

IDENTIFYING EQUILIBRIUM MODELS OF LABOR MARKET SORTING

BY MARCUS HAGEDORN, TZUO HANN LAW, AND IOURII MANOVSKII¹

We assess the empirical content of equilibrium models of labor market sorting based on unobserved (to economists) characteristics. In particular, we show theoretically that all parameters of the classic model of sorting based on absolute advantage in [Becker \(1973\)](#) with search frictions can be nonparametrically identified using only matched employer–employee data on wages and labor market transitions. In particular, these data are sufficient to nonparametrically estimate the output of any individual worker with any given firm. Our identification proof is constructive and we provide computational algorithms that implement our identification strategy given the limitations of the available data sets. Finally, we add on-the-job search to the model, extend the identification strategy, and apply it to a large German matched employer–employee data set to describe detailed patterns of sorting and properties of the production function.

KEYWORDS: Sorting, assortative matching, search frictions, identification.

1. INTRODUCTION

DOES THE MARKET ALLOCATE THE RIGHT WORKERS TO THE RIGHT JOBS? Are complementarities between workers and employers important in determining output, productivity, and wages? Do large employers pay higher wages because they employ better workers? What are the sources of inter-industry wage differentials? What is the allocation of workers to employers that maximizes total output? These classic questions are at the heart of current debates in many areas of economics. In business cycle research, there is an ongoing discussion on whether the slow productivity and employment recovery after the Great Recession is due to the mismatch between human capital of unemployed workers and skill requirements of potential employers. In the international trade literature, researchers attempt to determine whether the wage premium of exporting firms is due to them being more productive or having better workers, a question with important implications for understanding the effects of changes in trade regimes. The industry dynamics literature is interested in the role of effective labor input reallocation across producers for productivity dynamics at the micro level. Misallocation at the micro level is relevant for the macro literature as it typically reduces total factor productivity with a potentially important impact on, for example, income differences across time and across countries. The enhanced focus on this role of resource misallocation represents one of the most important recent developments in the economic growth literature.

It has been long recognized that to make progress in studying these issues, it is essential to move the analysis beyond relying on the observable worker and firm attributes that account for only some 30% of the observed variation in wages. This involves expanding the scope of the analysis to include the study of assortative matching between

¹We would like to thank the editor and numerous anonymous referees as well as seminar participants at Arizona State, Chicago Fed, Collegio Carlo Alberto, Columbia, Einaudi Institute, Indiana, Mannheim, MIT, Notre Dame, Oslo, UPenn, Toulouse, Yeshiva, Vienna Institute for Advanced Studies, Bank of France, Yeshiva, Search and Matching Workshop at the Philadelphia Fed, SED Annual Meetings, NBER Summer Institute, Econometric Society Meeting, Cowles Summer Conference on “Sorting in Labor Markets,” Konstanz Workshop on Labor Market Search and the Business Cycle, Canadian Macro Study Group Meetings, Sandjberg conference, Human Capital Conference at Washington University in St. Louis, and Barcelona GSE Summer Forum on “Sorting: Theory and Estimation” for their comments. Support from the National Science Foundation Grants SES-0922406 and SES-1357903 is gratefully acknowledged. We are grateful to Kory Kantenga for his dedicated research assistance.

workers and employers based on their unobservable characteristics, which account for much of the remaining variation. These unobserved characteristics are typically associated, following the lead of [Abowd, Kramarz, and Margolis \(1999\)](#), with worker and firm fixed effects in wages that are estimated using longitudinal matched employer–employee data sets. Unfortunately, the literature has recently established that the key identifying assumptions of this regression approach are inconsistent with the standard equilibrium sorting models and that the worker and firm fixed effects identified using this methodology have no economic interpretation in the context of these models.² The key problem is that the assumption underlying the fixed effect regression is that wages are monotone in firm’s productivity (fixed effect). This is inconsistent with an explicit sorting model, where a productive firm may agree to hire a relatively unproductive worker only if that worker accepts a sufficiently low wage to compensate the firm for the option value of waiting for a more productive potential hire.

Faced with the limitations of the fixed effect regression approach, one might hope that an approach more firmly grounded in the theory of sorting models might prove more fruitful. From the perspective of economic theory, a typical starting point for thinking about assignment problems in heterogeneous agent economies is the model of [Becker \(1973\)](#). In labor market applications, the current state-of-the-art formulation is due to [Shimer and Smith \(2000\)](#), who extended the competitive framework in [Becker \(1973\)](#) to allow for time-consuming search between heterogeneous workers and firms. This framework is then a natural choice to answer the empirical questions motivating this research agenda. However, the empirical content of this model is not well understood. As a consequence, existing quantitative work on assortative matching in the labor market has to rely on strong assumptions on technology to be able to take the model to the data. This is problematic as it is these assumptions on technology that determine the patterns and consequences of sorting in the model.

The first contribution of this paper is to theoretically prove nonparametric identification of the model primitives, including the production function, from standard matched employer–employee data on wages and labor market transition rates. In other words, we establish that, from these data alone, one can recover the output of any observed employer–employee match and the consequences for output, productivity, and wages of moving any individual worker to any firm in the economy (subject to some limitations that will be formally spelled out below). Importantly, the proof does not impose strong assumptions on the production function but allows to infer its properties from the data. Moreover, the proof is constructive and relies on statistics that are fairly easy to interpret and to compute in the data. The second contribution of this paper is to develop an implementation algorithm for the proposed identification strategy.

Our identification strategy consists of three main steps. First, we need to globally rank workers. To accomplish this task, the literature typically relies on extremum statistics such as workers’ highest or lowest observed wages that rank workers in theory. Given that workers are observed being employed in only relatively few firms in the data and with a

²[Gautier and Teulings \(2006\)](#) were the first to establish this in a model of sorting based on comparative advantage. This important class of models violates the underlying assumption of the fixed effect regression that workers and firms are globally rankable. [Eeckhout and Kircher \(2011\)](#) later made an even stronger point. They proved that even in a model of sorting on absolute advantage that allows for globally rankable workers and firms, the worker and firm fixed effects in wages have no relationship to underlying productivities. These theoretical insights have been confirmed quantitatively in a range of assortative matching models in [Lopes de Melo \(2013\)](#), [Lentz \(2010\)](#), and [Lise, Meghir, and Robin \(2016\)](#), among others.

plausible amount of measurement error in recorded wages, such statistics result in relatively noisy rankings. The key insight we offer is that comparisons of worker wages with wages of her coworkers, coworkers of her coworkers at other firms, etc., provide an enormous amount of information that can be used to infer the accurate working ranking. The precise way this information can be exploited is model-specific. However, as ranking workers is the foundational step in identification of this class of models, exploiting this information seems essential. The way we implement this idea in the [Shimer and Smith \(2000\)](#) model is as follows. In this model, workers can be ranked based on their wages within firms (potentially observed with an error). Workers who change firms provide links between the partial rankings inside the firms they work at. This enables us to solve a rank aggregation problem which effectively maximizes the likelihood of the correct global ranking. This problem is equivalent to the problem of how to aggregate rankings of candidates submitted by voters in the social choice literature. These problems are extremely computationally complex (they are NP-hard), but, fortunately, the computer science literature has recently made substantial progress in designing computational algorithms that can efficiently approximate their solution. We draw on these advances in algorithm research to develop a method that is fast and accurate for the applications we study.

The second key insight relates to ranking of firms. We show that the value of a vacant job, or the surplus a vacancy is expected to generate, is increasing in its productivity. We expect this property to hold in most empirically relevant models based on our empirical findings reported below. Standard assumptions on wage determination imply that both parties benefit from an increase in the match surplus. This implies that more productive firms expect to deliver higher surplus to the workers they hire. To operationalize this insight, we show that firms can be ranked based on the expected average difference between the wages they pay to each of their workers and the reservation wages of those workers. This is a simple statistic to compute, but it relies on having an accurate estimate of the reservation wage for each worker, which might be difficult to obtain in short samples. The ostensibly simple but crucial methodological insight we offer is that once workers are accurately globally ranked, similarly ranked workers must have similar reservation wages. Thus, we can estimate the reservation wage by considering a group of similar workers, despite the fact that each of those workers is observed for a relatively short period of time.

Being able to rank firms and workers allows us to recover the output of every match. In the model, wages, which are observed in the data, are a function of the output of the match as well as of two objects that our identification strategy allows to measure—the reservation wage of a worker, and the value of a vacancy. Thus, the wage equation can be solved for output as a function of three measurable variables. While the [Shimer and Smith \(2000\)](#) model is particularly convenient in that it implies an invertible wage equation, in many other models the inversion can be achieved using the equation for match surplus. We expect these insights to form the basis of any attempt to nonparametrically estimate the production function in models of labor market sorting. The key potential impediment to an accurate recovery of the production function in available data samples is the presence of measurement error. We show that this problem can be overcome by once again exploiting the insight that similarly ranked workers and firms can be binned and the production function estimated at the bin level.

We assess the performance of the proposed methods in a Monte Carlo study imposing the limitations (on sample size, frequency of labor market transitions, measurement error, etc.) of the commonly used matched worker-firm data sets. We find that the identification strategy and the implementation method that we develop are successful at measuring the relevant objects in the model.

Thus, in the first part of the paper, we develop all the theoretical and computational tools required to enable the empirical analysis using the [Becker \(1973\)](#) model with time-consuming search. We focus our theoretical analysis on its formulation in [Shimer and Smith \(2000\)](#) because of its well-understood theoretical properties. We also think it has considerable pedagogical merit to understand the sources of identification and to tackle the key implementation issues in the simplest possible but relevant model.

An important limitation of the model in [Shimer and Smith \(2000\)](#) is that it does not include search on the job, which is a key feature of the data. Thus, the third major contribution of the paper is to make the model empirically relevant by introducing on-the-job search. We prove nonparametric identification of that version of the model and verify the performance of the proposed methods in a Monte Carlo study. The key identification steps and insights are the same as in the baseline model, with some minor modification required by the change in the model structure.³

The fourth contribution of the paper is an empirical analysis, in which we nonparametrically estimate the model with on-the-job search using a large German matched employer–employee data set. We find a very strong degree of sorting with a rank correlation of 0.75 between workers and firms. Firms matching with more productive workers also have a much higher value of the vacancy. This finding is not hardwired by our estimation strategy, but indicates that firms cannot scale up production arbitrarily and drive the value of the vacancy down to zero at the firm level as is assumed in many macro models. While overall more productive workers tend to work in more productive firms, locally, the patterns of sorting are much more complicated. In particular, in contrast to the standard assumptions of the globally sub- or supermodular production function, the cross-partial derivative of the production function does not have a constant sign. This curvature is relatively well exploited by market participants. In particular, solving the optimal output maximizing assignment problem, we find that optimally assigning individual workers to individual firms increases output only by 1.83%. In contrast, reassigning workers to the main diagonal, as would be optimal given the typical assumption of a globally supermodular production function, would imply a 0.23% *decline* in output. This highlights the importance of a non-parametric recovery of the production function, especially for counterfactual analysis.

The paper is organized as follows. In Section 2, we describe the standard model with frictional labor market and assortative matching between workers and firms. Section 3 shows theoretically the identification of the model. In Section 4, we develop computational tools needed to implement our identification strategy and evaluate its performance in simulated data sets designed to mimic existing matched employer–employee data sets. In Section 5, we extend the model to include on-the-job search and show how to apply our identification strategy in this environment. Next, we use this methodology to measure the degree of sorting, identify the production function, and estimate the gains from eliminating search frictions in German data. Section 6 concludes. Most proofs and details of computations are in the Supplemental Material ([Hagedorn, Law, and Manovskii \(2017\)](#)).

2. THE ECONOMIC MODEL

The model description builds on [Shimer and Smith \(2000\)](#), who added time-consuming search to [Becker \(1973\)](#), with slight generalizations and some modifications. In particular, we do not impose symmetry between the two sides of the market, but have workers on the

³[Lamadon, Lise, Meghir, and Robin \(2014\)](#) showed identification of a different sorting model with search on the job.

one side and firms on the other, both sides with potentially different primitives. We also use a linear search technology instead of the quadratic search technology in [Shimer and Smith \(2000\)](#), which seems the better choice for labor market applications. None of our results hinge on this modification.

2.1. Environment

2.1.1. Basics

Time is discrete, all agents are infinitely-lived and maximize the present value of pay-offs, discounted with a common discount factor $\beta \in (0, 1)$. The unit mass of workers is either employed (e) or unemployed (u), while firms are either producing (p) or vacant (v). Workers and firms are heterogeneous with respect to their productivities, denoted by $x \in [0, 1]$ and $y \in [0, 1]$, respectively. To simplify the exposition, we treat each firm as having one job. All the results immediately generalize, however, to each firm having a mass of jobs sharing the same productivity y .⁴

Output of a match between worker x and firm y is given by the twice differentiable non-negative production function $f : [0, 1]^2 \rightarrow \mathbb{R}_+$. The existence proof in [Shimer and Smith \(2000\)](#) also requires that f has uniformly bounded first partial derivatives on $[0, 1] \times [0, 1]$. It is assumed that match output is increasing in worker and firm type, that is, $f_x > 0$ and $f_y > 0$.⁵ This assumption allows x and y to be measured as a worker's or a firm's rank in the corresponding productivity distribution. The rank of a worker (firm) is given by the fraction of workers (firms) who produce weakly less with the same firm (worker). In this paper, *productivity*, *rank*, or *type* have identical meanings. Therefore, the distributions of worker and firm types are both uniform. If the "original" (non-rank) distributions of worker and firm types are F and G , respectively, and the "original" production function is $\hat{f}(\hat{x}, \hat{y})$, then we transform the production function

$$f(x, y) = \hat{f}(F^{-1}(x), G^{-1}(y))$$

and the distributions are $F(\hat{x}) = x$, $G(\hat{y}) = y$.

We place no additional assumptions on the production function (except for mild technical conditions that ensure existence of an equilibrium). In particular, we do not assume that sorting is either positive or negative, but show how to recover this information from the data.

⁴This model of the firm, as simplistic as it is, represents the current state of the art in this literature. As [Lentz and Mortensen \(2010, pp. 593–594\)](#) put it, "all the analyses that we know of assume that output of any given job-worker match is independent of the firm's other matches. Furthermore, firm output is the sum of all the match outputs. Hence, the identification challenge reduces to that of identifying worker and firm contributions over matches and a common match production function. Of course, as the research frontier moves to improve our understanding of multiworker firms, it is likely and appropriately an assumption that will be challenged." We agree with this assessment and hope the identification results established here will continue to be relevant as more sophisticated and empirically implementable theories of the firm are developed.

⁵The assumption that economic agents can be globally ranked is standard in the models of sorting based on absolute advantage, such as [Becker \(1973\)](#) and [Shimer and Smith \(2000\)](#), and is implicit in the approach of [Abowd, Kramarz, and Margolis \(1999\)](#). In this paper, this assumption is only relevant for identifying rankings of workers and firms when they can be ranked. In [Hagedorn, Law, and Manovskii \(2014\)](#), we showed that if some agents cannot be ranked, for example, firms in the comparative advantage model of [Gautier and Teulings \(2012\)](#), our identification strategy will reveal this and it will continue to recover the production function correctly.

TABLE I
FUNCTIONS DESCRIBING DISTRIBUTIONS

Description	Density Function
Matches	$d_m(x, y)$
Employed workers	$d_e(x) = \int d_m(x, y) dy$
Unemployed workers	$d_u(x) = d_w(x) - d_e(x)$
Producing firms	$d_p(y) = \int d_m(x, y) dx$
Vacant firms	$d_v(y) = d_f(y) - d_p(y)$

2.1.2. Distributions

The measures characterizing the set of matched and unmatched workers and firms are assumed to be absolutely continuous, implying the existence of a density. Given our identification of types with ranks, the worker and firm time-invariant populations are given by $d_w = 1$ and d_f . The distribution of producing matches is described by $d_m : [0, 1]^2 \rightarrow \mathbb{R}_+$. The functions characterizing the employed and unemployed workers as well as the producing and vacant firms are denoted $d_e(x)$, $d_u(x)$, $d_p(y)$, and $d_v(y)$, respectively.⁶ Table I summarizes the relationships between these functions. Integrating the densities from Table I gives the time-invariant measures of aggregate employment, $E = \int d_e(x) dx$, of unemployment, $U = \int d_u(x) dx$, of producing firms, $P = \int d_p(y) dy$, and vacant firms, $V = \int d_v(y) dy$.

2.1.3. Timing

It is convenient to think of each period as consisting of two subperiods. In the first subperiod, a worker of type x matched with a firm of type y produces $f(x, y)$. Output of this match is exhausted by payments to the firm, $\pi(x, y)$, and the worker, $w(x, y)$. There is free entry of vacancies. Creating a vacancy costs a fixed cost c . After paying this cost, the vacancy learns its productivity y , which is a random draw from the uniform distribution on $[0, 1]$, implying that the time-invariant firm distribution is uniform, $d_f = 1/(V + P)$. In the second subperiod, new matches are formed when all unmatched workers and firms participate simultaneously in a single labor market subject to search frictions. After matching, existing matches (including newly formed ones) are destroyed with probability δ .⁷

2.2. Search and Matching

Only and all unmatched agents engage in random search.⁸ A function $m : [0, 1] \times [0, 1] \rightarrow [0, \min(U, V)]$ takes the masses of unemployed workers U and vacant firms V as

⁶Note that these functions do not integrate to 1, but to the mass of employed and unemployed workers and producing and vacant firms, respectively.

⁷The assumption that newly formed matches are also subject to job destruction shocks enhances the elegance of some expressions below but has no relevance for the substantive results.

⁸Random search means that workers and firms do not observe the types of their potential trading partners prior to meeting them, that is, they have the same information as is available to the econometrician (e.g., age, sex, education, occupation, etc. of a worker and industry, location, etc. of a firm). An alternative assumption is that workers (firms) know the type y of every firm (type x of every worker) and can direct their search to specific types (e.g., Moen (1997), Shi (2001), Shimer (2005), Eeckhout and Kircher (2010)), for example, workers direct their search to firms that are willing to accept them. In the analysis below, these informational assumptions matter only for the computation of the job filling probability for firms. These informational assumptions will not affect the analysis at all if the data allow to observe the number of vacancies at individual firms (as in, e.g.,

its inputs and generates meetings. The probability a worker meets a potential employer is given by $\mathbb{M}_u = \frac{m(U,V)}{U}$, while the probability of a vacant firm meeting a potential hire is $\mathbb{M}_v = \frac{m(U,V)}{V}$. These probabilities are time-invariant in the steady-state equilibrium we will consider. The probability for a worker to meet any firm $y \in Y \subseteq [0, 1]$ equals $\mathbb{M}_u \frac{\int_Y d_v(y) dy}{V}$. The probability for a firm to meet any worker $x \in X \subseteq [0, 1]$ equals $\mathbb{M}_v \frac{\int_X d_u(x) dx}{U}$. These probabilities reflect our assumption of a linear search technology. Using the quadratic search technology in [Shimer and Smith \(2000\)](#), these probabilities would be $\mathbb{M}_u \int_Y d_v(y) dy$ and $\mathbb{M}_v \int_X d_u(x) dx$, respectively. Since we obtain the same search technology by simply setting $U = V = 1$ in the matching process, it will become clear that our results do not depend on the returns to scale of the matching function.

Not all meetings necessarily result in matches. Some meetings are between workers and firms who are unwilling to consummate a match and who prefer to continue the search process.

2.3. Strategies, Acceptance Sets, and Surplus

The steady-state pure strategy of a worker of type x is to decide which firms to match with, taking all other strategies as given. This strategy is described by a Borel measurable *acceptance set* $A^w(x)$ of firms that a worker type x is willing to match with. Symmetrically for firms, the Borel measurable *acceptance set* $A^f(y)$ is composed of the workers that a firm of type y is willing to match with. Matching takes place when both the worker and the firm find it mutually acceptable. For a worker of type x , the *matching set* $B^w(x)$ consists of firms which accept worker type x and are accepted by worker type x . Similarly, for a firm of type y , $B^f(y)$ consists of workers who accept to match with firm type y and who are accepted by firms of type y . Specifically,

$$B^w(x) \equiv \{\tilde{y} : x \in A^f(\tilde{y}) \wedge \tilde{y} \in A^w(x)\},$$

$$B^f(y) \equiv \{\tilde{x} : y \in A^w(\tilde{x}) \wedge \tilde{x} \in A^f(y)\}.$$

$\overline{B^w}$ and $\overline{B^f}$ denote the complements of B^w and B^f , respectively. Define \mathcal{B} to represent all (x, y) pairs that form in equilibrium:

$$\begin{aligned} \mathcal{B} &\equiv \{(x, y) : y \in A^w(x) \wedge x \in A^f(y)\} \\ &= \{(x, y) : y \in B^w(x)\} \\ &= \{(x, y) : x \in B^f(y)\}. \end{aligned}$$

2.4. Bellman Equations and Surplus Sharing

Let $V_u(x)$ denote the value of unemployment for a worker of type x , $V_e(x, y)$ the value of worker x employed at a firm of type y , $V_v(y)$ the value of a vacancy for firm y , and $V_p(x, y)$ the value of firm y employing a worker of type x . The surplus of a match between worker x and firm y is defined as

$$S(x, y) \equiv V_p(x, y) - V_v(y) + V_e(x, y) - V_u(x).$$

the German LIAB data that we use below). In this case, one can compute the job filling rate directly without the need to make any informational assumptions. Without data on vacancies, the computation of the job filling rate is conditional on the specification of the matching process.

Shimer and Smith (2000) assumed that wages are determined by Nash bargaining over the match surplus $S(x, y)$ between workers and firms who have equal bargaining powers. We maintain this assumption in this paper, although it is not essential. First, we show below that the assumption of equal bargaining powers can be relaxed and the bargaining power can be identified in the data if the model incorporates either an idiosyncratic or an aggregate stochastic component affecting, say, firm productivity. In terms of notation, we allow for unequal bargaining powers by denoting workers' bargaining power $\alpha \in (0, 1)$ ($\alpha = \frac{1}{2}$ corresponds to the model in Shimer and Smith (2000)). Second, our method for identifying the sign and strength of sorting does not use the assumption of Nash bargaining but applies to any bargaining game whose solution implies that payoffs to both parties increase in match surplus. Finally, our method for the nonparametric identification of the production function only relies on specifying the bargaining protocol which yields a wage equation that can be inverted for output.

Generalized Nash bargaining over the match surplus with workers' bargaining power α implies

$$(1) \quad \begin{cases} \alpha S(x, y) = V_e(x, y) - V_u(x), \\ (1 - \alpha)S(x, y) = V_p(x, y) - V_v(y). \end{cases}$$

Following this rule, it is clear that $y \in A^w(x)$ if and only if $x \in A^f(y)$. Hence,

$$(2) \quad \begin{cases} A^w(x) = B^w(x) = \{y : S(x, y) \geq 0\}, \\ A^f(y) = B^f(y) = \{x : S(x, y) \geq 0\}. \end{cases}$$

Using the surplus sharing rule (1), we obtain the following steady-state value functions (the derivations of these equations are provided in Section A.1 of the Supplemental Material):

$$(3) \quad V_u(x) = \beta V_u(x) + \beta \alpha (1 - \delta) \mathbb{M}_u \int_{B^w(x)} \frac{d_v(\tilde{y})}{V} S(x, \tilde{y}) d\tilde{y},$$

$$(4) \quad V_v(y) = \beta V_v(y) + \beta (1 - \alpha) (1 - \delta) \mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} S(\tilde{x}, y) d\tilde{x},$$

$$(5) \quad V_e(x, y) = w(x, y) + \beta V_u(x) + \beta \alpha (1 - \delta) S(x, y),$$

$$(6) \quad V_p(x, y) = f(x, y) - w(x, y) + \beta V_v(y) + \beta (1 - \alpha) (1 - \delta) S(x, y).$$

Free entry requires

$$(7) \quad c = \int V_v(y) dy.$$

2.5. Stationary Distribution of Matches

In the stationary match distribution, for all worker and firm type combinations in the matching set, the numbers of destroyed and created matches are the same:

$$(8) \quad \forall (x, y) \in \mathcal{B}, \quad \underbrace{\delta d_m(x, y)}_{\text{destruction}} = \underbrace{(1 - \delta) d_u(x) \mathbb{M}_u \frac{d_v(y)}{V}}_{\text{new match formation}}.$$

The probability for a worker (of any type) to meet a firm of type y is the product of the probability to meet any firm, \mathbb{M}_u , and the probability that this firm is of type y , $\frac{d_v(y)}{V}$. This is multiplied by $(1 - \delta)$ because newly formed matches can get destroyed in the same period. Integrating over all matches yields that the total inflow into unemployment equals the total outflow out of unemployment:

$$\underbrace{\int_{\mathcal{B}} \delta d_m(x, y) dx dy}_{\text{inflow}} = \delta E = \underbrace{(1 - \delta) \int_0^1 d_u(x) \mathbb{M}_u \int_{B^w(x)} \frac{d_v(y)}{V} dy dx}_{\text{outflow}}.$$

2.6. Equilibrium

In a *steady-state search equilibrium* (SE), all workers and firms maximize expected pay-offs, taking the strategies of all other agents as given.⁹ The economy is in steady-state. A SE is then characterized by the density $d_u(x)$ of unemployed workers, the density $d_v(y)$ of vacant firms, the density of formed matches $d_m(x, y)$, and wages $w(x, y)$. The density $d_m(x, y)$ implicitly defines the matching sets, as it is zero if no match is formed and is strictly positive if a match is consummated. Wages are set to ensure the surplus sharing rule (1) and match formation is optimal given wages w , that is, a match is formed whenever the surplus is (weakly) positive (see (2)). The densities $d_u(x)$ and $d_v(y)$ ensure that the flow equations in (8) hold.

To prove existence, [Shimer and Smith \(2000\)](#) assumed that the production function is either globally supermodular or globally submodular.¹⁰ A stronger assumption would be to require that the production function induces either positive assortative matching (PAM) or negative assortative matching (NAM), defined as follows:

DEFINITION 1: Consider worker types $x_1 < x_2$ and firm types $y_1 < y_2$. There is PAM if $x_1 \in B^f(y_1)$ and $x_2 \in B^f(y_2)$ whenever $x_1 \in B^f(y_2)$ and $x_2 \in B^f(y_1)$. There is NAM if $x_1 \in B^f(y_2)$ and $x_2 \in B^f(y_1)$ whenever $x_1 \in B^f(y_1)$ and $x_2 \in B^f(y_2)$.

Whereas this stronger assumption is not necessary for the existence proof, it is commonly imposed in the literature as we discuss below. The equilibrium existence proof in [Shimer and Smith \(2000\)](#) also uses their assumption of a quadratic matching function. [Nöldeke and Tröger \(2009\)](#) extended the proof to a linear matching technology used in this paper and showed that if f is either supermodular or submodular, then a SE exists. [Shimer and Smith \(2000\)](#) suggested that the assumption of either super- or submodularity just avoids a more complicated existence proof and thus can be dispensed with. More specifically, this assumption rules out an atom of zero surplus matches, that is,

$$(9) \quad \forall x \neq x': \quad \mu(\{y : S(x, y) = S(x', y) = 0\}) = 0,$$

where μ is the Lebesgue measure. Imposing

$$\forall x \neq x', \forall y: \quad \mu(\{y' : f(x, y) + f(x', y') = f(x, y') + f(x', y)\}) = 0$$

⁹As in [Shimer and Smith \(2000\)](#), we assume that a match is formed if agents are indifferent.

¹⁰A production function is supermodular if the cross-derivative is positive and it is submodular if the cross-derivative is negative.

ensures this property. It thus avoids both the assumption of super- or submodularity and also a more complicated existence proof (see Step 1 of the proof of Lemma 3 in [Shimer and Smith \(2000\)](#)). This property is, for example, satisfied by the two production functions used in [Shimer and Smith \(2000\)](#) as examples which satisfy neither PAM nor NAM: $(x + y)^2$ and $(x + y - 1)^2$. It does not hold for modular production functions such as $x + y + k$ (k is a constant). However, for large enough k , every worker matches with every firm and thus (9) is trivially satisfied. Thus, a SE exists.

3. THE ECONOMETRIC MODEL: IDENTIFICATION

The description of the econometric model requires to determine which variables are observable and which are unobservable. The identity of a worker i and of a firm j are observed but their respective types $x(i)$ and $y(j)$ are not. The wage is only observed with mean zero measurement error ε_i , which is independent from all other variables, so that the observed wage of a worker i employed at firm j equals

$$(10) \quad w(x(i), y(j)) + \varepsilon_{i,j,t}.$$

All remaining variables or model primitives are unobserved.¹¹

The model is (fully) *identified* if a unique function from the joint distribution of observables to (all) the underlying elements of the model exists. In particular, different model primitives generate different joint distributions of observables, that is, they are observationally not equivalent.¹² In this section, we establish that the model is identified by providing a unique mapping from the joint distribution of wages, worker, and firm identity to the primitives of the model. The identification proof is constructive: We express the model parameters in terms of the observable distribution and we also use these expressions to recover the primitives of the model in the implementation in small samples. We proceed in three steps. First, we show how to identify the ranking of workers, that is, the mapping $x(i)$. Second, we identify the ranking of firms, that is, the mapping $y(j)$. Having identified the rankings of workers and firms, an investigation of the empirical matching patterns allows us to identify the presence and sign of sorting. Third, we identify the remaining primitives of the model, in particular, the output of every observed match between any worker and any firm.

Using the joint distribution of wages, workers, and firms, we can infer the conditionally expected wage

$$(11) \quad E(w(x(i), y(j)) + \varepsilon_{i,j,t} \mid i, j) = w(x(i), y(j)),$$

since the measurement error has mean zero and is independently distributed. To prove identification, we can therefore proceed under the assumption that we observe $w(x(i), y(j))$, which is free of measurement error. Measurement error is, however, obviously a potential impediment in small samples and we show how to deal with this issue successfully in the implementation section below.

¹¹In some data sets, such as the German LIAB data used in this paper, the number of vacancies $v(j)$ posted by firm j is observed. Adding this to the lists of observables is not necessary but simplifies the measurement.

¹²This is the standard definition of identification in the literature building on [Hurwicz \(1950\)](#) (see [Matzkin \(2013\)](#) for a recent survey). As [Matzkin \(2013\)](#) explained, “The analysis of identification is separate from statistical issues, which are dependent on sample size. Identification analysis assumes that the whole probability distribution of the observable variables, rather than a sample from it, is available.”

3.1. Ranking Workers

We now derive several statistics which are monotonically increasing in worker types. Such statistics naturally provide a way to rank workers.

The easiest such statistic is the value of unemployment. It is increasing in a worker's type because a more productive worker can always imitate the acceptance strategy of the less productive worker but produce more and consequently receive higher wages. This induces a more productive worker to set a higher reservation wage. As the production function and the value of unemployment increase in worker productivity, wages within firms are also increasing in worker type. This yields for every firm a correct ranking of workers in the matching set of that firm. If one firm were to match with all workers in the economy, the ranking of workers based on wages in that firm would automatically represent a global ranking of all workers in the economy. If no firm matches with all workers, we have to aggregate the partial within-firm rankings to a global one. To illustrate how this works, consider a firm A which hires workers $a_1 < a_2 < \dots < a_N$ and another firm B which hires workers $b_1 < b_2 < \dots < b_M$, where the ranking within each firm is denoted by " $<$." Now suppose there is an overlap in the matching sets of these two firms so that the best ranked workers in firm A are lowest ranked workers in firm B , that is, for some k , $a_{N-k} = b_1$, $a_{N-k+1} = b_2$, \dots , $a_N = b_{k+1}$. We can then combine the two rankings to rank all workers in the two firms to obtain $a_1 < a_2 < \dots < a_{N-k} = b_1 < a_{N-k+1} = b_2 < \dots < a_N = b_{k+1} < b_{k+2} < \dots < b_M$. Iterating yields a global ranking of workers under the mild assumption that the set of workers can be split into overlapping matching sets. As the matching sets cannot be guaranteed to be overlapping, we provide three further rankings—the highest wage of a worker, the lowest wage of a worker, and the adjusted average wage—which provide global rankings of workers.¹³ Thus, these rankings can be used to initialize the rank aggregation procedure and this ensures a resulting global ranking even in cases with non-overlapping matching sets.

Let $y^{\min}(x)$ be the firm that pays the lowest wage accepted by a worker of type x and $y^{\max}(x)$ be the firm that pays the highest wage to a worker of type x . In Section A of the Supplemental Material, we prove the following result.

RESULT 1:

- (i) $V_u(x)$, $V_e(x, y)$, and $w(x, y)$ are increasing in x .
- (ii) The lowest wage, given by $w(x, y^{\min}(x))$, is increasing in x .
- (iii) The highest wage, given by $w(x, y^{\max}(x))$, is increasing in x .
- (iv) The adjusted average wage, defined as

$$(12) \quad w^{av}(x) \equiv \left(1 - \mathbb{M}_u + \delta \mathbb{M}_u + \mathbb{M}_u (1 - \delta) \int_{B^w(x)} \frac{d_v(y)}{V} dy \right) w(x, y^{\min}(x)) \\ + \mathbb{M}_u (1 - \delta) \int_{B^w(x)} \frac{d_v(y)}{V} w(x, y) dy,$$

is increasing in x .

¹³Flinn and Heckman (1982) and Wolpin (1987) represent some of the earliest work on order statistic estimators such as the lowest or highest wages.

Note that while the adjusted average wage is increasing in x , the average wage (without the adjustment) is not.¹⁴ To see this, consider two workers with different productivities. A more productive worker might be matching with a wider set of firms (some of which do not accept the less able worker). However, the more able worker might be only marginally acceptable to those firms because they typically match with even better workers. As a consequence, those firms pay low wages to this worker. Thus, the average wage of the worker over his employment history might be lower than that of a less productive worker. The more productive worker still obtains higher utility because he spends a larger fraction of his lifetime employed. Result 1(iv) corrects for this effect by imputing the value of unemployment to unemployed workers and defining the average wage over the lifetime rather than of the portion of lifetime the worker spends employed.

We have derived a number of statistics that provide theoretically valid and equivalent rankings of workers. In Section 4, we discuss their implementation and assess their performance in small samples and in the presence of measurement error in wages. We find that the best way to rank workers is to use the global statistics to initialize the ranking and then refine it by aggregating within-firm rankings. For a realistic amount of worker mobility across firms, this yields a very accurate complete ranking of workers.

3.2. Ranking Firms

To rank firms, we derive a statistic which is monotonically increasing in firm type y .¹⁵ This is nontrivial since the wage of worker x , $w(x, y)$, is not always increasing in firm productivity. The same problem applies to the surplus of a match, $S(x, y)$. Our strategy is as follows. We first establish that the value of a vacancy is increasing in y . This implies that the surplus a vacancy is expected to generate is also increasing in y . Any bargaining game where both parties benefit from an increase in the surplus implies that the average surplus of workers employed by firm y is also increasing in y . Finally, we show that the average surplus of workers employed by firm y can be expressed as a function of wages, yielding a simple observable statistic that is increasing in y and thus allows to rank firms. In this section, we include some of the proofs in the main text as we consider them instructive (and surprisingly simple).

The foundation for our strategy of ranking firms is provided by the following result.

RESULT 2: $V_v(y)$ and $V_p(x, y)$ are increasing in y .

Our strategy is to relate these monotone statistics to observable statistics from the worker side. The next result is stated only in terms of workers' value functions.

¹⁴Since separation rates are identically δ at all firms a worker matches with, a worker's average wage is proportional to $\int_{B^w(x)} w(x, y) d_v(y) dy$. Assuming, for simplicity, that $B^w(x) = [\underline{\varphi}(x), \bar{\varphi}(x)]$, we get

$$\begin{aligned} & \frac{\partial}{\partial x} \int_{B^w(x)} w(x, y) d_v(y) dy \\ &= \int_{B^w(x)} \frac{\partial w(x, y)}{\partial x} d_v(y) dy + \bar{\varphi}'(x) w(x, \bar{\varphi}(x)) d_v(\bar{\varphi}(x)) - \underline{\varphi}'(x) w(x, \underline{\varphi}(x)) d_v(\underline{\varphi}(x)). \end{aligned}$$

Clearly, this equation is not necessarily increasing in x .

¹⁵Note that firms cannot be ordered based on data on average profits. This is because, just as average wages do not necessarily increase in x , average profits are not necessarily increasing in y . In addition to this theoretical obstacle, we are not aware of a convincing argument on how to overcome the well-known difficulties in measuring profits in the data in a way consistent with the model.

RESULT 3: *The expected surplus due to newly hired workers, given by*

$$(1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} (V_e(\tilde{x}, y) - V_u(\tilde{x})) d\tilde{x},$$

is increasing in y .

PROOF: Using (1),

$$\begin{aligned} (1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} (V_e(\tilde{x}, y) - V_u(\tilde{x})) d\tilde{x} \\ = \alpha(1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} S(\tilde{x}, y) d\tilde{x}. \end{aligned}$$

From (4), it follows that

$$(13) \quad \frac{V_v(y)(1 - \beta)}{\beta(1 - \alpha)} = (1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} S(\tilde{x}, y) d\tilde{x}.$$

From Result 2, both sides of (13) are increasing in y . Multiplying both sides of (13) by α yields the desired result. *Q.E.D.*

The proof used that the value of a vacancy is increasing in firm type y and then involved two steps. First, since the value of a vacancy is related to the expected surplus by an accounting identity (4), the expected surplus is also increasing in firm type (13). The next step uses that Nash bargaining implies that both the worker and the firm benefit from an increase in the surplus. Nash bargaining has an even stronger implication as the two parties benefit from an increase in the surplus in fixed proportions, determined by the bargaining power. This strong implication is, however, not used here and our results extend to other bargaining games where both parties benefit from an increase in the surplus.

Next, we relate this statistic to wages which are observable in the data.

RESULT 4: *The expected wage premium over the reservation wage of newly hired workers, given by*

$$(14) \quad \Omega(y) = (1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(x)}{U} (w(x, y) - w(x, y^{\min}(x))) dx,$$

is increasing in y .

Note that this expectation is taken when the vacancy is still unfilled. The proof uses three simple insights. Recall that $w(x, y^{\min}(x))$ denotes the lowest wage (the reservation wage) that worker x receives and $y^{\min}(x)$ is the firm type that pays this wage. The first insight is that the lowest wage is equal to the return of being unemployed,¹⁶

$$w(x, y^{\min}(x)) = (1 - \beta)V_u(x) = (1 - \beta)V_e(x, y^{\min}(x)).$$

¹⁶If a worker of type x is accepted by all firms, the lowest wage $w(x, y^{\min}(x))$ is not equal to the reservation wage and thus not equal to the return of being unemployed. In Section A.3 of the Supplemental Material, we show that this does not change the derivative of Ω . In particular, Ω is increasing in y independently from the properties of the matching set.

Second, the wage of a worker is a premium over the reservation wage (see (5)),

$$\begin{aligned} w(x, y) &= (1 - \beta)V_u(x) + (1 - \beta(1 - \delta))(V_e(x, y) - V_u(x)) \\ &= w(x, y^{\min}(x)) + (1 - \beta(1 - \delta))(V_e(x, y) - V_u(x)). \end{aligned}$$

Finally, this implies that the worker's surplus is proportional to the difference between the wage and the reservation wage,

$$w(x, y) - w(x, y^{\min}(x)) = (1 - \beta(1 - \delta))(V_e(x, y) - V_e(x, y^{\min}(x))).$$

Using Result 3 completes the proof.¹⁷

It is convenient to decompose $\Omega(y)$ into two factors that, as we show below, can be easily measured in the data. The first is the average wage premium of newly hired workers at firm y , $\Omega^e(y)$, and the second one is the probability to fill a vacancy, $q(y)$. The average wage premium equals

$$(15) \quad \Omega^e(y) = \frac{\int_{B^f(y)} \frac{d_u(x)}{U} (w(x, y) - w(x, y^{\min}(x))) dx}{\int_{B^f(y)} \frac{d_u(\tilde{x})}{U} d\tilde{x}}$$

The probability that a vacancy of type y is filled equals

$$(16) \quad q(y) = (1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} d\tilde{x}.$$

It then holds that

$$(17) \quad \Omega(y) = q(y)\Omega^e(y).$$

Computing $\Omega^e(y)$ requires knowing workers' reservation wages. While the reservation wage is clearly conceptually related to the lowest accepted wage by a worker (which is observable in the data), a more sophisticated measurement procedure is required in small samples and in the presence of measurement error. We develop such a procedure in Section 4. We also discuss the computation of the job filling rates $q(y)$ in Section 3.4.3 below.

3.3. Sign and Strength of Sorting

Having ranked workers and firms, we can compute Spearman's rank correlation between x and y in the data, which is just the Pearson correlation coefficient since both types are already ranked. The sign of this correlation is a natural indicator of the sign of sorting. For example, a value of 1 indicates perfect positive assortative matching and a value of -1 indicates perfect negative assortative matching.

¹⁷The strict monotonicity of Ω and thus the ranking of firms depends on $V(y)$ being increasing in y , as shown in Result 2. This result can be tested in the data using, for example, the monotonicity test by Hall and Heckman (2000) which we describe and apply in Hagedorn, Law, and Manovskii (2014) to discriminate between models with comparative and absolute advantage. In the empirical analysis below, we indeed reject the constant Ω hypothesis, which is perhaps not surprising in view of the large firm heterogeneity documented by, for example, Bagger, Christensen, and Mortensen (2011).

3.3.1. Relationship to the Literature

Note that $\Omega(y)$ is increasing in y regardless of whether the model features positive or negative assortative matching, or indeed neither. In particular, it does not require any assumptions on the production function f , that is, neither super- nor submodularity. Because of this, Result 4 enables us to identify the sign of sorting. This is in contrast to the recent results of [Eeckhout and Kircher \(2011\)](#) who used a simplified version of [Atakan \(2006\)](#) to show that the sign of sorting cannot be identified from wage data. More precisely, they demonstrated that in their two-period model, for every supermodular production function that induces PAM, there exists a submodular production function that induces NAM and generates identical wages. In Eeckhout and Kircher's model, workers and firms meet randomly in the first period. They can either form a match or pay a search cost and be paired up with their ideal partners in a perfectly competitive labor market in the second period. There is no discounting between the two periods. These specific modeling choices simplify the theoretical analysis substantially, but, unfortunately, they also ultimately prevent the identification of the model.

Consider first the role of discounting. One difference between [Shimer and Smith \(2000\)](#) and the models in [Atakan \(2006\)](#) and [Eeckhout and Kircher \(2011\)](#) is that the former uses discounting whereas search costs are explicit (and additive) in the latter two papers (as in [Chade \(2001\)](#)). This matters for ranking firms, as can be seen by rearranging the Bellman Equation (4) in our model:

$$V_v(y)(1 - \beta) = \beta(1 - \alpha)(1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} S(\tilde{x}, y) d\tilde{x}.$$

In the limit $\beta \rightarrow 1$, we get that the expected surplus is a constant,

$$0 = (1 - \alpha)(1 - \delta)\mathbb{M}_v \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} S(\tilde{x}, y) d\tilde{x},$$

the Constant Surplus Condition in Theorem 1 in [Atakan \(2006\)](#). If, instead, $\beta < 1$, then $V_v(y)(1 - \beta)$ is increasing in y and so is the expected surplus. This monotonicity (independent of the production function) of expected surplus is the key step in our ranking of firms. We then measure the surplus as being proportional to the wage premium of a worker resulting in our statistic $\Omega(y)$ that is expressed in terms of wages only. Constructing the same statistic in [Atakan \(2006\)](#) does not yield a function that is monotonically increasing in y but is, instead, a constant. The impossibility to rank firms in [Atakan \(2006\)](#) is thus due to the knife-edge assumption of no discounting. As soon as this assumption is relaxed, firms can be ranked.¹⁸

However, introducing discounting is not sufficient to achieve identification in [Eeckhout and Kircher \(2011\)](#) due to further simplification they make relative to [Atakan \(2006\)](#). In particular, their assumption of a frictionless second-period matching implies that it is the

¹⁸[Eeckhout and Kircher \(2011\)](#) noted that their non-identification proof breaks down in case of discounting. While this does not imply the possibility of identification, they conjectured that if one could somehow establish identification, it might be difficult to achieve in practice for plausible values of β . The strength of our method is that it is immune to such concerns. The statistic we use to rank firms is monotonically increasing in firm type regardless of how close β is to 1. In terms of the quantitative results reported below, we will show in Section F.5 of the Supplemental Material that even using the monthly discount factor as high as 0.999 does not measurably affect our ability to identify the objects of interest.

frictionless second-period outcome $\pi^*(y)$ that serves as the continuation value for firms, that is, the value of a vacancy equals

$$V(y) = \underbrace{\int S(x, y) dx}_{\text{expected surplus}} + \beta \pi^*(y).$$

Our statistic $\Omega(y)$ is monotone in y if and only if the expected surplus is. In [Shimer and Smith \(2000\)](#), we showed that this is the case because the value of a vacancy is increasing in y . In [Eeckhout and Kircher \(2011\)](#), such a simple relationship between the value of a vacancy and the expected surplus does not exist. Solving the above equation for the expected surplus yields

$$\int S(x, y) dx = V(y) - \beta \pi^*(y),$$

which is not necessarily increasing in y even with $\beta \neq 1$ since $\pi^*(y)$ is increasing in y and enters with a negative sign. As a result, our statistic $\Omega(y)$, which is proportional to expected surplus, is not necessarily monotonically increasing. Thus, the model in [Eeckhout and Kircher \(2011\)](#) is not identified due to the assumption that the continuation value is the frictionless allocation and not the value of a vacancy as in [Shimer and Smith \(2000\)](#) and in [Atakan \(2006\)](#).

3.4. Identifying Remaining Model Parameters

We now show how to identify the remaining objects in the model. Our primary interest is in identifying the production function $f(x, y)$. We recover it, at the end of this section, by inverting the wage equation. To accomplish that, we require the measures of the value of unemployment $V_u(x)$, the value of a vacancy $V_v(y)$, and the probability to fill a vacancy $q(y)$. Alongside with measuring these key objects, we also show how to measure the value of being employed $V_e(x, y)$, the value of producing for a firm $V_p(x, y)$, and the meeting probabilities for unemployed workers and vacant firms \mathbb{M}_u and \mathbb{M}_v .

3.4.1. Measuring $V_u(x)$, $V_e(x, y)$, and $S(x, y)$

The Bellman Equation (5) implies, using $V_e(x, y^{\min}(x)) = V_u(x)$, that $V_u(x)(1 - \beta) = w(x, y^{\min}(x))$. Thus, the reservation wage for workers of type x can be used to measure the (type-dependent) value of unemployment as¹⁹

$$V_u(x) = \frac{w(x, y^{\min}(x))}{1 - \beta}.$$

¹⁹Implementation of the measurement of (type-dependent) reservation wages is described in Section 4.3. Alternatively, one can, for example if no firm pays exactly the reservation wage, compute the value of unemployment using (3) and (5) as

$$V_u(x) = \frac{\frac{\beta(1 - \delta)\mathbb{M}_u}{1 - \beta(1 - \delta)} \int_{B^w(x)} \frac{d_v(\tilde{y})}{V} w(x, \tilde{y}) d\tilde{y}}{(1 - \beta) \left(1 + \frac{\beta(1 - \delta)\mathbb{M}_u}{1 - \beta(1 - \delta)} \right) \int_{B^w(x)} \frac{d_v(\tilde{y})}{V} d\tilde{y}}.$$

To measure $V_e(x, y)$, consider a worker of type x , who starts working at a firm of type y at time $t = 0$, becomes unemployed at time t_U , and receives wage $w_t = w(x, y)$ for all t between $t = 0$ and $t = t_U - 1$. We then define

$$\sum_{t=0}^{t_U-1} \beta^t w_t + \beta^{t_U} V_u(x),$$

where, of course, we use the measured value for $V_u(x)$. Averaging across all these sums for all types x starting at firm y yields the estimate $V_e(x, y)$.

We then also have a measure of surplus multiplied by the bargaining power

$$\alpha S(x, y) = V_e(x, y) - V_u(x).$$

Using that $\alpha = \frac{1}{2}$ in the model of [Shimer and Smith \(2000\)](#) yields the value of $S(x, y)$. In Section A.4 of the Supplemental Material, we follow [Hagedorn and Manovskii \(2008, 2013\)](#) and describe how the parameter α can be identified from the data in a more general version of the model.

3.4.2. Measuring $V_v(y)$ and $V_p(x, y)$

We next turn to the measurement of $V_v(y)$, which is related to our estimate $\Omega(y)$ through

$$V_v(y)(1 - \beta) = \beta \frac{1 - \alpha}{\alpha} (1 - \delta) \Omega(y).$$

Since, as discussed above, we can measure Ω in the data, and we can follow the standard approaches in the literature to estimate or calibrate δ and β , we obtain

$$V_v(y) = \frac{\beta}{1 - \beta} \frac{1 - \alpha}{\alpha} (1 - \delta) \Omega(y).$$

Using this, $V_p(x, y)$ can then be computed from

$$V_p(x, y) = V_v(y) + (1 - \alpha) S(x, y).$$

3.4.3. Transition Rates

The probability, $\tilde{q}(y)$, that a vacancy posted by firm j of type $y(j)$ is filled conditional on meeting a worker is simply the share of unemployed workers belonging to this firm's matching set in total unemployment. Indexing workers by their (estimated) rank x , denote by $u(x)$ this type's average unemployment rate. Using the law of large numbers, it holds that

$$(18) \quad \tilde{q}(y) \equiv \int_{B^f(y)} \frac{d_u(\tilde{x})}{U} d\tilde{x} = \frac{\int_{B^f(y)} u(\tilde{x}) d\tilde{x}}{\int_{[0,1]} u(\tilde{x}) d\tilde{x}}.$$

Note that $\tilde{q}(y)$ can be computed in the data and that its knowledge is sufficient to rank firms as $q(y)$ is proportional to $\tilde{q}(y)$ (see (16)). Next, we measure \mathbb{M}_v . Denote by $H_t(y)$

the observed number of new hires in firms of type y at time t , and by $V_t(y)$ the unobserved number of vacancies posted by these firms. Equation (16) and the law of large numbers imply that

$$\frac{H_t(y)}{(1 - \delta)\tilde{q}(y)} = \mathbb{M}_v V_t(y).$$

Adding up across all firms and time periods, and rearranging, yields \mathbb{M}_v (and \mathbb{M}_u as $\mathbb{M}_u U = \mathbb{M}_v V$):

$$\mathbb{M}_v = \frac{1}{1 - \delta} \frac{\int_{[0,1]} \frac{H_t(\tilde{y})}{\tilde{q}(\tilde{y})} d\tilde{y}}{\int_{[0,1]} V_t(\tilde{y}) d\tilde{y}},$$

where $\int_{[0,1]} V_t(\tilde{y}) d\tilde{y}$ is the total number of vacancies, which if unobserved, can be inferred by matching the wage share in output.

These computations simplify if data on firm-level vacancies are available.²⁰ In this case, we can directly measure the probability to fill a vacancy $q(y)$. For every worker type, we can measure the probability to leave unemployment $\lambda(x)$. With firm-level vacancy data (i.e., data for $d_v(y)$), we can then measure \mathbb{M}_u (and consequently \mathbb{M}_v) from

$$\lambda(x) = (1 - \delta)\mathbb{M}_u \int_{B^w(x)} \frac{d_v(\tilde{y})}{V} d\tilde{y}.$$

A more robust way is to integrate over all worker types and solve for \mathbb{M}_u from

$$\int_0^1 \lambda(x) dx = (1 - \delta)\mathbb{M}_u \int_0^1 \int_{B^w(x)} \frac{d_v(\tilde{y})}{V} d\tilde{y} dx.$$

3.4.4. Measuring Output $f(x, y)$

Using the equation for wages (A.2) in the Supplemental Material, the production function $f(x, y)$ on the set of matches that actually form then equals

$$f(x, y) = \frac{w(x, y) - \alpha(\beta - 1)V_v(y) - (1 - \alpha)(1 - \beta)V_u(x)}{\alpha}.$$

The output of a match is determined by inverting the wage equation, expressing the output $f(x, y)$ as a function of the observed wage $w(x, y)$ and the two measured outside options $V_v(y)$ and $V_u(x)$. For this to be feasible, the researcher has to know the exact wage equation. In the model of [Shimer and Smith \(2000\)](#), this is the case since Nash bargaining is imposed. Other wage determination mechanisms which imply an invertible wage equation would also allow for an identification of output.

An alternative strategy for recovering the output is to invert the equation for surplus (A.1) in the Supplemental Material. As $S(x, y)$, $V_v(y)$, and $V_u(x)$ have already been measured, this immediately implies $f(x, y)$.

²⁰The computation is also not affected if there is random noise leading to observing some firms hiring without having a posted vacancy.

4. IMPLEMENTATION AND QUANTITATIVE EVALUATION

In this section, we develop the key implementation steps of the proposed identification strategy and evaluate their performance in a Monte Carlo study over a range of parameter values that are likely to be encountered in empirical work. As our identification proof is fully constructive, the only challenge is to deal with the fact that available data sets are finite whereas theoretical identification assumes infinite data. This is particularly relevant to estimating a worker’s reservation wage which can only be imprecisely estimated from an individual worker’s few own wage observations. To obtain precise estimates, we propose an ostensibly simple but indeed very effective methodological innovation. After workers have been ranked, we bin similarly ranked workers. Given the large number of workers in available data sets, closely ranked workers are very similar. We then use wage observations for all workers in this bin as if they were a single worker’s observations and compute the relevant statistics accordingly. Analogously, we also bin firms after they have been ranked to compute statistics for a single firm using the observations for all firms in the respective bin. In Monte Carlo simulations, we find that binning workers and firms is the key to precisely recovering all model primitives in the data.²¹ We now describe the key steps of our implementation strategy, with the detailed implementation algorithm provided in Section B of the Supplemental Material.

4.1. *Parameterization*

We assume that a researcher has access to a matched employer–employee panel data set with a time dimension of 20 years. Most currently available and commonly used matched data sets (e.g., from Brazil, Denmark, Germany, France) have a similar or longer time span. We assume that the data include the information on wages, all employment and unemployment spells of the worker over the duration of the sample, and all hires and separations at the firm level. We simulate the model at a monthly frequency. The production functions commonly used in the literature belong to the constant elasticity of substitution (CES) family. We consider three such functions:

- (i) $f(x, y) = 0.6 + 0.4(x^{1/2} + y^{1/2})^2$, which induces positive assortative matching (PAM),
- (ii) $f(x, y) = (x^2 + 2y^2)^{1/2}$, which induces negative assortative matching (NAM), and
- (iii) $f(x, y) = 0.4 + \mathbb{1}_{\{x \leq 0.4\}}(x + 0.6)y + \mathbb{1}_{\{x > 0.4\}}((x - 0.4)^2 + y^2)^{1/2}$, which induces neither positive nor negative assortative matching (NEITHER). Instead, the pattern of sorting changes over its domain (PAM for $x \leq 0.4$ and NAM for $x > 0.4$).

The literature has largely restricted attention to identifying sorting assuming that the production function induces either globally positive or negative assortative matching. This motivates our choice of the first two production functions. Our method, however, does not rely on placing such restrictions. The choice of the third production function is designed to illustrate this point. The production functions are scaled to generate a realistic amount of wage dispersion.

We also consider three distributions of workers and firms (these are the “original” non-rank distributions F and G). Common choices in the literature are either uniform or normal distributions. We consider both and for the normal distribution we choose the mean of 0.5 and the variance of 0.25 (the distribution is truncated and normalized on $[0, 1]$ interval). We also consider a bimodal distribution constructed as the sum of two normals:

²¹Krasnokutskaya, Song, and Tang (2014), Krasnokutskaya, Terwiesch, and Tiererova (2016) pursued a similar approach in the context of online auctions where they grouped bidders on the basis of their quality.

$N(0.2, 0.25) + N(0.8, 0.25)$ truncated and normalized to integrate to one on $[0, 1]$. The distributions are discretized into 50 values on an evenly spaced grid. We simulate a small sample of 30,000 workers. The vacancy creation cost is such that there is the same number of jobs in the economy. After the productivity of the vacancies is learned, vacancy creators sell them in competitive markets to operating firms with the same productivity level. We use 300 firms in simulations, which implies in a symmetric equilibrium that operating firms have 100 jobs per firm (all sharing the same productivity level). As not all of these jobs are filled at a point in time, the actual size of employment at each firm varies across parameterizations, but is not more than 100 workers. We set the discount factor to 0.996 at monthly frequency to be consistent with the annual interest rate of 4%.

We assume the standard Cobb–Douglas form of the meeting function $m(s, v) = \kappa s^\nu v^{(1-\nu)}$. We set the elasticity parameter $\nu = 0.5$, as this parameter plays no interesting role in our stationary model. We consider a wide range for the scale parameter $\kappa = \{0.4, 0.7\}$ to generate job-finding probabilities ranging between a high of about 50% a month in the United States and a low of about 15% in some European countries. Similarly, we choose two values for the separation rate $\delta = \{0.01, 0.025\}$, roughly spanning the U.S. and European evidence. Bargaining weights of 0.5 are symmetric for workers and firms.

Finally, we allow for measurement error in wages. Hagedorn and Manovskii (2012) estimated that measurement error accounts for approximately 20% of the variance of residual wages in the U.S. NLSY data. This is likely an upper bound on the matched employer–employee data sets, as these data are typically based on administrative sources with highly reliable wage information. Nevertheless, to make the test of the proposed method more stringent, we add i.i.d. noise to every wage observation with a variance of 20% of the correctly measured wage variance. The error is simulated as draws from a normal distribution truncated at three standard deviations around the mean of zero.

The values of parameters used in simulations are summarized in Table II. All combinations of parameters result in 108 distinct parameterizations. Figure A-3 in the Supplemental Material summarizes the values that a number of variables of interest take across all parameterizations. Most tend to lie within empirically plausible ranges.

4.2. Ranking of Workers

Result 1 implies that wages within a firm are increasing in worker productivity x . This provides a way to rank workers according to their productivity within a firm, and these

TABLE II
PARAMETERIZATIONS

Parameter	Symbol	Option 1	Option 2	Option 3
Production function	$f(x, y)$	PAM	NAM	NEITHER
Worker distribution	d_w	Uniform	Normal	Bi-Modal
Firm distribution	d_f	Uniform	Normal	Bi-Modal
Discount factor	β	0.996		
Separation rate	δ	0.01	0.025	
Meeting function scale	κ	0.4	0.7	
Meeting function elasticity	ν	0.5		
Worker's bargaining weight	α	0.5		
Measurement error in wages	ε	20% of overall wage variance		

partial rankings can be aggregated to construct a more global ranking. However, the presence of measurement error in wages presents a complication. Within one firm, one worker could be ranked higher than another worker not because he is more productive (actually he is less productive) but just because of the measurement error. And the ranking between these two workers may not be consistent with the ranking from other firms. To solve this problem, we build on the insights from social choice theory, where voters rank candidates, potentially inconsistently with each other. In our application, voters correspond to firms and workers to voting alternatives. An aggregate ranking then minimizes the number of disagreements between individual votes which defines the *Kemeny–Young rank aggregation* problem first described in [Kemeny \(1959\)](#) and [Kemeny and Snell \(1963\)](#). We refine this procedure. We not only count the number of disagreements but also assign weights to the ranking of all worker pairs, determining how likely it is that the observed wages (with measurement error) indicate the true ranking. If, for example, the wage of worker i is much higher than the wage of worker j , we assign a high weight, whereas the weight is small if the wages are very similar. We use a Bayesian approach to compute these weights. The goal is then to find a ranking that maximizes the sum of weights in favor of a proposed ranking. To deal with the computational complexity of this problem, we build on insights from [Kenyon-Mathieu and Schudy \(2007\)](#), who provided a polynomial time algorithm that approximates the solution to this problem with arbitrary degree of accuracy. In practice, we found that implementing a portion of their algorithm achieves a high level of accuracy while being quite fast. A detailed description is provided in Section C of the Supplemental Material.

Figure 1 reports the distribution of the rank correlations of the true worker types and types recovered using the rank aggregation procedure across simulations. The results in-




FIGURE 1.—Distribution of the correlation between the true and estimated, using the rank aggregation procedure, worker ranks across parameterizations.

dicates that the procedure recovers the correct rankings quite well despite relatively small samples and in the presence of sizable measurement error in wages.²²

4.3. *Ranking of Firms*

Result 4 implies that, to rank firms, one simply needs to compute the expected average difference between the wages a firm pays to each of its newly hired workers and the reservation wage of those workers. The only challenge is to obtain an accurate estimate of the reservation wage for each worker, despite the limited time dimension of the available data. The key insight we use is that once workers are ranked, similarly ranked workers must have similar reservation wages. Thus, we can estimate the reservation wage by considering a group of similar workers, despite the fact that each of those workers is observed for a relatively short period of time. Since workers are ranked and ranks are uniformly distributed, we can group workers into bins of equal size. One can think of workers as ordered on a line and bins corresponding to intervals (without holes) on this line. What remains to be determined is the number of bins or, equivalently, the size of the bin. The answer to this question obviously depends on the size of the sample. If we had infinite observations for each worker, the choice would be easy, as each bin would consist of one worker only, as these infinite observations are sufficient to compute the worker's reservation wage from this worker's observed wages only. However, we only have a small sample available and the appropriate bin size has to be evaluated in Monte Carlo simulations. In these simulations, we find that 50 bins are sufficient to reliably compute all statistics, such as the reservation wage and the production function. Additional details of the implementation algorithm for firm ranking are relegated to Section B of the Supplemental Material.

Figure 2 plots the distribution of the correlations between the ranking of firms based on this procedure and true ranking across the parameterizations we consider. In all parameterizations, the ranking of firms is identified quite precisely.

After firms are ranked, we can also group similarly ranked firms (i.e., firms with very similar Ω values) into bins. As firms (ranks) are uniformly distributed, bins are of equal size and we find 50 bins to be sufficient to deliver reliable results. The measured value of a vacancy for a firm type is then the average value of a vacancy for all the firms assigned to this bin. The estimated value of a vacancy is found to be increasing in firm type as we also verified using the monotonicity test in [Hall and Heckman \(2000\)](#). While the production function can be estimated at the level of an individual firm, estimating it on worker and firm bins helps eliminate the effects of the sampling error in small samples.

4.4. *Sign and Strength of Sorting*

Figure 3 plots the correlation between identified worker and firm ranks against the true correlation (each black dot corresponds to a different parameterization). Here we separate the three production functions to illustrate that our identification strategy easily

²²As implied by Result 1, we initialize the rank aggregation procedure using global rankings based on the lowest accepted wage, the highest accepted wage, or the adjusted average wage. The performance of each of these methods is reported in Appendix Figure A-4 in the Supplemental Material. The correlations are relatively high for all these measures. In particular, the adjusted average wage dominates both the minimum and the maximum wage in its performance in ranking workers. The measure based on rank aggregation substantially outperforms any of the individual measures.




FIGURE 2.—Distribution of the correlation between the estimated and true firm ranks across parameterizations.

identifies the sign of sorting. In all cases, this relatively crude measure of the strength of sorting performs quite well (as all the black dots lie on a 45 degree line, the correlation between worker and firm types is identified nearly exactly).

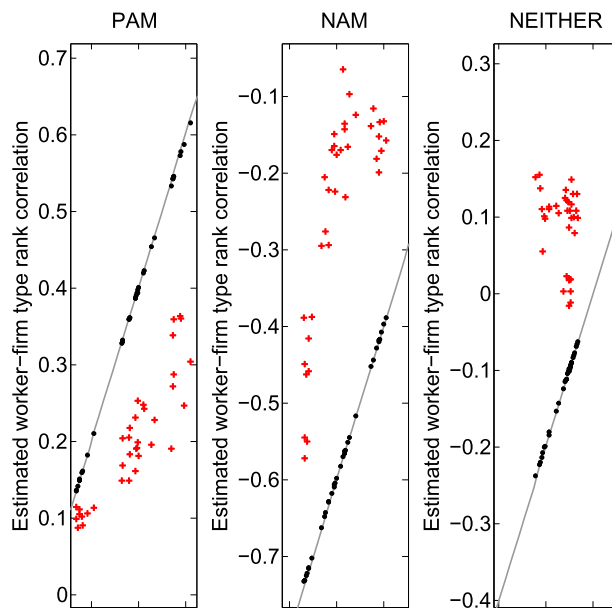


FIGURE 3.—Correlation between identified worker and firm ranks against true correlation.

For comparison, the (red, if viewed in color) crosses in the same figures correspond to the correlation between worker and firm fixed effects estimated using the exact least squares formulas provided by [Abowd, Creecy, and Kramarz \(2002\)](#). The results confirm the findings in the literature that such reduced-form estimates of the strength of sorting are severely biased towards zero.

4.5. Measuring $f(x, y)$

To evaluate how well our method recovers the production function, in Figure 4 we plot the distribution of the correlation between the true and the estimated production functions across all parameterizations (we have grouped similarly ranked firms and workers into types x and y as discussed above). The correlations are generally very high, with the median correlation of 0.9987.

To provide a better sense of the ability of our identification and implementation strategy to recover the production function, in Figures A-5–A-7 in the Supplemental Material, we plot the true and estimated production functions for three particular examples from the set of the parameterizations we considered. The three production functions induce positive, negative, or neither positive nor negative assortative matching and we use the same set of parameters for all cases: $\beta = 0.996$, $\kappa = 0.3$, $\delta = 0.01$, $\alpha = 0.5$, and $d_w = d_f = U[0, 1]$, with a measurement error equal to 20% of the variance of wages. Each figure contains the true production function (dark red with black lines if viewed in color) defined on the equilibrium matching set and the estimated one (transparent blue). As the functions are essentially on top of each other, to help appreciate the closeness of the fit, for each production function we provide four views with the red line representing the axis




FIGURE 4.—Distribution of correlation between true and estimated production functions across parameterizations.

of rotation. The estimated production functions are presented without any smoothing or filtering.

4.6. *Measuring the Effect of Search Frictions on Output*

Having obtained the nonparametric estimate of the production function and of the rankings of workers and firms, one can address a number of important economic questions. One such question is to assess the magnitude of output losses due to mismatch between workers and firms. We now evaluate the ability of our identification and implementation strategies to provide a reliable quantitative answer to this question without imposing functional form assumptions on technology.²³

To do so, we first derive the (counterfactual) allocation in a world without frictions. To solve for the frictionless assignment, we need to find a one-to-one assignment (bijection) $\mu : [0, 1] \rightarrow [0, 1]$ of workers to firms such that the total output $\sum_x f(x, \mu(x))$ is maximized. Our identification strategy identifies the production function only on the set of (x, y) matches observed in the data. Since our objective is to find an optimal assignment on this set, we assume that the output outside of the observed frictional matching set is zero.

This assignment problem is a well-studied combinatorial optimization problem and there are several existing algorithms that can solve it in polynomial time.²⁴ However, a complete solution is not required to approximate the effect of the elimination of search frictions on output. Instead, a much smaller scale assignment problem can be solved on a random sample of workers and firms. We choose the size of the sample so that the maximum total output of the sample scaled to the size of the total population of workers and firms becomes invariant to the sample size. Across our simulations, we found that a sample of about 5000 workers and 5000 jobs is sufficient. On a sample of that size, we can solve the problem in minutes using the [Jonker and Volgenant \(1987\)](#) algorithm without special hardware.

Denote by $\mathcal{E}^{\text{no fric}}$ the expectation of frictionless output $f(\cdot, \mu(\cdot))$:

$$(19) \quad \mathcal{E}^{\text{no fric}} = \int_0^1 f(x, \mu(x)) dx,$$

where we used that worker ranks are uniformly distributed. In the presence of frictions, let $\mathcal{E}^{\text{fric}}$ be the expectation of f :

$$(20) \quad \mathcal{E}^{\text{fric}} = \int_{\mathcal{B}} f(x, y) d_m(x, y) dx dy.$$

²³An alternative approach to measuring the cost of mismatch was pursued by [Eeckhout and Kircher \(2011\)](#). They noted (as did [Teulings and Gautier \(2004\)](#) and [Gautier and Teulings \(2006\)](#)) that differentiating the wage equation (such as (A.2) in the Supplemental Material), the cross-partial derivative of the wage equals the cross-partial derivative of the production function multiplied by the worker bargaining weight. They also showed that in their two-period model, where a period with search frictions is followed by a frictionless period and there is no discounting, and under the assumption that the production function globally induces PAM or NAM, the average cross-partial of the production function can be theoretically related to the output cost of mismatch. They did not explore the empirical properties of this estimator given the nature of the available data. Even maintaining their restrictive assumptions on the production technology, it is not clear whether it is possible to generalize this method to the fully dynamic model in [Shimer and Smith \(2000\)](#).

²⁴See [Burkard, Dell'Amico, and Martello \(2009\)](#) for a review.

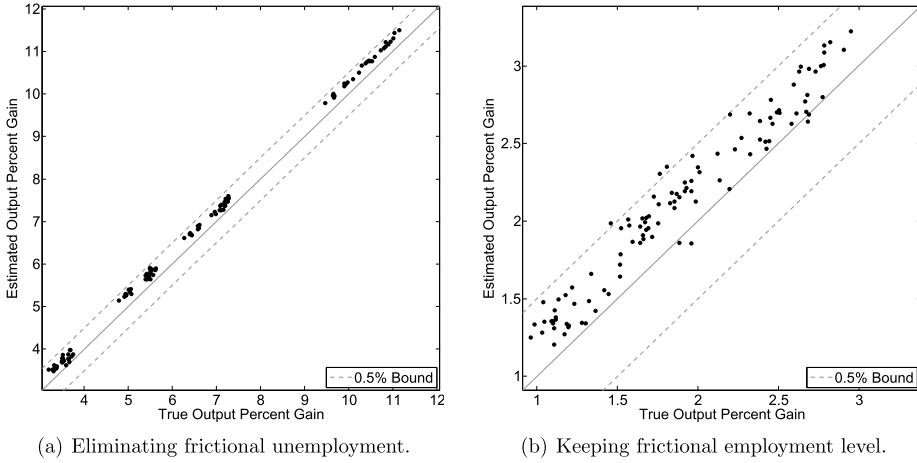


FIGURE 5.—Estimated gains from eliminating frictions.

Then, the output loss due to misallocation is the difference between the expected output without frictions, $\mathcal{E}^{\text{no fric}}$, and the expected output with frictions, $\mathcal{E}^{\text{fric}}$:

$$(21) \quad \Delta \mathcal{E} = \mathcal{E}^{\text{no fric}} - \mathcal{E}^{\text{fric}}.$$

In Figure 5, we plot the percent output gain from the optimal reallocation of workers $100 \cdot \frac{\Delta \mathcal{E}}{\mathcal{E}^{\text{fric}}}$. The true output gain as a percent of frictional output is on the horizontal axis, while the estimated gain as a percent of frictional output is on the vertical axis. To help visually interpret the results, the figure also includes two dotted lines that represent a mistake of plus or minus one half of one percent of output. In Panel 5(a), we plot the estimated gains from eliminating all frictions, including the elimination of frictional unemployment. In Panel 5(b), we only consider the effect from optimal reassignment of workers employed in the economy with frictions (i.e., we keep the employment of each worker type fixed at its level in the economy with frictions).²⁵ We interpret the results as indicating that the method performs quite well in estimating the gain from optimal worker reallocation.²⁶

4.7. Robustness

In this section, we conduct several robustness checks. First, we allow for a shorter time horizon and assume that data are available only for 10 years. Second, we look at smaller firms, and third, we allow for random match-specific productivity.

²⁵The relatively small gains from reallocation in this experiment are driven by the parsimonious CES production function. Larger gains from eliminating mismatch can be obtained with more cumbersome specifications of the production functions. We have explored a number of such specifications and found that our method continues to perform equally well.

²⁶Note that the high correlations between the true and estimated production functions reported in Figure 4 do not necessarily imply an accurate estimate of the gains from optimal assignment. The reason is that the solution to the optimal assignment problem exploits all imperfections of the estimated production function. Thus, it is possible to have a high correlation of the true and estimated production function and to substantially miss on the estimated gains from reallocation. The good fit of Figure 5 implies that our methodology does not introduce such significant exploitable imperfections.

4.7.1. *Shorter Time Horizon*

In the quantitative evaluation, we assumed, guided by the available matched employer–employee data sets, that we have data for 20 years. We now investigate how the results change if we have only a shorter data set of 10 years available, which may be the case for some developing countries or if an estimation on subsamples of available data is considered. We find that the quantitative performance of our method is not substantially affected by this restriction. Results reported in Section F.2 of the Supplemental Material imply that when using only 10 years of data, the median correlation between the true and the estimated production function is 0.9972, the correlation between the firm and the worker ranking is very well recovered, and so are the gains from eliminating search frictions.

4.7.2. *Small Firms*

In the benchmark, the maximum potential size of the firm, that is, the sum of employment and vacancies, was set to 100. Results reported in Section F.3 of the Supplemental Material suggest that the quantitative performance of the method does not worsen substantially if we instead restrict the maximum potential firm size to 20 workers. In particular, we find that the median correlation between the true and the estimated production function is 0.9962, the correlation between firm and worker ranking is very well recovered, and so are the estimated gains from eliminating search frictions.

4.7.3. *Match Quality*

The Monte Carlo experiments reported above illustrate that our implementation strategy can successfully overcome the effects of measurement error in wages. Another potential source of noise in identification is idiosyncratic match productivity. When a worker of type x and a firm of type y meet, they draw, before deciding whether to produce together, a random productivity realization z from a distribution Γ , so that production equals $f(x, y) + z$ for the duration of the match.

Thus, there are now two sources of match-specific productivity. The first depends on the fixed types x and y and the second one is z . Although z is a random draw from a distribution that is independent of either x or y , the accepted value of z depends on the worker and firm types. As a result, a worker and a firm either do not produce at all or they start producing for all realizations of match quality above some threshold level $\bar{z}(x, y)$.

Using this model, we repeat the Monte Carlo study with the same parameterizations. The choice of the offer distribution of match quality Γ is made consistent with the findings in [Low, Meghir, and Pistaferri \(2010\)](#). These authors found that the variance of the offer distribution for match quality Γ is about 0.22 of the variance of overall wages, where observables are removed. As we have an equilibrium model where the wage of a worker is determined endogenously given the production function, the “offer” distribution for output z is exogenous and not the offer distribution for wages as in the partial equilibrium model in [Low, Meghir, and Pistaferri \(2010\)](#). In particular, changes in z in the production function do not necessarily translate one-to-one into wage changes. We therefore pick the distribution Γ such that the endogenous wage offer distribution in the model has the same variance as the exogenous wage offer distribution in [Low, Meghir, and Pistaferri \(2010\)](#). The resulting median variance of Γ relative to the variance of wages in the model equals 0.2085. Similarly, we use the estimates in [Low, Meghir, and Pistaferri \(2010\)](#) to set the variance of measurement error relative to the variance of wages to 0.027. Thus,

our identification methodology is potentially complicated due to two sources of noise: idiosyncratic match productivity and measurement error. In Section F.4 of the Supplemental Material, we report that this is not quantitatively important, however, by conducting the same experiments as in the benchmark and using the same implementation strategy. We find that the median correlation between the true and the estimated production function is 0.9961, the correlation between firm and worker ranking is very well recovered, and so are the estimated gains from eliminating search frictions.

We should point out that adding match productivity allows for an easy extension of our methodology which renders it more powerful. As the support of the distribution of z is very wide, there exist values of z that induce a match between many worker types and many firm types. For example, we find that a match between a worker of type x and a firm of type y that would not form for low realizations of match-specific productivity, will form if the realization of z is sufficiently high. This implies that for each firm where some z 's but not all are accepted by a worker of type x , the lowest wage paid by this firm to this worker is the reservation wage of this worker. Given the wide support of z , we find that for the majority of firms, the lowest wage paid is the reservation wage. Clearly, this makes the estimation of the reservation wage much more precise.²⁷

5. IDENTIFICATION: DATA

In this section, we apply our methodology to a German matched employer–employee data set to identify all features of the model, in particular the ranking of workers, the ranking of firms, and the production function. A key feature of the data is a substantial movement of workers between jobs without intervening unemployment spells. We therefore first extend the basic model to include on-the-job search and then show that the identification and implementation methodology extends from the basic model to the model with on-the-job search.

5.1. *Model and Identification With On-the-Job Search*

An unemployed worker makes a take-it-or-leave-it offer to a firm and thus extracts the full surplus.²⁸ As in [Postel-Vinay and Robin \(2002\)](#) and [Cahuc, Postel-Vinay, and Robin \(2006\)](#), when a worker employed at some firm \tilde{y} meets a firm y which generates higher surplus, the two firms engage in Bertrand competition such that the worker moves to firm y and obtains the full surplus generated with firm \tilde{y} . Small costs of writing an offer prevent firms y which generate lower surplus than the current firm \tilde{y} from engaging in Bertrand competition. The remaining details of the model description are relegated to Section D of the Supplemental Material.

We now show that all the theoretical results which established identification in the benchmark model continue to hold in our model with on-the-search and establish identification in this model as well.

²⁷We thank an anonymous referee for drawing our attention to this improvement of our methodology. One can use this insight to also recover the threshold acceptance level $\hat{z}(x, y)$ and thus also decompose the output of a pair (x, y) into $f(x, y)$ and the contribution from match quality.

²⁸This assumption is special but, as we will see below, the model based on this assumption provides a very tight fit to the data. Alternative assumptions may break some of the monotonicity properties on which our identification strategy relies. Establishing identification of a more general model of on-the-job search is a promising avenue for future research.

Ranking Workers

The following result establishes that wages within a firm rank workers hired out of unemployment.

RESULT 5: *Within a firm, wages of workers hired from unemployment are increasing in worker type.*

We can therefore use the same algorithm as in the benchmark model with the restriction that it is applied to workers hired out of unemployment only. This reduces the number of usable observations relative to the benchmark model, but we show in Section D of the Supplemental Material that this reduction does not significantly worsen the performance of our methodology.

Ranking Firms

To rank firms, we derived a statistic—the value of vacancy—which is monotonically increasing in firm type y and which can be expressed as a function of wages observed in the data. The same result holds in the model with on-the-job search:

RESULT 6: *The value of a vacancy $V_v(y)$ is increasing in firm type y and can be expressed as a function of wages.*

We can therefore again use observable wages to rank firms in terms of their productivity. The actual statistic is slightly different from the one in the benchmark model, but it is even easier to implement as it does not involve estimating workers' reservation wages.

Measuring Output $f(x, y)$

Finally, we can also recover the output of every match from inverting the wage equation, which relates output to observable variables.

RESULT 7: *The wage equation can be inverted so that output can be expressed in terms of observables.*

5.2. Data Description

Our empirical analysis uses the model with on-the-job search and is based on the German LIAB data set, which is the primary source of data used in the research on the contribution of worker and firm effects to wage dispersion in Germany, for example, [Dustman, Ludsteck, and Schönberg \(2009\)](#). These data, spanning the 1993–2007 period, represent the linked employer–employee data set of the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB). The LIAB combines the Employment Statistics of the German Federal Employment Agency (Bundesagentur für Arbeit) with establishment level data from the IAB Establishment Panel. For detailed information on the LIAB, see [Alda, Bender, and Gartner \(2005\)](#).

We follow [Card, Heining, and Kline \(2013\)](#) in selecting the sample and in constructing the residual wages which form the basis of the analysis in both papers. The only adjustments we make are due to the fact that, in the data we have access to, all workers employed at the establishments from the representative IAB Establishment Panel can be




FIGURE 6.—Estimated match density.

observed, while we do not necessarily observe all workers employed in other establishments. Card, Heining, and Kline (2013) had access to the universe of workers and establishments. This does not limit our analysis as LIAB provides us with complete continuous (up to a day) earnings and employment histories of workers who worked in the establishments from the Establishment Panel for at least one day (even when these workers are employed in other establishments). One potential advantage of using the LIAB is that it contains information on vacancies at the establishment level which are not available for firms outside the Establishment Panel. Descriptive statistics and additional details are provided in Section E of the Supplemental Material.

5.3. Empirical Results

5.3.1. Sign and Strength of Sorting

In Figure 6, we plot the match density estimated by applying our methodology to the data. Table III provides the implied correlation between the identified worker and firm ranks. The table also provides the correlation implied by the methodology of Abowd, Kramarz, and Margolis (1999) estimated using the exact least squares formulas provided by Abowd, Creedy, and Kramarz (2002).

TABLE III
SIGN AND STRENGTH OF SORTING

	HLM	AKM
Corr(W-rank, F-rank)	0.7547	0.055

Our identification strategy reveals a very high degree of sorting. The correlation between worker and firm ranks is positive. The correlation obtained using the AKM methodology is very close to zero. Our AKM estimate is similar to that obtained by [Andrews, Gill, Schank, and Upward \(2008\)](#), who also used the LIAB data. Moreover, a positive sign of sorting (despite the negative AKM estimate) was recently found by [Bagger and Lentz \(2014\)](#), who used a different model of sorting from the one in this paper. This was also confirmed by [Bartolucci and Devicienti \(2013\)](#), although the theoretical basis for their empirical measures of sorting is not clear.

The strength of sorting implied by our estimate is also remarkably high. The correlation between worker and firm ranks of 0.7547 is much larger than the estimate obtained using the model in [Bagger and Lentz \(2014\)](#) in Danish data. [Lopes de Melo \(2013\)](#) proposed measuring the strength of sorting through the correlation of the wage with the wages of coworkers. We find the value of this correlation to be 0.556. This measure is somewhat difficult to interpret, but the fact that it is slightly higher than the range of estimate for other countries reported in [Lopes de Melo \(2013\)](#) seems suggestive.

One might expect the strong positive sorting to be indicative of a (log) supermodular production function. Such a conjecture would be premature, however, as we now illustrate.

5.3.2. Measuring $f(x, y)$

In Figure 7, we plot the estimated production functions presented without any smoothing or filtering. The estimated production function is increasing in worker and firm types. The substantial increase with firm type is due, in part, to a strongly increasing value of the vacancy (see (A.37) in the Supplemental Material), plotted in Figure 8(a). Applying the [Hall and Heckman \(2000\)](#) monotonicity test, we confirm that the value of the vacancy is monotonically increasing at any conventional significance level. The substantial variation in the value of a vacancy across firms implies that firms cannot easily scale up production. In particular, this finding is inconsistent with models in which firms face a constant cost of vacancy and scale up production until the value of the vacancy is driven down to zero.

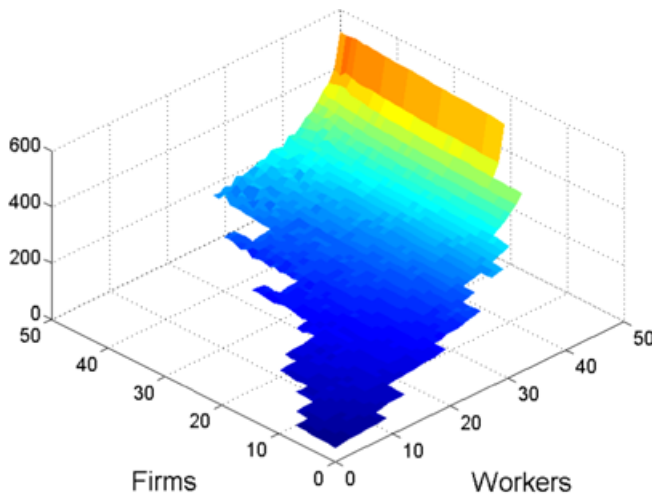


FIGURE 7.—Estimated production function.

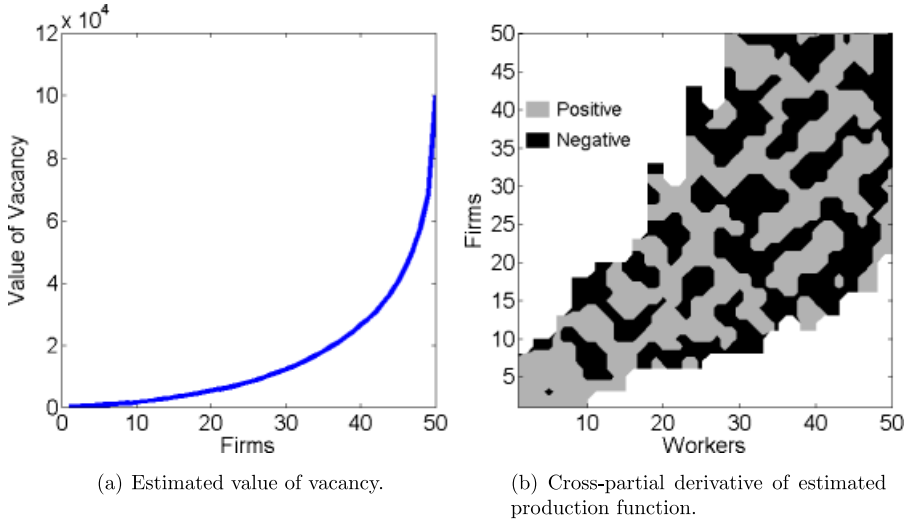


FIGURE 8.—Features of the estimated production function.

at the individual firm level.²⁹ In addition, the large systematic variation in the value of the vacancy implies that the variance of output is substantially higher than the variance of wages.

The cross-partial derivative of the production function is plotted in Figure 8(b).³⁰ We observe that the sign of the cross-partial is not the same everywhere.

5.3.3. *Fit of the Model*

An in-depth analysis of the ability of our model with on-the-job-search to fit various dimensions of the data was undertaken by [Kantenga and Law \(2015\)](#).³¹ We now summarize the relevant findings. [Kantenga and Law \(2015\)](#) first estimated all the parameters of the model, including the production function, using our methodology. Then they simulated the model economy with the estimated production function and asked whether it reproduces the matching sets and wages observed in the data. There is nothing in our methodology that would guarantee a good match if either the model is misspecified or the implementation strategy is not recovering the correct objects. Moreover, the estimation of the model with on-the-job search is based on wages of those hired from unemployment only, while the fit of the model is assessed through all wages and job transitions.

The results are as follows. First, the endogenously generated matching set in the model replicates the one in the data nearly exactly. Second, the model fits wages well. Considering the fit to individual-level wages, the R^2 is 0.92. For comparison, the R^2 obtained when implementing a much more flexible reduced-form AKM regression with fixed effects for

²⁹Of course, such models would also not generate the strong worker-firm sorting that we and others document in the data.

³⁰The cross-partial is quite volatile. To make the figure easier to interpret, we very slightly smooth the production function prior to computing the cross-partial using local quadratic least squares regressions (LOESS) with a narrow bandwidth of 0.05. In Monte Carlo simulations reported above, this bandwidth is sufficient to accurately recover the sign of the cross-partial everywhere. Even using a much higher bandwidth does not generate the cross-partial with a constant sign.

³¹They also used the LIAB data, although a different sample.

each worker and firm is 0.94. Given this fit, it is perhaps not surprising that estimating an AKM regression on predicted wages yields very similar estimates of the variance and covariances of worker and firm effects to those obtained when estimating the same regression on raw wage data.³² Third, a direct test rejects the linearity of log wages in worker and firm effects in the data, while the nonlinearity implied by our model is consistent with the data. Finally, even the patterns of job-to-job flows are replicated very well, resulting in a very good fit to the match density. Overall, these results lead us to conclude that our parsimonious structural model with its specific assumptions provides a very good fit to the data.

5.3.4. *Measuring the Effect of Search Frictions on Output*

Having obtained the nonparametric estimate of the production function and of the rankings of workers and firms, we can assess the magnitude of output losses due to mismatch between workers and firms. We use the same methodology as presented in the implementation section. In particular, we derive the (counterfactual) allocation in a world without frictions using the same one-to-one assignment algorithm. We find that the output loss due to misallocation, that is, the difference between the expected output without frictions and the expected output with frictions, equals 1.83% in the data. This is the effect from the optimal reassignment of workers employed in the economy with frictions (i.e., we keep the employment of each worker type fixed at its level in the economy with frictions). If we also eliminate frictional unemployment, the output gain is 8.47%, an increase smaller than the increase in the number of employed because unemployment and vacancies are concentrated among low productivity workers and firms, respectively.

Interestingly, forcing the matching to be on the main diagonal (i.e., matching the best worker to the best job, second best worker to the second best job, etc.) implies a 0.23% *decline* in output relative to the observed frictional allocation. This illustrates the danger of the common parametric assumption that high positive correlation between worker and firm types must be associated with a globally supermodular production function. In other words, our results indicate that overall positive sorting is generated by the production function with substantial local fluctuations in the cross-partial. Workers in the data exploit this local curvature relatively well, so that the optimal assignment does not deviate dramatically from the observed one and the associated output gains are relatively small. But it is important to take this curvature into account. The common but erroneous assumption that on average positive sorting should imply a globally supermodular production function with the output maximizing assignment on the main diagonal would lead to misleading conclusions. This illustrates the importance of the nonparametric identification of the production function proposed in this paper.

³²Importantly, this experiment replaces raw wages with wages predicted by our model but holds constant the observed transitions of workers across firms. Andrews, Gill, Schank, and Upward (2008, 2012) have pointed out a potentially important limited mobility bias in AKM regressions arising because many firms have only a few movers, whose wage observations are used to identify the firm effect. They argued that this likely leads to an overestimation of the variance of firm effects and underestimation of the correlation between worker and firm effects. Using our estimated model to eliminate this bias reveals that, in our sample, the bias is quite small for the correlation but is severe for the variance of firm effects, which drops from around 20% of overall wage variance to only about 1% (with a similar corresponding increase in the variance of worker effects). Bonhomme, Lamadon, and Manresa (2015) reported qualitatively similar findings using the data from Sweden. Of course, these experiments do not eliminate the biases due to the fundamental misspecification of AKM regression implied by our theory.

6. CONCLUSION

In this paper, we have developed an empirical methodology that allows to study assortative matching between employers and employees based on their unobserved (to economists) characteristics. In particular, we have shown theoretically that all parameters of the classic Becker model of sorting based on absolute advantage with search frictions, as analyzed in [Shimer and Smith \(2000\)](#), can be identified using only matched employer–employee data on wages and labor market transitions. For example, these data are sufficient to assess whether matching between workers and firms is assortative, whether sorting is positive or negative, and to measure the potential effect on output from moving any given worker to any given employer in the economy. We have also provided computational algorithms that allow to implement our identification strategy given the limitations (on sample size, frequency of labor market transitions, measurement error, etc.) of the commonly used matched employer–employee data sets, and found that they perform well in a Monte Carlo study. Next, we introduced on-the-job search into the model and proved that our identification and implementation strategies extend to this environment. Finally, we estimated the model with on-the-job search using a matched employer–employee data set from Germany.

Our theoretical analysis was based on relatively simple versions of the model which allowed us to obtain clear and transparent results. Nevertheless, we expect that the four key insights listed below will be of substantial practical importance for empirical research going forward. Each individual insight may or may not apply to every empirically relevant model. But we do expect the combinations of these insights to be at the foundation of a nonparametric identification and estimation strategy for the models in this class. These insights are:

(i) Within- and between-firm wage comparisons contain an extraordinary amount of information that can be exploited to identify relative worker productivities. Advances in algorithm research allow a number of approaches to extract this information. We exploit the one (rank aggregation) that applies to the models we consider. Other models may call for other methods, but they are readily available. We do not expect the ranking of workers based on, say, lowest or highest wages to achieve similar precision.

(ii) Firms can be ranked based on the value of the vacancy. This clearly will not be true in every model, but we do expect this to be true in most empirically relevant models. One piece of evidence for this is the empirical results in this paper. We very strongly reject the hypothesis that the value of the vacancy is constant (e.g., zero). Our measurement strategy is not biased towards finding this answer.

(iii) Wage or surplus equations can be inverted for output. Not every model will permit such an inversion, but many will. Once we know how to measure the value of unemployment and the value of the vacancy (even if it is zero), nonparametric identification of the production function immediately follows. We expect this to form the basis of any attempt at nonparametric identification of the production function.

(iv) The crucial element of our proposed implementation strategy that leads to very precise estimates is binning of workers and firms. Given the limited time dimension of any possible data, exploiting the cross-sectional variation will likely be the cornerstone of any future attempt at nonparametric estimation.

There are numerous important positive and normative questions that can be answered once the nonparametric identification of this class of models is established, in particular of the production function and of the rankings of workers and firms. For example, as we have shown, without any additional assumptions we can compute the optimal assignment of workers to jobs on the set where match productivity can be measured. A comparison

between the optimal assignment and the observed one reveals the extent of the output loss due to search frictions. A similar approach can be used to separately measure the extent of changes in technology and sorting frictions over time (e.g., to understand the reasons for the substantial rise in the college premium in the United States in the 1980s) and across, say, countries. The ability to separately identify changes in technology from changes in mismatch seems essential for understanding the effects of changes in many economic policies, for example, trade liberalizations. We can also determine the importance of complementarities in production and measure the role of frictions and sorting in determining the dispersion of output and productivity across establishments. It is also possible to measure the extent to which sorting on unobservables can account for wage differences across groups of employers (large or small, exporters and no-exporters, belonging to different industries, located in different geographic regions, etc.). Turning to wage dispersion, an application of our method allows to decompose wages into components due to workers, firms, and the assortative matching between them as well as to estimate the role of search frictions and sorting in driving the observed wage dispersion.

REFERENCES

- ABOWD, J. M., R. H. CREECY, AND F. KRAMARZ (2002): "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," LEHD Program Technical Paper TP-2002-06, U.S. Census Bureau. [52,58]
- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67 (2), 251–334. [30,33,58]
- ALDA, H., S. BENDER, AND H. GARTNER (2005): "The Linked Employer-Employee Dataset Created From the IAB Establishment Panel and the Process-Produced Data of the IAB (LIAB)," *Journal of Applied Social Science Studies/Schmollers Jahrbuch*, 125 (2), 327–336. [57]
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" *Journal of the Royal Statistical Society: Series A*, 171 (3), 673–697. [59,61]
- (2012): "High Wage Workers Match With High Wage Firms: Clear Evidence of the Effects of Limited Mobility Bias," *Economics Letters*, 117 (3), 824–827. [61]
- ATAKAN, A. E. (2006): "Assortative Matching With Explicit Search Costs," *Econometrica*, 74 (3), 667–680. [43, 44]
- BAGGER, J., AND R. LENTZ (2014): "An Empirical Model of Wage Dispersion With Sorting," NBER Working Paper No. 20031. [59]
- BAGGER, J., B. J. CHRISTENSEN, AND D. T. MORTENSEN (2011): "Wage and Productivity Dispersion: The Roles of Rent Sharing, Labor Quality and Capital Intensity," Report. [42]
- BARTOLUCCI, C., AND F. DEVICIENTI (2013): "Better Workers Move to Better Firms: A Simple Test to Identify Sorting," Report, Collegio Carlo Alberto. [59]
- BECKER, G. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81 (4), 813–846. [29,30,32, 33]
- BONHOMME, S., T. LAMADON, AND E. MANRESA (2015): "A Distributional Framework for Matched Employer Employee Data," Report. [61]
- BURKARD, R., M. DELL'AMICO, AND S. MARTELLO (2009): *Assignment Problems*. Philadelphia, PA: SIAM. [53]
- CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage Bargaining With On-the-Job Search: Theory and Evidence," *Econometrica*, 74 (2), 323–364. [56]
- CARD, D., J. HEINING, AND P. KLINE (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128 (3), 967–1015. [57,58]
- CHADE, H. (2001): "Two-Sided Search and Perfect Segregation With Fixed Search Costs," *Mathematical Social Sciences*, 42, 31–51. [43]
- DUSTMAN, C., J. LUDSTECK, AND U. SCHÖNBERG (2009): "Revisiting the German Wage Structure," *Quarterly Journal of Economics*, 124 (2), 363–376. [57]
- ECKHOUT, J., AND P. KIRCHER (2010): "Sorting and Decentralized Price Competition," *Econometrica*, 78 (2), 539–574. [34]

- (2011): “Identifying Sorting—In Theory,” *The Review of Economic Studies*, 78 (3), 872–906. [30,43, 44,53]
- FLINN, C., AND J. HECKMAN (1982): “New Methods for Analyzing Structural Models of Labor Force Dynamics,” *Journal of Econometrics*, 18 (1), 115–168. [39]
- GAUTIER, P. A., AND C. N. TEULINGS (2006): “How Large Are Search Frictions?” *Journal of the European Economic Association*, 4 (6), 1193–1225. [30,53]
- (2012): “Sorting and the Output Loss due to Search Frictions,” Discussion Paper TI 2011-010/3, Tinbergen Institute. [33]
- HAGEDORN, M., AND I. MANOVSKII (2008): “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” *American Economic Review*, 98 (4), 1692–1706. [45]
- (2012): “Search Frictions and Wage Dispersion,” Report, University of Pennsylvania. [48]
- (2013): “Job Selection and Wages Over the Business Cycle,” *American Economic Review*, 103 (2), 771–803. [45]
- HAGEDORN, M., T. LAW, AND I. MANOVSKII (2014): “Measuring Sorting on Comparative and Absolute Advantage,” Report, University of Pennsylvania. [33,42]
- (2017): “Supplement to ‘Identifying Equilibrium Models of Labor Market Sorting,’” *Econometrica Supplemental Material*, 85, <http://dx.doi.org/10.3982/ECTA11301>. [32]
- HALL, P., AND N. E. HECKMAN (2000): “Testing for Monotonicity of a Regression Mean by Calibrating for Linear Functions,” *The Annals of Statistics*, 28 (1), 20–39. [42,50,59]
- HURWICZ, L. (1950): “Generalization of the Concept of Identification,” in *Statistical Inference in Dynamic Economic Models*, Cowles Commission Research in Economics Monograph, Vol. 10, ed. by T. C. Koopmans. New York: Wiley, 245–257. [38]
- JONKER, R., AND A. VOLGENANT (1987): “A Shortest Augmenting Path Algorithm for Dense and Sparse Linear Assignment Problems,” *Computing*, 38, 325–340. [53]
- KANTENGA, K., AND T.-H. LAW (2015): “Sorting and Wage Inequality,” Report, University of Pennsylvania. [60]
- KEMENY, J. G. (1959): “Mathematics Without Numbers,” *Daedalus*, 88 (4), 577–591. [49]
- KEMENY, J. G., AND J. L. SNELL (1963): *Mathematical Models in the Social Sciences*. New York: Blaisdell. [49]
- KENYON-MATHIEU, C., AND W. SCHUDY (2007): “How to Rank With Few Errors,” in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing. STOC '07*. New York: ACM, 95–103. [49]
- KRASNOKUTSKAYA, E., K. SONG, AND X. TANG (2014): “The Role of Quality in Internet Service Markets,” Report, Johns Hopkins University. [47]
- KRASNOKUTSKAYA, E., C. TERWIESCH, AND L. TIEREROVA (2016): “Trading Across Borders in Online Auctions,” Report, Johns Hopkins University. [47]
- LAMADON, T., J. LISE, C. MEGHIR, AND J.-M. ROBIN (2014): “Matching, Sorting, and Wages,” Working Paper, University College London. [32]
- LENTZ, R. (2010): “Sorting by Search Intensity,” *Journal of Economic Theory*, 145 (4), 1436–1452. [30]
- LENTZ, R., AND D. T. MORTENSEN (2010): “Labor Market Models of Worker and Firm Heterogeneity,” *Annual Review of Economics*, 2 (1), 577–602. [33]
- LISE, J., C. MEGHIR, AND J.-M. ROBIN (2016): “Matching, Sorting and Wages,” *Review of Economic Dynamics*, 19 (1), 63–87. [30]
- LOPES DE MELO, R. (2013): “Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence,” Working Paper, University of Chicago. [30,59]
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2010): “Wage Risk and Employment Risk Over the Life Cycle,” *American Economic Review*, 100 (4), 1432–1467. [55]
- MATZKIN, R. L. (2013): “Nonparametric Identification in Structural Economic Models,” *Annual Review of Economics*, 5 (1), 457–486. [38]
- MOEN, E. R. (1997): “Competitive Search Equilibrium,” *Journal of Political Economy*, 105 (2), 385–411. [34]
- NÖLDEKE, G., AND T. TRÖGER (2009): “Matching Heterogeneous Agents With a Linear Search Technology,” Working Paper, Basel and Mannheim. [37]
- POSTEL-VINAY, F., AND J.-M. ROBIN (2002): “Wage Dispersion With Worker and Employer Heterogeneity,” *Econometrica*, 70 (6), 2295–2350. [56]
- SHI, S. (2001): “Frictional Assignment. I. Efficiency,” *Journal of Economic Theory*, 98 (2), 232–260. [34]
- SHIMER, R. (2005): “The Assignment of Workers to Jobs in an Economy With Coordination Frictions,” *Journal of Political Economy*, 113 (5), 996–1025. [34]
- SHIMER, R., AND L. SMITH (2000): “Assortative Matching and Search,” *Econometrica*, 68 (2), 343–370. [30-33, 35-38,43-46,53,62]
- TEULINGS, C. N., AND P. A. GAUTIER (2004): “The Right Man for the Job,” *Review of Economic Studies*, 71 (2), 553–580. [53]

WOLPIN, K. I. (1987): "Estimating a Structural Search Model: The Transition From School to Work," *Econometrica*, 55 (4), 801–817. [39]

Dept. of Economics, University of Oslo, Box 1095 Blindern, 0317 Oslo, Norway; marcus.hagedorn07@gmail.com,

Dept. of Economics, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, U.S.A.; tzuohann@gmail.com,

and

Dept. of Economics, University of Pennsylvania, 160 McNeil Building, 3718 Locust Walk, Philadelphia, PA 19104-6297, U.S.A.; manovski@econ.upenn.edu.

Co-editor Jean-Marc Robin handled this manuscript.

Manuscript received 20 December, 2012; final version accepted 31 March, 2016; available online 27 June, 2016.