) methodology to construct an aggregate recruiting intensity index. Using cross-sectional micro data, Davis, Faberman, and Haltiwanger (2013) estimate the elasticity between the job-filling rate and the gross hiring rate at the establishment level as $\phi = 0.82$. Applying this micro-based elasticity, they proxy aggregate recruiting intensity as $\rho_t^v = \left(\frac{H_t}{E_t}\right)^{\phi}$. We apply that proxy here.[11] It also is possible that, in addition to depending on the gross hiring rate, recruiting intensity may change over time due to composition effects (Davis, Faberman, and Haltiwanger 2012b) or to the endogenous responses of firms to overall labor market conditions, independent of their gross hiring rate (Gavazza, Mongey, and Violante 2018; Mongey and Violante 2020). We leave these possibilities for future exploration.

For time variation in relative job search intensities, we return to the specifications of Hall and Schulhofer-Wohl (2018). Their estimation allows not only for differences in average search intensities (what we have denoted as γ_i) but also for differences across job searcher types in the elasticity of job finding with respect to vacancy duration. We build on this insight to construct time-varying relative search intensities for our twenty-two groups, exploiting group-specific differences in the time variation in relative job-finding rates to infer what is happening to relative search intensities.[12] Details are reported in online appendix B. We do not use the common variation in job-finding rates in reweighting the different job searcher groups as we cannot distinguish changes in job-finding rates inherent in the matching function from those due to changes in search intensity.

Because the approach we have taken to constructing our time-varying recruiting intensity and relative job search intensity measures makes use of information related to the number of hires over time, some might be concerned that we have somehow built in the better performance of our model for explaining the job-filling and job-finding rates. Given the nature of our exercise, however, we do not believe this to be the case. We are

11. Online appendix figure A.12 shows how this recruiting intensity measure has moved over time.

12. See online appendix figure A.13 for the effective searcher measure using time-varying relative search intensities.

Figure 9. Actual versus Calibrated Job-Filling Rates with Standard and Generalized Tightness Measures including Time-Varying Recruiting (RI) and Search Intensity (SI)

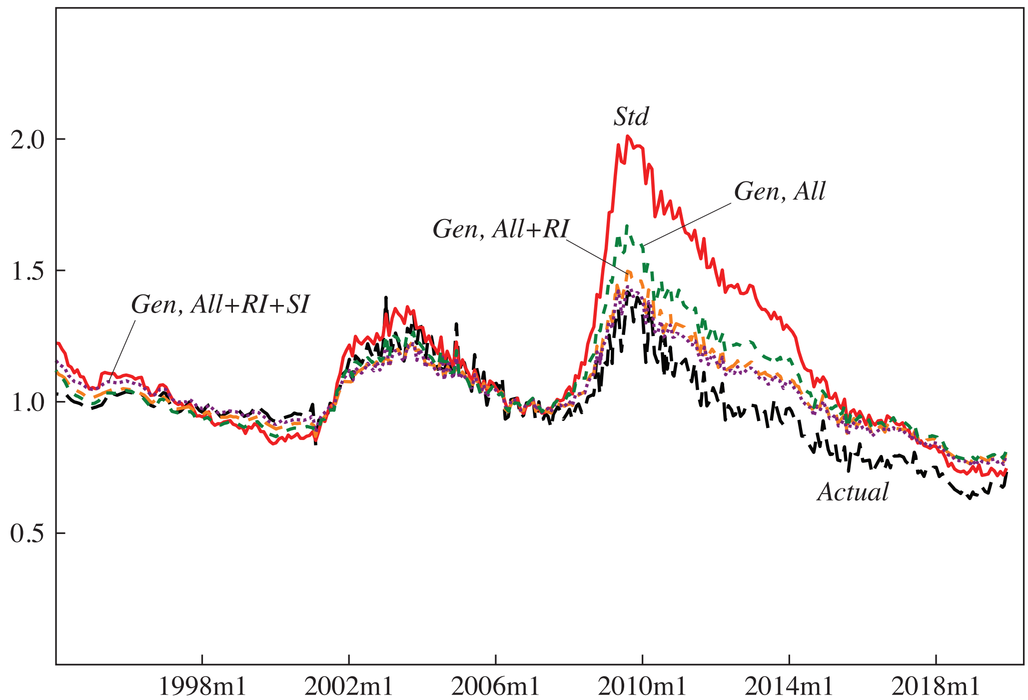

Source: Authors' calculations using CPS, JOLTS, and DFH (2012a) vacancies.

Notes: All rates normalized to 1 in 2006. Actual = job filling rate (H/V) from JOLTS; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and constant recruiting intensity; Gen, All + RI = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and time-varying recruiting intensity; Gen, All + RI + SI = calibrated using V/ES (all twenty-two groups) with time-varying search intensities and recruiting intensity.

evaluating alternative functional forms for the matching function by comparing the calibrated job-filling and job-finding rates based on each to the actual data.[13] There is no inherent reason that incorporating time-varying recruiting and search intensity into these calibrated rates in fact will reduce the unexplained residual variation in the outcome of interest. Indeed, as will become clear, we obtain mixed results when we extend the general model in these directions.

Figure 9 depicts the actual and calibrated job-filling rates from the standard model, the general model with fixed relative job search intensities, the extension of that specification to include time-varying recruiting intensity, and finally a specification that includes both time-varying recruiting

13. Online appendix B includes the functional forms of job-filling and job-finding rates in this extended version of the general model.

Table 4. Relative Performance for Job-Filling and Job-Finding Rates using Standard versus General Model: Time-Varying Recruiting Intensity (RI) and Relative Job Search Intensity (SI)

	Std dev	*Corr w/ actual*	*RMSE*	*Ratio RMSE to standard*
A. Job-filling rate				
Actual	0.17	1.00	NA	NA
Standard	0.29	0.75	0.25	1.00
General, constant SI, RI = 1	0.19	0.81	0.13	0.54
General, time-varying RI	0.15	0.88	0.09	0.38
General, time-varying SI	0.17	0.84	0.12	0.49
General, time-varying RI and SI	0.14	0.91	0.09	0.34
B. Job-finding rate for the unemployed				
Actual	0.16	1.00	NA	NA
Standard	0.20	0.31	0.26	1.00
General, constant SI, RI = 1	0.15	0.62	0.17	0.67
General, time-varying RI	0.18	0.77	0.14	0.54
General, time-varying SI	0.29	0.68	0.22	0.85
General, time-varying RI and SI	0.31	0.75	0.22	0.83

Source: Authors' calculations using CPS; JOLTS; Davis, Faberman, and Haltiwanger (2012a) vacancies; and Bureau of Labor Statistics Gross Flows.

Notes: Panel A shows statistics for calibrated job-filling rates using matching function elasticities specific to each. Panel B shows statistics for calibrated job-finding rate for the unemployed.

intensity and time-varying relative search intensities. As can be seen in table 4, adding time-varying recruiting intensity helps to reduce the residual unexplained variation in the job-filling rate, lowering the RMSE in the generalized model's projected job-filling rate from 0.13 to 0.09. Allowing in addition for time-varying relative job search intensity has a much smaller effect that is not apparent when the numbers are rounded to two digits to the right of the decimal point, though the fourth column of the table shows that the RMSE declines relative to that in the standard model.

Figure 10 depicts the analogous exercise for the job-finding rate for the unemployed. Just adding time-varying recruiting intensity yields a modest improvement in the RMSE for the residual unexplained variation relative to the baseline case, lowering it from 0.17 to 0.14, but incorporating the time series variation in relative job search intensities leads to a worsening in the model's performance for tracking the job-finding rate among the unemployed, raising the RMSE to 0.22.

We regard the results shown in figures 9 and 10 as more suggestive than conclusive. The proxies we use for time variation in recruiting intensity and job search intensities are relatively crude and indirect. Developing better measures of these intensities is an important direction for future research.

Figure 10. Actual versus Calibrated Job-Finding Rates for the Unemployed with Standard and Generalized Tightness Measures including Time-Varying Recruiting (RI) and Search Intensity (SI)

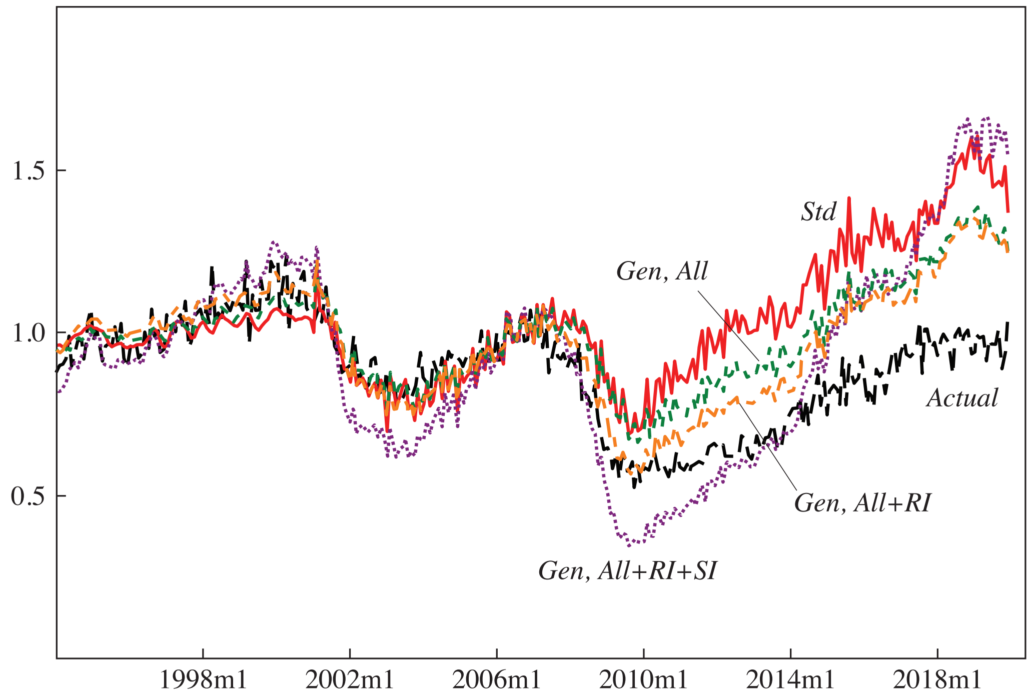

Source: Authors' calculations using CPS, JOLTS, DFH (2012a) vacancies, and BLS Gross Flows.
Notes: All rates normalized to 1 in 2006. Actual = job-finding rate for unemployed from published BLS Gross Flows; Std = calibrated using V/U; Gen, All = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and constant recruiting intensity; Gen, All + RI = calibrated using V/ES (all twenty-two groups) with constant relative search intensities and time-varying recruiting intensity; Gen, All + RI + SI = calibrated using V/ES (all twenty-two groups) with time-varying search intensities and recruiting intensity.

VII. Conclusions and Next Steps

The generalized measure of labor market tightness we have constructed based on the ratio of vacancies to effective searchers suggests that the US labor market was considerably less tight at the end of 2019 than implied by the standard ratio of vacancies to unemployment. The differing behavior of the two measures reflects the fact that the standard tightness measure does not account for important variation in search behavior on the part of workers. Job searchers include not only the unemployed but also those who are out of the labor force and the employed. In downturns, a more general index of effective searchers rises proportionally less than unemployment. The fact that the unemployed are only about 30 percent of all effective searchers contributes to this result, as any percentage increase in unemployment has a proportionally smaller effect on the overall number of effective searchers. Another contributing factor is that, during a protracted contraction such as the Great Recession, the distribution of unemployment shifts toward the long-term unemployed, meaning that effective searchers rise less than the simple unemployment count.

The central question motivating our analysis is whether substituting a generalized measure of effective searchers for the standard unemployment measure reduces the need to appeal to fluctuations in matching efficiency to explain what is happening to employment flows. We observe that the Beveridge curve constructed using effective searchers is much more stable than the standard Beveridge curve. Further, the matching function for hires with our generalized measure of labor market tightness as its argument outperforms the matching function based on the ratio of vacancies to unemployment. Specifically, the calibrated job-filling rate (hires per vacancy) using the generalized measure tracks the actual job-filling rate much more closely than the job-filling rate calibrated using the standard measure of labor market tightness. The calibrated job-finding rate among the unemployed (hires from unemployment per unemployed person) based on the generalized measure also comes closer to tracking the actual series than the calibrated rate based on the standard measure. These findings imply that our approach reduces the unexplained residual variation required in the matching function to be consistent with the real-world data.

Our baseline effective searcher measure is constructed using data for twenty-two separate population groups. We also have examined a number of alternatives that are less inclusive or based on more aggregated population groups. Taken together, the results make clear that the key to better matching the actual time series behavior of the job-filling and job-finding rates is having a broad-based measure of effective searchers that also distinguishes among core groups. Measures that include the employed and, especially, those out of the labor force do better than those limited to the unemployed, but the very detailed disaggregation of the broader groups we have adopted does little better than the simpler five-category breakout we examined as a sensitivity check. Given the large differences in base-period job-finding rates across the detailed categories combined in the five-category version of the generalized measure, it might seem surprising that the further breakouts we apply do not add more to the performance of the job-filling and job-finding models. Over our period, the changes in composition occurring within the five more aggregated groups are not large enough to make much difference, but this might not always be the case. As there is little cost to using more disaggregated groups to construct the generalized tightness measure, it seems preferable to us to do so.

Our baseline measure of labor market tightness undoubtedly could be improved and built upon. In a suggestive analysis, we find that incorporating proxies for time variation in relative job search intensities and also taking into account variation in recruiting intensity across employers further reduces the unexplained variation in job-filling rates. We find mixed evidence on using these proxies for explaining the fluctuations in job-finding rates among the unemployed. There clearly is more to be done to develop time series measures of search and recruiting intensity.

One topic that we have deliberately avoided but that is of critically important interest is whether and how the generalized labor market tightness approach could improve our understanding of wage and price pressures. It would be interesting to explore the estimation of Phillips curve–type relationships using a generalized measure of labor market tightness rather than the unemployment rate gap as the central explanatory variable. Even if it is true that labor market tightness is a better predictor of wage and price changes than the unemployment rate, however, there are other sources of instability in the Phillips curve relationship that seem likely to pose problems for the estimation of such relationships. Still, given that estimating and interpreting Phillips curves is an active area of research and highly relevant for policymakers, there would be value in exploring the role of generalized labor market tightness measures in this context.

ACKNOWLEDGMENTS We are grateful to Daniel Aaronson, Steven Davis, Jason Faberman, Robert Hall, Chris Nekarda, Ayşegül Şahin, Sam Schulhofer-Wohl, Jim Stock, Robert Valletta, and Justin Wolfers for helpful conversations and comments on an earlier draft; to colleagues at the University of Maryland for comments at an early stage of this research that helped to sharpen our thinking about these issues; and to participants in both the Federal Reserve System's June 2019 Conference on Monetary Policy Strategy, Tools and Communication Practices and the March 2020 Brookings Panel on Economic Activity meeting for their comments.

References

Aaronson, Daniel, Luojia Hu, Arian Seifoddini, and Daniel G. Sullivan. 2015. "Changing Labor Force Composition and the Natural Rate of Unemployment," Chicago Fed Letter 338.

Abraham, Katharine G. 2015. "Is Skill Mismatch Impeding U.S. Economic Recovery?" *ILR Review* 68, no. 2: 291–313.

Abraham, Katharine G., John C. Haltiwanger, Kristin Sandusky, and James R. Spletzer. 2019. "The Consequences of Long-Term Unemployment: Evidence from Linked Survey and Administrative Data." *ILR Review* 72, no. 2: 266–99.

Ahn, Hie Joo, and Ling Shao. 2017. "Precautionary On-the-Job Search over the Business Cycle." Finance and Economics Discussion Series 2017–025. Washington: Board of Governors of the Federal Reserve System.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109, no. 3: 175–78.

Barnichon, Regis, and Geert Mesters. 2018. "On the Demographic Adjustment of Unemployment." *Review of Economics and Statistics* 100, no. 2: 219–31.

Black, Matthew. 1980. "Pecuniary Implications of On-the-Job Search and Quit Activity." *Review of Economics and Statistics* 62, no. 2: 222–29.

Blanchard, Olivier, and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, no. 1, 1–76.

Blanchard, Olivier, and Peter A. Diamond. 1992. "The Flow Approach to Labor Markets." *American Economic Review* 82, no. 2: 354–59.

Blau, David M., and Philip K. Robins. 1990. "Job Search Outcomes for the Employed and Unemployed." *Journal of Political Economy* 98, no. 3: 637–55.

Broersma, Lourens, and Jan C. van Ours. 1999. "Job Searchers, Job Matching and the Elasticity of Matching." *Labour Economics* 6:77–93.

Casselman, Ben. 2018. "As Labor Pool Shrinks, Prison Time Is Less of a Hurdle." *New York Times*, January 13.

Crump, Richard K., Stefano Eusepi, Marc Gionnoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring U*." *Brookings Papers on Economic Activity*, Spring, 143–214.

Davis, Steven J. 2011. Comment on "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data," by Alan Krueger and Andreas Mueller. *Brookings Papers on Economic Activity*, Spring, 58–70.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012a. "Labor Market Flows in the Cross Section and Over Time." *Journal of Monetary Economics* 59, no. 1: 1–18.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012b. "Recruiting Intensity during and after the Great Recession." *American Economic Review* 102, no. 3: 584–88.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. “The Establishment-Level Behavior of Vacancies and Hiring,” *Quarterly Journal of Economics*, 128, no. 2: 581–622.

DeLoach, Stephen B., and Mark Kurt. 2013. “Discouraging Workers: Estimating the Impacts of Macroeconomic Shocks on the Search Intensity of the Unemployed.” *Journal of Labor Research* 34:433–54.

Diamond, Peter A. 1982. “Wage Determination and Efficiency in Search Equilibrium.” *Review of Economic Studies* 49, no. 2: 217–27.

Elsby, Michael, Ryan Michaels, and David Ratner. 2015. “The Beveridge Curve: A Survey.” *Journal of Economic Literature* 53, no. 3: 571–630.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2017. “Job Search Behavior among the Employed and the Non-Employed.” Working Paper 23731. Cambridge, Mass.: National Bureau of Economic Research.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2020. “The Shadow Margins of Labor Market Slack.” Working Paper 26852. Cambridge, Mass.: National Bureau of Economic Research.

Fujita, Shigeru, and Guiseppe Moscarini. 2017. “Recall and Unemployment.” *American Economic Review* 107, no. 12: 3875–916.

Fujita, Shigeru, Guiseppe Moscarini, and Fabien Postel-Vinay. 2019. “Measuring Employer-to-Employer Reallocation.” Working Paper.

Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante. 2018. “Aggregate Recruiting Intensity.” *American Economic Review* 108, no. 8: 2088–127.

Gomme, Paul, and Damba Lkhagvasuren. 2015. “Worker Search Effort as an Amplification Mechanism.” *Journal of Monetary Economics* 75:106–22.

Hall, Robert, and Sam Schulhofer-Wohl. 2018. “Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers.” *American Economic Journal: Macroeconomics* 10, no. 1: 1–32.

Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erica McEntarfer. 2018. “Cyclical Job Ladders by Firm Size and Firm Wage.” *American Economic Journal: Macroeconomics* 10, no. 2: 52–85.

Haltiwanger, John C., Henry R. Hyatt, and Erica McEntarfer. 2018. “Who Moves Up the Job Ladder?” *Journal of Labor Economics* 36, no. S1: S301–36.

Hornstein, Andreas, and Marianna Kudlyak. 2016. “Estimating Matching Efficiency with Variable Search Effort.” Working Paper 2016–24. San Francisco: Federal Reserve Bank of San Francisco.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange. 2014. “Measuring Resource Utilization in the Labor Market.” *Federal Reserve Bank of Richmond Economic Quarterly* 100, no. 1: 1–21.

Jones, Stephen R. G., and W. Craig Riddell. 1999. “The Measurement of Unemployment: An Empirical Approach.” *Econometrica* 67, no. 1: 147–62.

Kaitz, Hyman. 1970. “Analyzing the Length of Spells of Unemployment.” *Monthly Labor Review* 93, no. 11: 11–20.

Katz, Lawrence F. 1986. "Layoffs, Recalls and the Duration of Unemployment." Working Paper 1825. Cambridge, Mass.: National Bureau of Economic Research.

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* 41, no. 1: 45–72.

Krueger, Alan B., Judd Cramer, and David Cho. 2014. "Are the Long-Term Unemployed on the Margins of the Labor Market?" *Brookings Papers on Economic Activity*, Spring, 229–80.

Krueger, Alan B., and Andreas Mueller. 2010. "Job Search and Unemployment Insurance: New Evidence from Time Use Data." *Journal of Public Economics* 94:298–307.

Krueger, Alan B., and Andreas Mueller. 2011. "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High-Frequency Longitudinal Data." *Brookings Papers on Economic Activity,* Spring, 1–57.

Kudlyak, Marianna. 2017. "Measuring Labor Utilization: The Non-Employment Index." Federal Reserve Bank of San Francisco Economic Letter 2017–08, March 17.

Modestino, Alicia Sasser, Danial Shoag, and Joshua Balance. Forthcoming. "Upskilling: Do Employers Demand Greater Skill when Workers Are Plentiful?" *Review of Economics and Statistics*.

Mongey, Simon, and Gianlucia Violante. 2020. "Macro Recruiting Intensity from Micro Data." Working Paper 26231. Cambridge, Mass.: National Bureau of Economic Research.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61, no. 3: 397–415.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin. 2018. "Job Search Behavior over the Business Cycle." *American Economic Journal: Macroeconomics* 10, no. 1: 190–215.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity*, no. 3: 411–48.

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

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

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104, no. 11: 3529–64.

Sedlacek, Petr. 2016. "The Aggregate Matching Function and Job Search from Employment and Out of the Labor Force." *Review of Economic Dynamics* 21:16–28.

Shimer, Robert. 2001. "The Impact of Young Workers on the Aggregate Labor Market." *Quarterly Journal of Economics* 116, no. 3: 969–1008.

Shimer, Robert. 2004. "Search Intensity." Working paper. Chicago: University of Chicago.

Smialek, Jeanna. 2019. "The Economy Is Strong and Inflation Is Low. That's What Worries the Fed." *New York Times*, May 21.

Veracierto, Marcelo. 2011. "Worker Flows and Matching Efficiency." *Economic Perspectives* 35, no. 4: 147–69.

Comments and Discussion

COMMENT BY

AYŞEGÜL ŞAHIN Abraham, Haltiwanger, and Rendell developed a generalized measure of labor market tightness which takes into account intensive and extensive margins of search activity on both demand and supply sides of the labor market. Their measure captures the hiring process in the US economy better than the standard measure of labor market tightness. Their success is probably not surprising given that they build on two well-documented facts about the hiring process in the United States. The first is that the majority of jobs are filled by workers who are not unemployed. The second is that the number of vacancies is an imperfect proxy for firms' total recruiting effort since firms vary their recruiting intensity over the business cycle.

Assessment of labor market conditions is a fundamental issue in macroeconomics and a key input to implementation of monetary and fiscal policy. The authors' work is a valuable addition to the wealth of labor market indicators developed in the last decade to better evaluate labor market developments in light of ongoing secular trends. With the unemployment rate jumping from 3.5 percent in February 2020 to 14.7 percent in April 2020, the labor market will be our main focus of attention for years to come. Understanding how workers search for employment opportunities, how firms fill their open positions, and how search activity responds to macroeconomic conditions will help us in characterizing the adjustment path of the US labor market to the COVID-19 shock. This comment reviews and interprets the authors' findings and suggests new directions of research.

FRAMEWORK Abraham, Haltiwanger, and Rendell use the matching function specification—a key building block of the Diamond-Mortensen-Pissarides search and matching framework—to characterize the behavior of hiring in the US economy. In its basic form, the matching function takes the Cobb-Douglas form and specifies hires, h_t, as a function of two inputs: vacancies posted by firms looking to hire, v_t, and unemployed workers looking for jobs, u_t:

$$h_t = \Phi_t (u_t)^{\alpha} (v_t)^{1-\alpha}, \tag{1}$$

where $\alpha \in (0,1)$ is the unemployment share and Φ_t is the aggregate matching efficiency parameter. The job-finding and the job-filling rates can be written as a function of labor market tightness ($\theta_t = v_t/u_t$) as $\Phi_t\ \theta_t^{1-\alpha}$ and $\Phi_t (1/\theta_t)^{\alpha}$, respectively. This specification implies that θ_t is the only determinant of the hiring process if Φ_t does not vary over time. However, as discussed by the authors, this specification ignores the fact that most hires originate from employment or nonparticipation. Moreover, it does not take into account the intensive margins of firm and worker search effort. A generalized tightness measure, $\tilde{\theta}_t$, which incorporates these factors can be defined as

$$\tilde{\theta}_t = \frac{\sum_j \rho_t^{V_j} V_{jt}}{\sum_i \rho_t^{S_i} S_{it}}, \tag{2}$$

where $\rho_t^{V_j}$ is the recruiting intensity of vacancy group in firm type j and $\rho_t^{S_i}$ is the search intensity of searcher group i. The interpretation of $\tilde{\theta}_t$ is straightforward: it is the ratio of the recruiting-intensity-weighted sum of vacancies to the search-effort-weighted sum of searchers in the economy. While searchers include all individuals regardless of their labor force status, not all groups contribute equally to the total search effort in the economy. The contribution of each group depends on their relative search intensity, $\rho_t^{S_i}$.

Abraham, Haltiwanger, and Rendell do not explore the variation in the composition of firms by age, size, or industry over the business cycle. Instead, they use the direct measure of aggregate recruiting intensity from Davis, Faberman, and Haltiwanger (2013). With the aggregate recruiting intensity ρ_t^V in hand, they compute the effective vacancies as $\rho_t^V V_t$. The implementation of the generalized measure of effective searchers is more involved since it requires using detailed micro data from the Current Population Survey (CPS). First, the authors identify twenty-two distinct labor force states following Hall and Schulhofer-Wohl (2018) and compute the job-finding rates of these groups using the matched CPS data. They assume that the relative job-finding rate of each group corresponds to its relative search intensity. This assumption allows them to compute an aggregate measure of effective searchers by weighting each group i by its relative job-finding rate. Among these twenty-two detailed labor market groups, thirteen groups are unemployed; seven groups are those who are out of the labor force, and two are the employed. Table 1 of the paper shows that these groups vary considerably in their job-finding rates with recently temporarily laid-off workers having the highest job-finding rate.

ARE TWENTY-TWO LABOR MARKET STATES NECESSARY? Abraham, Haltiwanger, and Rendell differentiate between various detailed labor market states and end up using twenty-two distinct groups in their implementation. While this level of disaggregation allows them to capture even small variations in the relative job-finding rates of searchers, it restricts their analysis to the post-1994 period since the questions that allow them to differentiate between these twenty-two states are not available in the CPS before 1994. I show that focusing only on a few broad labor market states is sufficient to capture almost all the variation in the measure of effective searchers over time. This simplification is useful for two reasons. First, it makes it possible to compute the generalized tightness measure without using the CPS micro data. Second, it allows us to compute the generalized tightness measure starting in 1948.

Starting with five broad labor force states, the measure of effective searchers can be simplified as:

$$S_t = \underbrace{\rho^{UST} U_t^{ST} + \rho^{ULT} U_t^{LT}}_{\text{Unemployed}} + \underbrace{\rho^{W} W_t + \rho^{N} N_t}_{\text{Nonparticipants}} + \underbrace{\rho^{E} E_t}_{\text{Employed}} .$$

I focus on unemployment, nonparticipation, and employment and further distinguish between short-term (U^{ST}) and long-term unemployed (U^{LT}) among the unemployed and between those who want a job (W) and those who do not (N) among nonparticipants. I approximate the relative job-finding rates of these groups using the raw job-finding rates from the authors' table 1 and choose the following weights:

$$\rho^{UST} = 1, \rho^{ULT} = 0.48, \rho^{W} = 0.40, \rho^{N} = 0.09, \rho^{E} = 0.07.$$

I then compute the rate of effective searchers as

$$s_t = \frac{S_t}{P_t} = \frac{S_t}{U_t + N_t + E_t}$$

and normalize it to 1 in 2006. Figure 1 shows that the series computed using only five labor market states line up remarkably well with the authors'

Figure 1. Measures of Effective Searchers using Five Labor Force States

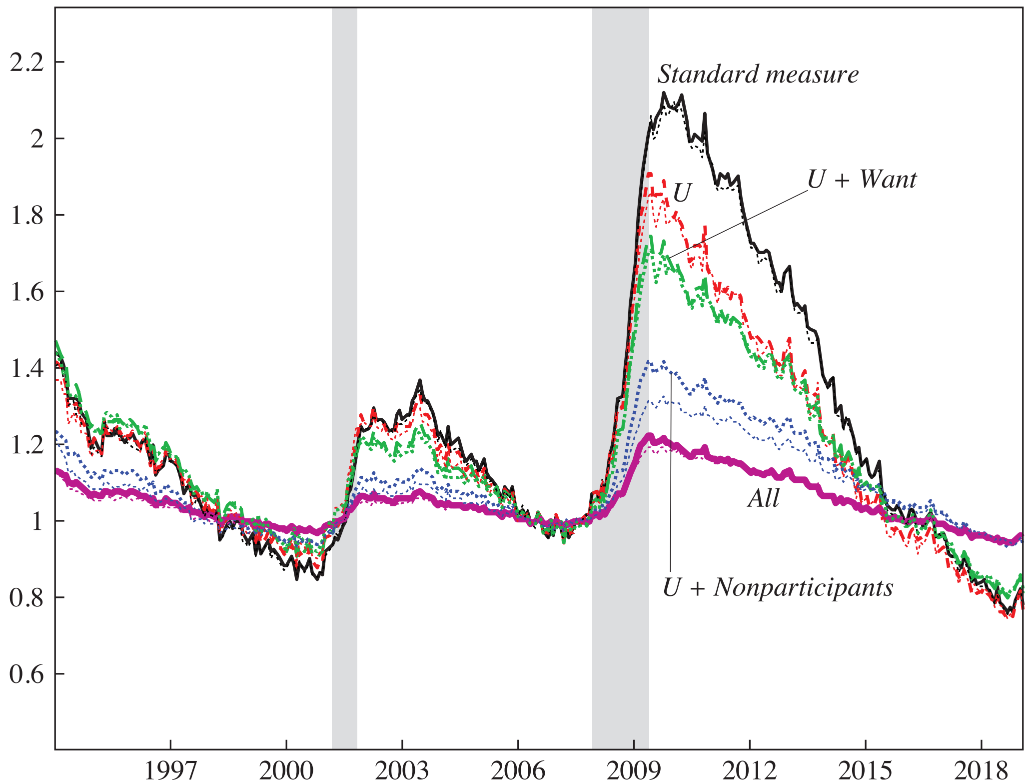

Source: Author's calculations based on the Current Population Survey (CPS).

Notes: AHR measures indicated by thin dotted lines. Standard measure: U_t/P_t; U: $(U_t^{ST} + 0.48U^{LT})/P_t$; U + Want: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t)/P_t$; U + Nonparticipants: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t)/P_t$; and all: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t + 0.07E_t)/P_t$. Series normalized to 1 in 2006.

measures.[1] The intuition is clear: while job-finding rates vary substantially by detailed labor market state, most of these groups are too small to affect the aggregate measure. Therefore I conclude that focusing on five broad labor market states does not change the essence of the authors' findings but simplifies their analysis considerably by making it possible to compute the index without using the CPS micro data.

HISTORICAL PERSPECTIVE Given that the five-state simplification does remarkably well, I now turn to historical data and calculate the measure of effective searchers starting in 1948. This requires consolidating the two groups of nonparticipants to only one, but even this abstraction does not change the behavior of the generalized measure. Figure 2 plots the generalized measure of effective searchers along with the standard measure

1. While these weights are not exact, following my discussion at the panel, the authors repeated their analysis with only five groups and found very similar results to mine.

Figure 2. Generalized Measure of Effective Searchers and the Standard Measure in the 1948–2018 Period

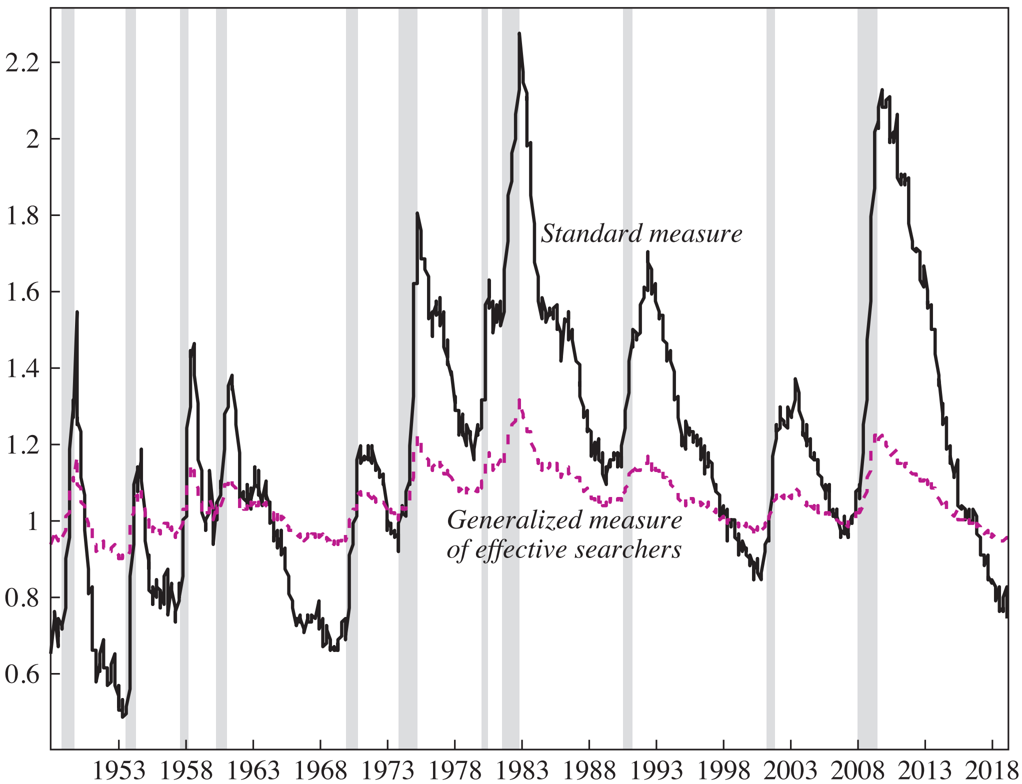

Source: Author's calculations based on the Current Population Survey (CPS).

Notes: AHR measures indicated by dotted lines. Standard measure: U_t/P_t; generalized measure of effective searchers: $(U_t^{ST} + 0.48U^{LT} + 0.09N_t + 0.07E_t)/P_t$.

starting in 1948 and shows that the deviation between the two measures is not specific only to the most recent expansion. The effective searchers measure exhibits more muted countercyclical behavior than the standard measure throughout the sample—a distinct feature of the generalized measure that I will discuss below in detail.

Interestingly, despite differences in the magnitude of their cyclicality, measures based on unemployment and effective searchers provide a very similar assessment of the historical business cycles. Both measures imply that the 1973–1975, 1980–1982, and 2007–2009 recessions were the deepest downturns and the second halves of the 1960s, 1990s, and 2010s were the tightest labor markets of the postwar period.

I also examine the historical Beveridge curve using the standard measure and the measure of effective searchers combined with Barnichon's (2010) composite help-wanted vacancy index in figure 3. The Beveridge

Figure 3. Historical Beveridge Curves using the Generalized Measure of Effective Searchers and the Standard Measure

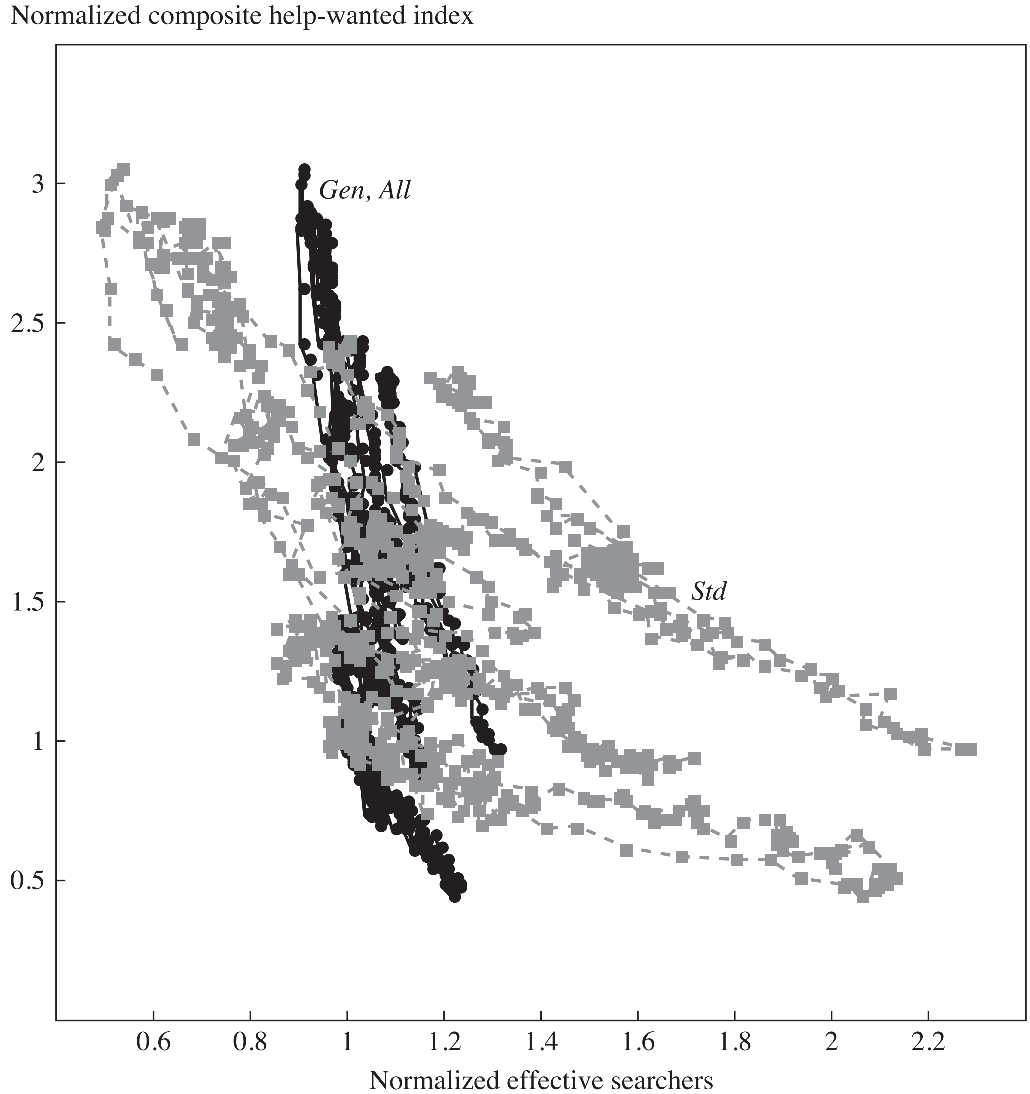

Source: Author's calculations based on the Current Population Survey (CPS) and composite help-wanted vacancy index developed by Barnichon (2010).

Note: Historical Beveridge curve using the unemployment-to-population ratio (gray dots) and the effective measure of searchers (black dots).

curve constructed using unemployment exhibits substantial horizontal shifts over time as is well documented in Diamond and Şahin (2015). The shifts are much less pronounced when the measure of effective searchers is used. However, the Beveridge curve also becomes much steeper when it is constructed with the new measure. The steepness of the Beveridge curve suggests that fluctuations in vacancies have very little effect on labor

Figure 4. Measures of Effective Searchers in Levels

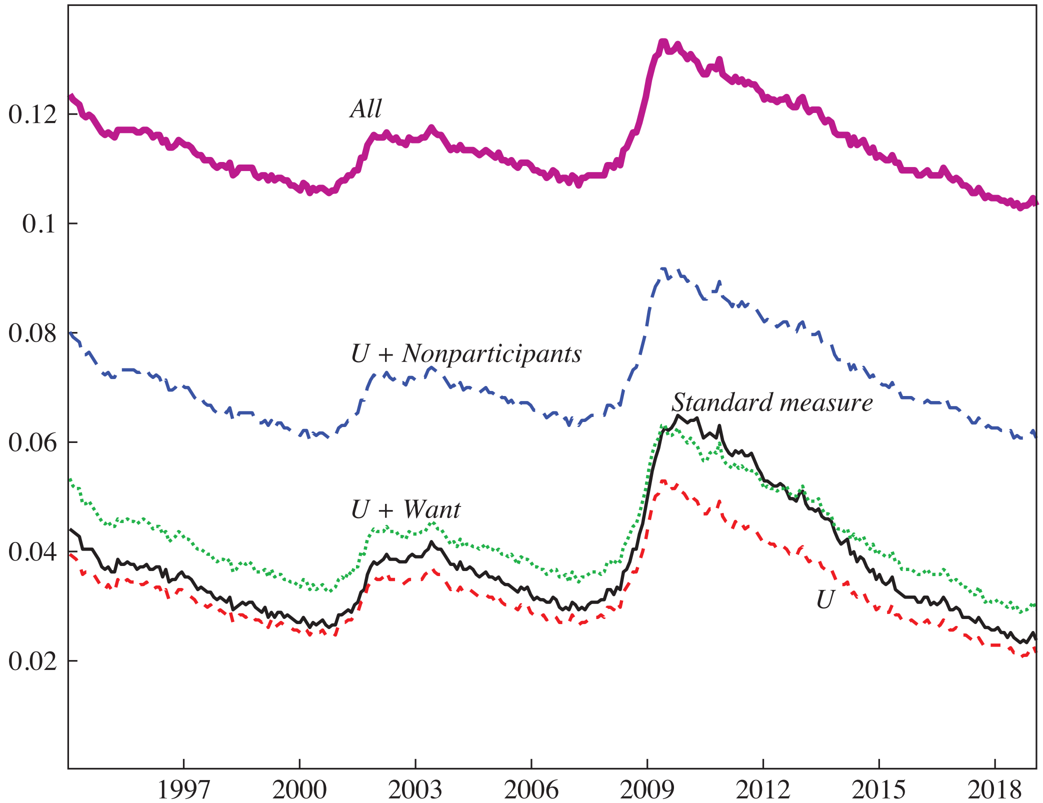

Source: Author's calculations based on the Current Population Survey (CPS).

Notes: Alternative searcher measures in levels. Standard measure: U_t/P_t; U: $(U_t^{ST} + 0.48U^{LT})/P_t$; U + Want: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t)/P_t$; U + Nonparticipants: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t)/P_t$; and all: $(U_t^{ST} + 0.48U^{LT} + 0.40W_t + 0.09N_t + 0.07E_t)/P_t$.

market underutilization. This interesting observation deserves further discussion and analysis since it suggests that the link between labor demand and labor underutilization has been considerably weaker than is typically assumed in the postwar period.

WHY IS THE COUNTERCYCLICALITY DAMPENED? Abraham, Haltiwanger, and Rendell focus only on normalized measures of effective searchers and do not report or interpret the *levels* of their measures. While this type of normalization is useful in the context of the matching function framework, I argue that the difference in the levels of alternative measures of searchers is important for the interpretation of the authors' findings.

Figure 4 plots different measures of searchers in levels without normalizing them to 1 in 2006. Comparison of figures 1 and 4 shows that measures plotted in levels look very different than their normalized counterparts. Expanding the pool of searchers by adding employed workers and nonparticipants increases the level substantially despite the low relative job-finding rates of these groups, and the normalization reduces cyclicality of the measures substantially.

This comparison is very helpful in understanding the reasons for dampened countercyclicality of the effective searchers measure relative to unemployment. First, the level of the generalized measure is higher than the level of the standard measure due to the addition of the high number of employed workers and nonparticipants. Second, there are offsetting changes in the composition of searchers over the business cycle. During recessions, unemployment increases but due to the decline in employment, the number of employed searchers declines, dampening the rise in unemployment. Similarly, during expansions as the number of unemployed searchers declines, the number of employed searchers increases, moderating the decline in the number of unemployed searchers. As a result, the generalized measure of effective searchers fluctuates less than unemployment over the business cycle. This finding has an important implication about the hiring incentives of firms: firms do not find it that much easier to fill jobs during recessions despite high unemployment since the number of employed searchers declines. This is why the generalized tightness measure does better in capturing the behavior of hires in the matching function framework.

HOW SHOULD WE ASSESS LABOR MARKET CONDITIONS? The Great Recession and the subsequent period of sluggish recovery in the labor market triggered an important line of research, with numerous studies developing labor market indicators as alternatives to the unemployment rate. Most of these measures exploit one of the following two approaches. The first focuses on estimating a time-varying natural rate of unemployment and uses the unemployment gap—the difference between the actual and the natural rate of unemployment—as a measure of labor market tightness. While this line of research typically focuses on estimating the natural rate of unemployment using only simple aggregate data, more recent work unifies the macro approach with rich labor market data (Crump and others 2019). The main advantage of this approach is that it provides a unified framework that takes into account secular demographic trends, wage and price inflation, and inflation expectations. Therefore it directly connects to maximum employment and price stability objectives.

The second approach is to develop broader measures which take into account additional margins of labor market underutilization. A common practice is to weight different groups of workers depending on their demographic characteristics, wages, search activity, or job-finding rates, such as in Perry (1970), Hornstein, Kudlyak, and Lange (2014), and this paper. While implementing this approach, choice of weights is an important issue since the interpretation of the aggregate measure depends heavily on the weights used. For example, the authors choose to use relative job-finding rates and show that their measure relying on realized transitions is informative about the hiring process in the economy. However, this does not necessarily mean that their measure would be preferable to other measures for other purposes. For example, one fundamental issue in macroeconomics is to estimate the potential output of the economy. In that case, an alternative measure which weights workers by their average idle hours (desired hours minus actual hours worked), as in Faberman and others (2020), would be informative about potential output.

To conclude, the authors' generalized tightness measure is a valuable addition to our arsenal of labor market indicators. They convincingly demonstrated that the measure they develop captures the essence of the hiring process well. Future work calls for broadening the metric for success by considering the usefulness of the new metric in capturing wage growth and inflation. This extension would help to connect the new measure more directly to monetary policy implementation. With the unemployment rate at its highest level in the postwar period, we need as much information as possible about the labor market for real-time assessment of the state of the labor market and its medium-term evolution. The work of Abraham, Haltiwanger, and Rendell undoubtedly will contribute to our understanding of the labor market in the years to come.

REFERENCES FOR THE ŞAHIN COMMENT

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109, no. 3: 175–78.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring U*." *Brookings Papers on Economic Activity*, Spring, 143–214.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring." *Quarterly Journal of Economics* 128, no. 2: 581–622.

Diamond, Peter A., and Ayşegül Şahin. 2015. "Shifts in the Beveridge Curve." In *Research in Economics*, edited by Federico Etro. Amsterdam: North-Holland.

Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2020. "The Shadow Margins of Labor Market Slack." Working Paper 26852. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w26852.

Hall, Robert, and Sam Schulhofer-Wohl. 2018. “Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers.” *American Economic Journal: Macroeconomics* 10, no. 1: 1–32.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange. 2014. “Measuring Resource Utilization in the Labor Market.” *Federal Reserve Bank of Richmond Economic Quarterly* 100, no. 1: 1–21.

Perry, George L. 1970. “Changing Labor Markets and Inflation.” *Brookings Papers on Economic Activity*, no. 3: 411–48.

COMMENT BY

JUSTIN WOLFERS This paper by Abraham, Haltiwanger, and Rendell follows a long tradition of Brookings papers in proposing a new measure of labor market tightness. The new measure aims to capture effective job search effort, and its key advantages are that it recognizes that the unemployed aren’t the only group who search and it accounts for the reality that different groups of workers search with different degrees of success. To be precise, the proposed new measure—which the authors call the “generalized measure of effective searchers”—is constructed as a weighted average of the share of the population in twenty-two different labor market states, with the weights reflecting each group’s baseline level of search effectiveness (measured by each group’s historical job-finding rate). As a weighted average rather than simple average, it generalizes more standard metrics like unemployment (which effectively puts a weight of one on the unemployed and zero on all others).

This paper was written during what now seems a distant time in early 2020 when the official unemployment rate was lower than it had been in about half a century and one of the central macro policy questions of the day was whether unemployment could go any lower. The final draft was submitted a few weeks later, when the coronavirus pandemic had pushed the unemployment rate to levels higher than at any time since the Great Depression. A superficial assessment might argue that a paper titled “How Tight Is the US Labor Market?” has been overtaken by events which have rendered it less relevant. But in reality, this research question may actually be more relevant than ever before, as the recession caused by the coronavirus—which includes elements of demand shocks and supply shocks—has led to arguably greater uncertainty about the extent of slack in the labor market than the United States has experienced at any point throughout the postwar period. The present-day relevance simply requires a reframing from this being a paper assessing how best to measure labor market tightness to one assessing how best to measure labor market looseness.

In proposing a new measure of labor market conditions, this paper enters an already quite crowded marketplace. Between the various measures of labor market slack published by the Bureau of Labor Statistics (BLS), alternatives tracked by the Federal Reserve banks, the preferred metrics of an array of private-sector forecasters, and new measures proposed in back issues of *Brookings Papers on Economic Activity*, it's no exaggeration to say that dozens of measures have been proposed and are closely tracked. Whether it's worth tracking one more measure depends on whether this new measure yields additional useful information.

The main test that the paper offers is to ask whether this new measure does a better job in explaining time series variation in the job-filling and job-finding rates than standard metrics like the official BLS measure of unemployment. The authors argue that their measure does add useful additional information. The analysis that I present below disagrees. To preview, I find little evidence that their proposed new measure outperforms standard measures. While the authors claim the generalized measure yields a better fit, the tests reported reflect arbitrary auxiliary assumptions that tilt the playing field against the standard measure of unemployment. When these assumptions are relaxed, the proposed new measures no longer appear to be more predictive. Indeed, the proposed new measure appears to largely track the official measure of unemployment. By this telling, there is no problem with the newly proposed measure, but there is also little to recommend it, as it adds little information beyond that in the standard measure.

EXPLAINING TIME SERIES VARIATION IN THE JOB-FILLING RATE The main test that the paper implements is to ask which measure of labor market slack—the proposed generalized measure of effective searchers or the widely used official measure of unemployment—does a better job explaining time series variation in the job-filling rate, which is the ratio of hiring to vacancies. To motivate their analysis, the authors start with a Cobb-Douglas matching function of the form:

$$H_t = \mu U_t^{\alpha} V_t^{1-\alpha}.$$

This can be understood as a production function, where the output is the number of people hired, H_t, and the production of these matches is a function of the aggregate effort workers put into searching for jobs (U_t) and the aggregate effort that firms put into searching for workers (V_t). In addition, μ is an index of the efficiency of this production function, and the empirical analysis in this paper holds this parameter constant, which is why it lacks a time subscript.

It is common for empirical analyses of the matching function to use the level of unemployment (which I'll denote U_t^s) as a proxy for total search effort by workers (so $U_t = U_t^s$) and the level of vacancies as a proxy for total search effort by firms (V_t). The authors propose that their generalized measure of effective searchers (which I'll denote U_t^g) may be a better proxy for the level of total search effort by workers (so $U_t = U_t^g$, instead). In a later extension, they also argue that accounting for recruiting intensity may yield a better proxy of the search effort by firms.

To test this, the paper asks which competing proxy measure of U_t—the standard measure, which is the number of unemployed people U_t^s, or their generalized measure of effective searchers U_t^g—better explains the observed time series movements in the job-filling rate. The specific empirical exercise it implements comes from a simple rearrangement of the matching function so that the dependent variable is now the job-filling rate:

$$\frac{H_t}{V_t} = \mu\left[\frac{V_t}{U_t}\right]^{-\alpha}.$$

The variable to be forecast, H_t/V_t, is readily measured as monthly hires divided by monthly vacancies, both from the Job Openings and Labor Turnover Survey (JOLTS) data. Likewise, the two candidate inputs into this forecasting equation—U_t/V_t, measured as either U_t^S/V_t, the ratio of the unemployed (from the Current Population Survey [CPS]) to the number of vacancies, or as U_g^g/V_t, the ratio of the generalized measure of effective searchers (constructed by the authors from CPS data) to the number of vacancies—are easily measured.

So far, this sounds like a standard forecast evaluation exercise. But there's a twist. No statistical agency publishes estimates of either μ (the efficiency of the matching function) or α (the Cobb-Douglas coefficient in the matching function). Think of the parameter μ as determining the position of the Beveridge curve, while the parameter α determines its slope. The values of these parameters will shape how well any measure of slack will fit the data. The values that are imposed reflect an auxiliary set of assumptions, and as I will show below, these assumptions largely drive the empirical findings that follow.

Unfortunately, the paper does not contain much discussion of where these parameters come from nor much of a defense of the specific parameter values that are imposed. In order to appropriately assess the quality of the evidence presented in the paper, it's worth being more transparent about these assumptions. First, μ reflects an arbitrary normalization, in which the time series of forecasts of the job-filling rate (and the inputs into that forecast) are indexed to be equal to 1, on average, during 2006. This effectively imposes

$$\mu = \frac{\overline{H_t/V_t}^{2006}}{\overline{(V_t/U_t)^{-\alpha}}^{2006}},$$

which would make sense if μ were constant and the matching function exactly fit the data (with an error term equal to zero) on average through 2006. When U_t is proxied by the level of unemployment, this yields $\widehat{\mu^s} = 0.80$, but when it is proxied by the generalized measure of effective searchers, it yields $\widehat{\mu^g} = 0.48$. Next, α comes from a structural estimation procedure described only in online appendix B of the paper, in which the evolution over time of the job-finding rates for each of the twenty-two groups are estimated as a function of a group fixed effect, a common time trend, and the aggregate vacancy-to-hires ratio. In this framework, α is recovered as a nonlinear function of the coefficient on the vacancy-to-hires ratio, and the various α's for each of the twenty-two groups are averaged to get an aggregate $\widehat{\alpha^g}$. (I append the superscript g to denote that this is the α estimated for the *g*eneralized measure). However, this equation is not estimated directly but rather using instrumental variables, where employment from the payroll survey is an instrument for the vacancy-to-hires ratio. For the generalized measure of effective searchers, this yields $\widehat{\alpha^g} = 0.57$. For the standard measure, this procedure collapses so that there is only one group of searchers (the unemployed), and so the dependent variable is the ratio of unemployment-to-employment flows to unemployment, and the dependent variables are a constant, a time trend, and the aggregate vacancy-to-hires ratio (which again is instrumented using payroll employment). This yields $\widehat{\alpha^s} = 0.49$ (where the superscript s denotes that this is the α that applies to the *s*tandard measure).

Imposing these assumptions yields the forecasts of the job-filling rate shown in figure 1 (which replicates exactly the corresponding lines from figure 5 in the paper). The top (dashed) line is the forecast generated using the standard measure of U_t, the next (dotted) line is the forecast generated using the generalized measure, while the actual job-filling rate, which is shown as a solid line, lies below both of these forecasts for most of the sample. The authors emphasize that the forecast generated using the standard

Figure 1. Predicting the Job-Filling Rate

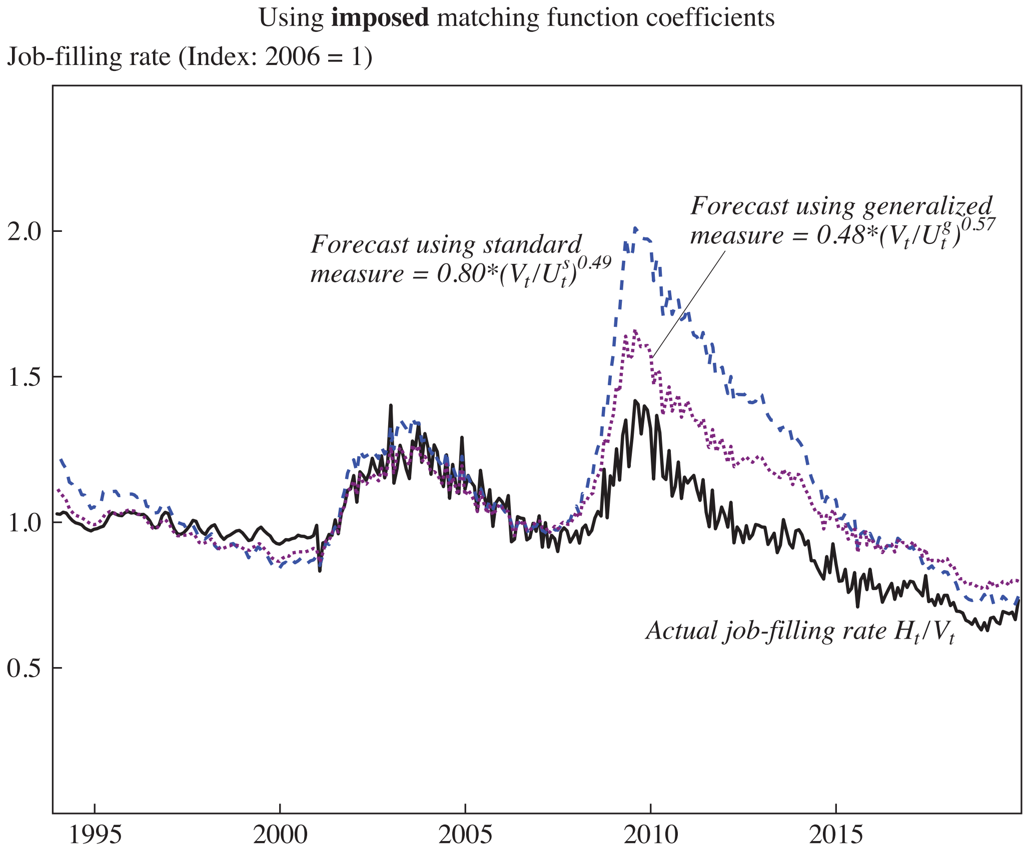

Source: Author's calculations, from data in the paper.

measure is further from the actual outcomes than the forecast generated using the generalized measure. This difference is the basis of the claim that the standard measure is outperformed by the authors' preferred generalized measure.

But eyeballing this figure reveals something more troubling: both lines look like problematic forecasts. Both are (substantially) higher on average than the actual job-filling rate. And both forecasts rise and fall substantially more over the business cycle than the actual job-filling rate. Neither forecast appears to fit the data well.

This is easy to confirm using standard forecast evaluation methods. In particular, figure 2 shows scatter plots of the actual job-filling rate against each of these forecasts. Neither fits the data well. Figure 2 also reports the results of regressions of the following form:

$$Actual_t = a + b \times Forecast_t + \epsilon_t.$$

Figure 2. Forecast Evaluation

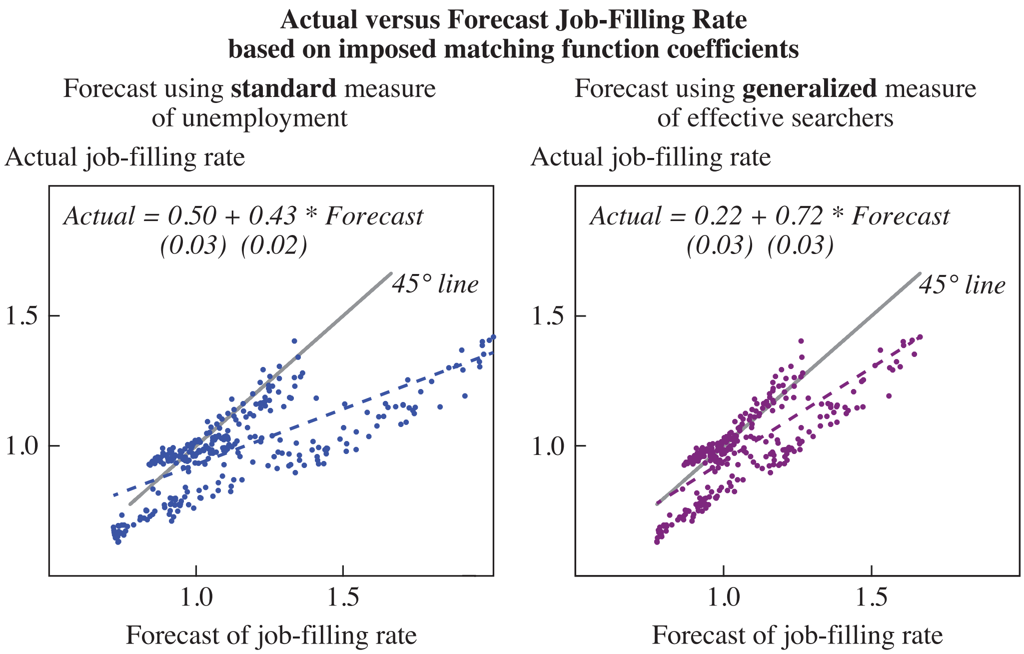

Source: Author's calculations.

An unbiased and efficient forecast would yield a constant term of zero and a coefficient on the forecast of one. (Failing this test implies that forecast errors are forecastable, and hence statistically inefficient.) The left panel shows that the forecast generated using the standard measure of unemployment fails both of these tests, and the right panel shows that the forecast generated using the generalized measure of effective searchers also fails the same test. (F-tests of the joint null that a = 0 and b = 1 are overwhelmingly rejected in both cases.)

These artifacts reflect the fact that the μ and α coefficients that were imposed appear to be inappropriate if one is interested in forecasting the job-filling rate. In addition, as we shall see in a moment, the differences between the two forecasts are largely attributable to these imposed coefficients. Correcting for these artifacts reveals that both the standard and the generalized measure yield very similar forecasts.

A more transparent (and arguably less arbitrary) approach to figuring out what values of μ and α to apply for this forecast evaluation exercise would be to simply estimate each of these parameters directly and then use those estimates to assess the accuracy of the ensuing forecasts. This alternative would be more consistent with the authors' claim that their analysis aims to estimate α in a way that would "give each set of matching function

Figure 3. Predicting the Job-Filling Rate Based on Estimated Matching Function Coefficients

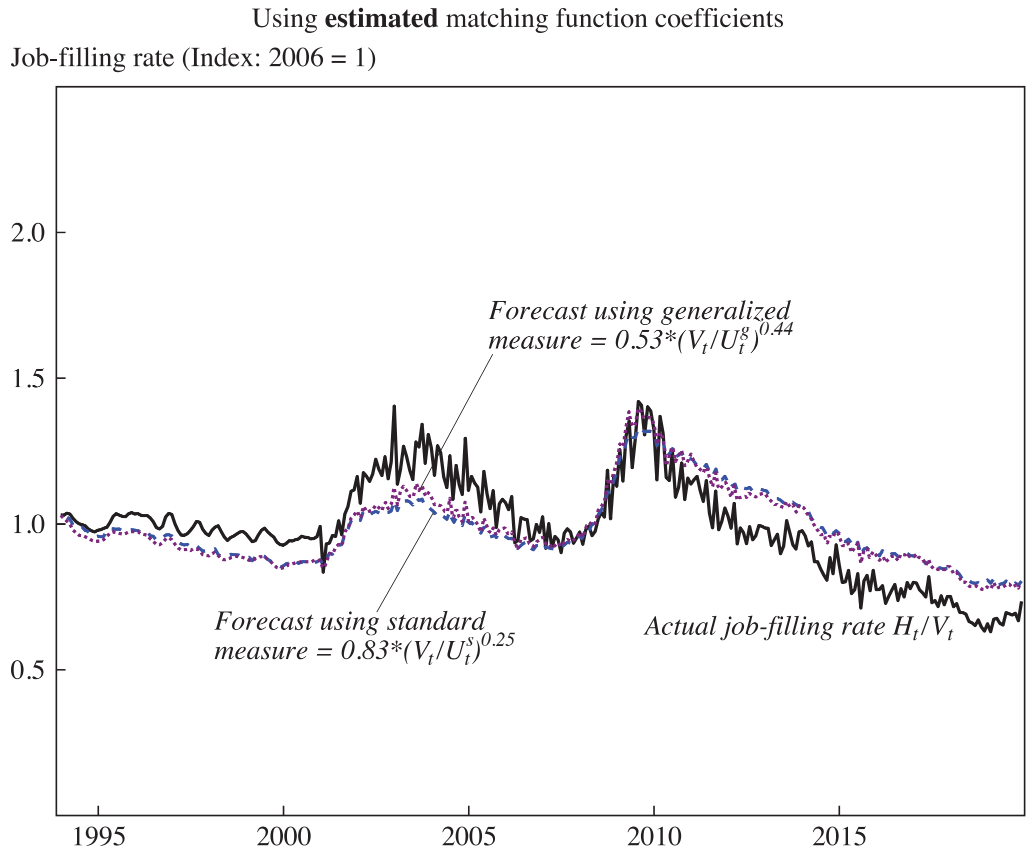

Source: Author's calculations.

arguments the best possible chance to fit the data well." As such, I ran the following simple nonlinear least squares regression:

$$\frac{H_t}{V_t} = \mu\left[\frac{V_t}{U_t}\right]^{-\alpha}.$$

Initially, I ran this regression using the standard measure of U_t^s (the number of unemployed people), which yielded estimates of $\widehat{\mu^s} = 0.83$ (se = 0.01) and $\widehat{\alpha^s} = 0.25$ (se = 0.01). Running this regression, but using the generalized measure of effective searchers U_t^s as the measure of worker search instead, yielded estimates of $\widehat{\mu^g} = 0.53$ (se = 0.014) and $\widehat{\alpha^g} = 0.44$ (se = 0.02). (Running the regression in logs: $\log(H_t) = \log(\mu) + \alpha\log(U_t) + (1-\alpha)\log(V_t)$ yielded relatively similar estimates of α in each case.)

Figure 3, which is drawn on the same scale as figure 1, plots each of the resulting forecasts of the job-filling rate, along with the actual job-filling

Figure 4. Forecast Evaluation Based on New (Estimated) Forecasts

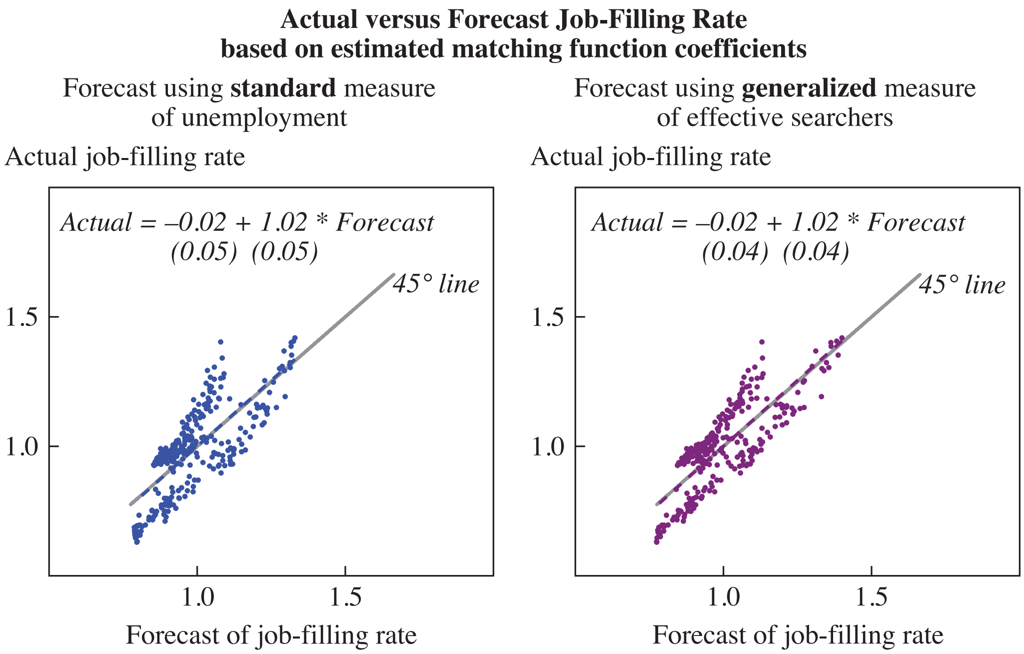

Source: Author's calculations.

rate. This yields a dramatically different picture than figure 1 (or figure 5 in the paper under discussion). Both sets of forecasts of the job-filling rate are now much more accurate.

Most importantly, the dashed and the dotted lines are now almost identical. Once one abandons the arbitrarily imposed matching function coefficients, the forecast for the job-filling rate generated using the standard measure of unemployment is almost identical to that forecast using the generalized measure of effective searchers. Indeed, the correlation between these two series is 0.991.

This analysis yields an interpretation of the paper's findings that is largely at odds with the interpretation offered by the authors. It suggests that any differences in the performance of these measures is not due to richer information embedded in the generalized measure of effective searchers but rather is due to the specific coefficients imposed on the matching function in generating forecasts. As such, the authors' claim that "the relatively poor performance of the standard model cannot be rescued with an alternative estimate of the matching function elasticity" appears to be wrong.

Not surprisingly, estimating the coefficients that we use to generate forecasts of the job-filling rate yields better behaved forecasts. Figure 4 plots each of our new forecasts of the job-filling rate against the actual

Table 1. Formal Forecast Evaluation

	Standard deviation	*Correlation with actual*	*RMSE*	*Ratio of RMSE to standard*
Panel A: Generating forecasts using imposed coefficients (from table 3 of the paper)				
Actual job-filling rate	0.17	1.00	—	—
Forecast using standard measure of unemployment	0.29	0.75	0.25	1.00
Forecast using generalized measure of search activity	0.19	0.81	0.13	0.54
Panel B: Generating forecasts using estimated coefficients (preferred alternative)				
Actual job-filling rate	0.17	1.00		
Forecast using standard measure of unemployment	0.12	0.76	0.11	1.00
Forecast using generalized measure of search activity	0.13	0.81	0.10	0.90

Source: Author's calculations.

job-filling rate based on the coefficients estimated above. Forecasts generated using either the standard or the generalized measure of labor market slack are unbiased, and they both rise and fall in proportion with the business cycle. Of course, this occurs largely because these forecasts are generated using coefficients that were estimated with an eye to closely matching the actual outcomes. (In the regression of actual outcomes on forecasts, while the constant is close to zero, it is not precisely zero—and the slope is close to one, but not precisely one—because these forecasts were generated from an equation that is not linear in the parameters.)

At this point, it is worth revisiting the authors' evaluation of the relative performance of each measure. Table 1 presents a formal forecast evaluation. Panel A replicates the findings the authors present in their table 3, showing that using the imposed matching function coefficients leads to the conclusion that the root-mean-square error (RMSE) of forecasts generated using the standard measure is much higher (that is, worse) than that of the forecasts generated using the generalized measure of search activity. Indeed, the authors' preferred measure yields forecast errors that tend to be nearly half that when using the standard measure. A Diebold-Mariano test reveals that this difference in RMSE is (highly!) statistically significant.

But this finding no longer survives when one is no longer willing to impose the specific values that the authors impose on the matching function coefficients. To see this, panel B has a parallel structure to panel A, but it reports on the accuracy of the forecasts generated using each measure of labor market slack, but using the matching function coefficients that I estimated above. The fit of both sets of forecasts—generated using either measure of labor market slack—is now much better. Moreover, the performance of the forecasts generated using the competing measures of slack—whether measured as the correlation between the forecast and the actual outcome or as the RMSE—are now remarkably similar. The RMSE of the two alternative forecasts evaluated in panel B are no longer economically meaningful, and a Diebold-Mariano test fails to reject the null that they're equal. All told, the evidence in panel B is that there's little to recommend one measure over the other. In turn, this reflects the reality that the alternative forecasts are remarkably similar.

THE REMARKABLE SIMILARITY OF ALTERNATIVE MEASURES OF SLACK The most striking finding so far is that using the estimated matching function coefficients yields forecasts of the job-filling rate that are almost identical whether using the standard measure of unemployment or the generalized measure of effective searchers as an input. It's worth pausing a bit to see where this comes from.

The left panel of figure 5 plots the forecast of the job-filling rate generated using the generalized measure of effective searchers versus the forecast generated using the standard measure of unemployment. This panel uses my estimated matching function coefficients, and as reported above, it reveals a correlation between these measures of 0.991. The middle panel turns to a similar plot, but this time the forecasts are based on the matching function coefficients imposed by the authors. The correlation remains very high. This suggests that the high correlation of the forecasts generated by the two competing measures of labor market slack is not driven by my preference for estimating the matching function coefficients rather than imposing them, as this high correlation between these two sets of forecasts is a feature even in the authors' preferred measures. The third panel shows the source of this similarity, plotting the basic measure of labor market tightness, V_t/U_t, when estimated using the two alternative measures of U_t. The measure of slack based on the generalized measure of effective searchers, V_t/U_t^s, is remarkably similar to that generated using the standard measure of unemployment, V_t/U_t^s.

This, in turn, implies that the differences in the forecasts of these two measures aren't driven by one measure being more informative than the other. Both rise and fall almost in tandem. Rather, they are driven by difference in the scale and location of the two measures of labor market slack. Indeed, this is the key to reconciling the evidence in figure 1 which shows that the authors' forecasts of the job-filling rate based on their generalized

Figure 5. Comparing Generalized and Standard Measures

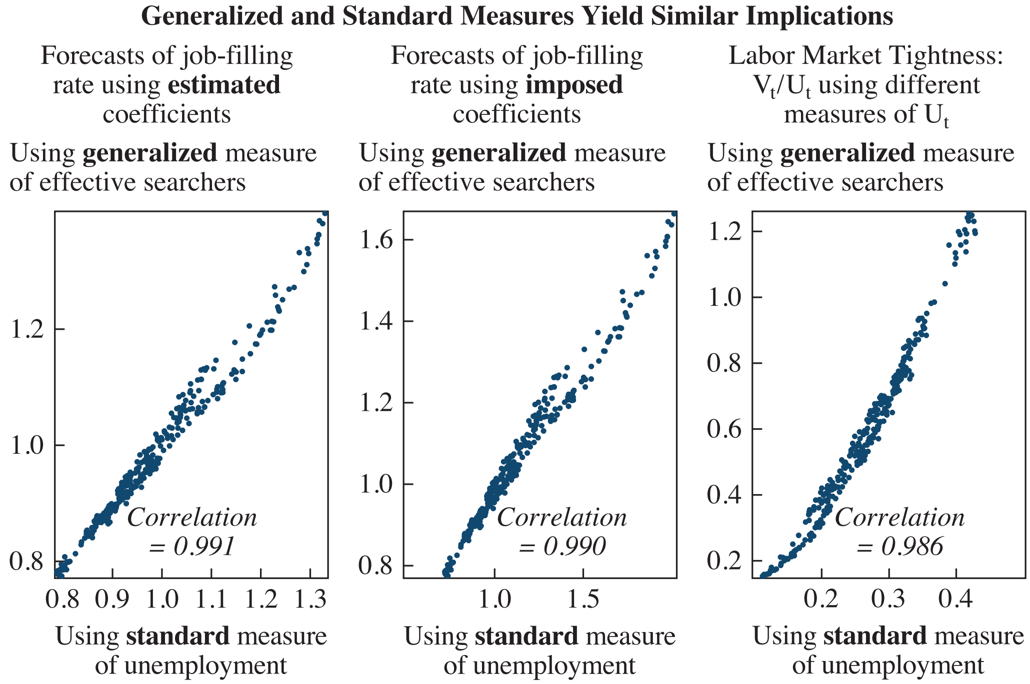

Source: Author's calculations.

measure of effective searchers is quite different from their forecasts based on the standard measure of unemployment, with the evidence in figure 5 that these two sets of forecasts are still very highly correlated. As figure 1 shows, both measures rise and fall in lockstep, but using the coefficients imposed by the authors yields a series of forecasts of the job-filling rate based on the standard measure that are typically located higher than those based on the generalized measure, and they are scaled so that they are more variable.

The location and scale of each series is determined by the coefficients α and μ. Thus, nearly *all* of the difference in the forecasts generated by these series reflects the choices of these auxiliary parameters. Remarkably little of the difference is due to the generalized measure bringing extra information about the state of the labor market.

Figure 6 illustrates the role that the matching function coefficients play in determining the scale and location of each series. It should be read in concert with the right panel of figure 5, which illustrates that the inputs into these forecasts are remarkably highly correlated. Yet as the left panel of figure 6 shows, the matching function coefficients imposed by the authors yields a series of forecasts based on the standard measure of unemployment

Figure 6. Scale and Location of Alternative Forecasts of the Job-Filling Rate

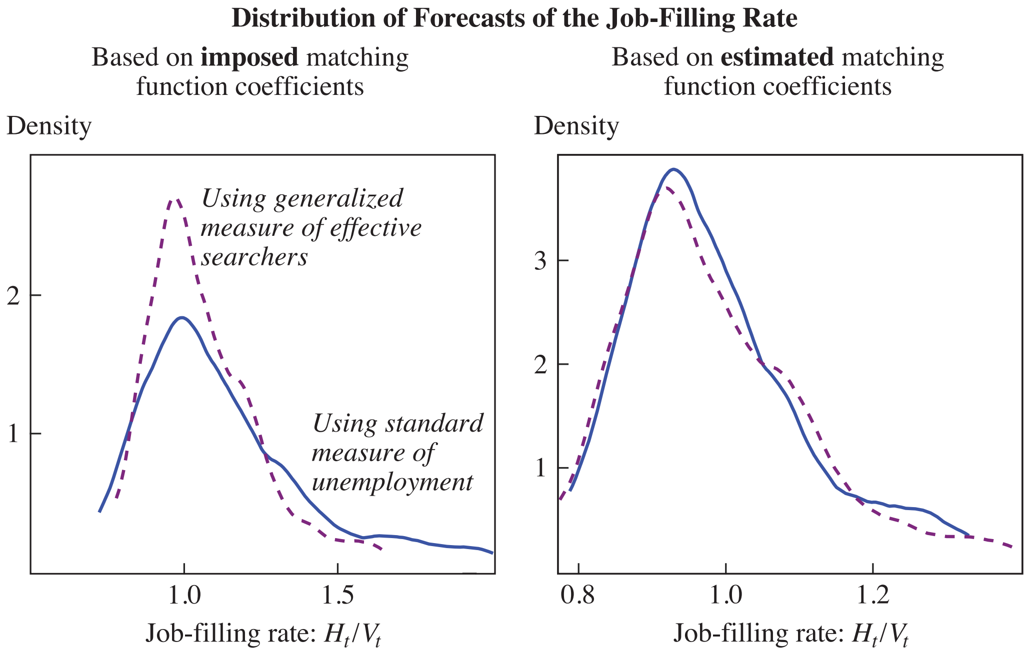

Source: Author's calculations.

that is both larger on average and much more variable than the forecasts based on the generalized measure of effective searchers. The right panel shows that these differences in scale and location are no longer meaningful when using estimated matching function coefficients.

DO THE IMPOSED COEFFICIENTS MAKE SENSE? At this point, it is worth assessing whether there is any reason to prefer the tests based on the imposed matching function coefficients presented in the paper versus the estimated coefficients emphasized here. On this score, there are two key points to make, one methodological and the other quantitative.

First, the methodological point yields a prima facie case that the matching function coefficients imposed in the paper are not appropriate for this sort of analysis. The exercise the paper performs is effectively a forecast evaluation exercise. But the source of the imposed α's comes from an instrumental variables regression, designed to recover the "structural" matching function coefficient. While structural coefficients are useful for constructing policy counterfactuals, in the presence of measurement error (and this paper is motivated by the idea that there is measurement error in our measures of labor market slack) a structural coefficient will not generally yield an efficient or unbiased forecast. (Recall, while ordinary least squares may yield biased coefficients, those coefficients are still the best linear unbiased estimates.) Indeed, figure 3 illustrates that the imposed coefficients yield forecasts of the job-filling rate that are both biased and inefficient. Even if policymakers were to rely on the generalized measure of effective searchers, it is hard to believe that they would rely on these imposed coefficients to evaluate the likely implications for the job-finding rate.

To further muddy the methodological waters, the μ coefficients that the authors impose are neither structural nor reduced form, but the result of an arbitrary normalization in which it is essentially assumed that the unemployment and vacancy data for 2006 determine the location of the Beveridge curve. As far as I can tell, the authors never defend this choice, but it turns out not to be an innocuous normalization, as it also helps determine how strongly the job-filling rate responds to changes in labor market tightness.

Second, there is a quantitative point to be made. The authors are correct to argue that they need to use different estimates of α and μ when generating forecasts based on the two different measures of slack. Different estimates are needed because the generalized measure of effective searchers is scaled in a way that does not permit direct comparison to the standard measure, which is the level of unemployment. (The generalized measure is the sum over twenty-two population groups of the product of their share in the population at that point in time and their relative job-finding rate in 2006. Its precise scale depends on the normalization of these job-finding rates, or γ_i, in the authors' terminology.) As such, it makes sense to use different estimates of α_s and α_g to generate forecasts based on either the standard or generalized measures of labor market tightness.

But while the precise numeric values of α_s and α_g used to generate each forecast may differ, for this to be an apples-to-apples comparison, they should have a similar *economic* interpretation. They should each predict a similar response of the job-filling rate (H_t/V_t) to business cycle fluctuations in labor market tightness ($\theta_t = V_t/U_t$). After all, only by ensuring that both comparisons embed a similar sensitivity to the state of the business cycle can we ensure that any observed differences are due to the extra information embedded in one measure rather than the other.

To assess this, start by taking the first derivative of the earlier job-filling rate equation to obtain:

$$\frac{d\frac{H}{V}}{d\theta} = -\alpha\mu\theta^{-(\alpha+1)}.$$

Next, recall that the measure of labor market tightness based on the standard measure of unemployment ($\theta_t^s = V_t/U_t^s$) is scaled differently from that based on the generalized measure of effective searchers ($\theta_t^g = V_t/U_t^g$). To be precise, the former has a mean of 0.58 and a standard deviation of 0.26, while the latter has a mean of 0.26 and a standard deviation of 0.71. So rather than comparing the responsiveness of the job-filling rate to a one-unit change in each measure, it would be more of an apples-to-apples comparison to evaluate how the job-filling rate responds to a one standard deviation change in each measure of labor market tightness, σ_θ. I call this the "cyclical sensitivity of the job-filling rate" and evaluate it at the mean level of tightness, $\bar{\theta}$, as follows:

$$\textit{Cyclical sensitivity of job-filling rate} = \sigma_\theta \times \left.\frac{d\frac{H}{V}}{d\theta}\right|_{\theta=\bar{\theta}} = -\sigma_\theta \alpha \bar{\theta}^{-(\alpha+1)}.$$

Table 2 shows how I calculate the cyclical sensitivity of the job-filling rate. The first four columns show the inputs into my calculations, and the final column shows the calculated sensitivity.

Importantly, notice that in panel A, the estimated sensitivity of the forecasts generated using the standard measure of unemployment is quite different from that calculated using the generalized measure of search activity. The point is that the relationship between job-filling, workers' search, and firm search used to compare these two measures is quite different in each case. And that in turn is the sense in which the comparisons calculated by the authors are not apples-to-apples comparisons.

By contrast, the panel B shifts the analysis to using the estimated matching function coefficients that I focus on in this comment. While the precise α's are quite different (that is α_s is different from α_g—reflecting differences in the scaling of the two competing measures of worker search, their implications), if instead we evaluate their consequences in terms of the effect of a one standard deviation change in search, the effects are quite similar. This suggests that the comparison based on the estimated coefficients is more of an apples-to-apples comparison. That comparison, in turn, found that the forecast of the job-filling rate based on the generalized measure of search effectiveness was almost identical to the forecast based on the standard measure of unemployment.

Finally, I should add that I'm not quite sure why the process by which the authors arrived at the α coefficients they impose are so different when

Table 2. Evaluation of the Cyclical Sensitivity of the Job-Filling Rate

	α *coefficient*	μ *coefficient*	*Average tightness* $\bar{\theta} = \overline{(V/U)}$	*Standard deviation of tightness* σ_θ	*Cyclical sensitivity of job-filling rate* $= -\sigma_\theta \alpha \bar{\theta}^{-(\alpha+1)}$
Panel A: Generating forecasts using imposed coefficients					
Using standard measure of unemployment	0.49	0.80	0.58	0.26	–0.23
Using generalized measure of search activity	0.57	0.48	0.26	0.71	–0.16
Panel B: Generating forecasts using estimated coefficients					
Using standard measure of unemployment	0.25	0.83	0.58	0.26	–0.10
Using generalized measure of search activity	0.44	0.53	0.26	0.71	–0.11

Source: Author's calculations.

using a measure of worker search based on the standard versus the generalized measure. The estimation is relatively opaque, and it is surely worth exploring why highly correlated series yield such different estimates of the cyclical sensitivity of the job-filling rate.

CONCLUSION This comment has dived pretty deep into the weeds, so it's worth panning back to the big picture. Abraham, Haltiwanger, and Rendell propose a new measure of labor market search that at a conceptual level has a lot to recommend it. The reality, however, is that it yields a measure that is remarkably highly correlated with a more standard measure like unemployment. The particular empirical exercise carried out in the paper asks which measure does a better job at predicting the job-filling rate. (It also asks which does a better job in matching the job-finding rate. For the sake of space, I've not dug into that measure here, but similar conceptual issues arise in analyzing competing predictions of that measure.)

Of course, one needs not just a measure of labor market slack but also a model if one is to predict the job-filling rate. And any model brings with it a set of auxiliary assumptions This comment has argued that these auxiliary assumptions are responsible for much of the difference in the forecasts generated by the two competing measures of labor market slack. My alternative set of tests are based on a different—and I would argue simpler, more transparent, more statistically coherent, and economically meaningful—set of auxiliary assumptions. And my alternative tests reveal that both measures of labor market slack yield almost identical forecasts of the job-filling rate. As such, there is little evidence that the new generalized measure of active searchers includes much useful information that's not already embedded in a more standard measure like unemployment. The new measure is neither meaningfully better nor worse than a more standard alternative.

GENERAL DISCUSSION Olivier Blanchard noted that the authors clearly made progress on the matching function but he was struck by the lack of discussion around the Phillips curve. He also commented that the analysis directly speaks to the paper in this volume by Del Negro, Lenza, Primiceri, and Tambalotti, which finds that wages are less responsive to measures of slack.[1] However, Abraham, Haltiwanger, and Rendell's paper suggests that slack is likely mismeasured. Blanchard continued by noting the importance of capturing the heterogeneity of workers in the matching function. It may also be that wage bargaining associated with these different groups of searchers is not the same, resulting in a more complex relationship between slack and wage setting.

Giorgio Primiceri commented that to understand if this new measure of slack performs better with the standard Phillips curve correlation, one needs to examine data from the 1960s to the 1980s. Primiceri continued by noting that Ayşegül Şahin's discussion showed that the movement of the authors' new measure of slack across the business cycle is likely very similar before and after the 1990s. So it is unlikely that this new measure of slack would solve the puzzle of the post-1990s Phillips curve correlation.

James Stock said he is unsure about the paper's argument that the Beveridge curve is more stable and not shifting up with this new measure of slack. While roughly vertical in the way it is graphed, it is not obvious whether it has less instability and appears to shift to the right during the relevant period.

Ricardo Reis wondered if the authors could relate their paper to the Stansbury and Summers paper in this volume, which finds a decline in worker bargaining power.[2] In a strict Diamond-Mortensen-Pissarides model,

1. Marco Del Negro, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti, "What's Up with the Phillips Curve?," *Brookings Papers on Economic Activity*, Spring (2020).

2. Anna Stansbury and Lawrence H. Summers, "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy," *Brookings Papers on Economic Activity*, Spring (2020).

one would think of the Beveridge curve separately from the bargaining of the surplus.[3] However, in a McCall model of the labor market where a change in the bargaining power of workers can impact their search decisions, one may find, as a result, a different relationship between the amount of vacancies and the amount of job seekers.[4]

Katharine Abraham and John Haltiwanger made several points in response to Justin Wolfer's discussion of the relative performance of forecasts based on generalized versus standard measures of labor market tightness. (Further details regarding the authors' response to these comments can be found in online appendix D.) Abraham and Haltiwanger observe, first, that Wolfers's analysis does not recognize the trade-off in how well the standard model tracks the job-filling rate versus how well it tracks the job-finding rate among the unemployed. Using Wolfers's ordinary least squares (OLS) estimates of the matching function elasticity, the standard model performs almost as well as the generalized model with regard to the job-filling rate. But, using those same elasticity estimates, the relative performance of the standard model for tracking the job-finding rate among the unemployed is significantly *worse*—applying Wolfers's OLS estimates to tracking the job-finding rate among the unemployed yields an RMSE for the standard model that is three times as large as that for the generalized model.

Another issue Abraham and Haltiwanger note with Wolfers's analysis is that it uses the same data to estimate model parameters and then evaluates their performance through the lens of a forecasting perspective. As discussed in more detail in online appendix D, when Wolfers's approach is modified so that models are fit with data through 2007 and then used to forecast outcomes in later years, the result is a generalized matching function that significantly outperforms the standard matching function both for predicting the job-filling rate and for predicting the job-finding rate among the unemployed. Finally, Abraham and Haltiwanger argue, the OLS approach Wolfers suggests for estimating the elasticity of the matching function yields estimates that are inherently biased. This is not true of the instrumental variables (IV) estimates developed in the paper.

3. Peter A. Diamond, "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies* 49, no. 2 (1982): 217–27; Dale T. Mortensen, "The Matching Process as a Noncooperative Bargaining Game," in *The Economics of Information and Uncertainty*, ed. J. J. McCall (Chicago: University of Chicago Press, 1982); Christopher A. Pissarides, "Short-Run Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review* 75, no. 4 (1985): 676–90.

4. John Joseph McCall. "Economics of Information and Job Search," *Quarterly Journal of Economics* 84, no. 1 (1970): 113–26.

Using the IV estimates, the generalized model again significantly outperforms the standard model in tracking both the job-filling and the job-finding rates.

Haltiwanger addressed Şahin's and Blanchard's comments noting no incorporation of the Phillips curve in the analysis. Haltiwanger said that the paper produced an alternative measure to better capture slack in the labor market. Hopefully, there is consensus that the standard matching function is mis-specified. According to the data, a large portion of job seekers come out of the labor force and from among the employed, which is not captured by the standard matching function. Given this fact, their paper builds an alternative generalized matching function where the heterogeneity of job seekers is captured. Haltiwanger continued by saying that they are sympathetic to working out the details around their new measure and taking the next step to incorporate it into the Phillips curve framework. The new measure may have important implications for wage and price dynamics that should be explored in the future.

Haltiwanger responded to Reis's question by saying that if one were to pursue a Diamond-Mortensen-Pissarides model in the context of the Phillips curve, one would need to take a generalized approach like they do in the paper.

Haltiwanger continued by addressing Stock's question and part of Şahin's comment around normalizing measures. The authors normalized their measures because they think that the intercept of the hiring function should be different across specifications. The standard model should not have the same intercept as the generalized model. The authors are interested in removing those intercept differences, so they normalized their measures.

Haltiwanger then turned to another part of Şahin's comment around the number of subgroups needed to calculate a generalized matching function. He acknowledged that it's not clear whether one needs twenty-two groups for the sample period covered by the paper. The most important subgroups in this sample period are decomposing the unemployed into short-term and long-term unemployed and including and decomposing the group of those out of the labor force into want a job and other. However, over a longer sample period (e.g., including the entire post-WWII period), it is likely that additional subcategories matter. In considering this issue, it is noteworthy that there are large differences in relative search intensities across the detailed groups but for this to matter there also need to be changes in composition across the groups over time.