American Economic Journal: Macroeconomics 2015, 7(4): 222–249
http://dx.doi.org/10.1257/mac.20140116

Labor Market Heterogeneity and the Aggregate Matching Function†

By Regis Barnichon and Andrew Figura*

We estimate an aggregate matching function and find that the regression residual, which captures movements in matching efficiency, displays procyclical fluctuations and a dramatic decline after 2007. Using a matching function framework that explicitly takes into account worker heterogeneity as well as market segmentation, we show that matching efficiency movements can be the result of variations in the degree of heterogeneity in the labor market. Matching efficiency declines substantially when, as in the Great Recession, the average characteristics of the unemployed deteriorate substantially, or when dispersion in labor market conditions—the extent to which some labor markets fare worse than others—increases markedly. (*JEL* E24, E32, J41, J42)

The search and matching model (Mortensen and Pissarides 1994) has become the canonical framework to introduce equilibrium unemployment in macroeconomic models. One of its building blocks is the aggregate matching function that relates the flow of new hires to the stocks of vacancies and unemployment. Like the aggregate production function, the matching function is a convenient device that "partially captures a complex reality […] with workers looking for the right job and firms looking for the right worker" (Blanchard and Diamond 1989).

An important feature of the labor market is its matching efficiency, i.e., the market's ability to match unemployed workers to jobs. However, in a standard specification of the matching function, matching efficiency is akin to a Solow residual; a parameter that adjusts to capture any hiring behavior that cannot be explained by the observed levels of unemployment and vacancy posting. We estimate such a matching function over 1967–2012, and we find that the regression residual, or movements in

* Barnichon: Centre de Recera en Economia, Internacional (CREI), Ramon Trias Fargas 25, 08005 Barcelona, Spain, Universitat Pompeu Fabra, and Barcelona GSE (e-mail: rbarnichon@crei.cat); Figura: Board of Governors of the Federal Reserve System, 20th Street and Constitution Avenue, NW, Washington, DC 20551 (e-mail: andrew.figura@frb.edu). We thank Jan Eeckhout, Bruce Fallick, Shigeru Fujita, Jordi Gali, Robert Hall, John Haltiwanger, Bart Hobijn, Philipp Kircher, Rob Valletta, Thijs van Rens, William Wascher, Yanos Zylberberg, two anonymous referees and seminar participants at the 2012 AEA Annual Meetings, the 2012 CEPR European Summer Symposium in International Macroeconomics (ESSIM), the CREI-CEPR Conference "Understanding Jobless Recoveries," the 2014 Essex Search and Matching workshop, the Chicago Fed, the New York Fed, the Norges Bank, and the San Francisco Fed. We thank Peter Chen for excellent research assistance. The views expressed here do not necessarily reflect those of the Federal Reserve Board or the Federal Reserve System. Barnichon acknowledges financial support from the Ministerio de Ciencia e Innovacion (ECO2011-23188) and the Agència de Gestió d´Ajuts Universitaris i de Recerca (AGAUR) through the Beatriu de Pinós fellowship (2011 BP00152). Any errors are our own.

† Go to http://dx.doi.org/10.1257/mac.20140116 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

matching efficiency, displays nontrivial procyclical fluctuations. In particular, over 2008–2012, matching efficiency experienced an unprecedented decline that lowered the aggregate job finding rate by 30 percent.

In this paper, we aim to better understand fluctuations in matching efficiency. To do so, we take a "look into the [matching function] black-box" (Petrongolo and Pissarides 2001), and we construct an aggregate matching function that explicitly incorporates (i) heterogeneity across workers and (ii) labor market segmentation. We incorporate worker heterogeneity by allowing for different levels of search efficiency across workers, i.e., we allow for the possibility that some individuals have a higher propensity to form a match than others. We incorporate labor market segmentation by allowing the labor market to be segmented in submarkets, where each submarket is described by a matching technology. This setup captures the idea that because of geographic distance, skill mismatch, or degree requirements, a worker can only match with the vacancies opened in his submarket.

In this framework, matching efficiency is not a residual. Instead, matching efficiency is a function of worker and submarket heterogeneity, and matching efficiency moves over the cycle because of variations in the average characteristics of the labor market. We highlight the role of two effects. The first one is a *composition effect*, due to the fact that the average search efficiency of the unemployment pool can vary. For instance, if composition changes, and a group with a lower than average search efficiency becomes more represented among the unemployed, matching efficiency will decline. The second effect is a *dispersion effect*, in which dispersion in labor market conditions, the fact that tight submarkets coexist with slack ones, drives down matching efficiency because of the concavity of the matching function.[1] When the degree of heterogeneity across workers and labor markets is constant, the two effects are constant and matching efficiency is constant.

Estimating our framework requires data on worker characteristics as well as labor market characteristics, in particular, local labor market conditions. We use matched CPS micro data over 1976–2012 to control for worker characteristics. Controlling for local labor market conditions (i.e., labor market tightness at the segment level) is difficult because highly disaggregated vacancy data start being available only in 2006, just one year before matching efficiency began its unprecedented decline. To address this data limitation, we propose a two–stage estimation procedure that overcomes the need for job openings data before 2006. The method combines CPS micro data over 1976–2012 with Conference Board online help wanted ads data available since 2006.

We find that our aggregate matching function does a very good job at capturing movements in the aggregate job finding rate over 1976–2012, including the post-2007 period, and we conclude that explicitly allowing for heterogeneity across workers and labor markets is important to understand labor market fluctuations. Aggregate matching efficiency is procyclical because both the composition effect and the dispersion effect are procyclical. First, in recessions, dispersion in labor

[1] The effect of labor misallocation on matching efficiency in the context of the matching function is similar to the effect of capital misallocation on aggregate TFP in the context of the production function and emphasized in recent studies (e.g., Hsieh and Klenow 2009).

market tightness across segments rises—some segments fare much worse than others—and aggregate matching efficiency declines.[2] Second, in recessions, composition changes and the average quality, or employability, of the unemployment pool worsens leading to a decline in matching efficiency. The two key individual characteristics responsible for the composition effect are reason of unemployment (e.g., job loser versus job leaver)—likely capturing unobserved heterogeneity across workers—and unemployment duration—capturing unobserved heterogeneity across workers and/or the fact that workers' employability declines with the length of the unemployment spell (Kaitz 1970). In recessions, the share of long-term unemployed and the share of job losers go up, leading to a decline in aggregate matching efficiency. Since 2007, both dispersion and composition—in particular, a large increase in the share of long-term unemployed—have driven down aggregate matching efficiency to exceptionally low levels.

While there is a large literature studying the aggregate matching function,[3] this paper is the first to propose, and estimate with micro data, a framework in which labor market segmentation and heterogeneity across workers and jobs affect aggregate matching efficiency. Our matching function framework encompasses two separate strands of the literature. The first strand, related to our composition effect, has studied the individual determinants of unemployment duration, although without specific concern for the underlying matching technology.[4] More recently, Hall and Schulhofer-Wohl (2013) extend our analysis of the composition effect to on-the-job and out-of-the-labor-force jobseekers.[5] The second strand, related to our dispersion effect, has focused on measuring the extent of mismatch in the labor market.[6] Recently, Herz and van Rens (2011) propose an approach to disentangle various potential sources of mismatch, and Şahin et al. (2014) construct mismatch indices based on a theoretical model of mismatch. Şahin et al.'s mismatch measure and our dispersion measure are related, both relying ultimately on the concavity of the matching function.

The next section estimates a standard matching function. Section II presents the empirical framework underlying our aggregate matching function. Section III uses micro data to estimate that framework. Section IV presents the results. Section V interprets the movements in matching efficiency over time and Section VI concludes.

[2] Different mechanisms could explain the procyclicality of dispersion. For instance, changes in the location or nature of jobs can lead to more misallocation of jobs and workers in recessions and, hence, to a higher level of dispersion. Alternatively, different cyclical sensitivities to aggregate shocks across labor markets could also generate procyclical dispersion (Abraham and Katz 1986).

[3] See, e.g., Pissarides (1986); Blanchard and Diamond (1989); Bleakley and Fuhrer (1997); the review of Petrongolo and Pissarides (2001); Davis, Faberman, and Haltiwanger (2010); and Borowczyk-Martins, Jolivet, and Postel-Vinay (2012). Since the first draft of this paper, a number of papers have studied the behavior of the matching function during the last recession. See, e.g., Barlevy (2011); Veracierto (2011); Barnichon et al. (2012); Hobijn (2012); Ghayad and Dickens (2012); and Sedláček (2012).

[4] See Devine and Kiefer (1991) and Baker (1992). Lindeboom, van Ours, and Renes (1994) and Petrongolo (2001) are two noteworthy exceptions (although with a different focus) that exploit the link between a matching function and workers' job finding to estimate matching functions from micro data.

[5] Two recent papers have also emphasized the importance of variations in the composition of the unemployment pool over the cycle. Mueller (2012) shows that in recessions the pool of unemployed shifts toward workers with high wages in their previous job. Kroft et al. (2013) investigate whether composition can explain the high level of long-term unemployment since the Great Recession.

[6] See Padoa Schioppa (1991) and Layard, Nickell, and Jackman (2005).

I. The Aggregate Matching Function

The matching function relates the flow of new hires to the stocks of vacancies and unemployment. In a continuous time framework, the flow of hires is typically modeled with a Cobb-Douglas matching function with constant returns to scale, and we can write

$$m_t = \mu_t U_t^{\sigma} V_t^{1-\sigma} \tag{1}$$

with m_t, the number of new hires at instant t; U_t, the number of unemployed; V_t, the number of vacancies; and μ_t denoting matching efficiency.[7]

Since the job finding rate f_t is the ratio of new hires to the stock of unemployed, we have $f_t = \frac{m_t}{U_t}$, so that

$$f_t = \mu_t \theta_t^{1-\sigma} \tag{2}$$

with $\theta = \frac{V}{U}$ the aggregate labor market tightness, and we can estimate the matching function in the log-linear form with

$$\ln f_t = (1 - \sigma) \ln \theta_t + \varepsilon_t. \tag{3}$$

We measure the job finding rate f_t from unemployment-employment transitions from the Current Population Survey (CPS) over the period 1976–2012 and from the worker flows data tabulated by Joe Ritter for the period 1968–1975. We use the composite help wanted index presented in Barnichon (2010) as a proxy for vacancy posting. We use nondetrended quarterly data and estimate (3) over 1968–2007. Table 1 presents the results. Using OLS, the elasticity is estimated at 0.33. Using lagged values of v_t and u_t as instruments gives similar results, and the elasticity is little changed at 0.34.[8]

Figure 1 plots the empirical job finding rate, its fitted value, and the regression residual ε_t, which captures movements in matching efficiency. While aggregate labor market tightness does a good job at capturing movements in the aggregate job finding rate up until 2007, the residual shows a spectacular decline after 2007, and as of late 2012, the observed value of the job finding rate is 30 percent lower than implied by the level of the vacancy-unemployment ratio alone.[9] In other words, matching efficiency has dropped markedly since 2007.

Interestingly, even before 2007, the matching function residual displays a puzzling cyclical pattern; increasing in the later stages of expansions, peaking in the late stages of recessions or the early stages of recoveries, and declining thereafter.

[7] The Cobb-Douglas matching function is used in almost all macroeconomic models with search and search and matching frictions (e.g., Pissarides 2000). Allowing for nonconstant returns to scale or using a more general CES matching function $m_t = \mu[\sigma U_t^{\rho} + (1 - \sigma) V_t^{\rho}]^{1/\rho}$ gives very similar results.

[8] As argued by Borowczyk-Martins, Jolivet, and Postel-Vinay (2012), OLS may suffer from an endogeneity bias because of agents' endogenous behavior.

[9] Elsby, Hobijn, and Şahin (2010) report a similar finding using the unemployment outflow rate, and Davis, Faberman, and Haltiwanger (2010) also report a dramatic decline in the vacancy yield using JOLTS data.

Table 1—Estimates of the Matching Function Elasticity

Dependent variable	f_t	f_t	F_{jit}	F_{jit}
Sample (quarterly frequency)	1968–2007	1968–2007	1976–2007	2006–2012
Regression	(1)	(2)	(3)	(4)
Estimation	OLS	GMM	BF	MLE
$1 - \sigma$	0.33***	0.34***	0.18***	0.21***
	(0.01)	(0.01)	(0.02)	(0.02)
R^2, 1976–2012	0.78	—	0.88	—

Notes: Standard errors are reported in parentheses. Regressions 1 and 2 are the aggregate regressions of (log) f on (log) tightness as described in Section II. In regression 2, we use three lags of v and u as instruments. Regression 3 is the two-stage procedure (labeled BF) described in the main text. Regression 4 estimates all model parameters in one stage with MLE using HWOL vacancy data available over 2006–2012.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

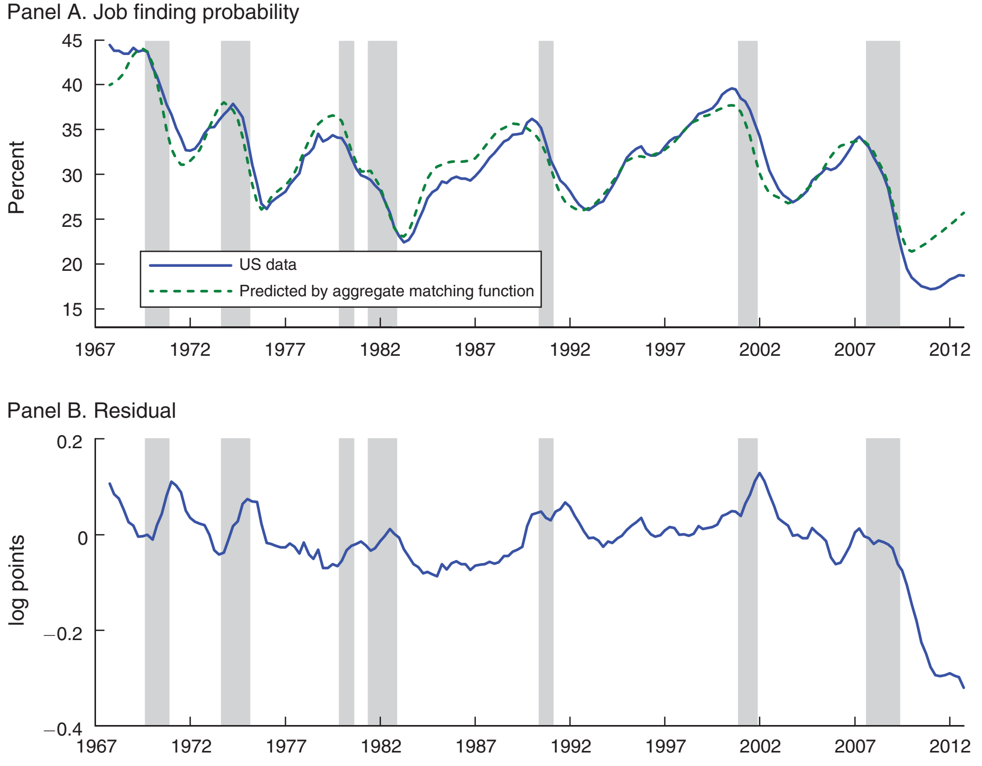

Figure 1. Empirical Job Finding Rate

Notes: Job finding rate predicted by an aggregate matching function estimated over 1968–2007 and (log) residual, the (log) difference between the empirical and the predicted job finding rate, over 1968–2012. The plotted series are the four-quarter moving averages. Grey bars indicate National Bureau of Economic Research (NBER) recession dates.

II. A Matching Function Framework with Labor Market Heterogeneities

In this section, we construct an aggregate matching function that explicitly incorporates labor market heterogeneities across workers and labor markets. We then show how, in this framework, aggregate matching efficiency moves over the cycle because of variations in the average characteristics of the labor market.

A. *An Aggregate Matching Function*

We first show how an aggregate matching function can arise out of the aggregation of segmented labor markets populated by heterogeneous workers.

There are I labor market segments and J worker types. The labor market segment $i \in \{1,..,I\}$ of individual type $j \in \{1,..,J\}$ is the labor market in which individual j can look for work and find a job. Each labor market segment i has a matching technology that depends on V_{it}, the number of job openings in segment i; U_{it}, the number of unemployed in segment i; and μ_i, the constant matching efficiency of segment i. Heterogeneity in matching efficiency captures the idea that some occupations or locations have a higher rate of matching than others.[10] The matching technology in each segment is described by a CRS Cobb-Douglas matching function.[11]

Each worker type j in segment i is characterized by his search efficiency s_{jit}, which depends on characteristics that make him more or less likely to form a match. We do not take a stand on the mechanism behind the different search efficiencies, but simply allow for the presence of heterogeneity in that dimension. Without loss of generality, we normalize average search efficiency to 1 by appropriately rescaling the μ_is (the matching efficiency levels of the segments).

The number of new hires in segment i at time t, m_{it}, is thus given by

$$m_{it} = \mu_i V_{it}^{1-\sigma}(s_{it} U_{it})^{\sigma}, \tag{4}$$

with s_{it}, the average search efficiency in segment i, given by

$$s_{it} \equiv \sum_{j=1}^{J} \frac{U_{jit}}{U_{it}} s_{jit}, \tag{5}$$

with U_{jit}, the number of unemployed workers of type j in segment i at time t, so that $U_{it} = \sum_j U_{jit}$.

[10] For instance, hiring for high-skill occupations may be more time consuming than hiring for low-skill occupations. As a result, low-skill occupations may display a higher number of new matches per unit of time (for a given number of job seekers and job openings), i.e., a higher matching efficiency.

[11] While we relax the standard matching function apparatus by considering a segmented labor market, we still make a number of simplifying assumptions. The matching function elasticity σ is constant across segments, and matching efficiency μ_i is constant across time in each segment. These two assumptions are common in the mismatch literature (Jackman and Roper 1987, Padoa Schioppa 1991, Şahin et al. 2014).

The total number of matches in the economy, $m_t \equiv \sum_{i=1}^{I} m_{it}$, is then given by an aggregate matching function

$$m_t = \mu_t V_t^{1-\sigma} U_t^{\sigma}, \tag{6}$$

with aggregate matching efficiency given by

$$\mu_t = \sum_{i=1}^{I} \frac{U_{it}}{U_t} \mu_i s_{it}^{\sigma} \left(\frac{\theta_{it}}{\theta_t}\right)^{1-\sigma}, \tag{7}$$

with $V_t \equiv \sum_{i=1}^{I} V_{it}$ and $U_t \equiv \sum_{i=1}^{I} U_{it}$ the total number of vacancies and unemployed in the economy, $\theta_{it} \equiv \frac{V_{it}}{U_{it}}$ the labor market tightness in segment i, and $\theta_t \equiv \frac{V_t}{U_t}$ the aggregate labor market tightness.

Expression (7) generalizes the standard matching function by explicitly allowing (i) for worker heterogeneity and (ii) segmentation in the labor market. Thanks to (7), we can link movements in aggregate matching efficiency to observable characteristics of the labor market, and movements in aggregate matching efficiency μ_t can be decomposed into a composition effect and a dispersion effect.

B. *A Decomposition of Aggregate Matching Efficiency*

With some manipulation of (7) left for the Appendix, the aggregate job finding rate $f_t = \frac{m_t}{U_t}$ can be approximated as

$$\ln f_t = \ln \mu_t + (1 - \sigma) \ln \theta_t \tag{8}$$

$$\text{with } \mu_t \simeq \mu_0 \left(1 + \underbrace{\mu_t^s + \mu_t^m}_{\text{Composition}} - \underbrace{\frac{\sigma(1-\sigma)}{2} \operatorname{var}\left(\frac{\theta_{it}}{\theta_t}\right)}_{\text{Dispersion}} \right) \tag{9}$$

to a second-order in the degree of heterogeneity across worker characteristics and across labor market tightnesses, with μ_0 the average matching efficiency level across segments.

This decomposition of the aggregate job finding rate highlights how, with worker heterogeneity and concavity in the matching technology, changes in composition and dispersion can lead to movements in aggregate matching efficiency μ_t. For small variations in the degree of labor market heterogeneity, the terms on the right-hand side of (9) move little, and we have $\mu_t \simeq \mu$,[12] and the aggregate matching

[12] We have $\mu = \mu_0 E\left(1 + \mu_t^s + \mu_t^m - \frac{\sigma(1-\sigma)}{2} \operatorname{Var}\left(\frac{\theta_{it}}{\theta_t}\right)\right)$.

function—$m_t = \mu_t V_t^{1-\sigma} U_t^{\sigma}$—can be approximated by a matching function with constant matching efficiency—$m_t = \mu V_t^{1-\sigma} U_t^{\sigma}$.

Looking into the components of (9), the first term in (9) captures the aggregate job finding rate $\mu_0 \theta_t^{1-\sigma}$ absent worker heterogeneity and absent dispersion in labor market tightness across segments.

The second term in (9), $\mu_t^s + \mu_t^m$, describes the composition effect coming from:

- $\mu_t^s = \sigma \sum_{i,j} \frac{U_{jit}}{U_t} \left(s_{jit} - 1\right)$ capturing the effect of changes in the composition of the unemployment pool. For instance, if the share of a group (e.g., long-term unemployed) with a lower than average job finding probability increases in recessions, then the *average* job finding probability will decline without any change in individuals' job finding probabilities.
- $\mu_t^m = \sum_i \frac{U_{it}}{U_t} \left(\frac{\mu_i}{\mu_0} - 1\right)$ capturing the effect of changes in the distribution of the unemployed across segments with different average matching efficiency. For instance, if a higher fraction of the unemployed becomes concentrated in a segment with higher matching efficiency, the average job finding probability will increase even if the aggregate numbers of vacancy and unemployed remain constant.

The third term in (9) captures the effect of dispersion in labor market conditions on aggregate matching efficiency. Intuitively, dispersion in labor market tightness across segments negatively affects the *average* job finding rate because the segment-level job finding rate $f_{it} = \frac{m_{it}}{U_{it}}$ is a concave function of labor market tightness θ_{it} (because the matching function $m_{it} = \mu_i V_{it}^{1-\sigma} (s_{it} U_{it})^{\sigma}$ is a concave function of U_{it} and V_{it}). As a result, if some segments (such as health care) display a relatively tight labor market and some segments (such as manufacturing) display a slack labor market, the average job finding probability will be lower than in an economy where labor market tightness is identical across segments.

C. *Discussion of Empirical Framework*

Before bringing our aggregate matching function to the data, we briefly discuss the economic rational behind our setup. Our accounting framework rests on two premises: (i) workers differ in their search efficiency, and (ii) the labor market is segmented and the matching process in each segment is described by a matching function.

Starting with search efficiency, two mechanisms could generate variations in search efficiency across workers. First, the intensity with which an individual searches for a job influences the probability of receiving a job offer. Search intensity can vary across workers because of worker heterogeneity in the disutility cost of search or in the utility of market production relative to home production (e.g., Pissarides 2000, chapter 5). Second, conditional on a worker meeting a firm, a match may or may not be viable depending on the worker's reservation wage. As with search intensity, with worker heterogeneity in the disutility cost of search or in the net utility of market work, the reservation wage will vary across workers, and this can generate variation in matching rate across workers.[13]

Turning to the segmentation of the labor market, we follow the mismatch literature (Jackman and Roper 1987; Padoa Schioppa 1991, and more recently Şahin et al. 2014) and do not impose the existence of a unified labor market, but instead allow the labor market to be segmented into (i) *distinct* and (ii) *frictional* submarkets.

Within each segment, the labor market is *frictional*, and the matching process is governed by a matching function with constant matching efficiency. Underlying this assumption lies the existence of coordination frictions (as captured for instance by a simple urn-ball matching process (Butters 1977) that affect the matching rate of job seekers and vacancies.

The segments are *distinct*, and we rule out that workers or firms spread out their search effort over several segments. Underlying this assumption lies the existence of large costs of moving and searching for jobs across submarkets, either across large geographic distances or across different occupation or industry groups. Naturally, the validity of this assumption depends on the size of a labor market segment, a topic to which we will return in the empirical section.

III. Estimation Procedure

In this section, we present our approach to bring our aggregate matching function to the data.

To map the continuous time aggregate matching function to the data, we consider a continuous time environment in which data are available only at discrete dates. For $t \in \{0, 1, 2 \ldots\}$, we refer to the interval $[t, t+1[$ as 'period t.' We assume that during period t, the instantaneous flow of new matches in island i is constant and given by m_{it}. Given the matching technology (4), the job finding rate of an individual type j in segment i is constant during period t and satisfies

$$f_{jit} = \frac{s_{jit}}{s_{it} U_{it}} m_{it}$$

$$= \mu_i \frac{s_{jit}}{s_{it}} s_{it}^{\sigma} \theta_{it}^{1-\sigma},$$

and the job finding probability over period t is given by

$$F_{jit} = 1 - e^{-\mu_i \frac{s_{jit}}{s_{it}} s_{it}^{\sigma} \theta_{it}^{1-\sigma}}. \tag{10}$$

To estimate (10), we use matched monthly data from the Current Population Survey (CPS) covering January 1976 to December 2007 to measure

[13] Alternatively, one could also think of a stochastic job matching model (e.g., Pissarides 2000, chapter 6) in which the output of a match involving a worker of type j is drawn from a distribution Γ_j . Heterogeneity across workers in the distribution Γ_j will generate heterogeneity in matching rates, i.e., in search efficiency.

unemployment-to-employment transitions (Nekarda 2009) and to control for worker characteristics. A major data limitation is the absence of data on job openings, and hence labor market tightness, at the segment level over a long time sample.[14] In particular, the Help Wanted OnLine (HWOL) dataset provided by the Conference Board provides information on the number of job openings by geographic location, occupation and/or industry at a very disaggregated level, allowing researchers to measure labor market tightness at a high level of disaggregation (as recently used by Şahin et al. 2014). However, the sample period covered by HWOL starts only in 2006 and covers precisely the period in which matching efficiency displayed an unprecedented decline. Since our intention is to see how far taking labor market heterogeneities into account can go in accounting for the movements in the aggregate job finding rate after 2007, we cannot rely on post-2007 data to estimate our model.[15] Instead, we want to estimate our model over a long sample period that excludes the post-2007 data.

To get around this data limitation, we propose a two-stage estimation procedure that overcomes the need for job openings data over a long time sample. In the first stage, we use the fact that each individual is atomistic in his labor market segment,[16] so that we can use the segment-specific average job finding rate (measurable from CPS micro data) to control for market tightness at the segment level. This first stage allows us to measure the effect of worker characteristics—the composition effect—while controlling for local labor market conditions. In the second stage, we combine HWOL data and CPS micro data with our first-stage estimate of the composition effect to estimate σ, the elasticity of the matching function, from time series variation over 1976–2007.[17]

We define a labor market segment by its geographic location and occupation group, and we disaggregate the labor market into 36 segments defined by 9 geographic locations (the US Census divisions) and 4 occupation groups: professional, services, sales, and production.[18]

The appropriate size of a labor market segment, i.e., the definition of the labor market unit, is an open question in the literature (Petrongolo and Pissarides 2001). As stressed by Abraham (1991), a segment should be small enough to accurately capture the relevant labor market faced by individuals, but not too small so that it

[14] In the United States, the two public data sources with vacancy posting data are the Job Openings and Labor Turnover Survey (JOLTS) and the Help-Wanted OnLine series from the Conference Board. The JOLTS measure of job openings can be disaggregated into about 15 industry groups, but the series only start in 2000.

[15] In particular, since the matching function elasticity parameter σ is estimated with information from the time dimension only, using post-2006 data would bias our σ estimate and bias our results into fitting the large decline in matching efficiency during the recent recession.

[16] This approach is thus valid as long as the labor market segment is not too tightly defined.

[17] In addition, and as a robustness check to our two-stage procedure, in the Appendix, we report the results of a direct maximum likelihood estimation of all model parameters using HWOL vacancy data over 2006–2012. If our model is well specified, and worker heterogeneity and dispersion do indeed explain movements in matching efficiency, estimating the model with post-2006 data should give estimates similar to the ones obtained with pre-2006 data. We find that this is indeed the case.

[18] Specifically, we use the nine US Census Divisions and four high-level occupation groups: professional, services, sales and office, and production. At this level of disaggregation, an individual is clearly atomistic. The nine census divisions are New England, Middle Atlantic, East North Central, West North Central, South Atlantic, East South Central, West South Central, Mountain and Pacific. The occupation groups are taken from the SOC high-level groups: professional (management, business, science, and arts), services (personal services), sales (sales and office), and production (construction, maintenance, production, and transportation).

can still be described by a standard matching technology with no interactions with outside segments. Moreover, and in addition to these theoretical considerations, data availability limits the level of labor market disaggregation that can be studied. In the case of the United States, we are limited by the sample size of the CPS.

We strived to strike a balance between all these constraints. By defining a segment at the intersection of a US Census division and a high-level occupation group, the segments are sufficiently different (e.g., production in the New England division versus professional in the Mountain division) to be considered approximately distinct, while at the same time provide a more reasonable description of the relevant labor market faced by an individual. Moreover, at this level of disaggregation, the CPS sample size is still large enough to ensure a good signal-to-noise ratio.

A. *Stage 1: Estimating the Effect of Workers' Characteristics*

In the first stage, we estimate the vector β—capturing the effect of individual characteristics on job finding probabilities—while controlling for local labor market conditions.

To capture the effect of individual characteristics on search efficiency, we posit that s_{jit}, the search efficiency of worker type j at time t, is given by

$$s_{jit} = e^{\beta X_{jit}}, \tag{11}$$

with $X_{jit} = \left[1, x_{jit}^1, \ldots, x_{jit}^K\right]$ a vector of worker characteristics (detailed below) for type j in segment i at time t.

To estimate β without data on local market tightness θ_{it}, we use the fact that, given (4), the average job finding rate in segment i is

$$f_{it} \equiv \frac{m_{it}}{U_{it}} = \mu_i s_{it}^{\sigma} \theta_{it}^{1-\sigma}, \tag{12}$$

so that an individual job finding probability can be written as

$$F_{jit} = 1 - e^{-\frac{s_{jit}}{s_{it}} f_{it}}, \tag{13}$$

with $\frac{s_{jit}}{s_{it}} = \frac{e^{\beta X_{jit}}}{\sum_j \frac{U_{jit}}{U_{it}} e^{\beta X_{jit}}}$ and f_{it} independent of F_{jit} since individuals are atomistic in their labor market segment. We can then estimate β by maximizing the likelihood of the sample of unemployment-employment transitions given that the monthly job finding probability F_{jit} is given by (13).[19]

[19] We provide more details about the maximum likelihood estimation in the Appendix. Note that because the job finding rate of an individual depends on his search efficiency *relative* to the average search efficiency in the labor market at a given time $\left(\text{i.e., } \frac{s_{jit}}{s_{it}}\right)$, the effect of characteristics on an individual job finding rate is estimated only in the cross section.

We use three main types of information from the CPS to capture worker characteristics: demographics, reason for unemployment, and duration of unemployment.[20]

Demographic information includes the age, sex, and education level of the unemployed individual. We use 10 bins of 5 years to capture the effect of age on the job finding probability: less than 20, 20–25, …, 55–60, and over 65.

We distinguish between four main reasons for unemployment: permanent layoff, temporary layoff, re-entering the labor force, and quit job. We use dummy variables for each reason. Reason for unemployment likely captures unobserved heterogeneity across individuals.

The CPS records the duration (in weeks) of individuals' current spells of unemployment. Prior research (e.g., Kaitz 1970; Machin and Manning 1999) found that the job finding probability declines with duration, and we include unemployment duration as an explanatory variable. To capture the effect of duration, we use ten bins of equal size (in terms of number of unemployed).

In 1994, a major redesign of the CPS survey was implemented and introduced breaks in many important variables, such as reason for unemployment and duration of unemployment (Polivka and Miller 1998). To control for these breaks, we estimate separate coefficients for the preredesign and postredesign periods.

Finally, we assign each job seeker to his/her location-occupation submarket from CPS information on current state of residence and previous occupation.[21]

B. *Stage 2: Estimating the Elasticity of the Matching Function*

We still have two parameters to estimate: σ, the elasticity of the matching function, and μ_i the segment-specific matching efficiency.

Although data on θ_{it} are not available before 2006, our aggregate matching function framework provides just enough structure on the data to allow us to estimate σ from time series variation in θ_t and f_{it} (both of which available back to 1976), as well as recover the time series for θ_{it} consistent with our model.[22]

Specifically, our approach proceeds in two steps:

- **Step 1:** We consider a grid over $[0, 1]$ of possible values of σ, and for each value of σ on this grid, we do two things: (i) estimate the μ_is and (ii) construct series of local labor market tightness, the θ_{it}s.

[20] We also experimented with race/ethnicity but found that these characteristics play little role in the cyclicality of matching efficiency, consistent with the findings of Baker (1992). We thus omitted them for clarity of exposition. We also include a set of monthly dummies to control for seasonality in job finding probabilities.

[21] Less than 10 percent of unemployed are missing occupation information. They are almost exclusively new entrants to the labor force and comprise mostly individuals younger than 20. Restricting our analysis to individuals older than 20 gives very similar results.

[22] Specifically, we can circumvent the absence of data on θ_{it} over 1976–2007, thanks to two assumptions in the model: (i) the matching function elastically, σ, is identical across segments, and (ii) μ_i, segment-specific matching efficiency, is constant over time. As stated previously, these assumptions are standard in the mismatch literature.

To first estimate the μ_is, we use Conference Board HWOL job openings data to measure θ_{it} over 2006–2007.[23] Since β was estimated in the first-stage, $\frac{s_{jit}}{s_{it}}$ is known, and we can estimate μ_i for a given σ from

$$F_{jit} = 1 - e^{-\frac{s_{jit}}{s_{it}} s_{it}^{\sigma} \mu_i \theta_{it}^{1-\sigma}} \tag{14}$$

by maximum likelihood.

Second, given μ_i, we can rearrange (12) to construct the θ_{it}s implied by our model given the observed segment-specific average job finding rates f_{it}. Specifically, we construct

$$\hat{\theta}_{it} = \left(\frac{f_{it}}{\mu_i s_{it}^{\sigma}}\right)^{\frac{1}{1-\sigma}}. \tag{15}$$

- **Step 2:** The first step allowed us to construct series of local tightness as functions of σ, i.e., to construct functions $\hat{\theta}_{it}(\sigma)$. Using the definition of aggregate tightness, the implied aggregate tightness is then also a function of σ with $\hat{\theta}_t(\sigma) = \sum \frac{U_{it}}{U_t} \hat{\theta}_{it}(\sigma)$. Although local tightness θ_{it} is not directly measurable, aggregate tightness θ_t is. Thus, using the time series variation in aggregate labor market tightness θ_t over 1976–2007, we can estimate σ from

$$\min_{\sigma} \sum_t \left(\theta_t - \hat{\theta}_t(\sigma)\right)^2. \tag{16}$$

  Specifically, we do a grid search over $[0, 1]$ to find the σ that minimizes the sum of squared differences between observed aggregate tightness θ_t and implied aggregate tightness $\hat{\theta}_t(\sigma)$.[24]

IV. Estimation Results

In this section, we present the results of our estimation and then analyze the behavior of the aggregate job finding rate since 1976 through the lens of our aggregated matching function.

A. *Coefficient Estimates*

Column 3 of Table 1 reports the results of our two-stage estimation procedure. At 0.18, the elasticity is substantially lower than when using only aggregate labor market tightness as an explanatory variable. This indicates that the effect of labor

[23] Although HWOL data are also available after 2007, we restrict our time sample to 2006–2007 to avoid using data from a period with unusual movements in aggregate matching efficiency. The 2006–2007 period is a period *before* the dramatic decline in matching efficiency, which allows us to estimate the μ_is without biasing our results in favor of explaining the behavior of the job finding rate after 2007.

[24] The grid covers [0,1] in increments of 0.01. The standard error is computed by Monte Carlo methods. Specifically, we model the residuals $\theta_t - \hat{\theta}_t(\sigma)$ with an AR(1) to allow for serial correlation. We then sample from the residuals of this AR(1) to generate a series ε_t of model residuals. We then generate a new series $\hat{\theta}_t(\sigma) + \varepsilon_t$ to which we apply our procedure and estimate a new value for σ. We repeat this exercise 1,000 times. The standard error is then the standard deviation of these estimated σ across all draws.

market heterogeneities is on average procyclical, and that failing to control for heterogeneity biases estimates of the aggregate matching function elasticity upward.

Figure 2 presents the coefficients for the determinants of search efficiency, expressed in units of job finding rate for ease of comparison.[25] The most important individual characteristic is unemployment duration. Search efficiency (i.e., the propensity to form a match) is decreasing in unemployment duration, consistent with previous findings on the existence of duration dependence (e.g., Kaitz 1970; Machin and Manning 1999; Shimer 2008; Kroft, Lange, and Notowidigdo 2013).[26]

We find that the effect of duration on an individual's employment probability is large: for instance, an individual unemployed for 6 months is 50 percent less likely to find a job than an individual who just entered the unemployment pool. This estimate, based on workers' actual job finding rates, is remarkably similar to Kroft, Lange, and Notowidigdo's (2013) result based on field experiment data on employers' callback rate. Moreover, consistent with Kroft, Lange, and Notowidigdo (2013), we find that workers' search efficiency drops sharply over the first six months of the unemployment spell and then stabilizes.[27]

The second most important characteristic is reason for unemployment. The estimates reveal that it is more difficult for permanent job losers and entrants to the labor force to find employment. Not surprisingly, workers on temporary layoff are the most likely to find a job, i.e., have the highest search efficiency.

Turning to demographics, the coefficients on the age variables indicate that search efficiency decreases with age. Quantitatively, a 60-year old individual is 10 percent less likely to find a job than a 20-year old individual. The coefficient on the male dummy indicates that males are slightly more likely to find jobs than females.

Finally, more educated workers have higher search efficiency: a college graduate is 8 percent more likely to find a job than a high-school dropout.

B. *Accounting for Movements in the Aggregate Job Finding Rate*

Using our estimated coefficients, we now evaluate whether our aggregate matching function can account for movements in the aggregate job finding rate. Figure 3 plots the movements in the job finding rate unexplained by our aggregate matching function—the difference $\ln f_t - \ln\left(\mu_t \theta_t^{1-\sigma}\right)$ with μ_t given by (7)—along with the movements in the job finding rate unexplained by an aggregate matching

[25] Figure 2 presents the coefficients estimated over 1994–2007. Recall that because of a break in 1994, we allowed for a break in the coefficients in 1994. The coefficients estimated over 1976–1993 (available upon request) are very similar.

[26] A contribution to that literature is that we estimate the strength of the duration dependence phenomenon after controlling for worker characteristics as well as local labor market conditions, in a manner fully consistent with the matching function framework.

[27] Theoretically, duration dependence can arise through two channels. First, the "accumulation" of unemployment duration could have a causal effect on workers' search efficiency and job finding probability (Kaitz 1970), for instance through skill deterioration (e.g., Pissarides 1992). Second, duration dependence could arise out of a dynamic selection process driven by unobserved worker heterogeneity: workers with high search efficiency leave unemployment faster than those with low search efficiency, thereby generating a negative correlation between duration and job finding rates (Salant 1977). While discriminating between these two channels is outside the scope of this paper, the fast decline in workers' search efficiency over the first months of unemployment (Figure 2) suggests that gradual loss of skill is unlikely to be the sole factor and points toward some role for unobserved heterogeneity, in line with the recent findings of Kroft, Lange, and Notowidigdo (2013).

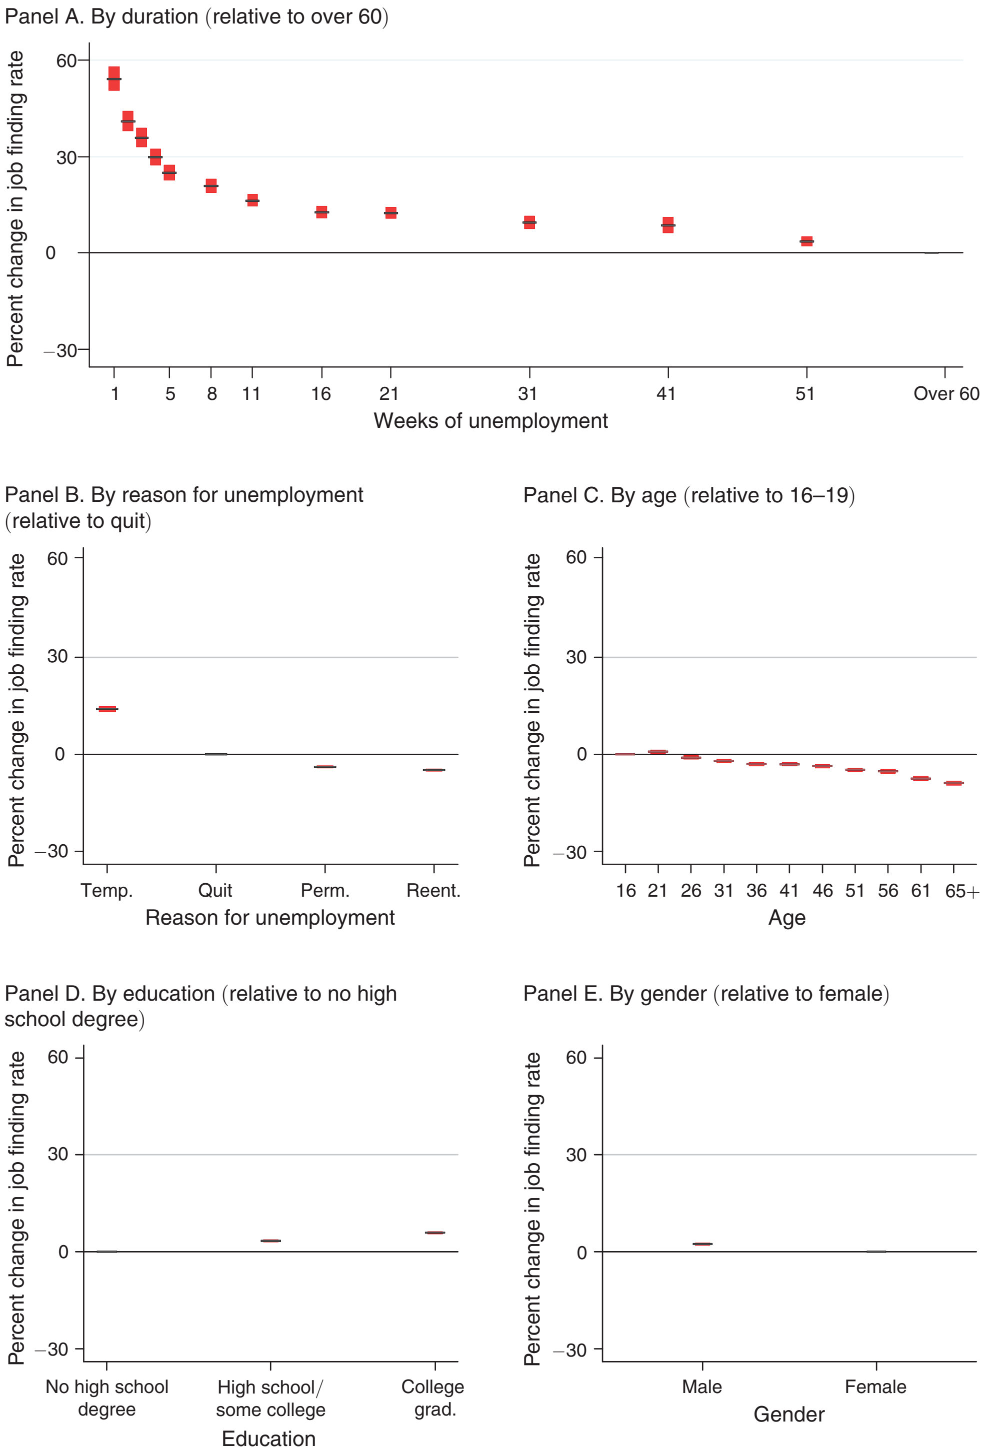

FIGURE 2

Notes: Coefficient estimates, 1994–2007. The black bars denote the point estimates and the grey bars denote ±2 standard errors.

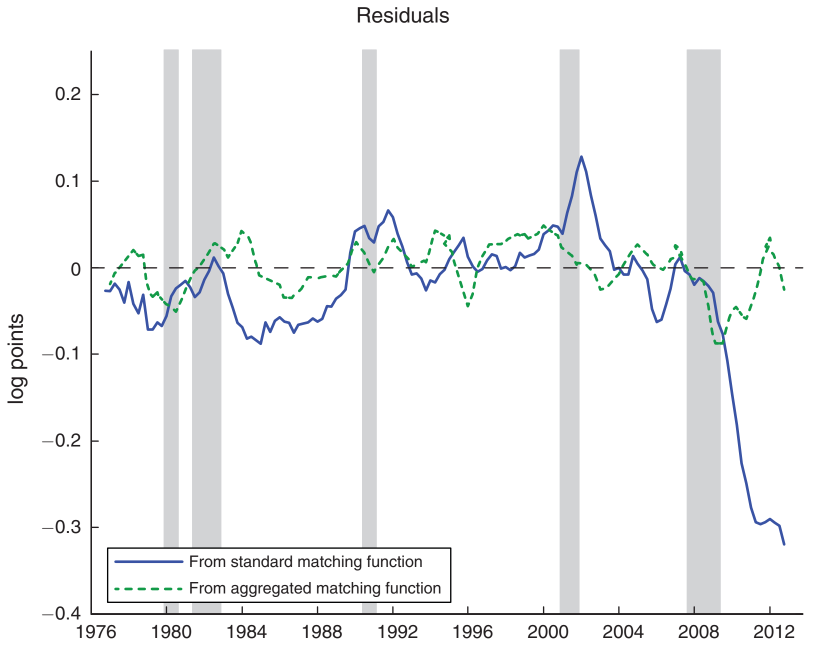

FIGURE 3. UNEXPLAINED MOVEMENTS IN THE AGGREGATE JOB FINDING RATE, 1976–2012

Notes: Residuals from models estimated over 1976–2007. The "aggregated matching function" refers to our matching function framework with worker heterogeneity and market segmentation. All series are four-quarter moving averages. Grey bars indicate NBER recession dates.

function with constant matching efficiency—the residual from regression (3), $\ln f_t - \ln\left(\mu\theta_t^{1-\sigma}\right)$, estimated in Section I.

Our aggregate matching function, *estimated with data prior to* 2008, does a very good job of explaining the dramatic and prolonged decline in the job finding rate since 2008. Even before the last recession, our aggregate matching function substantially improves the fit of the data, reducing by more than 50 percent the volatility of the (already small) residual of the aggregate matching function regression (3). Calculating the coefficient of determination for both models over 1976–2012, we find that the R^2 increases from 0.78 using the standard matching function to 0.88 using our aggregate matching function. Moreover, the cyclical pattern that was apparent in the residual from the aggregate regression (3) is absent in the residual from our aggregate matching function framework.[28]

[28] The correlation between the unemployment rate and the residual from the aggregate regression (3) is −0.41 (with a p-value for the null of no correlation of 0.00), whereas the correlation between the unemployment rate and the residual from our aggregate matching function framework is −0.15 (with a p-value of 0.09).

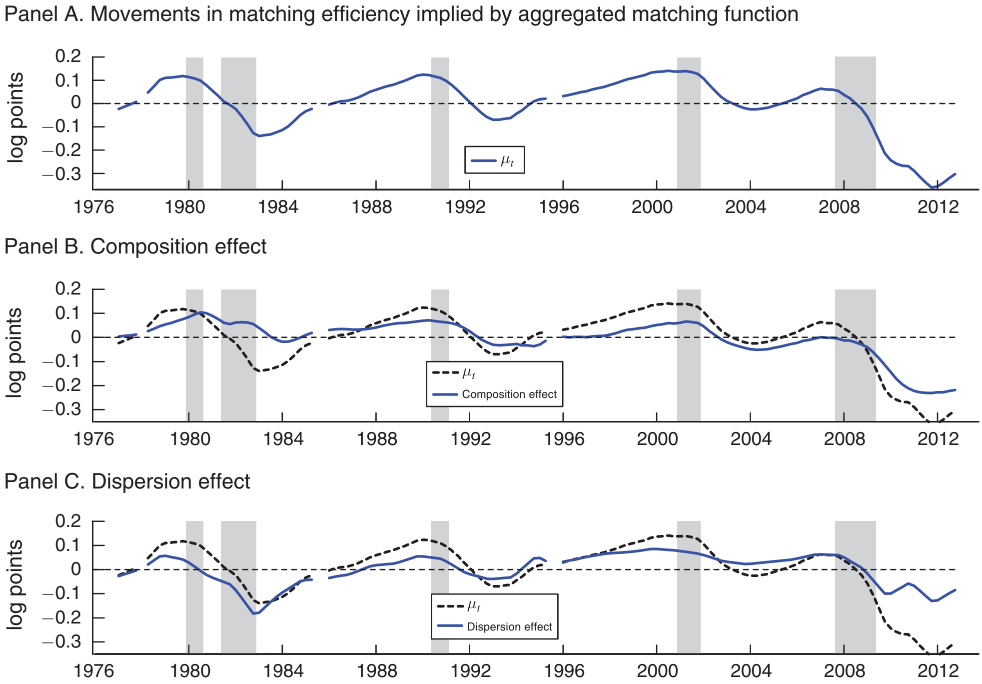

Figure 4. Decomposition of Movements in Aggregate Matching Efficiency

Notes: Panel A: movements in aggregate matching efficiency, μ_t, implied by the "aggregated matching function," i.e., by our matching function framework with worker heterogeneity and market segmentation. Panel B: movements in μ_t due to composition effect. Panel C: movements in μ_t due to dispersion effect. All series are four-quarter moving averages. Grey bars indicate NBER recession dates.

V. Movements in Aggregate Matching Efficiency

Having shown that our aggregate matching function can successfully capture movements in the aggregate job finding rate, we now analyze the cyclical properties of aggregate matching efficiency, μ_t, through the lens of our framework, and we discuss the reasons for the dramatic decline in matching efficiency after 2007.

A. *Aggregate Matching Efficiency over the Cycle*

Figure 4 plots μ_t and its two components: the composition effect and the dispersion effect.

First, we can see that composition and dispersion contribute roughly equally to movements in matching efficiency μ_t over the business cycle.

Second, dispersion appears to be a countercyclical phenomenon—rising during recessions and abating during expansions (Figure 5). The countercyclicality of dispersion is particularly interesting in the context of the literature on mismatch, where data availability constrained researchers to assess the cyclicality of mismatch from five to ten years of data only (Şahin et al. 2014).

Third, the composition effect is procyclical. In recessions, the average quality or employability of the unemployment pool worsens leading to a lower aggregate

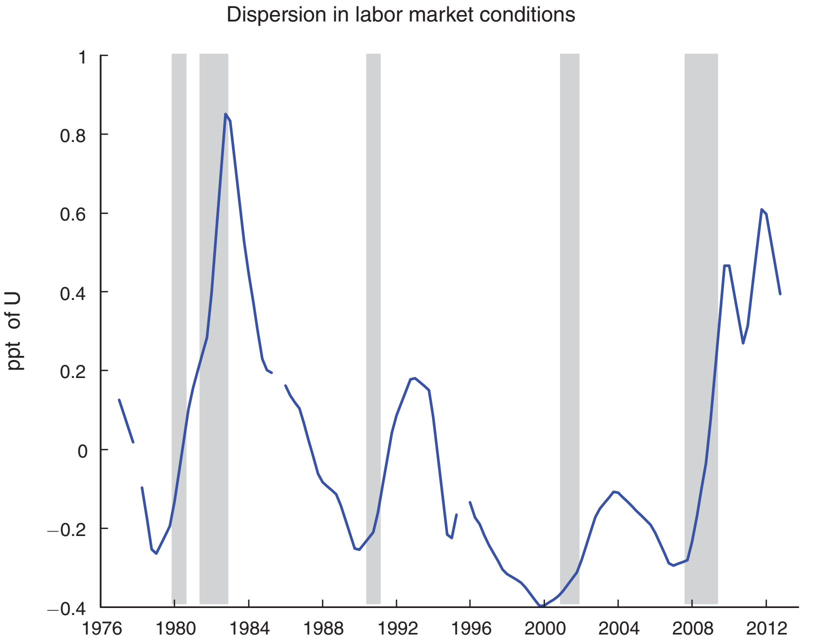

FIGURE 5

Notes: Dispersion in labor market conditions (expressed in percentage points of unemployment, see Appendix), 1976–2012. Four-quarter moving averages. Grey bars indicate NBER recession dates.

matching efficiency. We explore the procyclicality of the composition effect in more details in the next section.[29]

B. *The Composition Effect over the Cycle*

In order to better understand how the composition of the unemployment pool affects μ_t, we isolate the contributions of the different characteristics behind the composition effect. We find that the two key characteristics responsible for the composition effect are unemployment duration and reason of unemployment. In recessions, the share of long-term unemployed and the share of job losers go up, leading to a decline in aggregate matching efficiency.

Figure 6 graphs the contributions of individual characteristics μ_t^s—unemployment duration, reason for unemployment, demographics (grouping together the contributions of age, sex and education)—and the contribution of μ_t^m capturing the effect of changes in the distribution of the unemployed across segments with different average matching efficiency. The sum of these four components equal the contribution of the composition effect to μ_t.

[29] Recall that the effect of characteristics on an individual job finding rate was estimated only in the cross section. As a result, the composition effect does not mechanically adjust to capture the movements in matching efficiency over time. Instead, the composition effect only captures the movements in matching efficiency that are implied by changes in the distribution of worker across characteristics and by the effect (estimated with pre-2008 data) of each characteristic on search efficiency.

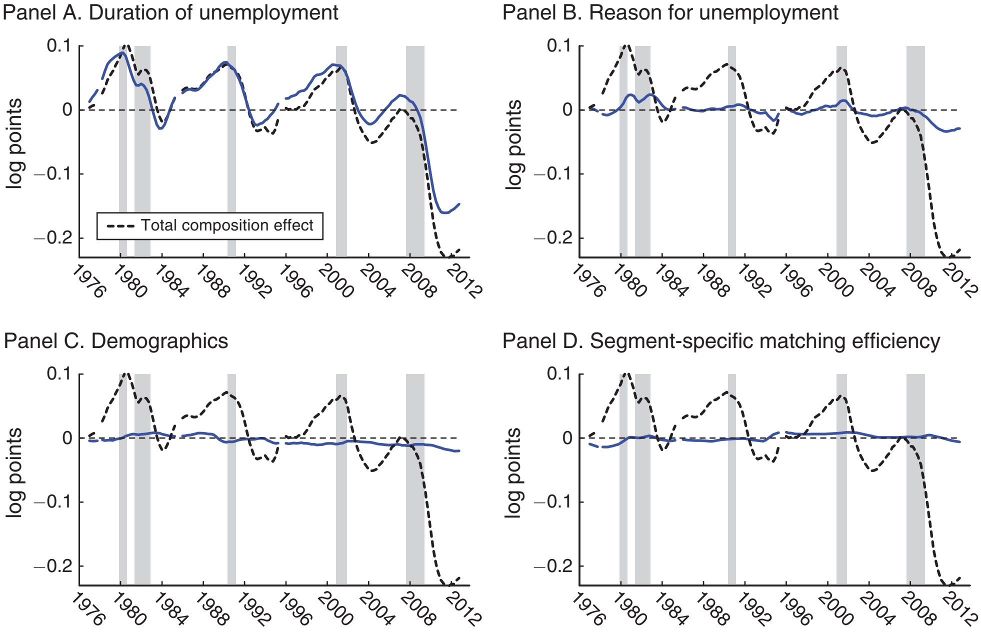

FIGURE 6

Notes: Decomposition of the composition effect into the contributions of unemployment duration, reason for unemployment, demographics (age, sex, and education), and segment-specific matching efficiency. The dashed line plots the total composition effect. The model was estimated over 1976–2007. All series are four-quarter moving averages. Grey bars indicate NBER recession dates.

Unemployment duration accounts for a large fraction of the composition effect, a perhaps not surprising result given the strength of duration dependence, and duration depresses matching efficiency in the aftermath of recessions. Reason for unemployment also lowers matching efficiency in recessions. This happens because recessions coincide with sharp increases in the fraction of permanent job losers (Figure 7), i.e., individuals with lower propensity to find a job, which worsens the employability of the unemployment pool.[30]

Interestingly, unemployment duration and reason for unemployment generate some inertia in the behavior of μ_t and thus in the behavior of the aggregate job finding rate. By definition, unemployment duration is an inertial variable, and average unemployment duration lags the cycle. As a result, the component of μ_t^s driven by duration also lags the cycle; peaking at the end of expansions and bottoming a few years into the recovery. Similarly, the fraction of permanent job losers in the unemployment pool is a persistent variable. Permanent job losers have a low search

[30] Note also that reason for unemployment tends to lift the job finding rate in recessions. This pattern owes to an increasing share of temporary job losers during recessions (especially before 1985, Figure 2). At the onset of recessions, bursts of temporary layoffs lift the job finding rate because job losers on temporary layoffs have a higher search efficiency than average. This was especially the case in the 1970s, and probably explains the sharp increases in the residual of the aggregate matching function regression (3) during the 1970 and 1974 recessions (Figure 1).

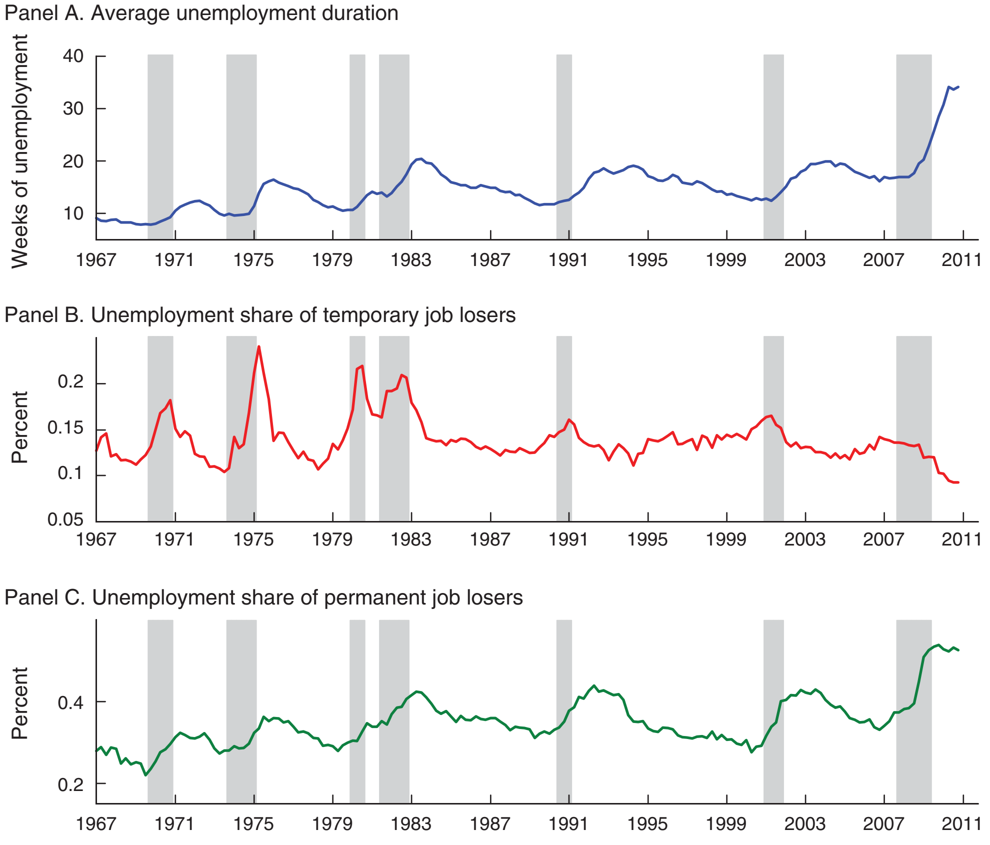

Figure 7

Notes: Average unemployment duration in weeks (panel A), share of job losers on temporary layoffs (panel B), and share of job losers on permanent layoffs (panel C), 1967–2010.

efficiency (Figure 2), and many years of expansion are necessary to bring their share back to prerecession levels (Figure 7).

Other characteristics play only a marginal role. Demographic characteristics have little effect on the cyclical behavior of aggregate matching efficiency, in line with Baker (1992),[31] and changes in the distribution of the unemployed across segments (μ_t^m) have virtually no effect.

C. *The Decline in Matching Efficiency since 2007*

We now turn to the behavior of matching efficiency since 2007.

As shown in Figure 4, both composition and dispersion drove down aggregate matching efficiency to exceptionally low levels. Moreover, since 2009—the end

[31] Demographics generated a downward trend in average search efficiency over the sample period because the labor force got older (search efficiency declines with age, Figure 2) and because the share of women in the labor force increased (women have lower search efficiency than men, Figure 2).

of the recession according to the NBER—both dispersion and composition have remained at high levels, keeping aggregate matching efficiency low and preventing unemployment from going down faster and participation from going up. Note also the very large contribution of duration during the recent recession (Figure 6). As average duration reached record highs (Figure 7), the average search efficiency of the unemployment pool deteriorated substantially, leading to a large decline in aggregate matching efficiency.[32]

It is interesting to contrast the recent large recession with the large recession of the early 1980s. While dispersion reached high levels in both recessions (Figure 5), the composition effect was much stronger in the 2008–2009 recession than it was in the 1980–1982 recession.[33] This is due to two effects: (i) a much larger increase in the share of long-term unemployed over 2008–2009 than over 1980–1982, and (ii) opposite contributions of reason for unemployment. In the 1980–1982 recession, reason for unemployment raised aggregate matching efficiency because of a sharp increase in the share of workers on temporary layoffs in the unemployment pool (Figures 6 and 7). In contrast, in the recent recession, the fraction of workers on temporary layoffs went down (firms rely less on temporary layoffs than in the early 1980s), while the fraction of workers on permanent layoffs went up, which lowered aggregate matching efficiency.

One last interesting difference between the 1980–1982 and 2008–2009 recessions is the behavior of aggregate matching efficiency *into the recovery*. Since 2009, both dispersion and composition have remained at high levels, keeping aggregate matching efficiency low and preventing unemployment from going down faster. This is in contrast to the early 1980s, where both dispersion and composition mean-reverted quickly after the end of the recession (Figures 4 and 5).

VI. Conclusion

This paper takes a "look into the [matching function] black-box" (Petrongolo and Pissarides 2001). We construct an aggregate matching function that explicitly takes into account worker heterogeneity as well as market segmentation. In this framework, and different from standard specifications of the aggregate matching function, matching efficiency is not a residual but is explicitly determined by the average characteristics of the labor market. We show how matching efficiency can move through a composition effect, due to changes in the composition of the unemployment pool, and through a dispersion effect, in which dispersion in labor market conditions drives down aggregate matching efficiency.

Our aggregate matching function can successfully capture the evolution of the aggregate job finding rate over 1976–2012, and we find that matching efficiency declined markedly after 2007 because the average characteristics of the unemployed worsened and because dispersion rose substantially.

[32] This last finding is consistent with the recent work of Kroft et al. (2013) who show that a search and matching model augmented with duration dependence in unemployment can account for a substantial fraction of the rise in long-term unemployment and the outward shift of the Beveridge curve during the Great Recession.

[33] Note also that, compared to 1980, μ_t entered 2008 at a much lower level, because both duration and the fraction of permanent job losers were not back to their pre-2001 level when the recession started.

An implication of our results is that heterogeneities across workers and labor markets are key aspects of unemployment fluctuations. As such, explicitly incorporating heterogeneity across agents and labor markets in search models are important research projects.[34]

Our empirical framework rests on the idea—often used in the mismatch literature—that the labor market is segmented into distinct submarkets. By not allowing for any interaction across segments, our framework shares the limitations of the mismatch literature regarding the appropriate definition of a labor market segment. As discussed by Abraham (1991), the size of a segment has consequences for the measurement of dispersion (as in our framework) and mismatch (as in Şahin et al. 2014). Since dispersion and mismatch come out of the concavity of the matching function, a higher level of disaggregation (i.e., a smaller definition of a segment) will mechanically generate a higher level of dispersion and thus a higher effect on matching efficiency. However, and counteracting the first effect, the smaller the segment, the more likely are workers to find a job outside of their local segment, leading the effect of mismatch on unemployment to be exaggerated (Abraham 1991). While our definition of a labor market segment appears reasonable and the empirical success of the framework is encouraging, an important task for future research is to model mobility decisions across segments and to better understand the link between the size of a segment and the matching process across segments.

Appendix

A. *A Decomposition of Movements in Aggregate Matching Efficiency*

Recall that aggregate matching efficiency is given by

$$\text{(A1)} \qquad \mu_t = \sum_{i=1}^{I} \frac{U_{it}}{U_t} \mu_i s_{it}^{\sigma} \left(\frac{\theta_{it}}{\theta_t} \right)^{1-\sigma},$$

with $s_{it} \equiv \sum_{j=1}^{J} \frac{U_{jit}}{U_{it}} s_{jit}$ the average search efficiency in segment i, $V_t \equiv \sum_{i=1}^{I} V_{it}$, and $U_t \equiv \sum_{i=1}^{I} U_{it}$ the total number of vacancies and unemployed in the economy, $\theta_{it} \equiv \frac{V_{it}}{U_{it}}$ the labor market tightness in segment i, and $\theta_t \equiv \frac{V_t}{U_t}$ the aggregate labor market tightness.

Without loss of generality, we normalize average search efficiency to 1, so that $s_0 \equiv \frac{1}{T} \sum_{t=1}^{T} \sum_{i,j} \frac{U_{ijt}}{U_t} s_{ijt} = 1$ by appropriately rescaling the μ_is.

Denote μ_0 the average matching efficiency level across segments with $\mu_0 \equiv \frac{1}{T} \sum_{t,i} \frac{U_{it}}{U_t} \mu_i$.

[34] For recent work in this direction, see Alvarez and Shimer (2011); Birchenall (2011); Merkl and Van Rens (2012); and Carrillo-Tudela and Visscher (2014).

Taylor expanding (A1) with respect to s_{ijt} around 1, μ_i around μ_0, and θ_{it} around θ_t to a second-order, the aggregate matching efficiency can be written

$$\text{(A2)} \qquad \ln f_t = \ln \mu_t + (1 - \sigma) \ln \theta_t,$$

$$\text{(A3)} \qquad \text{with } \mu_t \simeq \mu_0 \Bigg(1 + \underbrace{\mu_t^s + \mu_t^m}_{\text{Composition}} - \underbrace{\frac{\sigma(1 - \sigma)}{2} \mathrm{Var}\left(\frac{\theta_{it}}{\theta_t} \right)}_{\text{Dispersion}} \Bigg),$$

with

$$\begin{cases} \mu_t^s = \sigma \sum_{i,j} \dfrac{U_{jit}}{U_t} \big(s_{jit} - 1 \big) \\ \mu_t^m = \sum_i \dfrac{U_{it}}{U_t} \left(\dfrac{\mu_i}{\mu_0} - 1 \right) \\ \mathrm{Var}\left(\dfrac{\theta_{it}}{\theta_t} \right) = \sum_i \dfrac{U_{it}}{U_t} \left(\dfrac{\theta_{it}}{\theta_t} - 1 \right)^2, \end{cases}$$

where the second-order terms in μ_t^s and μ_t^m, as well as the cross-order terms, have been omitted for clarity of exposition, since they are in practice negligible.

Finally, using our specification to capture the effect of workers' characteristics on search efficiency

$$s_{jit} = e^{\beta X_{jit}},$$

with $X_{jit} = \big[1, x_{jit}^1, \ldots, x_{jit}^K \big]$ a vector of worker characteristics for type j in segment i at time t and $\beta = [\beta_0, \ldots, \beta_K]$ the corresponding vector of coefficients, we can decompose the composition effect as follows:

$$\mu_t^s = \sigma \sum_{i,j} \frac{U_{jit}}{U_t} \sum_{k=0}^{K} \beta_k \big(x_{jit}^k - \overline{x}^k \big),$$

with $\overline{x}^k = \frac{1}{T} \sum_{t=1}^{T} \sum_{i,j} \frac{U_{ijt}}{U_t} x_{ijt}^k$.

B. *Maximum Likelihood Estimation*

Recall from Section III that the job finding rate of an individual j in segment i, f_{jit}, is constant during "period t" and is given by

$$f_{jit} = \frac{s_{jit}}{s_{it}} \frac{m_{it}}{U_{it}}.$$

Thus, within period t, the duration of an individual's unemployment spell is characterized by an exponential distribution with parameter f_{jit}.

With data available at a monthly frequency, we set the period to a month, so that the probability that individual j in segment i finds a job within one month is given by

$$F_{jit} = 1 - e^{-\frac{m_{it}}{U_{it}}} \tag{A4}$$

$$= 1 - e^{-\frac{s_{jit}}{s_{it}} f_{it}},$$

with $\frac{s_{jit}}{s_{it}} = \frac{e^{\beta X_{jit}}}{\sum_j \frac{U_{jit}}{U_{it}} e^{\beta X_{jit}}}$. The job finding rate in segment i, f_{it}, can be measured from CPS data on worker transitions.[35]

For each individual j in segment i at time t, we observe whether he/she found a job within month t. Denoting $y_{jit} = \{1, 0\}$ the outcome of job search and treating the observations as independent and identically distributed across individuals and time, the log-likelihood function $\ell(\beta)$ is given by

$$\ell(\beta) = \sum_{t=1}^{T} \sum_{i=1}^{I} \sum_{j=1}^{J_i} \left[y_{jit} \ln\left(1 - F_{jit}\right) + (1 - y_{jit)} \ln F_{jit} \right],$$

where $y_{jit} = 1$ if individual j in segment i finds a job and $y_{jit} = 0$ otherwise, F_{jit} is given by (A4), and J_i is the number of individual observations in segment i. We estimate β by minimizing $\ell(\beta)$.

C. *Robustness Check: Estimation Using HWOL Data over 2006–2012*

The sample period covered by HWOL vacancy data starts only in 2006 and covers precisely the period in which matching efficiency displayed an unprecedented decline. Since the intention of the paper is to see how far taking labor market heterogeneities into account can go in accounting for the movements in the aggregate job finding rate after 2007, we cannot rely on post-2007 data to estimate our model.

Nonetheless, as a robustness check, it is instructive to estimate the model using HWOL data over 2006–2012. If our model is well-specified, and worker heterogeneity and dispersion do indeed explain movements in matching efficiency, estimating the model with post-2006 data should give relatively similar estimates to our baseline ones with pre-2006 data.[36] We find that this is indeed the case. As shown in Table 1, the matching function elasticity is estimated at 0.21 (only slightly higher than our baseline estimate of 0.18 despite very different sample periods), and the β estimates (capturing the effects of worker characteristics) are very similar (Figure 8), although duration dependence is estimated to be slightly stronger over 2006–2012 (more on this below).

[35] The monthly job finding probability in segment i, F_{it}, can be calculated from individual transitions between unemployment and employment in segment i from the law of large numbers, and the hazard rate f_{it} can be recovered from $f_{it} = -\ln(1 - F_{it})$.

[36] Another advantage of estimating the model with HWOL data is that we can estimate all model parameters simultaneously using a standard (one stage) maximum likelihood estimation.

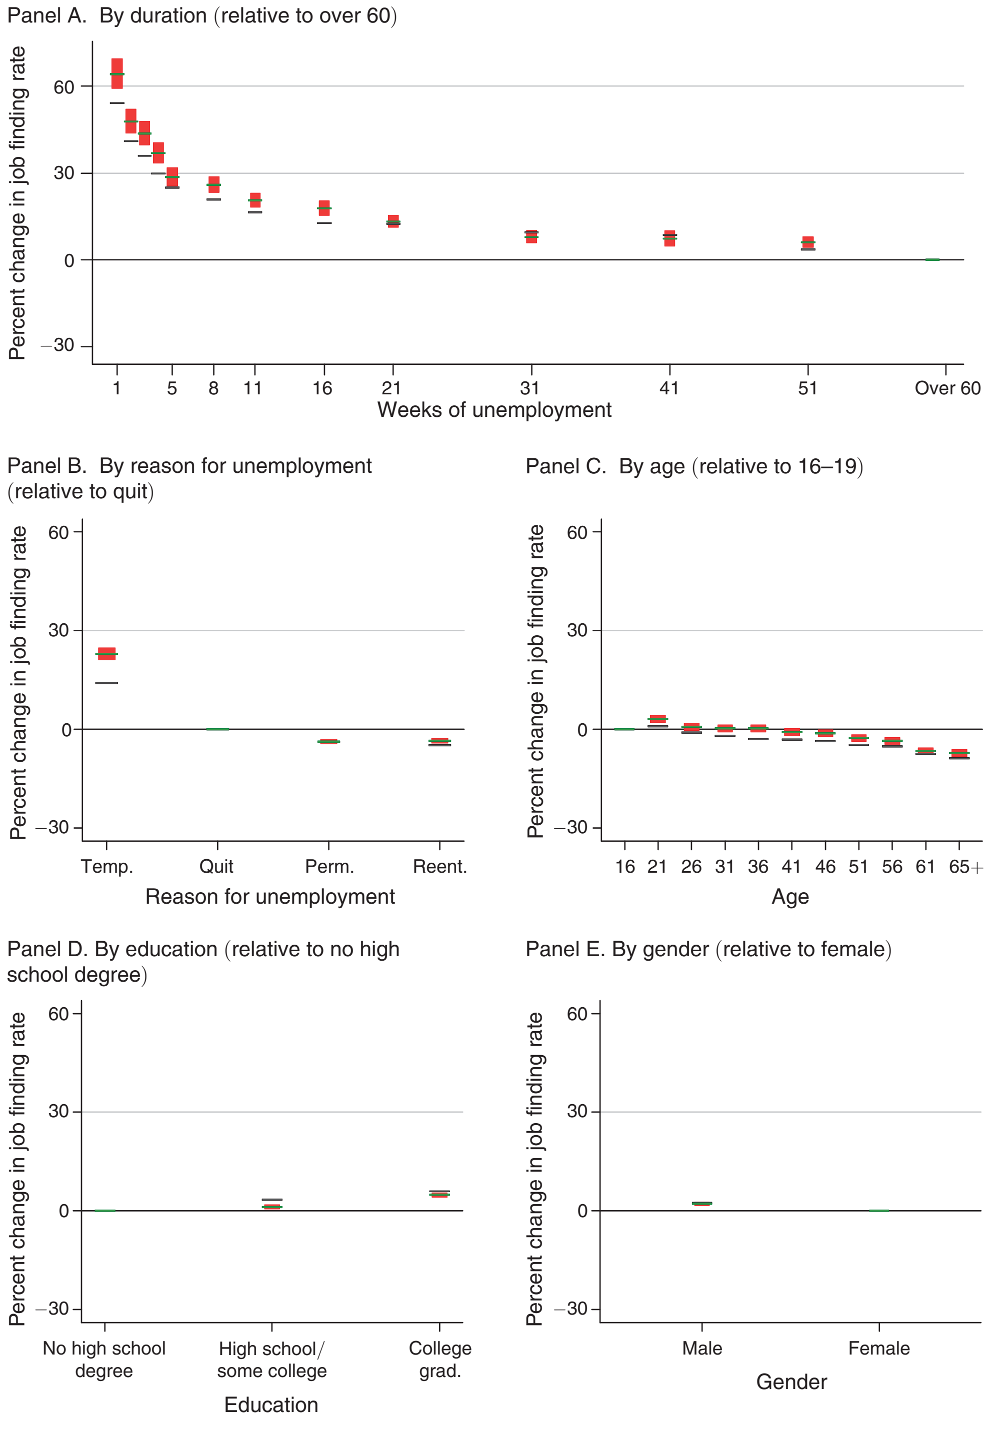

FIGURE 8. COMPARISON OF COEFFICIENT ESTIMATES

Notes: The dark grey lines refer to point estimates using HWOL data over 2006–2012 and the associated light grey bars denote ±2 standard-errors. The black lines refer to baseline point estimates using the two-stage procedure over 1994–2007.

Thus, the fact that the estimates are similar using pre-2006 or post-2006 data suggests that our framework provides an empirically successful characterization of the matching process.

Another advantage of doing a separate estimation with data covering only the Great Recession period is that it allows us to explore the robustness of our results to one possible critique—that duration dependence may be time-varying, leading us to possibly overestimate the contribution of duration.

While we imposed the effect of duration on an individual's job finding probability to be constant over time, recent research has shown that the effect of duration may actually vary over the cycle, leading us to possibly overestimate the contribution of duration in the Great Recession. Kroft, Lange, and Notowidigdo (2013) found that the effect of duration on the job finding rate is weaker in more depressed labor markets. By not allowing the strength of duration dependence (the slope of the duration dependence relationship) to vary over the business cycle, we could be overstating the contribution of duration to the decline in matching efficiency. Thus, one could worry that the large contribution of duration to the recent decline in aggregate matching efficiency is overstated because we did not allow the strength of duration dependence to vary over the business cycle (and become weaker during the Great Recession).

By comparing estimates obtained over 1976–2007 and over 2006–2012, Figure 8 shows that this worry is not warranted. We can see that the effect of duration on an individual's job finding rate was actually *stronger*, not *weaker*, during the last recession. As a result, the contribution of duration to the recent decline in matching efficiency could actually be even stronger, not weaker, than reported in Section V.

REFERENCES

Abraham, Katherine G. 1991. "Mismatch and Labour Mobility: Some final remarks." In *Mismatch and Labor Mobility,* edited by F. Padoa-Schioppa, 453–81. Cambridge: Cambridge University Press.

Abraham, Katharine G., and Lawrence. F. Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94 (3, Pt. 1): 507–22.

Alvarez, Fernando, and Robert Shimer. 2011. "Search and Rest Unemployment." *Econometrica* 79 (1): 75–122.

Baker, Michael. 1992. "Unemployment Duration: Compositional Effects and Cyclical Variability." *American Economic Review* 82 (1): 313–21.

Barlevy, Gadi. 2011. "Evaluating the Role of Labor Market Mismatch in Rising Unemployment." *Economic Perspectives* 35 (3): 82–96.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109: 175–78.

Barnichon, Regis, Michael Elsby, Bart Hobijn, and Ayşegül Şahin. 2012. "Which Industries are Shifting the Beveridge Curve?" *Monthly Labor Review* 135 (6): 25–37.

Barnichon, Regis, and Andrew Figura. 2015. "Labor Market Heterogeneity and the Aggregate Matching Function: Dataset." *American Economic Journal: Macroeconomics.* http://dx.doi.org/10.1257/mac.20140116.

Birchenall, Javier A. 2011. "A Competitive Theory of Mismatch." http://econ.ucsb.edu/~jabirche/Papers/mismatch.pdf.

Blanchard, Olivier Jean, and Peter Diamond. 1989. "The Beveridge Curve." *Brookings Paper on Economic Activity* 19 (1): 1–60.

Bleakley, Hoyt, and Jeffrey C. Fuhrer. 1997. "Shifts in the Beveridge Curve, Job Matching, and Labor Market Dynamics." *New England Economic Review* 10 (9): 3–19.

Borowczyk-Martins, Daniel, Grégory Jolivet, and Fabien Postel-Vinay. 2012. "Accounting For Endogeneity in Matching Function Estimation." *Review of Economic Dynamics* 16 (3): 440–51.

Butters, Gerard R. 1977. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies* 44 (3): 465–91.

Carrillo-Tudela, Carlos, and Ludo Visschers. 2014. "Unemployment and Endogenous Reallocation over the Business Cycle." http://www.bi.edu/InstitutterFiles/Samfunns%C3%B8konomi/Papers/Spring%202014/Visschers.pdf.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2010. "The Establishment-Level Behavior of Vacancies and Hiring." National Bureau of Economic Research (NBER) Working Paper 16265.

Devine, Theresa J., and Nicholas Kiefer. 1991. *Empirical Labor Economics: The Search Approach*. Oxford: Oxford University Press.

Elsby, Michael, Bart Hobijn, and Ayşegül Şahin. 2010. "The Labor Market in the Great Recession." *Brookings Papers on Economic Activity* 40 (1): 1–48.

Ghayad, Rand, and William Dickens. 2012. "What Can We Learn by Disaggregating the Unemployment-Vacancy Relationship?" Federal Reserve Bank of Boston Public Policy Brief 12-3.

Hall, Robert E., and Sam Schulhofer-Wohl. 2013. "Measuring Matching Efficiency with Heterogeneous Workers." Unpublished.

Herz, Benedikt, and Thijs van Rens. 2011. "Structural Unemployment." https://www.aeaweb.org/aea/2012conference/program/retrieve.php?pdfid=98.

Hobijn, Bart. 2012. "The Industry-Occupation Mix of U.S. Job Openings and Hires." Federal Reserve Bank of San Francisco Working Paper 2012-09.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124: 1403–48.

Jackman, R., and S. Roper. 1987. "Structural Unemployment." *Oxford Bulletin of Economics and Statistics* 49 (1): 9–36.

Layard, Richard, Stephen J. Nickell, and Richard Jackman. 2005. *Unemployment: Macroeconomic Performance and the Labour Market*. Oxford: Oxford University Press.

Lindeboom, Maarten, Jan van Ours, and Gusta Renes. 1994. "Matching Employers and Workers: An Empirical Analysis on the Effectiveness of Search." *Oxford Economic Papers* 46 (1): 45–67.

Kaitz, Hyman. 1970. "Analyzing the Length of Spells of Unemployment." *Monthly Labor Review* 93 (11): 11–20.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment." *Quarterly Journal of Economics* 128 (3): 1123–67.

Kroft, Kory, Fabian Lange, Matthew J. Notowidigdo, and Lawrence F. Katz. 2013. "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Non-Participation." http://korykroft.com/wordpress/KLNK_LTU_and_Great_Recession_june23_2014.pdf.

Machin, Stephen, and Alan Manning. 1999. "The Causes and Consequences of Long Term Unemployment in Europe." In *Handbook of Labor Economics,* Vol. 3, edited by Orley Ashenfelter and David Card, 3085–3139. Amsterdam: North-Holland.

Merkl, Christian, and Thijs van Rens. 2012. "Selective Hiring and Welfare Analysis in Labor Market Models." Institute for the Study of Labor (IZA) Working Paper 6294.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61 (3): 397–415.

Mueller, Andreas I. 2012. "Separations, Sorting and Cyclical Unemployment." Institute for the Study of Labor (IZA) Discussion Paper 6849.

Nekarda, Christopher J. 2009. "A Longitudinal Analysis of the Current Population Survey: Assessing the Cyclical Bias of Geographic Mobility." http://chrisnekarda.com/papers/lpd-geomob_20090527.pdf.

Padoa Schioppa, Fiorella. 1991. *Mismatch and Labor Mobility*. Cambridge: Cambridge University Press.

Petrongolo, Barbara. 2001. "Reemployment Probabilities and Returns to Matching." *Journal of Labor Economics* 19 (3): 716–41.

Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature* 39 (2): 390–431.

Pissarides, Christopher A. 1986. "Unemployment and Vacancies in Britain." *Economic Policy* 1 (3): 499–559.

Pissarides, Christopher A. 1992. "Loss of Skills during Unemployment and the Persistence of Employment Shocks." *Quarterly Journal of Economics* 107 (4): 1371–91.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. Cambridge: MIT Press.

Polivka, Anne, and Stephen Miller. 1998. "The CPS after the Redesign: Refocusing the Economic Lens." In *Labor Statistics and Measurement Issues,* edited by John Haltiwanger, Marilyn E. Manser, and Robert Topel, 249–89. Chicago: University of Chicago Press

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104 (11): 3529–64.

Salant, Stephen. 1977. "Search Theory and Duration Data: A Theory of Sorts." *Quarterly Journal of Economics* 91 (1): 39–57.

Sedláček, Petr. 2011. "Match Efficiency and the Cyclical Behavior of Job Finding Rates." http://www1.fee.uva.nl/toe/content/people/content/sedlacek/Documents/Working/Sedlacek_MatchEfficiency.pdf.

Shimer, Robert. 2007. "Mismatch." *American Economic Review* 97 (4): 1074–1101.

Shimer, Robert. 2008. "The Probability of Finding a Job." *American Economic Review* 98 (2): 268–73.

Shimer, Robert. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics* 15 (2): 127–48.

Veracierto, Marcelo. 2011. "Worker flows and matching efficiency." *Economic Perspectives* 35 (4): 147–69.