

Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency

Fabian Lange* and Theodore Papageorgiou^{†‡}

April 14, 2020

Exploiting results from the literature on non-parametric identification, we make three methodological contributions to the empirical literature estimating the matching function, commonly used to map unemployment and vacancies into hires. First, we show how to non-parametrically identify the matching function. Second, we estimate the matching function allowing for unobserved matching efficacy, without imposing the usual independence assumption between matching efficiency and search on either side of the labor market. Third, we allow for multiple types of jobseekers and consider an “augmented” Beveridge curve that includes them. Our estimated elasticity of hires with respect to vacancies is procyclical and varies between 0.15 and 0.3. This is substantially lower than common estimates suggesting that a significant bias stems from the commonly-used independence assumption. Moreover, variation in match efficiency accounts for much of the decline in hires during the Great Recession.

*McGill University, Department of Economics, 855 Sherbrooke St West, Room 511, Montreal, QC, H3A 2T7, Canada, fabian.lange@mcgill.ca

[†]Boston College, Department of Economics, Maloney Hall 343, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

[‡]We are thankful to Jesper Bagger, Alan Collard-Wexler, Steve Davis, Mike Elsby, Jason Faberman, and Andreas I. Mueller for useful comments. Codes for estimating matching functions following the procedure described in this paper can be found on the authors’ websites, fabianlange.ca and theodorepapageorgiou.com

1 Introduction

The matching function explains hiring as a function of search effort that job seekers and recruiters exert in the labor market. It is central in the literature on frictional labor markets where it serves as a modeling device to capture a costly trading process (Pissarides, 2000). The results of a substantial empirical literature investigating frictional labor markets, whether it focuses on cyclical fluctuations or on the (re-)allocation of workers across firms, occupations, industries or locations, depend on the specifics of the matching function. For instance, its properties determine whether search effort expended by job seekers and recruiters are constrained efficient.¹ Motivated by these observations, a sizeable literature has focused on estimating matching functions (see Petrongolo and Pissarides (2001) for an early survey).

Two major problems beset the literature on estimating the matching function. First, it is difficult to measure search effort. To start, it is not clear who exactly is searching. Traditionally search is proxied by unemployment and vacancy counts, but these proxies are incomplete. The pool of job seekers is not limited to the unemployed, but includes many that are currently out of the labor force (OLF) or employed but searching on-the-job. In fact, more than half of all transitions from non-employment to employment in the Current Population Survey (CPS) come from OLF rather than unemployment.² Moreover, while proxies of search effort of job seekers and recruiters do exist, it is unlikely that they fully capture variation in search effort or recruiting intensity.³

Compounding these difficulties is that both observed and unobserved search are likely to respond endogenously to labor market conditions. For instance, an unemployed worker may vary her search effort when the job finding rate changes (Hornstein and Kudlyak, 2016). Similarly, resources invested into hiring might well vary with search conditions (Davis et al., 2013).⁴ Just as unobserved search effort might be endogenous, the number of unemployed workers or vacancies might also respond to matching efficiency. Not accounting for these endogenous responses will bias estimates of the elasticity of the matching function with respect to search by job seekers and recruiters.⁵

¹See for instance the Hosios (1990) condition that governs efficiency in a large class of labor market models.

²Kudlyak and Lange (2017) and Elsby et al. (2015a) emphasize how important search among the OLF is. Similarly, Fallick and Fleischman (2004) and Fujita et al. (2019) demonstrate the importance of employment-to-employment flows in the dynamics of the labor market.

³Notably, Mukoyama et al. (2018) use time-use data combined with data on the search methods of individuals in the CPS to construct a proxy for individual search effort. Davis et al. (2013) provide a proxy for recruiting intensity. In our empirical analysis, we make use of both of these proxies to separate variation in search effort and recruiting intensity from variation in the efficiency of the matching function over time.

⁴As Stigler (1961) and Shimer (2004) point out, whether the hiring rate increases or reduces the incentives for search effort is ambiguous. See also chapter 5 of Pissarides (2000) and Gomme and Lkhagvasuren (2015).

⁵See also Borowczyk-Martins et al. (2013) and Hornstein and Kudlyak (2016) for recent attempts at addressing some of the problems that arise due to endogenous search. For direct evidence on the cyclical nature of search effort using data from American Time Use Survey (ATUS) and CPS see Mukoyama et al. (2018).

The second problem is that most of the literature relies on strong functional form assumptions that are neither empirically, nor theoretically grounded. The common practice is to assume that the matching function takes the Cobb-Douglas form. However, there is little beyond convenience that is offered in support of this assumption. This assumption is important for both normative and positive reasons. On the normative side, Hosios (1990) shows that how the elasticity of the matching function relates to wage bargaining determines whether a search equilibrium is constrained efficient.⁶ The Cobb-Douglas functional form imposes the matching elasticity to be constant over the entire support and thus restricts our ability to study the efficiency of the matching equilibrium over the cycle. On the positive side, the efficacy of policies to stimulate vacancy creation will vary with the matching elasticity. If the Cobb-Douglas specification is overly restrictive, then this will lead to biased estimates of the efficacy of such policies under different labor market conditions.

This paper addresses both the problem of unobserved, potentially endogenous search effort and the problem of overly restrictive functional forms imposed on the matching function. To address the former problem, we first incorporate rich measures of search effort and recruiting based on the work by Mukoyama et al. (2018) and Davis et al. (2013) and we allow for multiple types of job seekers. Beyond this, we use our framework to estimate unobserved search effort, which in our setup is captured by variation in matching efficiency. To address the latter problem, we relax the functional form restrictions on the matching function and instead non-parametrically identify how matches depend on total search effort among job seekers and recruiters.

Using our methodology we are able to quantify shocks that have shifted the Beveridge curve. The high unemployment rate that persisted well after the end of the Great Recession has troubled both economists and policymakers.⁷ Using our approach we consider an “augmented” Beveridge curve that takes into account not only the stocks of unemployed workers and vacant firms, but also search by employed workers, as well as those out of the labor force, recruiting effort by firms and time-varying (observed) search effort by workers. Our methodology allows us to quantify the shifts in this curve, i.e. changes in unobserved matching efficiency. In that regard, our results extend a literature starting from Blanchard and Diamond (1989) who first measured these types of shocks, albeit with a very different methodology. Elsby et al. (2015b) offer a recent review of the related literature.

Our empirical application using US data from 2001 through 2017 finds that aggregate search effort has

⁶Hosios (1990) focuses on markets with homogeneous agents. However the same condition is necessary (but not sufficient) in the case of (one-sided) heterogeneity (see Brancaccio et al., 2020b).

⁷See for instance discussion in Sahin et al. (2014) and references therein.

not changed substantially, but this masks a substantial decline in unobserved search effort or matching efficiency, conditional on observed search. In fact most of the decline in hires during the Great Recession is driven by the decline in matching efficiency. In addition, we estimate an elasticity of the matching function with respect to vacancies that ranges between 0.15 and 0.3. This elasticity varies substantially over time and is negatively correlated with labor market tightness. Our estimates differ from those obtained from the commonly used Cobb-Douglas specification that implicitly imposes that matching efficiency is independent from market tightness. However since matching efficiency is strongly procyclical this induces a positive bias in the estimates of the vacancy elasticity when matching efficiency is not controlled for resulting in estimates that are 3 times larger than the one produced from our methodology.

Of course, these gains in estimating unobserved match efficiency and relaxing functional form assumptions are not for free. Our approach relies on two assumptions. First, we require that the matching function exhibits constant returns to scale (or that the returns to scale are known to us). This is a strong assumption, yet one that many researchers impose as a matter of course when modeling the labor market. Our contribution is to show how to more fully exploit this assumption to identify the matching function. Second, we assume that match efficiency and vacancies are independent, *conditional* on observed worker search. Given these two restrictions, our model is overidentified. We can therefore test the independence assumption.

We proceed as follows. In the next section, we use the Cobb-Douglas function to illustrate the identification problems that need to be confronted when estimating the matching function and review the recent literature on estimating the matching function. Section 3 contains our main identification result based on Matzkin (2003) for a general class of functions. We show that, assuming that unobserved search effort among one group of workers is independent of vacancies, conditional on observed worker search, it is possible to recover the distribution of unobserved search effort(s) and the matching function non-parametrically up to a normalization. Our approach here is related to Brancaccio et al. (2020a) who also rely on Matzkin’s identification proofs to show how to estimate a matching function in a trade model with matching between ships and exporters.⁸ Our contribution is to extend these results to the context of matching job seekers to vacancies. This is discussed in Section 4, where we specify the structure that we employ in our empirical work, which allows for multiple types of job searchers.⁹ Section 5 describes the data from Job Openings and Labor Turnover Survey (JOLTS) and from the CPS that we use in our

⁸See also Bajari and Benkard (2005) for an application of the identification results from Matzkin (2003) in demand estimation.

⁹We follow the literature in treating search of different individuals as perfect substitutes.

empirical implementation. Section 6 describes the details of the empirical implementation, as well as the results. The Appendix contains additional results and figures, including an extension of our setup to allow for dependence between vacancies and the unobserved component of search effort of our reference group of workers and how it can be estimated.

Before we begin, a few words on nomenclature and notation. At times, we will use the terms unobserved search effort as well as matching efficiency interchangeably (and we have done so above). Both manifest themselves as variation in the rate at which matches are formed conditional on observed measures of search. Unobserved search effort emphasizes that some of this variation is due to endogenous choices on the part of job seekers. The term “match efficacy” indicates that variation in the rate of matching job seekers to vacancies can also arise for reasons that are not commonly thought of as effort. For instance, match efficacy includes variation due to the technology of matching job seekers to vacancies, as well as variation that arises because of mismatch between job seekers and available vacancies.

Throughout the paper, we will use unsubscripted letters (H, U, V, A) to refer to random variables representing hires, unemployed, vacancies, and match efficacy. Variables subscripted by t are realizations in period t . We use small caps (h_t, u_t, v_t, a_t) to denote the logarithms of (H_t, U_t, V_t, A_t). We denote distributions of random variables that are observed directly in the data using $G(\cdot)$. For distributions that involve unobserved random variables we use $F(\cdot)$.

2 The Cobb-Douglas Matching Function

In this Section, we fix ideas imposing the Cobb-Douglas functional form on the matching function $m_t(\cdot)$ that maps period- t unemployed U_t , per-capita search efficacy/matching efficiency of the unemployed A_t , and vacancies V_t into hires H_t .

$$H_t = m_t(A_t U_t, V_t) = (A_t U_t)^{1-\gamma} V_t^\gamma \tag{1}$$

The objective of the researcher is to estimate γ using data on (H_t, U_t, V_t) , while search efficacy A_t is unobserved. The variation used for identification is across time t , but the data could also come from multiple markets or markets interacted with time. For now, we assume that the underlying data generating process is stationary and that we observe a long enough time-series so that we can treat the joint distribution $G : \mathbb{R}_+^3 \rightarrow [0, 1]$ of (H_t, U_t, V_t) as observed. We also, for simplicity, assume that there is only one type of job seeker.

The Identification Problem

The identification problem arises because both observed and unobserved search effort depend on the rate at which job seekers and vacancies match.

The starting point for many an estimator summarized in Petrongolo and Pissarides (2001) is to divide equation (1) by U_t and take logs to obtain the log job finding rate $\lambda_{u,t} = \log\left(\frac{H_t}{U_t}\right)$ as a linear function of log market tightness $\theta_t = \log(V_t/U_t)$:

$$\lambda_{u,t} = \log\left(\frac{H_t}{U_t}\right) = (1 - \gamma)a_t + \gamma\theta_t, \quad (2)$$

where $a_t = \log A_t$, captures unobservable variation in match efficiency which might arise because of technological changes in the matching function or because unobserved search effort on the part of the unemployed varies over time.

Borowczyk-Martins et al. (2013) stress the identification problem that arises when U_t and V_t respond to variation in search efficiency a_t . When vacancy creation is governed by a zero profit condition, then (all else equal), periods of high a_t will also be periods in which V_t is large. High a_t might of course also encourage more unemployed to enter the market. Thus, θ_t is likely to correlate with a_t and naive estimators based on equation (2) will be biased.

Borowczyk-Martins et al. (2013) (and following their lead Barnichon (2012)) solve this problem by assuming that a_t follows a stochastic process with innovations ε_t that are uncorrelated with θ_{t-k} at some lag. Then, θ_{t-k} represent valid instruments for θ_t as long as tightness is correlated across time-periods. This approach solves the endogeneity problems that arise when observed measures of market tightness θ_t depend on a_t .

The approach championed by Borowczyk-Martins et al. (2013) however fails if a_t responds to conditions in the labor market, maybe because search effort expended by job seekers responds to how easy it is to find employment. To fix ideas, assume (for now) that a_t depends linearly on the log job finding rate $\lambda_{u,t}$:

$$a_t = \beta_0 + \beta_1\lambda_{u,t} \quad (3)$$

Substituting equation (3) into equation (2) and rearranging, we obtain

$$\lambda_{u,t} = \frac{1 - \gamma}{1 - (1 - \gamma)\beta_1}a_t + \frac{(1 - \gamma)}{1 - (1 - \gamma)\beta_1}\beta_0 + \frac{\gamma}{1 - (1 - \gamma)\beta_1}\theta_t$$

This expression shows that estimates of the matching elasticity γ using the relation between job finding rates and market tightness will be biased even if efficacy a_t is orthogonal to market tightness θ_t . An estimator based on an instrument for θ_t that is orthogonal to a_t will identify $\frac{\gamma}{1-(1-\gamma)\beta_1}$ rather than the search elasticity γ . If for instance search effort a_t increases with the job finding rate ($\beta_1 > 0$), then the estimated elasticity of the matching function with respect to vacancies exceeds the true elasticity. Unfortunately, it is equally plausible that search effort decreases with the job finding rate.¹⁰ Thus, estimates of $\frac{\gamma}{1-(1-\gamma)\beta_1}$ cannot serve as upper or lower bounds of the elasticity of the matching function with respect to market tightness.¹¹

In conclusion, no instrument can both simultaneously have power (affect θ_t) *and* be exogenous as long as match rates determine search effort of job seekers and recruiters. Any instrument that influences θ_t also affects match formation because it affects the effort market participants expend on finding matches.¹² What is there to be done? To attack this problem we devise an estimator that relies heavily on assuming that the matching function has constant returns to scale.¹³ We also assume that unobserved search efficacy A and vacancies V are independent, conditional on unemployment, U . We show that these assumptions identify both the matching function and the distribution of A . We propose an estimator that is constructive in the sense that it follows the steps of the identification proof.

At first glance, it might seem as if little has been gained since we still require A and V to be independent, conditional on U . However, our framework is overidentified so that we can test this independence assumption. Below, we first impose the independence assumption to estimate the matching function. Then, we use the estimated matching function together with data on (H_t, U_t, V_t) to back out estimates of unobserved search efficacy A_t . Importantly, these estimates are not by construction orthogonal to V_t so that we can use them to test whether A and V are indeed independent conditional on U . We therefore do not rely on untestable assumptions about the joint distribution of unobserved search effort and observed measures of market tightness. In Appendix A, we also show how to relax the independence assumption when estimating the matching function.

¹⁰It might increase, because the “yield” of effort increases, but on the other hand it might decrease because increases in the job finding rates imply that even a little effort will land a job (Shimer, 2004).

¹¹Of course, for some purposes $\frac{\gamma}{1-(1-\gamma)\beta_1}$ might well be the parameter relevant for policy.

¹²To our knowledge, Hornstein and Kudlyak (2016) are the only existing attempt to address this issue. Their work allows for multiple types of job seekers whose search effort endogenously responds to the type-specific match rate. Readers familiar with their paper will remember that they propose that the elasticity of the matching function and the search effort elasticity with respect to the matching efficiency are not separately identified. In our setup, as we discussed above, we do not differentiate between matching efficiency and unobserved search effort. Their result therefore does not apply to our case.

¹³We show below how to relax the assumption of constant returns to scale when the degree of returns to scale is known.

3 Identifying $m(A_t U_t, V_t)$ when A_t is unobserved

We now discuss how to identify the matching function as well as unobserved, time-varying matching efficiency, A .¹⁴ We assume that V and A are independent conditional on U and that the matching function $m(\cdot, \cdot) : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ has constant returns to scale. We do not impose additional functional form assumptions on $m(\cdot)$.¹⁵ This discussion is closely based on Matzkin (2003).

In this Section, we assume that there is a single type of vacancies and of unemployed job searchers. In Section 4 and in our empirical application, we relax this assumption and allow for multiple types of job seekers all of which might exert unobserved effort in job seeking.¹⁶

Proposition 1 states the main identification result, namely that the distribution $G(H, U, V)$ of hires, unemployment, and vacancies identifies the distribution of matching efficiency A as well as the function $m(\cdot)$.¹⁷

Proposition 1. *Assume that vacancies, V , are independent of matching efficiency, A , conditional on unemployment: $A \perp V|U$. Denote by $F(A, U)$ the joint distribution of A and U . Furthermore, assume that $m(AU, V) : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ is a continuous constant returns to scale function that is strictly increasing in AU . Then, the joint distribution $G(H, U, V)$ identifies $F(A, U)$ and $m(AU, V) : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ up to a normalization of A at one point of the support of (A, U, V) .*

Proof. Matching efficiency, A , does not have a natural unit. Thus, we normalize matching efficiency at (U_0, V_0) to A_0 .

Our starting point is the following condition

$$\begin{aligned}
 F(A|U) &= F(A|U, V) && \text{by independence of } A \perp V|U \\
 &= Pr(H \leq m(AU, V) | U, V) && \text{by monotonicity} \\
 &= G_{H|U, V}(H|U, V) && (4)
 \end{aligned}$$

The first equality holds because vacancies, V , are independent of matching efficiency, A , conditional on

¹⁴Recall our notational convention that unscripted capital letters (A, U, V) denote random variables while realizations are subscripted by t . Further, recall that distribution functions involving unobserved random variables are denoted by $F(\cdot)$, while those involving only observed random variables are denoted by $G(\cdot)$.

¹⁵Below, we show how one can relax the assumption that the matching function $m(\cdot)$ has constant returns to scale (see Corollary 3) as long as the returns to scale of the matching function is known.

¹⁶It is conceptually easy to apply the same approach to different types of vacancies. However, to implement this empirically requires data of vacancy-filling rates conditional on type of vacancy. Such data is rare.

¹⁷A corollary of proposition 1 is that we identify the joint distribution of U, V, A , and H . This follows because V is independent of A conditional on U . Thus, once we identify $F(A, U)$ we can obtain $F(A, U, V)$ using the marginal distribution of V . Together with the matching function $H = m(AU, V)$ we obtain the joint distribution of (U, V, A, H) .

unemployment, U , by assumption. The second equality holds, because the matching function is strictly increasing in its first argument, AU . In what follows we use this condition repeatedly.

Step 1: Obtain the distribution function $F(A|U)$ at point (A_0, U_0) :

$$F(A_0|U_0) = G_{H|U,V}(H_0|U_0, V_0)$$

where $G_{H|U,V}(\cdot|U_0, V_0)$ is the observed conditional distribution function of H given (U, V) evaluated at (U_0, V_0) .

Step 2: Identify $F(A_0|U)$ over the entire support of U . For this, consider an arbitrary scaling factor ψ .

Then

$$\begin{aligned} F(A_0|\psi U_0) &= F(A_0|\psi U_0, \psi V_0) &= Pr(H < m(\psi A_0 U_0, \psi V_0) | \psi U_0, \psi V_0) & (5) \\ &= Pr(H < \psi H_0 | \psi U_0, \psi V_0) &= G_{H|U,V}(\psi H_0 | \psi U_0, \psi V_0) \end{aligned}$$

Since $G_{H|U,V}$ is observed, we identify $F(A_0|\psi U_0)$ and as ψ varies, it traces out $F(A_0|U)$ over the entire support of U .

Step 3: Identify $F(A|U)$. Consider an arbitrary value of u in the support of U . For instance, U_0 will do fine. Employing again a scaling factor $\psi > 0$, we obtain for any arbitrary value of ψ :

$$\begin{aligned} F(\psi A_0|U_0) &= F(\psi A_0|U_0, \psi V_0) &= Pr(A < \psi A_0 | U_0, \psi V_0) \\ &= Pr(H < m(\psi A_0 U_0, \psi V_0) | U_0, \psi V_0) &= Pr(H < \psi m(A_0 U_0, V_0) | U_0, \psi V_0) & (6) \\ &= Pr(H < \psi H_0 | U_0, \psi V_0) &= G_{H|U,V}(\psi H_0 | U_0, \psi V_0) \end{aligned}$$

The first line again exploits the conditional independence assumption. The second line uses the monotonicity (strictly increasing) and homogeneity of $m(\cdot)$. Since $G_{H|U,V}(\cdot|U_0, \psi V_0)$ is observed, we identify the entire function $F(A|U_0)$ on the support of A by letting ψ vary across \mathbb{R}^+ . We can then repeat this procedure for all values in the support of U . Thus, steps 1-3 identify $F(A|U)$. Since $G(U)$, the distribution of U , is observed, we have identified the full joint distribution $F(A, U) = F(A|U)G(U)$.

Step 4: Identify the matching function, $m(\cdot, \cdot)$. For this, again we go back to equation (4):

$$G_{H|U,V}(m(AU, V)|U, V) = F(A|U) \tag{7}$$

Equation (7) immediately implies

$$m(AU, V) = G_{H|U,V}^{-1}(F(A|U)) \quad (8)$$

and since $G_{H|U,V}(\cdot)$ is observed and we already identified $F(A|U)$, we have thus identified m . In other words, we can back out the implied number of matches for any triplet (A, U, V) . \square

It is possible to provide some intuition for this identification result. Because of the independence assumption, match efficiency A_t does not systematically vary with vacancies V_t as long as U_t is held constant. Thus, variation in H_t as V_t varies and U_t is fixed identifies the elasticity of the matching function with respect to vacancies. Furthermore, because of constant returns to scale, the elasticity of $m(\cdot)$ with respect to V and AU respectively sum to one. Thus, by identifying the elasticity of the matching function with respect to V_t , we also obtain the elasticity of the matching function with respect to $A_t U_t$. Independence and constant returns to scale together thus imply that the function $m(\cdot)$ can be identified using observations of V_t and U_t only. With this function in hand, it is possible to identify the distribution of $A|U$ because conditional on (U_t, V_t) the distribution of hires maps one-to-one into the distribution of matching efficiency. This implies that the observed distribution of hires H conditional on U, V identifies the distribution of matching efficiency, A .

Since $m(\cdot)$ is monotone, we have that conditional on (U, V) , there is a one-to-one relationship between A and the realized number of hires H . This gives rise to the following corollary.

Corollary 2. *Under the conditions of Proposition 1, A_t is observed whenever (H_t, U_t, V_t) are observed.*

Proof. Immediate since $A_t = \frac{m^{-1}(H_t; V_t)}{U_t}$. Here $m^{-1}(H_t, V_t)$ denotes the inverse of $H_t = m(A_t U_t, V_t)$ with respect to the first argument. \square

To close this Section, we also establish that the matching function is identified when the matching function does not have constant returns to scale but the returns to scale are known. As long as the returns to scale are known, the intuition that identifying the elasticity of the matching function with respect to one argument identifies the elasticity of the other as well, goes through.

Corollary 3. *Assume that the matching function does not exhibit constant returns to scale, but can be represented by $\tilde{m}(AU, V) = z(m(AU, V))$ where $z : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a known monotone increasing function and $m : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$ exhibits constant returns to scale. Then, under the conditions of Proposition 1,*

the matching function \tilde{m} and the joint distribution $F(A, U)$ are identified and Corollary 2 holds (A is observed).

Proof. Since z is known, we can transform the random variable H and write $\tilde{H} = z^{-1}(H) = m(AU, V)$. Then, the joint distribution $G(U, V, \tilde{H})$ is observed and the conditions of proposition 1 are satisfied by $(m(\cdot), U, A, V)$. Thus, all results go through. \square

To summarize, in this Section we assumed that $A \perp V|U$ and the returns to scale of the matching function are known. Given these assumptions, we assumed that observing the joint distribution of (H, U, V) identifies (i) the matching function (ii) the joint distribution of unemployment, matching efficiency, and vacancies and (iii) the realizations of A in different time-periods. Looking ahead, we note that the estimation algorithm will assume, but not impose the independence $A \perp V|U$. This implies that we can (and will) use the estimated A_t in conjunction with the observed measures of search to test the independence assumption.

4 A Setting with Multiple Types of Job Seekers

So far we have allowed only for a single type of job seeker. Nonetheless, as discussed in the Introduction, more than half of all new matches are formed by workers who are either already employed or nominally out of the labor force. In this Section, we develop a specification which allows for multiple job seekers, all of which might exert unobserved effort in job seeking. It is this specification that we subsequently bring to the data.

Denote by Q_t the N-vector of counts of different groups of job seekers (e.g. those unemployed, employed, and OLF, perhaps also subdivided by gender, education etc.) and by A_t the N-vector of search efficacy associated with each type. Individual entries into these N-vectors are subscripted by $x \in \{1, \dots, N\}$ and when we take logarithms of these, we render them in small caps. In keeping with prior established notation, the random vectors associated with $\{Q_t, A_t\}$ are denoted by unsubscripted capitalized letters $\{Q, A\}$.

$A_{x,t}$ captures the search efficacy of group x at time t . This includes not just technological variation in the matching process between individual job seekers and jobs, but also search effort exerted by individuals.¹⁸

¹⁸For some purposes, it might be helpful to associate search efficacy explicitly to search effort and to think of search effort $A_{x,t}$ expended by group x at t as the average of individual search effort $A_{i,t}$ expended by those in group x . Then, $A_{x,t} = \frac{1}{Q_{x,t}} \sum_{i \in (x,t)} (A_{i,t})$ where this notation is meant to indicate that we sum individual effort across individuals i that are in group x at time t .

Assume that the overall number of hires is determined by a matching function $m(\sum_{x=1}^N A_{x,t}Q_{x,t}, V_t)$. We further assume that the average job finding rate of group x at t is proportional to effort $A_{x,t}$ and given by

$$p_{x,t} = \frac{m(\sum_{x=1}^N A_{x,t}Q_{x,t}, V_t)}{\sum_{x=1}^N A_{x,t}Q_{x,t}} A_{x,t} \quad (9)$$

Here $\sum_{x=1}^N A_{x,t}Q_{x,t}$ can be interpreted as the total effective search exercised by all the different job seekers and $\frac{m(\sum_{x=1}^N A_{x,t}Q_{x,t}, V_t)}{\sum_{x=1}^N A_{x,t}Q_{x,t}}$ is the search productivity per unit of effective search at time t .

Taking group $x = 1$ as our reference group, we can identify search of group x relative to group 1 using the ratio of job finding rates:

$$\frac{p_{x,t}}{p_{1,t}} = \frac{\frac{m(\sum_{x=1}^N A_{x,t}Q_{x,t}, V_t)}{\sum_{x=1}^N A_{x,t}Q_{x,t}} A_{x,t}}{\frac{m(\sum_{x=1}^N A_{x,t}Q_{x,t}, V_t)}{\sum_{x=1}^N A_{x,t}Q_{x,t}} A_{1,t}} = \frac{A_{x,t}}{A_{1,t}} \quad (10)$$

Then,

$$\tilde{S}_t \equiv \sum_{x=1}^N \frac{p_{x,t}}{p_{1,t}} Q_{x,t} \quad (11)$$

defines an index of overall job search effort relative to type 1's job search effort.¹⁹ In our empirical specification, we set group $x = 1$ to the unemployed. In addition, we also have access to an index of search effort among the unemployed, the Mukoyama-Patterson-Sahin (MPS) Index $I_{MPS,t}$ which we describe in more detail below. Using this index, we can now write total observed search as

$$S_t \equiv I_{MPS,t} \tilde{S}_t.$$

It is important to note that S_t contains only objects that can be directly observed from the data, namely relative job finding rates, counts of different groups of job seekers, and the MPS index.²⁰

¹⁹ \tilde{S}_t is analogous to the Hornstein-Kudlyak-Lange Non-employment index that summarizes search of different types of job seekers measurable in the CPS using fixed weights.

²⁰The fact that S_t is directly observed also implies that the assumption that search of different groups are perfect substitutes is less restrictive than it seems at first. For instance, it allows for variable search efficacy over time. This provides degrees of freedom that can for instance accommodate time-variation in the composition of vacancies that favor one type of job seeker

Thus we now have

$$H_t = m(A_t S_t, V_t) \tag{12}$$

where A_t denotes unobserved matching efficiency among the unemployed that is not accounted for by the MPS index. Contributing to A_t are search effort not captured by $I_{MPS,t}$ but also other unobserved reasons that lead to variation in matching over time, such as for instance recruiting effort by firms or geographical and/or occupational mismatch between job and worker characteristics.

This structure maps back into the identification arguments of the previous Section. Given S_t and as long as we are willing to assume that $V \perp A|S$, we can now directly apply the identification argument of the previous section and identify the functional form of $m(\cdot)$, as well as unobserved matching efficiency, A_t .

We now turn to empirically implement the approach developed in Sections 3 and 4. In the next two Sections, we first describe the data used in the empirical analysis as well as the approach taken to estimate the non-parametric matching function. We then present what our estimates imply about how the economy’s matching efficiency has changed over the recession, as well as the implied elasticity of the matching function with respect to vacancies.

5 The Data

We showcase the above methodology using up-to-date data sources, as well as some of the best available measurements from the literature on the search and recruiting efforts of workers and firms. The two main datasets we rely on are the Current Population Survey (CPS) and the Job Openings and Labor Turnover Survey (JOLTS). In addition, we use two indices of search and recruiting intensity. The first is based on Mukoyama et al. (2018) and exploits the American Time Use Survey (ATUS) jointly with the CPS to obtain an index of time spend search among the unemployed. We extend this series, following their methodology, to cover the period up to Dec. 2017. Our index for firm recruiting intensity is based on Davis et al. (2013) available on Steve Davis’s webpage.²¹

The CPS is the most commonly used data for analyzing trends in the labor market in the United States. We rely on basic monthly data from the CPS from January 2001 to December 2017.²² The CPS

over another. Note however that to make counterfactual predictions on search effort also requires taking a stand on what values S takes in the counterfactual world.

²¹We use the 2018-11 release.

²²The data was obtained from IPUMS-CPS on April 30th, 2019. Ruggles et al. 2020

is a very large data-set which allows measuring the stock of unemployed, employed, and OLF as well as their job finding rates across time. To estimate the job finding rates, we match observations in the CPS panel across two months. The CPS surveys addresses rather than individuals. To ensure that observations matched across months are from the same respondents we require that the variables age, race, and gender are consistent over time.

We limit ourselves to months-in-sample 1 and 5 as well as the short two months panels starting in those months when estimating job finding rates. While this restriction reduces the sample size substantially, the large size of the CPS still generates precise estimates of job finding rates and stocks by labor force status and age. The advantage of limiting ourselves in this way is that we reduce problems arising from the tendency of respondents to the CPS to report lower rates OLF and higher rates unemployed in later months of the short four-months panels produced by the CPS. All our estimates are obtained by weighing the data using the weights provided by the CPS.

JOLTS is a monthly survey providing monthly estimates of job openings, hires, and separations since December 2000. It is based on a sample of 16,400 establishments from the population of non-farm establishments including public employers. We use the data on vacancies and hires collected by JOLTS. Vacancies includes job postings for all positions that can start within 30 days for which the employer is actively recruiting outside the firm. Hires includes all additions to the payroll.

The CPS and the JOLTS are familiar to researchers in this literature and we therefore do not describe them in more detail here. The raw data and code used to produce our estimates are available upon request. Our final dataset covers the period from January 2001 to December 2017.²³

6 Empirical Analysis

6.1 Components of Observed Search and Recruitment

We next describe how we estimate the matching function and matching efficiency given in equation (12). This matching function depends on search from both sides of the market. Individuals seeking jobs represent the supply of labor and recruiters represent the demand for labor. As noted before, Mukoyama et al. (2018) measure search effort among the unemployed. They exploit the CPS jointly with the ATUS to provide an index representing search effort of the unemployed. Using their methodology, we extend

²³The JOLTS (and the Davis et al. 2013) data start on January 2001. Our data ends in December 2017 because the ATUS microfiles necessary to extend the search index of Mukoyama et al. (2018) were available only until then at the time of the analysis.

their index to cover the time-period up to Dec. 2017. In Figure 9 in the Appendix we show how the updated index compares to the original one. To capture variation in recruiting effort, we rely on the recruiting index provided by Davis et al. (2013) (denoted $I_{D,t}$). Thus, we have $V_t = I_{D,t}V_{JOLTS,t}$ where $V_{JOLTS,t}$ is the vacancy measure available from JOLTS.

To summarize, our task is to estimate $m(A_t S_t, V_t)$ as well as the unobservable A_t using the restrictions that the matching function is constant returns to scale and that V_t is independent of A_t conditional on S_t . Before we present the results from this estimation exercise, we present first how observed labor supply $S_t = I_{MPS,t}\tilde{S}_t$ and labor demand $V_t = I_{D,t}V_{JOLTS,t}$ and their components vary over time.²⁴

Recruitment

Figure 1 shows overall labor demand V_t and its components over the study period ranging from January 2001 to December 2017. The components are the vacancy count reported by JOLTS and the recruiting index provided by Davis et al. (2013).

The Davis et al. recruitment index ranges in a narrow band between 0.85 in the throes of the Great Recession in 2009, and 1.15 just prior to the 2001 recession and towards the end of our sample period. By contrast, $V_{JOLTS,t}$ varies by a factor of about 3 over the study period. Thus, the overall variation in labor demand over the cycle is dominated by the number of posted vacancies from JOLTS, so much so that it is difficult to distinguish between $V_t = I_{D,t}V_{JOLTS,t}$ and $V_{JOLTS,t}$ in the above graph.

Labor Supply

There are 3 sources of observable heterogeneity in the supply of labor. First, the number of unemployed varies over time. This variation has traditionally played a big role in estimates of the matching function. Second, for each unemployed individual in the market there are others that are not unemployed that are searching for jobs. These other job seekers might be OLF or currently employed. In recent years, a number of contributors showed that these job seekers that are not unemployed contribute significantly to new job relationships that are being formed and that their search has important implications for labor market dynamics (see, among others: Hornstein et al. (2014); Kudlyak and Lange (2017); Elsby et al. (2011); Kroft et al. (2016)). Third, search per individual might vary over time. Mukoyama et al. (2018) measure this variation in search effort for unemployed workers using data from the CPS and the ATUS.

²⁴We remove the strong seasonal effects present in this type of data. For this, we regress all variables on month dummies and remove the part predicted by these dummies.

Figure 1





Figure 2

Labor Supply



All series are deseasonalized and indexed to September 2007

We obtain counts of individuals by labor force status from the CPS. We aggregate these counts into unemployed-equivalency units using the relative job finding rates as weights, as in equation (11). Thus, each job seeker of type x in period t represents $\omega_{x,t} = \frac{Pr(job|x,t)}{Pr(job|unemployed,t)}$ unemployed individuals. Because there are substantially more individuals engaged in job-to-job search or OLF, the total stock of these job seekers expressed in unemployed equivalency units, \tilde{S}_t , typically exceeds the number of unemployed U_t by a factor of 4 to 8. Finally, we arrive at the total measure of job search by multiplying the total stock of job seekers, \tilde{S}_t , with the MPS index. In Appendix B we plot the observed search effort of employed workers (Figure 10), as well as those out of the labor force (Figure 11) of age 25-70.


Figure 2 shows in the top left panel how observed labor supply, $S_t = I_{MPS,t}\tilde{S}_t$, varies over our study period. The other panels show the three component parts of observed labor supply. These are (i) the MPS Search Index (ii) the number of unemployed age 25-70, and (iii) the factor $\frac{\tilde{S}_t}{U_t}$ which adjusts the number of unemployed to account for the other types of job seekers. Each of these is normalized to the Sept 2007 value, just prior to the Great Recession.

The top left panel shows that overall the supply of labor doubled during the Great Recession. It has since slowly fallen back and returned to its pre-recession values at some point during 2017. Traditional measures of labor supplied to the market based on the number of unemployed overstate how much the supply of labor increased during the Great Recession by about 25 percent (bottom left). This is because during the Great Recession (as during prior recessions) significant numbers of individuals flowed from OLF into unemployment, implying that search from these other sources declined relative to search supplied by the unemployed. The bottom right panel indicates the magnitude of this effect. Ignoring this composition effect means overestimating how countercyclical search in the market is. However, we also observe (top right) that time spend searching as measured using the MPS index increases when the labor market is weak, which partially but not fully offsets the composition effect.

Augmented Beveridge Curve

Finally we can also plot an “augmented” Beveridge curve. While the traditional Beveridge curve includes only search from unemployed workers, our augmented one also captures search from workers out of the labor force, as well as employed workers. In addition, it allows for variations in search per individual for all three groups of workers, as well as variation in vacancies’ recruiting intensity. Figure 3 presents the augmented Beveridge curve. Just like the traditional Beveridge curve, this curve also shifted during the Great Recession and as of 2017 has not returned to its pre-recession levels. In our framework this shift is

Figure 3



captured by decline in matching efficiency, which we quantify in our estimation next.

6.2 Estimation Algorithm

Our goal is to recover the matching function, $m(A_t S_t, V_t)$, as well the matching efficiency realizations, A_t . To do so, our estimation procedure follows closely the constructive identification arguments in Section 3. We normalize matching efficiency, A_t , in the period when vacancies V_t attain their median value throughout our time-period. This happens to be April of 2008. Denote this time-period by $t = 0$ and normalize A_0 to 1.

We first estimate $F(A_0|S)$ across the support of S . For this, we use the distribution of hires conditional on observed search effort, S , and observed recruiting effort, V , based on equation (6) which provides for $F(A_0|\psi S_0) = G_{H|S,V}(\psi H_0|\psi S_0, \psi V_0)$ for any arbitrary scalar ψ . We also have (see (6)) that $F(\psi A_0|\lambda S_0) = G_{H|S,V}(\psi H_0|\lambda S_0, \psi V_0)$ where $\lambda > 0$ is a scaling factor. Thus, by varying (ψ, λ) , we can trace out $F(A|S)$ across the entire support of (A, S) .

Of course, we have to rely on an estimate of $G_{H|S,V}$ obtained from our finite data to implement our constructive estimator. Consider then an arbitrary point (H_τ, S_τ, V_τ) . To obtain $G(H_\tau|S_\tau, V_\tau)$ we compute the proportion of observations with less than H_τ observed hires among observations close to (S_τ, V_τ) in (S, V) -space. In practice this is done by averaging across all observations in the data, penalizing observations with values (S_t, V_t) using a kernel that weighs down observations distant from (S_τ, V_τ) . That is, our estimate is

$$\hat{F}(\psi A_0 | \lambda S_0) = \sum 1(H_t < \psi H_0) \kappa(S_t, V_t, \lambda S_0, \psi V_0)$$

where $\kappa(\cdot)$ is a bivariate normal kernel.

Having recovered the distribution function $F(A|S)$, we now invert $F(A_t|S_t) = G(H_t|S_t, V_t)$ to obtain $A_t = F^{-1}(G(H_t|S_t, V_t)|S_t)$ for all observations (H_t, S_t, V_t) observed in the data. Finally, we recover the matching function as $m(A_t S_t, V_t) = G^{-1}(F(A_t|S_t)|S_t, V_t)$.

Our estimation procedure results in a set of fitted hires which the left panel of Figure 4 plots against the observed hires. It shows that we fit matches very closely.

The right panel shows the residuals from regressing fitted on actual hires. Besides some outliers, we do not observe that the fit varies systematically across the support of actual hires observed in JOLTS. Overall, the reason that our procedure fits the data so well is simply because we do not impose parametric restrictions. What Figure 4 demonstrates is that the restrictions we do impose are not obviously rejected by the data. These restrictions are mainly that the matching function is constant returns to scale and the regularity imposed by the kernel estimation procedure.

In addition, we are able to test the assumption that vacancies are independent of matching efficiency, conditional on overall labor supply, i.e. $V \perp A|S$. In particular, we use the residuals from a regression of vacancies, V on labor supply, S and similarly the residuals of matching efficiency, A on S . The correlation between these two residuals is close to zero (0.08) and as shown in Figure 5 there is no systematic relationship between these two residuals.

6.3 The elasticity of the matching function

The elasticity of the matching function with respect to vacancies is a crucial parameter influencing the dynamics of the labor market over the business cycle. In our framework, we cannot provide a single estimate of this elasticity, since we allow it to vary across (S, V) space rather than restrict it to be

Figure 4

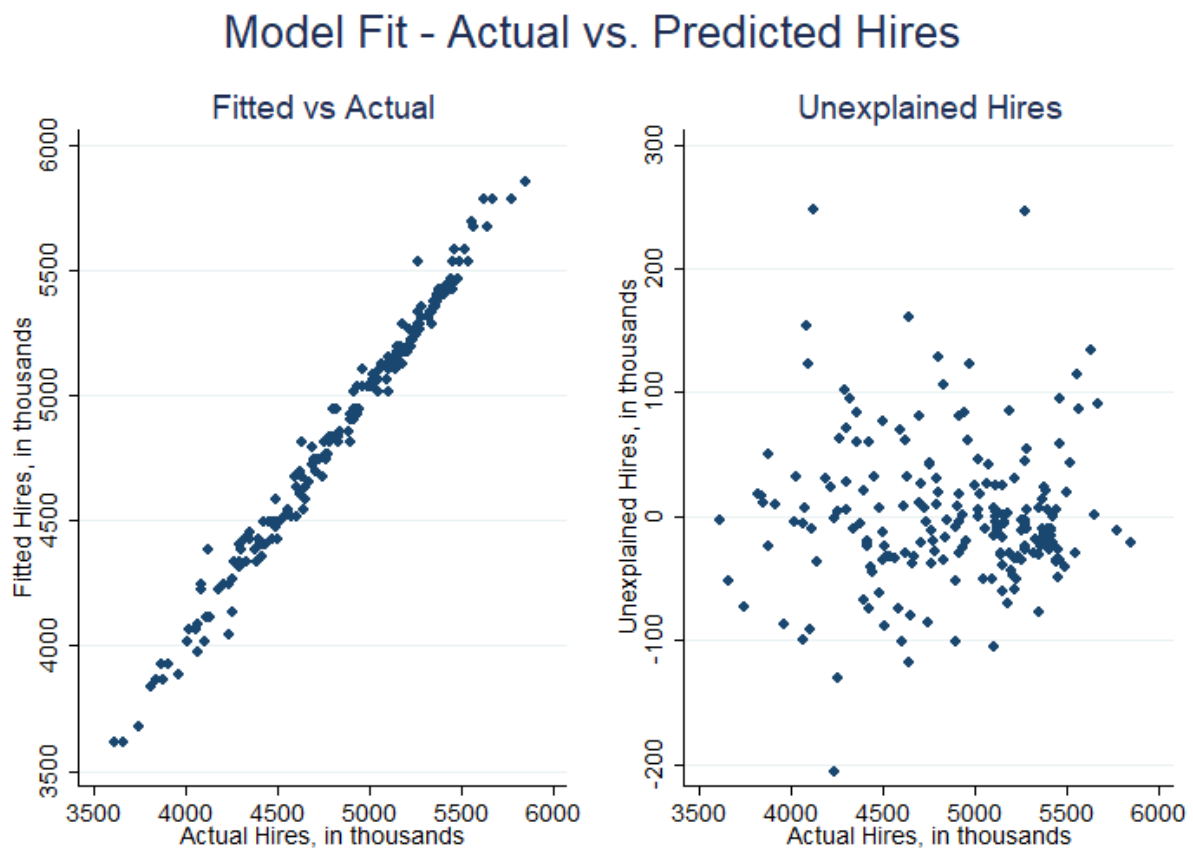


Figure 5





Figure 6




constant. As we argued in the introduction, relaxing this restrictive assumption has important normative and positive implications. Our non-parametric estimates of the matching function allow us to estimate how the elasticity of the matching function varies with market tightness. Here market tightness is defined as the ratio of total recruitment effort V_t to search effort by jobseekers S_t .²⁵

To obtain estimates of the responsiveness of hires to vacancies over the support of V_t and $A_t S_t$, we run a lasso regression projecting matches on fully interacted 3rd order polynomials of total recruiting effort, V_t and $A_t S_t$, the recovered matching efficiency multiplied by the total search by job seekers. The lasso estimates allow us to analytically differentiate with respect to vacancies and thus obtain an estimate of the elasticity of the matching function as market tightness varies. Figure 6 shows this elasticity of the matching function for 2001/1 - 2017/12 (Panel A) and also against market tightness (indexed against 9/2007.)

²⁵As discussed in Section 6.1, the former is the product of the vacancies observed in the JOLTS with the Davis recruitment index, while the latter is the product of the number of job seekers weighted by the relative job finding rates multiplied by the MPS index of observed search effort.

Figure 7



Our estimates are below those typically obtained in the literature, which often range above 0.5.²⁶ By contrast, our estimates range between 0.15 and 0.3. During “normal” labor market conditions, we tend to find an elasticity of approximately 0.22. We will return to the question why our estimates are below those in the literature (or conversely why previous estimates are so high) in Section 7 below.

In addition, we observe that the elasticity of the matching function is procyclical, so that during tight labor markets such as late in the recovery after the Great Recession the matching function is substantially more responsive to vacancies than during the Great Recession itself.²⁷

6.4 Observable Search and Matching Efficiency over the Great Recession

Figure 7 presents the implied matching efficiency A_t and how it changed over the Great Recession. Next, we consider how much variation in observed measures of labor demand (vacancies and recruiting effort), observed measures of labor supply, S_t , and our estimated measure of matching efficiency A_t contribute to variation in hires over the Great Recession. Figure 8 shows the actual number of matches over this period as well as the counterfactual number of matches obtained if we vary observed measures of search and recruiting, holding match efficiency fixed at the pre-recession level (October 2007) or, alternatively, hold observed measures fixed and vary match efficiency. Over the Great Recession, we see that, on net, the combined observed search and recruiting effort would in fact have increased the number of matches formed.²⁸ Therefore the decline in matches formed, is accounted for by a decline in matching efficiency A . Conversely, had observed search and recruiting effort stayed fixed at its pre-recession levels, then the decline in matches implied by the deterioration of the matching function would have exceeded the observed decline. In short, for reasons that are not yet well understood, the established relation between observed search and recruiting effort and matches broke down during the Great Recession.

In a way, this finding that aggregate search efficacy declined during the Great Recession is not new; the decline in search efficacy manifests itself in the much commented-upon outward shift of the Beveridge curve during the Great Recession. Using data covering 1976 to 2012, Barnichon and Figura (2015) provide evidence that matching efficacy is indeed procyclical. And, Sahin et al. (2014) provide evidence that this shift is only partially due to mismatch in the labor market across geography or industry. Our contribution is to quantify this decline in matching efficiency and show that this finding survives relaxing the functional form of the matching function as well as accounting for observed measures of recruiting intensity (Davis et al., 2013), search effort (Mukoyama et al., 2018), and type heterogeneity among job seekers.

7 Contrast to Cobb-Douglas as an Alternative Matching Function


By far the most common matching function employed in the literature is the Cobb-Douglas matching function. It is therefore worth comparing our estimates of the elasticity with respect to market tightness

²⁶Footnote 26 of Sahin et al. (2014) lists the most prominent estimates in the literature. These are 0.28 (Shimer, 2005), 0.5 (Davis et al., 2013), 0.54 (Mortensen and Nagypal, 2007), and 0.66-0.72 (Barnichon and Figura, 2015).

²⁷An exception to this pattern seems to be 2017, when our estimates of the elasticity of the matching function are somewhat low, while market tightness is very high. It is possible that this is due to boundary effects. We also note from panel B of Figure 6 that during no other time in our study-period were markets equally tight, contributing to instability in our estimates in this range. To draw more confident conclusions, we will use additional data gathered over time.

²⁸As discussed in Section 6.1, observed search effort increased during the Great Recession, while recruiting effort fell. Figure 8 shows that the net impact on match formation was positive.

Figure 8



overall with those obtained using the Cobb-Douglas specification.

Table 1 presents two regressions with log hires as the dependent variable and measures of log search and log recruitment as independent variables. Both specifications impose constant returns to scale, so that the coefficients on log search and log recruitment sum to one. The first column of estimates contains the regression specification using our measures of observed search, S_t and observed recruitment, V_t . Using these comprehensive measures of observed search and recruiting intensities, we obtain an estimate of the elasticity w.r.t. vacancies of 0.62, in line with estimates reported in the literature, but well above those we obtain using the non-parametric estimation approach, as discussed in Section 6.3. It is worth emphasizing that since in this regression we use our comprehensive measures of search and recruiting intensities, the difference in the results is not driven by the omission of search by employed workers or those out of the labor force.

Table 1: Cobb-Douglas Specifications

	Using Observed Search S	
Log(Obs. Recruiting Effort) (V)	0.617 (0.015)**	0.195 (0.005)**
Log(Obs. Search) (S)	0.384 (0.015)**	
Log(Matching Efficiency*Search) (A*S)		0.805 (0.005)**
Constant	-4.189 (0.174)**	-12.474 (0.085)**
N	204	204

* $p < 0.05$; ** $p < 0.01$

The table reports linear regressions of log hiring (JOLTS) on log search and recruiting effort. Recruitment effort equals the number of vacancies adjusted using the Davis et al. index. Both specifications impose that the coefficients sum to 1.

In column 1, observed search equals the total number of job seekers expressed in unemployed equivalencies (see Section 6.1) and the MPS index as described in Section 4. In column 2, the independent variable adjusts observed search using the matching efficiency which is estimated as described in Section 6.2.

The second column of estimates relies on the same recruitment measure but augments the measure of search using the imputed matching efficiency measure, A_t , obtained during the first step of our estimation algorithm, so that the right hand side term becomes $A_t S_t$. Using $A_t S_t$ as our measure of labor supply, we obtain a much lower estimate of 0.2 for the elasticity of hiring with respect to vacancies. Perhaps not surprisingly, this estimate of 0.2 is right in the range of estimates we report in figure 6.

Why then do we obtain such different estimates when we rely on observed measures of unemployed search compared to when we account for unobserved matching efficiency? The reason is that matching efficiency correlates with the traditional measure of tightness. Periods with many vacancies per observed job seeker are also periods when match efficiency tends to be high: Figure 7 shows that matching efficacy does vary systematically over the cycle - in particular, it is highest prior to the Great Recession and after the labor market recovered during the Great Recession. Indeed our estimated correlation is 0.88 (see also Figure 12 in Appendix B). This correlation between market tightness and search efficacy induces a positive bias in the estimates of the vacancy elasticity whenever unobserved matching efficacy is not controlled for, as is the case in traditional estimators of the matching function.

The Cobb-Douglas specification thus generates overall elasticities with respect to the vacancies that are comparable to those obtained from the non-parametric approach as long as we account for the changes in matching efficiency that is imputed from the first step of the estimation algorithm. Moreover, by construction, it cannot match how the elasticity varies with market tightness, shown in Figure 6.

We also considered a CES matching function, with limited success. Again, we find elasticities with respect to vacancies around 0.6 if we use observed search S_t and vacancies V_t as the arguments of the matching function. When using effective search $A_t S_t$ in lieu of S_t , we obtain an elasticity of 0.19. However, the cyclical variation implied by the CES runs counter to that obtained non-parametrically. In particular, the CES estimates predict that the elasticity with respect to vacancies declines with market tightness - the opposite of what we found non-parametrically.

8 Conclusion

This paper revisits how to estimate matching functions. Our proposed methodology allows for time-varying search of unemployed, as well as workers out of the labor force or currently employed. It relaxes both the strong independence assumption typically imposed between matching efficiency and search on either side of the labor market, as well as the functional form restrictions. We fully relax the functional

form of the matching function except for imposing constant returns to scale. Important welfare results (cf. Hosios 1990) related to search frictions hinge on the functional form of the matching function. For instance, the commonly used Cobb-Douglas form imposes strong restrictions on how search frictions affect welfare over the business cycle. Our non-parametric estimates of the matching function allow for a richer characterization of how search frictions affect welfare as market conditions vary.

Our results indicate an elasticity of the matching function with respect to vacancies of about 0.2. This elasticity varies significantly over the cycle, co-moving positively with market tightness. Our estimate is significantly lower than those obtained by other studies. This is because we account for bias induced because unobserved matching efficacy A_t covaries with market tightness over the cycle. An example is the Great Recession. Our estimates show that during this period the majority of the decline in hires can be attributed to declines in matching efficiency, which may partly capture unobserved search effort.

Appendix

A Estimating the Dependence Between E and V

Proposition 1 establishes the identification result for $m(AU, V)$ if $A \perp V|U$ and $m(\cdot)$ exhibits constant returns to scale. Under those conditions, the model is over-identified and we can test the independence assumption. This is the approach we have taken in the paper.

Here we show how to relax the conditional independence assumption and rather allow for A and V to be conditionally dependent. In particular, let $A = \tilde{A}V^\xi$ and $\tilde{A} \perp V|U$. We thus assume that systematic relation between matching efficiency A and vacancies V has a constant elasticity and that the stochastic component is multiplicative.

Assume for the moment that ξ is known. Then we can write $m(AU, V) = m(\tilde{A}UV^\xi, V) = \tilde{m}(\tilde{A}U, V)$ where \tilde{m} is a function with known returns to scale in $(\tilde{A}U, V)$. Corollary 3 implies that \tilde{m} and the distribution of \tilde{A} are known. We can also obtain realizations \tilde{A} for all observations of (H, V) . The latter then allow us to test whether $\tilde{A} \perp V|U$.

In the main text, we estimate $m(\cdot)$ assuming that $\xi = 0$. By allowing ξ to vary, we can obtain a estimated set for ξ that is consistent with the assumption that $A = \tilde{A}V^\xi$ and $\tilde{A} \perp V|U$ and for each point in this set, we obtain an estimate of $m(\cdot)$.

B Additional Figures

Figure 9




Figure 10




Figure 11

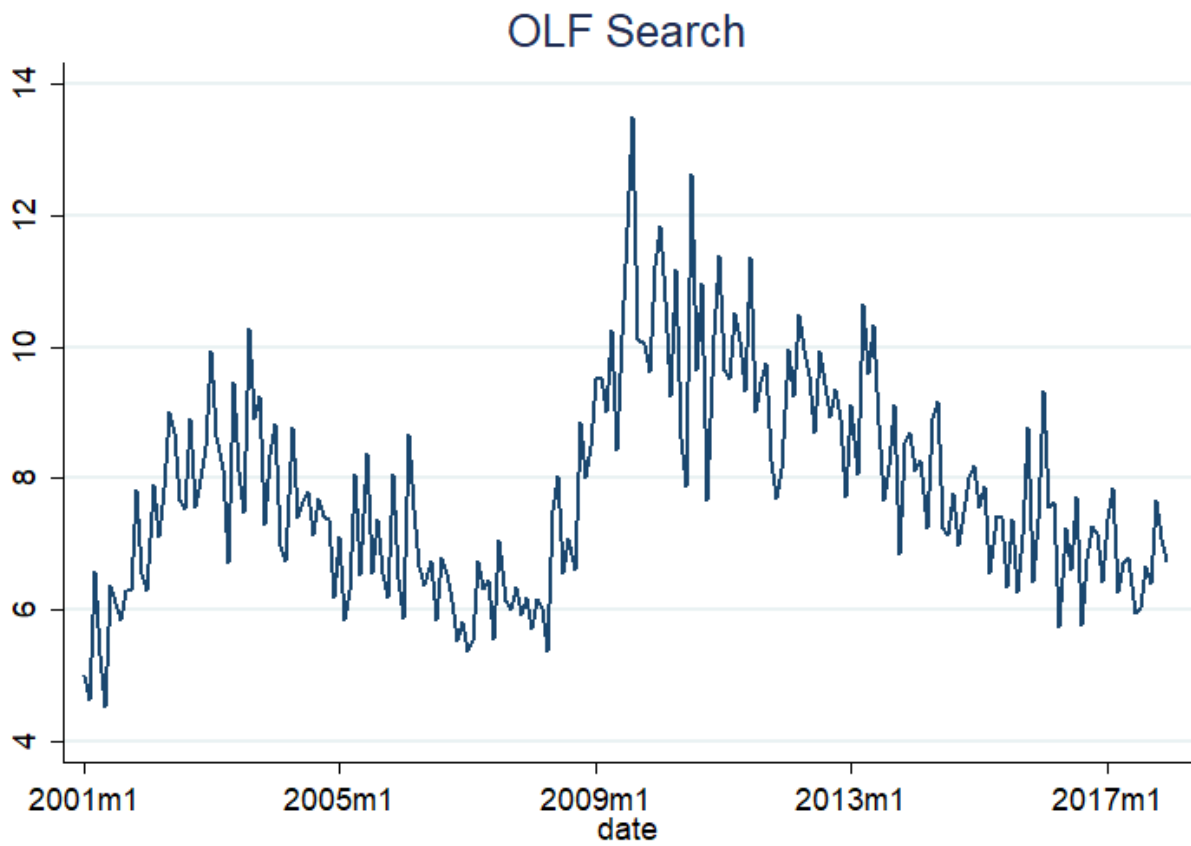



Figure 12



References

- BAJARI, P. AND C. L. BENKARD (2005): “Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach,” *Journal of Political Economy*, 113, 1239–1276.
- BARNICHON, R. (2012): “Vacancy Posting, Job Separation, and Unemployment Fluctuations,” *Journal of Economic Dynamics and Control*, 36, 315–330.
- BARNICHON, R. AND A. FIGURA (2015): “Labor Market Heterogeneity and the Aggregate Matching Function,” *American Economic Journal: Macroeconomics*, 7, 222–249.
- BLANCHARD, O. J. AND P. DIAMOND (1989): “The Beveridge Curve,” *Brookings Papers on Economic Activity*, 1989, 1–76.
- BOROWCZYK-MARTINS, D., G. JOLIVET, AND F. POSTEL-VINAY (2013): “Accounting for Endogeneity in Matching Function Estimation,” *Review of Economic Dynamics*, 16, 440–451.
- BRANCACCIO, G., M. KALOUPTSIDI, AND T. PAPAGEORGIU (2020a): “Geography, Transportation, and Endogenous Trade Costs,” *Econometrica*, 88, 657–691.
- BRANCACCIO, G., M. KALOUPTSIDI, T. PAPAGEORGIU, AND N. ROSAIA (2020b): “Search Frictions and Efficiency in Decentralized Transport Markets,” *mimeo, Harvard University*.
- DAVIS, S. J., J. R. FABERMAN, AND J. C. HALTIWANGER (2013): “The Establishment-Level Behavior of Vacancies and Hiring,” *Quarterly Journal of Economics*, 128, 581–622.
- ELSBY, M. W., B. HOBIJN, AND A. SAHIN (2015a): “On the Importance of the Participation Margin for Labor Market Fluctuations,” *Journal of Monetary Economics*, 72, 64–82.
- ELSBY, M. W., B. HOBIJN, A. SAHIN, AND R. G. VALLETTA (2011): “The Labor Market in the Great Recession - An Update to September 2011,” *Brookings Papers on Economic Activity*, 353–371.
- ELSBY, M. W. L., R. MICHAELS, AND D. RATNER (2015b): “The Beveridge Curve: A Survey,” *Journal of Economic Literature*, 53, 571–630.
- FALLICK, B. AND C. A. FLEISCHMAN (2004): “Employer-to-Employer Flows in the US Labor Market: The Complete Picture of Gross Worker Flows,” *Federal Reserve Board of Governors, Working Paper 2004-34*.

- FUJITA, S., G. MOSCARINI, AND F. POSTEL-VINAY (2019): “Measuring Employer-to-Employer Reallocation,” *mimeo*, *Yale University*.
- GOMME, P. AND D. LKHAGVASUREN (2015): “Worker Search Effort as an Amplification Mechanism,” *Journal of Monetary Economics*, 75, 106–122.
- HORNSTEIN, A. AND M. KUDLYAK (2016): “Estimating Matching Efficiency with Variable Search Effort,” Tech. Rep. WP 2016-24, Federal Reserve Bank of San Francisco Working Paper Series.
- HORNSTEIN, A., M. KUDLYAK, AND F. LANGE (2014): “A New Measure of Resource Utilization in the Labor Market,” *Federal Reserve Bank of Richmond*, *mimeo*.
- HOSIOS, A. J. (1990): “On the Efficiency of Matching and Related Models of Search and Unemployment,” *Review of Economic Studies*, 57, 279–298.
- KROFT, K., F. LANGE, M. J. NOTOWIDIGDO, AND L. F. KATZ (2016): “Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Non-Participation.” *Journal of Labor Economics*, 34(S1), S7–S54.
- KUDLYAK, M. AND F. LANGE (2017): “Measuring Heterogeneity in Job Finding Rates among those Out of the Labor Force,” *mimeo*, *McGill University*.
- MATZKIN, R. L. (2003): “Nonparametric Estimation of Nonadditive Random Functions,” *Econometrica*, 71, 1339–1375.
- MORTENSEN, D. T. AND E. NAGYPAL (2007): “More on Unemployment and Vacancy Fluctuations,” *Review of Economic Dynamics*, 10, 327–347.
- MUKOYAMA, T., C. PATTERSON, AND A. SAHIN (2018): “Job Search Behavior over the Business Cycle,” *American Economic Journal: Macroeconomics*, 10, 190–215.
- PETRONGOLO, B. AND C. A. PISSARIDES (2001): “Looking Into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390–431.
- PISSARIDES, C. A. (2000): *Equilibrium Unemployment Theory*, Cambridge, MA: MIT Press, 2nd ed.
- RUGGLES, S., S. FLOOD, R. GOEKEN, J. GROVER, E. MEYER, J. PACAS, AND M. SOBEK (2020): “IPUMS-USA: Version 10.0,” *Minneapolis, MN: IPUMS*.

SAHIN, A., J. SONG, G. TOPA, AND G. L. VIOLANTE (2014): “Mismatch Unemployment,” *American Economic Review*, 104, 3529–3564.

SHIMER, R. (2004): “Search Intensity,” *mimeo*, *University of Chicago*.

——— (2005): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25–49.

STIGLER, G. J. (1961): “The Economics of Information,” *Journal of Political Economy*, 69, 213–225.