

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

Fabian Lange^y and Theodore Papageorgiou^z

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 efficiency, 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 matching 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

^yBoston College, Department of Economics, Maloney Hall 343, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

^zWe 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 cyclicity 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 efficiency” 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 efficiency 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 efficiency. 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 effort/matching efficiency of the unemployed A_t , and vacancies V_t into hires H_t .

$$H_t = m_t(A_t U_t; V_t) = (A_t U_t)^{\alpha} V_t \quad (1)$$

The objective of the researcher is to estimate α using data on $(H_t; U_t; V_t)$, while search effort 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 $u_{j,t} = \log \frac{H_t}{U_t}$ as a linear function of log market tightness $\tau_t = \log(V_t/U_t)$:

$$u_{j,t} = \log \frac{H_t}{U_t} = (1 - \alpha) a_t + \tau_t; \quad (2)$$

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, τ_t is likely to correlate with a_t and naive

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 $m(;;) : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ 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

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 $F(A_jU)$ at point $(A_0; U_0)$:

$$F(A_0jU_0) = G_{HjU;v}(H_0jU_0; V_0)$$

where $G_{HjU;v}(:jU_0; V_0)$

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, matching 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 $m(AU; V) = z(m(AU; V))$ 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 \mathfrak{m} and the joint distribution $F(A; U)$

Assume that the overall number of hires is determined by a matching function $m(x)$

Thus we now have

$$H_t = m(A_t S_t; V_t) \quad (12)$$

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 $V \neq A$

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

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 $S_t = I_{MPS,t} \bar{S}_t$ 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

Total Labor Supply


MPS Search Index



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 $n_{x;t} = \frac{\text{Pr}(\text{job};x;t)}{\text{Pr}(\text{job};\text{unemployed};t)}$ 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, 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, 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, $S_t = \text{panel no. 396} - 2.758 T_d$;

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

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 $(H_0; S_0; V_0)$. To obtain $G(H_0 | S_0; V_0)$ we compute the proportion of observations with less than H_0 observed hires among observations close to $(S_0; V_0)$ 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 $(S_0; V_0)$. That is, our estimate is

$$\hat{F}(H_0 | S_0; V_0) = \frac{1}{n} \sum_{t=1}^n 1(H_t < H_0) K((S_t; V_t) - (S_0; V_0))$$

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

Havingwhere

Figure 4

Model Fit - Actual vs. Predicted Hires




Figure 5


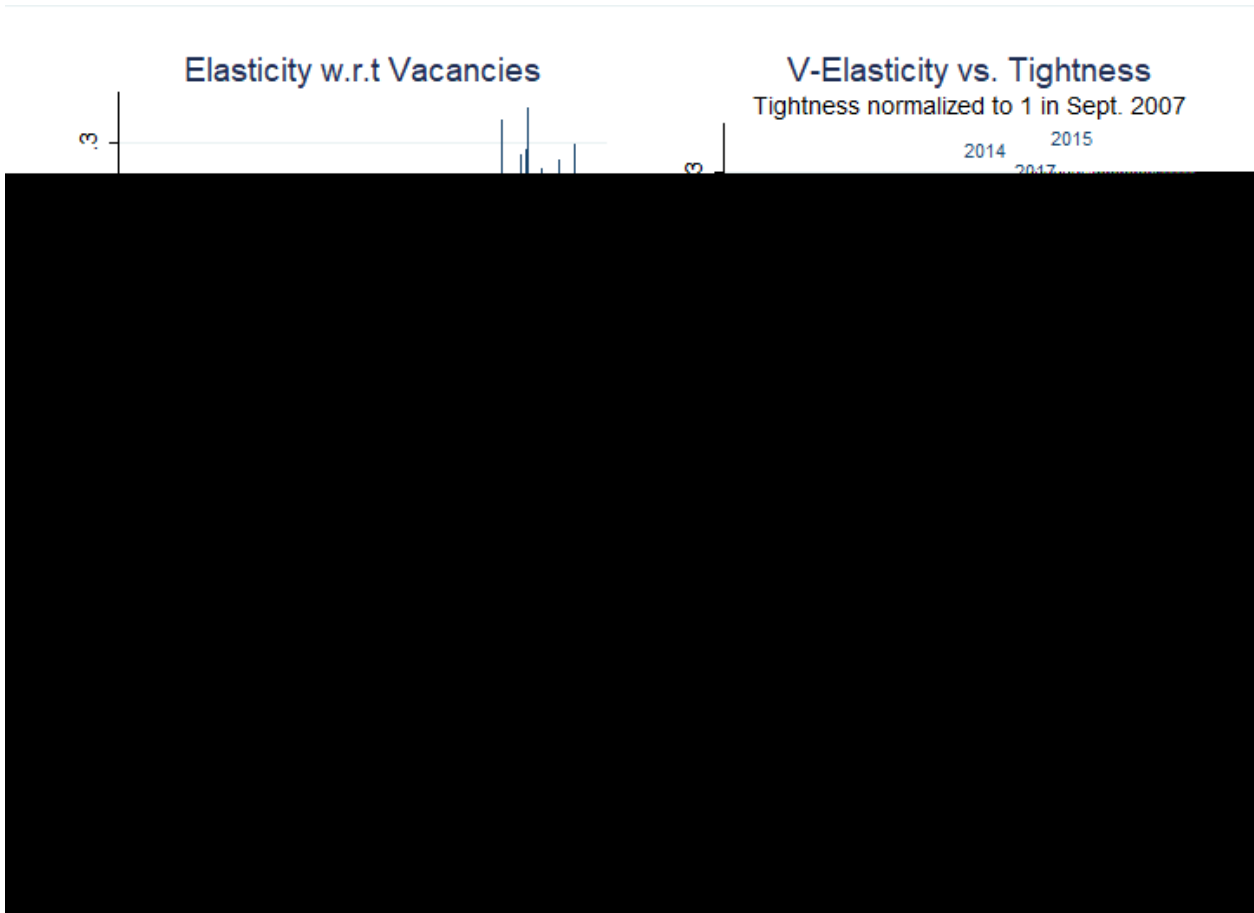


Figure 6




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

Matching Efficiency



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.

6.4 Observable Search and Matching Efficiency over the Great Recession

Figure 8

Decomposing the Variation in Matches Observed and Unobserved Components

5000



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.

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 efficiency 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 efficiency induces a positive bias in the estimates of the vacancy elasticity whenever unobserved matching efficiency 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 efficiency A_t covaries with market tightness over the cycle. An example is

Appendix

A Estimating the Dependence Between E and V

Proposition 1 establishes the identification result for $m(A; V)$ 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 $A = \alpha V$ and $A \perp V | U$: 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 α is known. Then we can write $m(A; V) = m(\alpha V; V) = \bar{m}(A; V)$ where \bar{m} is a function with known returns to scale in $A; V$. Corollary 3 implies that \bar{m} and the distribution of A are known. We can also obtain realizations A for all observations of $(H; V)$. The latter then allow us to test whether $A \perp V | U$.

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

B Additional Figures

Figure 9

Original and Updated MPS Measure





Figure 12

Matching Efficiency and Market Tightness



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. (2015): "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 123, 1239–1276.

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.