

Schulz, Bastian; Lochner, Benjamin

Conference Paper

Labor Market Sorting in Germany

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel -
Session: Heterogenous Firms and the Labor Market, No. G07-V3

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Schulz, Bastian; Lochner, Benjamin (2016) : Labor Market Sorting in Germany,
Beiträge zur Jahrestagung des Vereins für Socialpolitik 2016: Demographischer Wandel - Session:
Heterogenous Firms and the Labor Market, No. G07-V3, ZBW - Deutsche Zentralbibliothek für
Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/145902>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Labor Market Sorting in Germany*

Benjamin Lochner¹ and Bastian Schulz²

¹*University of Erlangen-Nuremberg*

²*Ifo Institute at the University of Munich*

August 31, 2016

Abstract

This paper analyzes the allocation of workers to jobs and the wage distribution in Germany. Our main contribution is to reconcile prominent empirical models of wage dispersion (Abowd et al., 1999; Card et al., 2013) with theoretical sorting models (Shimer and Smith, 2000; Eeckhout and Kircher, 2011; Hagedorn et al., 2016). We find that empirical fixed effect models provide a valid approximation of observed wages and matching patterns for a large part of the data. For low-type workers, however, wages are decreasing in the type of the firm a worker is matched with. This prediction of theoretical sorting models is at odds with the monotonicity assumption of fixed effect models. After ranking both workers and firms, we show that low-type workers have become increasingly sorted into low-type firms over time, especially out of unemployment. This increase is driven by selection into wage-maximizing matches at the bottom of the firm type distribution. It can be linked to increased domestic outsourcing of low-type workers to business service firms.

Keywords: Assortative Matching, Wage Dispersion, Unobserved Heterogeneity, Optimal Rank Aggregation, Job Mobility, Matched Employer-Employee Data, Firm Performance, Stochastic Frontier Analysis

JEL Classifications: J24, J31, J40, J62, J64, L25

*We thank Michele Battisti, Anja Bauer, Bernd Fitzenberger, Timo Hener, Christian Holzner, Philipp Kircher, Patrick Kline, Thibaut Lamadon, Jonas Maibom, Christian Merkl, Helmut Rainer, Jean-Marc Robin, Johannes Schmieder, Heiko Stüber, Coen Teulings, and Rune Vejlin for their comments and suggestions. Feedback from seminar and conference participants at the Universities of Aarhus, Cambridge, Nuremberg, and Regensburg, at the Ifo Institute and the Institute for Employment Research (IAB), and at SMYE 2016 (Lisbon), SaM 2016 (Amsterdam), ES-NASM 2016 (Philadelphia), and SED 2016 (Toulouse) has greatly improved the paper. All remaining errors are our own. Benjamin Lochner gratefully acknowledges financial support provided by the Graduate Program of the IAB and the School of Business and Economics of the University of Erlangen-Nuremberg (GradAB). Contact: benjamin.lochner@fau.de; schulz.b@ifo.de

1 Introduction

Increasing wage inequality is a topic of high interest for both policymakers and academics. For Germany, Dustmann et al. (2009) show that between 1990 and 2000 real wage growth for full-time working men was negative below the 18th percentile of the wage distribution and increasingly positive above.¹ Using by now widely available matched employer-employee data, Abowd et al. (1999) (henceforth AKM) first demonstrate how to quantify the respective contributions of unobserved worker and firm heterogeneity to wage dispersion. AKM models provide a very good fit and typically find that most of the observed wage dispersion is explained by unobserved heterogeneity. Card et al. (2013) (henceforth CHK) apply the AKM methodology to the universe of German social security records.² They decompose the increase of wage dispersion into three elements: rising worker heterogeneity, more dispersion of wage premiums paid by employers, and increased sorting based on unobservables, measured by a rising correlation of worker and firm-fixed effects.

This paper is motivated by the discussion about the potential misspecification of the AKM two-way fixed effect model. As pointed out by, among others, Gautier and Teulings (2006), Eeckhout and Kircher (2011), and Lopes de Melo (2015), the assumption that log wages are additively separable into worker and firm-fixed effects leaves, by construction, no role for match-specific effects.³ In theoretical sorting models, however, output and wages are match-specific; they are determined by a production complementarity between workers' skill and firms' productivity types. If this match-specific component of wages was quantitatively important, the AKM model would be inapplicable to quantify wage dispersion. CHK argue that the abstraction from match-specific effects is defensible because deviations in terms of wage residuals appear to be small for most, but not all, combinations of worker and firm types.

In this paper, we build a bridge between these two opposing views on the sources of wage dispersion. We analyze German matched employer-employee data⁴ through the lens of a structural sorting model featuring heterogeneity on both sides of the market, search frictions, and search on the job. Our main analytical tool is the identification technique proposed by Hagedorn et al. (2016) (henceforth HLM). To identify the sign

¹Remarkably, the yearly increase of wage dispersion in Germany has an order of magnitude comparable to the United States; this trend can be traced back well into the 1980s/70s. Dustmann et al. (2009) find that the gap between the 85th and 50th percentiles of the German wage distribution has increased by about 0.6 log points per year between 1975 and 2004. This is comparable to Autor et al. (2006), who report that the gap between the 90th and 50th percentile in the United States has increased at a rate of roughly 1 log point per year during the same period.

²Having access to the universe of data is crucial for the AKM methodology because the observation of all workers per firm reduces the potential impact of the so-called "limited mobility bias", a problem emphasized by Andrews et al. (2008, 2012).

³In the remainder of the paper, we will sometimes refer to this assumption as the "AKM assumption".

⁴German social security register data are provided by the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB) of the German Federal Employment Agency (Bundesagentur für Arbeit).

and strength of sorting, they estimate economy-wide rankings of both workers and firms using data on wages and labor market transitions only. In their model, the match-specific output is non-parametrically identified. HLM report a rank correlation coefficient, a natural measure for sorting, of 0.76. The respective correlation of worker and firm-fixed effects in CHK is much lower, about 0.21 for a comparable period.⁵ First, we make a methodological contribution by using the German labor market as a laboratory to reconcile these seemingly incompatible approaches. Using the worker ranking procedure proposed by HLM and a firm ranking which does not depend on wages, we confirm that matching patterns are positive assortative in Germany. We find an overall rank correlation coefficient of 0.24, which is closer to the CHK-AKM benchmark than to HLM.

Our second contribution is a thorough quantitative investigation of labor market sorting in Germany. We find that sorting increased significantly over time. It rose particularly for matches formed by workers out of unemployment. Low-skill workers are increasingly sorted into low-productivity firms. This development can be linked to increasing domestic outsourcing of firms.⁶ Sorting of high-skill workers into high-productivity firms has, if anything, slightly decreased during our period of observation. Interestingly, the increased sorting of workers at the bottom of the type distribution is related to non-monotonic wage patterns. These workers maximize their wages in matches with low-type firms and the increased sorting leads to wage gains. We are the first to provide direct empirical evidence that for some worker types wages are not monotonically increasing in the productivity-type of the firm they are matched with. This non-monotonicity of wages is a key feature of theoretical sorting models like Shimer and Smith (2000), Atakan (2006), and Eeckhout and Kircher (2011), but it is at odds with the AKM two-way fixed effect approach, because in the reduced-form model wages mechanically increase in the firm effect. Accordingly, the residuals of the CHK-AKM model are substantial for some combinations of worker and firm types, specifically those including low-type workers.⁷ We find that increased sorting and the non-monotonicity of wages are most pronounced for these high-residual workers at the bottom of the type distribution. For medium and high-type workers, however, wages increase monotonically in the firm type, consistent with CHK and AKM. This explains why two-way fixed effect models generate a good fit overall. Nevertheless, it is important to emphasize that the AKM model cannot explain wage determination for low-type workers. The outcomes for this group of employees are the ones with which labor market policy is typically most concerned.

⁵For simplicity, this is the arithmetic mean of the correlations reported by CHK for the two subperiods which span the same period as HLM.

⁶Goldschmidt and Schmieder (2015) analyze domestic outsourcing in Germany. They use an AKM model and find that displaced workers suffer wage losses. Using our more flexible framework, we observe that the wage of low-type workers increases on average when they move to firms that offer services which are typically outsourced (cleaning, security, temporary work agencies).

⁷In the AKM context, worker and firm types are equivalent to their estimated fixed effects. Regarding residuals in CHK, see Figure VI on p. 996 and the discussion on pp. 989-991 in Card et al. (2013).

In the related literature, a number of recent papers explores ways to identify the sign and strength of sorting and its contribution to wage dispersion without relying on the AKM assumption. Bagger and Lentz (2015), Gautier and Teulings (2015), and Lise et al. (2016) develop structural search models and take them to the data. Bonhomme et al. (2016) propose a flexible empirical framework with a discrete number of types (finite mixture model) which also allows for unrestricted interactions of worker and firm heterogeneity. Our analysis is inspired by HLM, which also belongs to this group of papers. HLM show how match-specific output, wages, and the sign and strength of sorting are non-parametrically identified with standard matched employer-employee data available for many countries. The key challenge for the empirical study of sorting is to construct credible global rankings of workers and firms. While studies in the AKM tradition implicitly rank workers and firms by their fixed effects, HLM solve a Kemeny-Young rank aggregation problem to merge intra-firm wage rankings into a global ranking of workers. This procedure uses the largest connected set of workers. A computational algorithm running on this graph can effectively maximize the likelihood of the correct global worker ranking, as proven by Kenyon-Mathieu and Schudy (2007). Once workers are globally ranked, HLM show how to use the model structure to rank firms based on the value of a job vacancy. This approach delivers a rank correlation which lies considerably above the correlation for Germany estimated by CHK and above related results for other countries in the studies mentioned before. The German case is well-suited to take the HLM approach to a test. We have access to very detailed information for a large number of firms from the IAB Establishment Panel. If all information necessary to rank both workers and firms globally was indeed contained in wages alone, additional firm data would be redundant and not alter the results in any meaningful way. We construct an efficiency-based firm ranking using the distance to an estimated production frontier. Our ranking is independent of wages and produces lower rank correlations, which are more in line with the empirical sorting literature and with CHK for Germany in particular. We also construct a profit-based firm ranking inspired by Bartolucci et al. (2015), which yields similar results.

A paper closely related to ours is Kantenga and Law (2016). They use the full HLM model to structurally decompose increasing overall wage dispersion into the contributions of worker heterogeneity, firm heterogeneity, search frictions, and sorting. They reject the assumption of additive separability statistically by following pairs of workers that jointly change employers. We do not have to impose the full structure of the HLM model on the data because we rank firms independently of wages. This prevents us from performing a structural decomposition of wage dispersion. We go beyond Kantenga and Law (2016), however, by providing an explanation for why additive separability is rejected in some parts of the type space. Thus, we see our results as largely complementary.

This paper is organized as follows: Section 2 describes our data. Section 3 briefly

discusses the theory of sorting that guides our thinking. Section 4 explains our approach to identifying the sign and strength of sorting. Section 5 presents our results: correlations of our estimated worker and firm rankings, the dynamics of sorting over time, and how this relates to wages, wage inequality, and increased domestic outsourcing. Section 6 concludes.

2 Data

We use matched employer-employee data for Germany provided by the Institute for Employment Research (Institut für Arbeitsmarkt- und Berufsforschung, IAB) of the German Federal Employment Agency (Bundesagentur für Arbeit). We use the “LIAB Mover Model”⁸ and restrict our analysis to the years 1998-2008.⁹ This data set is ideal for our purposes as it provides information about a large number of workers moving between firms and it can be linked directly to firm-level survey data from the IAB Establishment Panel. We provide a short overview of our data preparation procedures in the following; more details on sample selection and imputation procedures can be found in Appendix A.

2.1 Data Preparation

One major advantage of German matched employer-employee data is the high quality of the wage data due to plausibility checks carried out by the social security institutions followed by sanctions for non- and misreporting. In our raw data, we observe nominal gross daily wages, which we deflate using the consumer price index from German national accounts with 2005 as the base year. Every wage observation corresponds to one employment spell, which can last from one day up to one year due to the reporting rules of the German social security system.¹⁰ On average, the workers in our sample have 6.4 job spells between 1998 and 2008 (with a standard deviation of 5.0) and the average spell lasts 289 days (with a standard deviation of 114). We drop workers with more than 150 job spells (less than 0.01%). Observed spells are sometimes identical except for the ending date and the corresponding wage, which could be the result of multiple reports in case of a changing contract end date. In such cases the longer spell is typically associated with a higher wage, possibly due to Christmas bonuses or other salary supplements. We always keep the spell with the higher wage.

The education variable in German social security data suffers from missing values and inconsistencies, essentially because misreporting has no negative consequences. We

⁸File: LIAB_MM_9308. See Heining et al. (2012) for a detailed description of the data set.

⁹We chose this period of time because it is roughly split in half by the German labor market reforms passed and implemented between 2002 and 2005 (Hartz I-IV).

¹⁰The reporting rules require employers to file a report whenever an employee joins or leaves the establishment or, in the event of no change in an ongoing employment relationship, on December 31 each year.

impute missing and inconsistent observations using the methodology proposed in Fitzenberger et al. (2006). Missing values cannot be imputed for about 2% of the data and we drop these spells. A second limitation of the wage data is that the German social security system tracks earnings only up to a certain threshold, the contribution assessment ceiling (“Beitragsbemessungsgrenze”).¹¹ We follow the procedure suggested by Dustmann et al. (2009) and impute the upper tail of the wage distribution by running a series of Tobit regressions, allowing for a maximum degree of heterogeneity by fitting the model separately for years, education levels, and eight five-year age groups.¹²

We restrict our data set to West German male employees between 20 and 60 years of age who were liable for social security contributions. Self-employed workers, civil servants, and students are not included in the sample. Although it would be very interesting to analyze the sign and strength of sorting for the excluded subgroups, we put a higher priority on using a dataset which is, on the one hand, homogeneous and, on the other, comparable to data used in previous studies of wage dispersion in Germany, CHK in particular. To minimize working-time effects in our wage data, we further exclude part-time and marginal employment.¹³

After the initial data preparation, our sample consists of 16,361,068 employment spells, 1,824,580 workers, and 472,869 establishments for the years 1998-2008. Note that the employers we observe are not firms in the legal sense, but establishments or local production units. These do not necessarily coincide with the legal entity to which they belong. We use the terms “firm”, “establishment”, and “employer” interchangeably throughout the paper. The standard deviation of log wages in this sample is 0.455. This value is close to the measures of wage dispersion reported by CHK, indicating that our data preparation procedure indeed generates a comparable sample.¹⁴

2.2 Residual Wages

This paper analyzes how unobserved worker and firm heterogeneity as well as their potential interaction influence wages and the selection of workers into jobs, i.e sorting. Therefore, we use residual wages—wages net of the effects of observable worker characteristics—as the main input to rank workers. To compute them, we use a simple empirical model

¹¹The average yearly censoring rate is 13.6% of wage observations. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold.

¹²An alternative to imputing the censored part of the wage distribution would be to simply drop top-coded wages. Although we would lose the ability to analyze sorting patterns of workers earning very high wages, our main findings in this paper would not be affected as they pertain to workers with low and hence non-censored wages. Moreover, the calculation of residual wages, the main input for our analysis, is largely unaffected by the imputation: Table A.2a shows that a wage variance decomposition delivers very similar results with and without the imputed part of the wage distribution.

¹³To reliably identify spells of marginal or part-time employment we use both indicator variables in the data and, additionally, drop spells with wages below the time-varying marginal employment threshold, which is on average 12.2 Euro per day across years in our sample (“Geringfügigkeitsgrenze”).

¹⁴CHK report a standard deviation of log wages of 0.432 for 1996-2002 and 0.499 for 2002-2009.

Table 1: Decomposition of the Variance of Log Wages

	$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	\hat{r}_{it}
$\ln w_{it}$	0.207			
$x'_{it}\hat{\gamma}$	0.014	0.008		
$\hat{\alpha}_i$	0.158	0.006	0.152	
\hat{r}_{it}	0.035	0.000	0.000	0.035

Notes: Variance-Covariance matrix of regression model 1. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual (\hat{r}_{it}). Rounded to three decimal places.

of wage dispersion of the AKM-type and regress, following CHK, log wages on a person-fixed effect, an unrestricted set of year dummies, and quadratic and cubic terms in age fully interacted with educational attainment:

$$\ln w_{it} = x'_{it}\gamma + \alpha_i + r_{it}. \quad (1)$$

$\ln w_{it}$ denotes the log real daily wage of a worker i in year t , x'_{it} includes the time-varying observable characteristics, α_i is a worker-fixed effect, and r_{it} is the error term. Typical of an AKM-type wage regression, the explanatory power of this model is very high. The adjusted R^2 is 81%, slightly below the around 90% reported by CHK. This difference is easily explained by the lack of an establishment-fixed effect in our specification, which is, however, inconsequential for our ranking purposes.¹⁵ Table 1 shows the decomposition of the variance of log wages. The unobservable components of wages explain the vast majority of wage variation in our data, namely, 91%. The person-fixed effect alone explains 74% of log wage variance; the residual absorbs another 17%. To compute the residual wages, we simply subtract the share of the wage predicted by observable characteristics from the observed wage for every individual:

$$\ln \tilde{w}_{it} = \ln w_{it} - x'_{it}\hat{\gamma}. \quad (2)$$

The log residual wage has a standard deviation of 0.433 (variance 0.187). It is thus only slightly less dispersed than observed wages in our sample, underlining the small role that

¹⁵In contrast to CHK, who have access to the universe of German social security records, we cannot reliably estimate establishment-fixed effects because we do not observe the full workforce of the establishments in our sample. This is not a problem for our analysis. The ranking procedure proposed by HLM relies on pairwise wage comparisons of workers who are employed at the same establishment (see Section 4.2). The unobserved firm effect affects the wages of two workers at the same establishment by exactly the same amount. The ranking of a pair of workers employed by the same firm is thus not affected by the firm effect.

observable characteristics play for wage dispersion. The correlation between observed and residual wages is very high (0.98).

To address concerns about the influence of occupations on wages which we do not account for in the specification above, we augment Equation (1) by controlling for 32 occupational categories.¹⁶ We interact them with education and year dummies. Reassuringly, controlling for occupations does not significantly increase the explanatory power of the wage regression. The adjusted R^2 stays virtually the same (81%) and the portion of variance explained by observable characteristics rises only slightly from 3.9% in the baseline regression to 6.7% when including occupational controls.¹⁷ The correlation between baseline residual wages and residual wages net of occupational effects is very high, above 0.99. Taking into account the workers' occupation to calculate residual wages would not lead to a significantly different worker ranking and hence not affect our results. We are confident that the residual wages \tilde{w}_{it} are the appropriate input for our analysis.¹⁸

2.3 Firm Data

We use the IAB Establishment Panel to calculate measures of firm performance which we develop and use in Section 4.3. The IAB Establishment Panel draws a stratified random sample of establishments from the register data.¹⁹ We restrict our attention to establishments that employ at least 10 workers on average during our period of interest.²⁰ The data contains no direct information about the capital stock on which an establishment operates. To circumvent this shortcoming, we use the perpetual inventory method proposed by Müller (2008) to approximate the capital stock of the establishments in our sample. This method uses information about the average economic lives of different capital goods (buildings, IT, production machinery, transport equipment), which is available from national accounts, and the firms' net investments in the different categories of capital goods. We exclude the public sector of the economy and firms which do not report revenues as their primary measure of output.²¹ Finally, we merge our firm sample with the data on employment spells. This leaves us with 4,901 establishments for which we observe 1,714,450 employment spells of 234,800 workers.

¹⁶In accordance with Bundesagentur für Arbeit (1988) the classification in our data consists of about 330 occupational codes on the 3-digit level. We use 32 2-digit values ("Berufsabschnitte") in order to be able to estimate a large set of interaction terms.

¹⁷Table A.2b shows the decomposition of wage variance including occupational controls.

¹⁸CHK provide a comprehensive analysis of the explanatory content of additional controls in their setting. In line with our finding, they report that occupational (and industry) controls do not significantly increase the explanatory power of AKM-type models of wage dispersion.

¹⁹For details on the IAB Establishment Panel see Kölling (2000) and Fischer et al. (2009).

²⁰We lose only 0.2% of employment spells due to this restriction.

²¹The public sector includes health care, education, non-profit organizations. We exclude these firms because they typically do not seek to maximize profits or minimize costs, what is at odds with the rationale of our firm ranking procedure. The most important group of establishments which do not report revenues are banks.

2.4 Matches

The main unit of our empirical analysis in the following is a match between a worker and a firm. For our study of the allocation of workers to jobs, we focus on the first person-year observation of a new employment spell. Hence, we abstract from the spell length. This ensures that certain match types are not over-represented (under-represented) in the distribution of matches simply because they last longer (shorter) on average, leading to more (less) person-firm-year observations. This is the most conservative way to study sorting. Technically speaking, we define a new match as the first worker-id establishment-id combination we observe in the data. We further distinguish between two types of matches: matches out of unemployment and matches resulting from job-to-job switches of an employee. A match out of unemployment occurs whenever a particular worker is employed after a period of registered unemployment or an uncompensated time gap between two consecutive jobs which is longer than 1 month. Job-to-job matches are defined as job switches with no time gap or a time gap smaller than 1 month. Whenever a worker and an establishment match twice, we only consider the first match and exclude job recalls.

The worker ranking procedure we apply relies on the HLM model with search on the job.²² In this framework, identification of the worker ranking is based on wage information from matches formed by workers out of unemployment only. These wages can be shown to be monotonically increasing in the unobserved worker type in the context of the model. As a result, we lose all workers in our sample who are never unemployed and only switch jobs. Workers unemployed at some point in time, however, can be ranked and we follow them as they switch jobs.²³ Our final sample consists of 183,156 matches, 75,831 arise out of unemployment and 107,325 from workers switching jobs.

3 Sorting Theory

To support our analysis of labor market sorting in Germany, this section briefly summarizes the theory of sorting that guides our thinking about the allocation of workers to jobs and the determination of wages. Thinking about sorting conceptually requires heterogeneity of agents in a two-sided market. Since this is a paper about the German labor market, we denote our heterogeneous types *workers* and *firms*. Firms are heterogeneous in terms of their productivity and workers have different skill levels. Workers and firms carry identifiers i and k , respectively. We assume for now that unambiguous

²²Details are discussed in Section 4.2.

²³All our conclusions are unaffected by this choice. Ranking workers based on all their wages, not just out of unemployment, leads to a very similar worker ranking with a correlation of above 0.99. In this case, the final sample contains about three times as many matches. Even though the difference between the two rankings is very small, we use the ranking out of unemployment for the sake of methodological cleanness.

rankings of both worker and firm types exist and are observable to the researcher. In practice, rankings are unobservable. Establishing these rankings empirically is thus the key challenge in measuring the sign and strength of labor market sorting.

Let the true rank of worker i be denoted $x(i)$ while the true rank of firm k is $y(k)$. For simplicity, heterogeneity across workers and firms is assumed to be one-dimensional.²⁴ The starting point for the theory of labor market sorting is the neoclassical optimal assignment model proposed by Becker (1973). In the absence of search frictions, every worker finds his optimal firm assignment instantaneously and matches are formed. This leads to a Walrasian first-best allocation of workers to firms. The specific nature of the optimal assignment depends on the production structure of the economy. One possible assignment pattern is positive assortative matching (PAM). In Becker's theory, a complementarity between worker and firm types gives rise to PAM. Technically, the production function, which takes worker and firm types as arguments, must be supermodular, that is, the cross-derivatives must be strictly positive. To maximize output, the complementarity then requires the most productive firm to employ the most skilled worker, the second most productive firm to employ the second most skilled worker, and so forth. The worker with the lowest skill level is optimally matched with the least productive firm. This allocation implies a perfect positive correlation of worker and firm ranks: $\rho = \frac{\text{Cov}[x(i), y(k)]}{\sigma_{x(i)} \sigma_{y(k)}} = 1$, where ρ is Spearman's rank correlation coefficient, computed by dividing the covariance of the rank variables by the product of their standard deviations.²⁵ Spearman's ρ is a natural measure for the sign and strength of sorting. A negative rank correlation, $\rho < 0$, implies negative assortative matching (NAM) with a submodular production function. NAM features high-productivity firms employing low-skilled workers (and vice versa) to maximize output. In turn, $\rho = 0$ indicates the absence of sorting patterns in allocating workers to jobs, corresponding to a modular production technology.

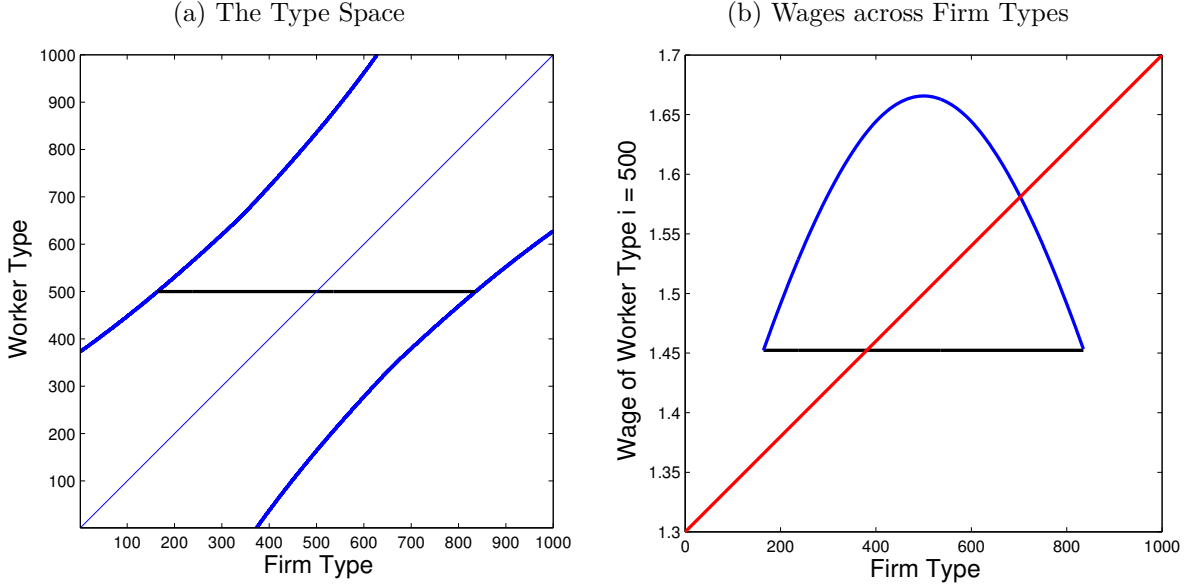
A major step forward for the theory of labor market sorting has been the incorporation of search frictions into the Beckerian model. Shimer and Smith (2000) show existence and the properties of an equilibrium in a two-sided assignment model with frictions. With a continuum of types and random search, frictions imply that the first-best allocation cannot be realized because the probability of meeting any specific partner type on the other side of the market is zero. Therefore, a subset of firm and worker types in the vicinity of the optimal allocation must be acceptable. These subsets of types, the matching sets, are determined by the option values of all potential matches.²⁶ Figure 1 depicts the equilibrium of a simple Shimer and Smith (2000) economy. This illustrative economy is

²⁴Of course, this is a simplifying assumption. However, it is not obvious ex-ante that having more than one dimension of heterogeneity would meaningfully increase the explanatory power of the model. For some recent explorations of the empirical content of multi-dimensional sorting models see Lise and Postel-Vinay (2016) and Lindenlaub and Postel-Vinay (2016).

²⁵Spearman's ρ is equal to a standard Pearson correlation coefficient applied to the rank variables.

²⁶Matches are acceptable if their option value is higher than the value of continued search.

Figure 1: A Simple Sorting Model



Note: The blue curves depict the equilibrium of a simple Shimer and Smith (2000) economy. The black line represents the equilibrium matching set of a worker of type 500. The model is a simplified version of the structural search model presented in HLM. It is solved numerically using standard parameter values. For the derivation and solution procedure see also Schulz (2015). The red line in Panel (b) is a simple stylized representation of a monotonic relation between worker's wages and the type of the firm as it is assumed in two-way fixed effect models of wage dispersion.

populated by a population of 1,000 discrete and heterogeneous worker and firm types, each of them having an equal mass normalized to one. Since we abstract from the ranking problem for now, the indices $(i, k) \in \{1, \dots, 1000\}$ are directly interpretable as the workers' and firms' ranks. The production function exhibits a complementarity which induces PAM as explained above. Panel (a) of Figure 1 shows the space of all worker and firm type combinations, the type space. The 45-degree line represents the Beckerian first best allocation. The blue curves depict the cutoffs of the matching sets of all worker and firm types. The black line is the matching set of one specific worker type, $i = 500$. This worker type is willing to match with, roughly, all firm types above 180 and below 820. Within these boundaries, the surplus is (weakly) positive and matches are formed. Matching with firm types outside of the boundaries, in turn, is precluded due to a negative surplus. According to the model, no matches are formed in the upper-left and lower-right corners of the type space, indicating the absence of matches between high-type workers and a low-type firms as well as low-type workers and high-type firms.

Panel (b) shows how the model-generated wage of the exemplary worker type, $i = 500$, varies across firm types. In the model, wages are determined by splitting the match surplus according to the Nash bargaining solution. Note that due to the production complementarity in this economy, the wage of a given worker type is non-monotonically related to the firm type. For worker type $i = 500$, it is maximized in a match with firm $k = 500$ (first best allocation). Apart from the wage-maximizing allocation, the worker has to accept lower wages because the production complementarity is not fully exploited.

The worker needs to compensate the firm for the foregone option value of waiting for a better hire. This non-monotonicity lies at the heart of the critique of the AKM two-way fixed effect approach (Gautier and Teulings, 2006; Eeckhout and Kircher, 2011; Lopes de Melo, 2015). In the log-linear AKM model, the wage must monotonically increase in the estimated firm-fixed effect, which is simply being equated with the unobservable firm type. Such a monotonic relation between wages and firm type might look like the stylized red line in Panel (b) of Figure 1. The higher the firm type, the higher is the wage of every worker the firm employs, a simple wage premium. The theoretically predicted non-monotonicity is clearly at odds with the AKM model, which assumes monotonicity. This potential misspecification of the AKM model led to a number of alternative strategies to quantify the contributions of worker and firm heterogeneity as well as their potential interaction to overall wage dispersion. The HLM identification strategy, which we rely on to a significant extent, makes direct use of the model structure described in this section. We discuss it in relation to alternative identification strategies in the next Section.

4 Identifying Sorting

We identify the sign and strength of sorting in the labor market by combining the wage-based worker ranking procedure (global rank aggregation) proposed by HLM with firm rankings which are constructed independently of wages. The next subsections explain our motivation for this approach and our ranking procedures in detail before we turn to the results in Section 5.

4.1 Motivation

The main contribution of HLM is to show how to construct global worker and firm rankings and identify the sign and strength of sorting using data on wages and labor market transitions only. Their identification strategy is viable in the class of structural models discussed in Section 3, building upon Shimer and Smith (2000). The quality of wage and transition data is typically high in commonly available matched employer-employee data sets. Detailed information about the firm, however, is not always available. It is thus an asset that the firm data requirements of the HLM approach are small. For Germany, we have access to very detailed information for a large number of firms from the IAB Establishment Panel.²⁷ The German case is therefore well-suited to take the HLM approach to a test: if all necessary information to rank both workers and firms globally were contained in wages, additional firm data would be redundant and not alter the results in any meaningful way. We show below that using additional firm data leads

²⁷HLM use data from the IAB Establishment Panel as well. Vacancy data at the firm level greatly simplify the computation but is not necessary for their approach to work in principle.

to a different firm ranking and different results.

Once workers are ranked and binned using the global rank aggregation algorithm (see details below), the value of unemployment for a specific worker type can be estimated by interpreting the lowest observed wage of all workers of the same type as this type’s reservation wage. Subsequently, HLM use the model structure, primarily Nash bargaining, to show that firms can be ranked based on the value of a vacant job.²⁸ It can be constructed as a statistic of wages and transition rates. The value of a vacant job is shown to be monotonically increasing in the firm type and it is identified from the observed differences of wages paid to different workers of the same type and their estimated reservation wage. This wage premium is proportional to the expected surplus of a vacancy due to the assumption of Nash bargaining.²⁹ Thus, the bargaining assumption is crucial to identify the value of a vacancy using wage data. Once the worker-specific value of unemployment and the firm-specific value of a vacancy are known, the model structure allows HLM to simply invert the wage equation for match-specific output. Thus, the production function is identified non-parametrically on the match level in this setting. As a consequence, it is possible to analyze the sign and strength of sorting without making functional form assumptions about the production function.

The assumption that wages are determined by contemporaneous bargaining made by HLM is critical for identification as it enables them to rank firms based on wages and invert the wage equation for match-specific output. Empirically, the assumption that wages are determined by splitting the surplus of a match period-by-period is backed up by evidence presented in Hagedorn and Manovskii (2013). For the U.S., they show that wages are robustly not history dependent but driven by current aggregate labor market conditions and idiosyncratic match-specific productivities, consistent with period-by-period bargaining.³⁰ For Germany, new evidence suggests that history dependence is, in contrast to the U.S., a prevalent feature of wage determination. Bauer and Lochner (2016) use the on-the-job-search model proposed by Hagedorn and Manovskii (2013) and estimate it on German social security data (SIAB-7514). Following Hagedorn and Manovskii (2013), they explicitly control for unobserved match-quality and even allow for heterogeneity across occupations. Unlike in the U.S., wages turn out to depend on the initial labor market conditions or on the best outside offer a worker receives while being in a specific job.³¹ The influence of current labor market conditions is thus smaller

²⁸Technically, Nash bargaining is not specifically required. It is sufficient that both parties benefit from an increase of the surplus, the proportions do not have to be fixed.

²⁹See Hagedorn et al. (2016) p. 15 ff. for more details and proofs.

³⁰In contrast, earlier studies for the U.S., most notably Beaudry and DiNardo (1991), find history dependence in wages and interpret it as evidence for infrequent bargaining or implicit contracts. This result disappears after controlling for unobserved match quality as Hagedorn and Manovskii (2013) show.

³¹The labor market conditions at a specific point in time are measured by the unemployment rate. Empirically, these papers use the unemployment rate at the begin of an employment spell to control for initial conditions. The likelihood of receiving outside offers, which the firm may match to retain the worker, is captured by the lowest measured unemployment rate during an employment spell.

than in the U.S. This suggests that the bargaining assumption necessary to rank firms and invert the wage equation might be less appropriate for Germany than it would be in the U.S. context.

The baseline results in HLM are derived using an extended version of the model with search on the job. As in Postel-Vinay and Robin (2002) and Cahuc et al. (2006), in case an employed worker meets another firm, the firms engage in Bertrand competition and depending on the potential surplus at the new firm the worker might switch jobs. HLM show that their constructive proof of identification extends to this richer framework under the assumption that unemployed workers extract the full surplus of the match when they are hired. In this case, the wage is monotonically increasing in the unobserved worker type, the key requirement to rank workers. A disadvantage of identifying the model from out of unemployment wages only is that workers who are never observed in an unemployment spell cannot be ranked.³²

Using German data, HLM find a very high correlation of estimated worker and firm ranks, it is 0.76 for the years 1993-2007. This value suggests a high degree of PAM and severe misspecification of AKM models. It is much higher than the correlation of estimated worker and firm-fixed effects reported by CHK using the AKM methodology on German data. CHK report correlations of 0.17 (1996-2002) and 0.25 (2002-2009).³³ The HLM result also lies considerably above rank correlations reported in related studies using alternative identification strategies and data from other countries. Bagger and Lentz (2015) use Danish data and estimate a structural model with endogenous search intensity and on the job search. Job-switchers are used to rank firms by their poaching rank. They report a correlation of only 0.11. Bonhomme et al. (2016) propose a clustering technique (finite mixture model) to identify a discrete number of firm types based on the similarity of their within-firm wage distributions. This very flexible empirical model allows for unrestricted interactions between worker and firm heterogeneity. Using Swedish data, they find correlations around 0.44. Bartolucci et al. (2015) rely on high-quality Italian firm data (balance sheets). They develop a refined reduced-form approach to identify the sign and strength of sorting and find a correlation of worker and firm types of 0.52. Finally, evidence for the U.S. is presented by Lise et al. (2016). They make a parametric assumption about the production function (CES) and directly estimate the elasticity of substitution. They also find evidence for PAM, but the magnitude of estimated substitution elasticity is not readily comparable to a rank correlation coefficient.

The sign of sorting found in all these studies is unambiguously positive, suggesting that

³²We find that ranking workers based on all their wages, not just out of unemployment, leads to a very similar worker ranking with a correlation of above 0.99 for the workers in both samples. However, we stick to the ranking out of unemployment for the sake of methodological cleanness.

³³Recall that CHK have access to the universe of German social security records. Using the AKM model on those data does not suffer from the limited mobility bias emphasized by Andrews et al. (2008, 2012).

some degree of PAM is indeed a prevalent feature of labor markets in developed economies. The variation of estimated correlation coefficients, however, is huge.³⁴ Using our method, we shed light on the large spread of estimated correlation coefficients for Germany. As mentioned before, the benchmark studies are CHK using the AKM model (0.17-0.25) and HLM using a structural framework (0.76). Both studies use German social security records and comparable time periods. Different correlation coefficients must therefore stem from methodological differences. Using additional firm information from the IAB Establishment Panel is instructive because we can construct firm rankings independently of wages and thus test whether the HLM approach of ranking both workers and firms on wage data alone is a reliable way of capturing both worker and firm heterogeneity.

We find confirmatory evidence for PAM in the German labor market. Interestingly, our results are much closer to the CHK-AKM benchmark than to HLM. Using the HLM worker ranking and our alternative firm ranking, we find an overall rank correlation of 0.24 (1998-2008). Thus, the specific structural assumptions made to rank firms based on wages, particularly wage bargaining, appear to severely increase the measured rank correlation. Overall, we find that the AKM model applied by CHK seems to deliver a satisfactory approximation of the data. However, we do not conclude from this observation that match-specific effects play no role for wage determination whatsoever. Low-type workers have large estimated residuals in CHK.³⁵ For these workers, the empirical wage profiles we observe are at odds with the monotonicity assumption. The sorting of low-type workers indeed appears to be driven by match-specific interaction effects, possibly as a result of production complementarities as suggested by theory. Applying the AKM model shrouds these effects: the monotonicity assumption may be met for large fraction of the data, leading to the well-known high explanatory power of the model. Deviations for smaller groups of workers, however, remain undetected. Whether AKM or a more flexible model is applicable therefore depends on the specific research question: the AKM model can be a valid approximation to study the extent of wage dispersion in the labor market as a whole, jointly analyzing all worker and firm types. In a study that focuses on low-type workers, the group that labor market policy typically is most concerned with, the AKM model would fail because wage patterns are highly non-monotonic in this group. Below, we show in detail for which worker types wages across firms are indeed non-monotonic (in line with the theory depicted in Figure 1) or monotonically increasing in the firm type (in line with the AKM assumption of additively separable log wages). Before we turn to the results, we describe our ranking procedures in more detail.

³⁴To some extent the large variation of estimated correlation coefficients is probably driven by cross-country differences rather than methodological issues. Also, possible cyclical fluctuations of the degree of labor market sorting are not taken into account in the aforementioned studies. Both topics are exciting avenues for future research.

³⁵CHK show residuals for type combinations, measured in worker and establishment effect deciles, in Figure VI on p. 996. In fixed effect models, the estimated fixed effects of both workers and firms are implicitly interpreted as their cardinally rankable type.

4.2 Ranking Workers

HLM present a computational algorithm that merges intra-firm wage rankings into a global ranking of workers by solving a Kemeny-Young rank aggregation problem.³⁶ The procedure makes use of the fact that matched employer-employee data allows wage comparisons across firms because workers switch jobs: co-workers at one firm move from firm to firm over time and form a graph (or connected set) of workers with comparable wage observations. A computational algorithm running on this graph can effectively maximize the likelihood of the correct global ranking, as proven by Kenyon-Mathieu and Schudy (2007). The input of the algorithm are the workers' residual wages, \tilde{w}_{it} , net of observable effects as presented in Section 2.2. The algorithm is initialized by ranking workers according to a simple wage statistic which needs to be monotonically increasing in the unobserved worker type.³⁷ Using a Bayesian approach with a normal prior, HLM show how to compute the probability of worker i being ranked higher than worker j given wage histories at firm k in the presence of measurement error:

$$c(i, j) = P(\tilde{w}_{i,k} > \tilde{w}_{j,k}) = \Phi \left(\frac{\bar{\tilde{w}}_{i,k} - \bar{\tilde{w}}_{j,k}}{\frac{\sigma^2}{n_{i,k}} + \frac{\sigma^2}{n_{j,k}}} \right). \quad (3)$$

Φ is the standard Normal CDF. Observed (residual) wages are assumed to follow a noisy process: $\tilde{w}_{i,k,t} = \tilde{w}_{i,k} + \epsilon_t$, with σ^2 being the variance of ϵ . Intuitively, the difference of the average residual wages $\bar{\tilde{w}}_{i,k} - \bar{\tilde{w}}_{j,k}$ at firm k is weighted by the wage variance σ^2 in proportion to the number of wage observations for workers i and j at firm k , $n_{i,k}$ and $n_{j,k}$. The more available observations, the smaller is the potential impact of measurement error on the average wage of worker i at firm k and the more plausible is the ranking implied by the wage observations at this firm, resulting in a higher value of $c(i, j)$.³⁸ Note that σ^2 is the overall wage variance and not firm-specific because HLM make the assumption that all variation in wages for a specific job stems from measurement error only.³⁹ The probability $c(i, j)$ is defined for worker pairs employed at the same firm.

³⁶Rank aggregation is an ancient problem that originated in social choice theory. Kemeny-Young rank aggregation solves this problem by minimizing the number of disagreements between potentially inconsistent rankings of voting alternatives by different voters, see Kemeny and Snell (1962). In the HLM application to the labor market, the firms are the voters and the workers the voting alternatives.

³⁷HLM prove that, in the context of their model, the reservation wage, the maximum wage, and the adjusted average wage of a worker are monotonically increasing in the unobserved type. Importantly, average wages, sometimes used to rank workers in empirical applications, are not monotonically increasing in the type because they do not factor in the values of workers' interjacent unemployment spells.

³⁸For details of the derivation of $c(i, j)$, see Appendix III.1 in Hagedorn et al. (2016).

³⁹While this assumption is consistent with period-by-period wage bargaining in their model, from an empirical perspective it could be desirable to allow for heterogeneity of the within-firm wage distributions beyond the mean. Imagine a firm using different contracts to discriminate between worker types: two workers could have different slopes in their wage profile over time because tenure is remunerated differently. Such patterns could be due to history dependence, as evidenced by Bauer and Lochner (2016), or due to the coexistence of wage bargaining and wage posting, as evidenced by Gartner and Holzner (2015) (both for Germany). Ranking the two workers based on their mean wage in this setting might

In case a pair of workers is observed at more than one firm, the wage observations are considered to be independent and the probabilities are simply multiplied. By comparing the initial ranking with the ranking implied by the posterior probabilities $c(i, j)$, the algorithm iteratively increases the value of the following objective function and, hence, maximizes the likelihood of the global ranking:

$$\sum_{i>j} [c(i, j) \Pi(i, j) + c(j, i) \Pi(j, i)]. \quad (4)$$

$\Pi(i, j) \Pi(j, i)$ is an indicator function that takes on the value 1 in case i (j) is ranked higher than j (i) and 0 otherwise. Whenever $c(i, j) > c(j, i)$ but $\Pi(i, j) = 0$ and $\Pi(j, i) = 1$, the values of the indicator functions are swapped and the value of the objective rises. The procedure continues until no further swap of workers increases the value of the objective. It runs on the set of worker pairs who are employed by the same firm at some point in time. The employment spells do not have to overlap.⁴⁰ We choose the “LIAB Mover Model” version of German matched employer-employee data because the sampling procedure maximizes the numbers of observed coworker pairs in our data, an ideal environment for the outlined computational procedure to run.⁴¹ Importantly, we do not need to observe all workers of a given establishment to compute $c(i, j)$. The pairwise comparison of residual wages of two workers at the same firm is not affected by a potential wage premium (or firm-fixed effect) because both workers receive it.⁴² We arrive at a final ranking which gives an estimate of the unobserved type, $\hat{x}(i)$, for every individual worker in our data. We group workers into 100 bins of equal size.⁴³ In the following, workers within one bin should be thought of as workers with the same estimated type $\hat{x}(i)$.

To understand the properties of our worker ranking $\hat{x}(i)$, we compare it to other wage statistics commonly used to rank workers in the empirical literature: mean wages, person-fixed effects (AKM), and rankings based on the observable characteristics age and education. Table 2 shows correlations of those statistics with our final ranking. Notably, the rank aggregation procedure produces a ranking which is markedly different

not yield the correct ranking. In contrast, the k-means clustering technique proposed by Bonhomme et al. (2016) allows for heterogeneity of within-firm wage distributions even beyond the second moment. However, the computational complexity of this method increases quickly with the number of moments to be estimated, hence the number of clusters/types is limited. The HLM method, in turn, allows for (almost) unique worker and firm ranks. The researcher faces a trade-off: to allow for more heterogeneity of the within-firm wage distributions, the number of types to be identified must be smaller. Our results indicate that it can be insightful to have a large number of types: it allows us to detect non-monotonic wage patterns for small worker groups.

⁴⁰Recall that residual wages are deflated and net of time effects.

⁴¹See Appendix A for sampling details.

⁴²This is a big advantage over the AKM model with respect to data requirements. In a reduced-form fixed effect model, the full workforce of a firm needs to be observed in order to reliably estimate the firm-fixed effect. CHK meet this data requirement by using the universe of German employment records.

⁴³The number of individual workers in every bin must be sufficiently large. We find that 100 bins/types is a good compromise between observations per bin and fineness of the type space which allows us to detect wage non-monotonicity even for small groups of workers.

Table 2: Properties of Worker Ranking

	\bar{w}_i	$\hat{\alpha}_i$	age	education
Correlation with $\hat{x}(i)$	0.75	0.87	0.19	0.48

Note: The table shows correlations of our estimated worker ranking ($\hat{x}(i)$) with other statistics used to rank workers in the literature: individual mean wages (\bar{w}_i) and estimated person-fixed effects (extracted from running the wage regression 2, $\hat{\alpha}_i$), as well as workers' observable characteristics, the individual means of age and education.

from rankings based on the alternative statistics. The correlations of our final ranking with individual mean wages (\bar{w}_i) and estimated person-fixed effects ($\hat{\alpha}_i$) are, naturally, positive but markedly different from 1. The correlation with the estimated worker-fixed effect (0.87) can be interpreted as a measure for how different the worker rankings in CHK-AKM and HLM are. Moreover, our final ranking is only weakly correlated with age and education. This is reasonable since the ranking procedure uses residual wages, which are net of the effect of observables.

To understand how the binning of workers modifies the ranking, Table 3 shows a decomposition of the respective variances of workers' observed wages, residual wages, age, and education into the shares explained within and between the bins. A relatively homogeneous distribution of a variable within the bins indicates that the binning provides a meaningful summary of the underlying heterogeneity in the respective dimension. Since our worker ranking is based on wages, the share of variance explained between the bins is relatively high, roughly two-thirds both for log wages ($\ln w_{it}$) and log residual wages ($\ln \tilde{w}_{it}$). Hence, our bins are internally homogeneous in terms of wage variation. Since our ranking is based on the residual wage (net of observables), the bins are much less homogeneous for the covariates age and education with, respectively, 95% and 66% share of the overall variance within the bins. We find workers of almost all ages (20-60 in our sample) in every bin. A high-type worker is not necessarily old and low-type workers are not simply young workers without experience. For our six education categories, the same is true, albeit to a lesser extent. Plots showing the distribution of age and education across worker bins are shown in Appendix B.2, Figure B.1.

4.3 Ranking Firms

As reasoned before, we seek to rank heterogeneous firms independently of wages. This approach sets a counterpoint to both CHK-AKM and HLM, who use wage data and worker mobility to capture firm heterogeneity. The downside of our approach is that the measures we use to rank firms are by construction firm-specific and not match-specific (like wages). Ideally, the researcher would like to directly observe the firms' share of

Table 3: Properties of Worker Bins

	$\ln w_{it}$	$\ln \tilde{w}_{it}$	age	education
Overall Variance	0.182	0.163	96.927	1.910
Between bins	0.119 (65%)	0.107 (65%)	4.725 (5%)	0.666 (35%)
Within bins	0.064 (35%)	0.057 (35%)	92.202 (95%)	1.244 (65%)

Note: The table decomposes the overall variance of log wages ($\ln w_{it}$), residual wages ($\ln \tilde{w}_{it}$), age, and education in our sample into the respective shares explained within and between the worker bins. The age of individual workers in our sample ranges from 20 to 60. There are 6 education categories: 1 = “no degree”, 2 = “vocational training”, 3 = “high school”, 4 = “high school and vocational training”, 5 = “technical college”, 6 = “university”.

rents obtained from matches with each individual worker.⁴⁴ In the absence of these data, we have to be content with the argument that match-specific influences are most likely largely integrated out of firm-level statistics, particularly in large firms.

We interpret firm heterogeneity in terms of efficiency. Conceptually, this idea builds upon stochastic frontier analysis (SFA).⁴⁵ An individual firm’s distance to the production frontier is a natural measure of relative performance. Given its inputs, capital and labor, the most productive firm in the economy is located on the frontier. Less efficient firms generate less output relative to their inputs. They have a positive distance to the frontier in output space. In this setting, each individual firm’s distance to the production frontier is inversely related to the firm-fixed effect, which is equivalent (in absolute terms) to the technical efficiency residual in SFA terminology. A firm on the frontier has a distance of zero and thus the highest estimated fixed effect.

To rank firms indexed by k based on the distance to the production frontier, we estimate a rich empirical model of firm performance with a large number of controls from the IAB Establishment Panel. We use a flexible translog specification with firm and time effects. Log value added is our measure of output, $\ln v_{kt}$. x'_{kt} contains time-varying explanatory variables. In the translog setting, these include capital, labor, capital and labor squared as well as their interaction.⁴⁶ z'_{kt} includes dummies for 32 sectors of the economy which we include as additional controls.⁴⁷ Time effects are captured by ω_t , ϕ_k

⁴⁴This would be equivalent to directly observing match-specific output. In this hypothetical case, it would not be necessary to make a structural assumption on bargaining (like HLM do) to identify the value of a firm’s vacancy and rank firms according to it.

⁴⁵SFA has been developed in the context of cross-sectional data, see Aigner et al. (1977). It was then extended to the panel data context by, among others, Schmidt and Sickles (1984) and Cornwell et al. (1990).

⁴⁶We approximate the capital stock using a perpetual inventory method, see Section 2.3. The labor input is measured by the size of the workforce.

⁴⁷We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel. The WZ classification of the German Federal Statistical Office is compatible to the common international classifications of industries, NACE and ISIC.

Table 4: Properties of Firm Ranking

	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	$\hat{\phi}_k$	$\ln \bar{N}_k$	$\bar{\pi}_k$
Correlation with $\hat{y}(k)$	0.47	0.75	0.94	0.17	0.60

Note: The table shows correlations of our firm ranking ($\hat{y}(k)$) with other statistics that could be used to rank firms: log mean value added (\bar{v}_k), log mean value added per worker ($\ln \bar{v}_k/\bar{N}_k$), and estimated firm-fixed effects (extracted from running regression 5, $\hat{\phi}_k$), the log of the average size of a firm's workforce, \bar{N}_k , as well as average profits per worker, $\bar{\pi}_k$.

is the firm-fixed effect, and r_{kt} is the residual.

$$\ln v_{kt} = \phi_k + \omega_t + x'_{kt}\beta + z'_{kt}\gamma + r_{kt}. \quad (5)$$

After running this regression, we rank firms based on their estimated fixed effect, $\hat{\phi}_k$. This is our estimate of each individual firm's type. We group the firms into 30 bins of equal size, denoting the estimated rank of all firms within one group $\hat{y}(k)$.

Table 4 shows the correlations of our firm ranking with other firm-level statistics: log value added, log value added per worker, the log of the firms' workforce size, and profits per worker (all in firm-level means), as well as the estimated firm-fixed effect. It is not surprising that the production frontier based ranking is highly correlated with the estimated firm-fixed effects due to the close link between the two. The moderately high correlations with log value added and log value added per worker, however, are not mechanic and show that the various controls included in the regression lead to a very different firm ranking as compared to, for instance, ranking on value added alone. Moreover, controlling for the size of the workforce, its square and the interaction with the capital stock in the translog setting leads to a low but positive correlation with the log of the mean workforce size, \bar{N}_k . The correlation of the production frontier based ranking with our measure of profits per worker (discussed below) is of the same order of magnitude as the ranking's correlation with value added and value added per worker.

As we did with the binned worker ranking, we decompose the variance of some key firm variables in our data into the shares explained between and within our firm bins to show in which dimension the bins are internally homogeneous and in which they are not. Table 5 shows the decomposition. The bins are internally homogeneous in terms of average profits per worker and log value added per worker with the majority of variance between the bins. This is not true for log value added alone, underlining the importance of controlling for size effects in the ranking exercise. Moreover, the variable indicating in which sector a firm operates has a very high share of within variance. Thus, the sector, which we also control for in the production function estimation, is not a strong determinant of a firm's position in the final ranking. In every bin, we find firms from

Table 5: Properties of Firm Bins

	$\bar{\pi}_k$	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	Sector
Overall Variance	4.734E+09	3.584	0.733	66.825
Between bins	2.715E+09 (57%)	1.084 (30%)	0.475 (65%)	6.321 (9%)
Within bins	2.019E+09 (43%)	2.500 (70%)	0.258 (35%)	60.504 (91%)

Note: The table decomposes the overall variance of profits ($\bar{\pi}_k$), log value added ($\ln \bar{v}_k$), log value added per worker ($\ln \bar{v}_k/\bar{N}_k$), and of the sectors the firms operate in into the respective shares explained within and between the firm bins in our sample. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = “Agriculture & Mining”, 3-18 = “Manufacturing”, 19-20 = “Construction”, 21-23 = “Retail Trade”, 24-32 = “Service Sector”.

almost all sectors. Additionally, we show in Figure B.2 the plotted distributions of firm size, the sector variable, as well as collective bargaining and employee representation dummies. There is no clear relation of our firm bins to collective bargaining schemes or employee representation, the dispersion of these attributes within the firm bins is huge.


The results we present in this paper do not hinge on a specific way of ranking firms. To illustrate this, we use an alternative measure of firm performance, the average profit per worker, and check the robustness of our findings. The average profit per worker is a simple and transparent statistic to rank firms. We build on Bartolucci et al. (2015), who use very detailed firm data (balance sheets) to study labor market sorting in the Italian region of Veneto.⁴⁸ Their main argument in favor of profits is that all firms share a similar objective: maximizing profits. In addition, as argued before, firms tend to be matched to a large number of workers and match-specific noise should thus be integrated out of firm-level profits.⁴⁹ On the other hand, in the context of the model presented in Section 3, average profits are not necessarily increasing in the unobserved type of the firm. Similar to the argument presented for the non-monotonicity of workers’ wages, profits could be non-monotonically related to the unobserved firm type. This would invalidate a firm ranking based on profit data. Since the production frontier ranking is unaffected by this theoretical obstacle, we chose it to generate our baseline results.⁵⁰ The correlation between the profit-based ranking and the production frontier ranking is very high, 0.79, indicating that both rankings do an almost similarly good job in summarizing firm heterogeneity. Hence, all our main conclusions are unaffected by the ranking choice. The next Section presents our baseline results, the robustness check using the profit ranking is relegated to Appendix C.

⁴⁸They test a variety of potential measures of firm performance based on the firms’ balance sheets and find that economic profits per worker are well suited to rank firms.

⁴⁹Conversely, workers are typically matched only to a small number of employers throughout their career, creating noisy wage histories.

⁵⁰In an earlier version of this paper, we used the profit-based ranking as baseline and checked robustness with the production frontier ranking. We exchanged the two rankings solely for the purpose of methodological cleanness, none of our main results changed.

Figure 2: Empirical Bivariate Match Density in Germany (1998-2008, $\rho = 0.24$)

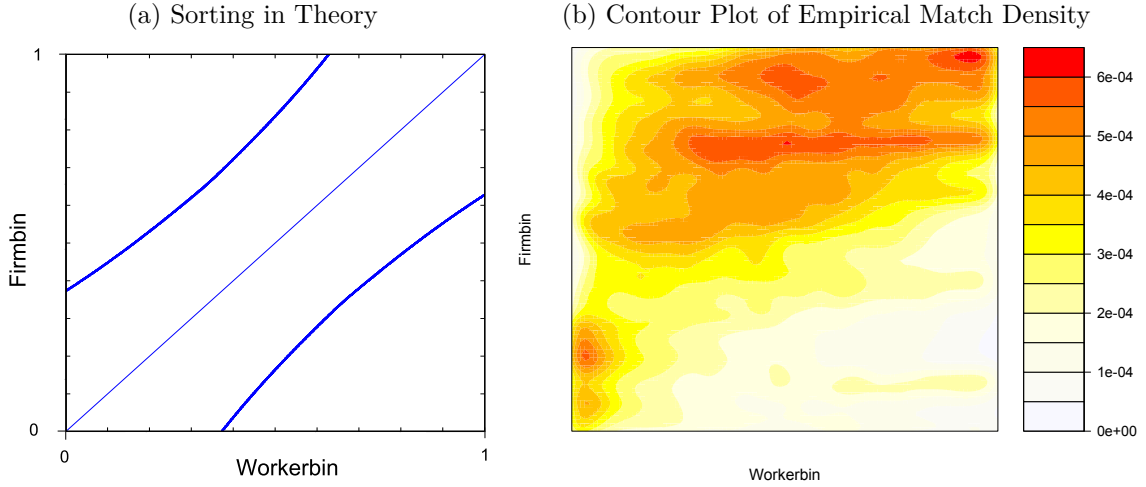


Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 (#worker types \times #firm types).

5 Labor Market Sorting in Germany

Having ranked both workers and firms, we are now in a position to construct and analyze the empirical bivariate density of matches across all possible combinations of worker and firm types. We use 100 worker and 30 firm bins based on the types estimated in the preceding section. To compute the density of matches, we use our definition of a match described in Section 2.4. A match is always the first new person-year observation of an employment spell between a worker and an establishment in our data, so we count every individual worker-establishment combination only once to be as conservative as possible when estimating rank correlations. This also excludes recall. Our match definition is well-suited for the analysis in this paper because our primary interest is the allocation of workers to jobs. The abstraction from the spell length ensures that certain match types are not over-represented (under-represented) in the distribution of matches simply because they last longer (shorter) on average, leading to more (less) person-firm-year observations. Figure 2 plots the bivariate density of matches to illustrate the sorting patterns in our data. There is a distinct tendency of low-type workers to be matched with low-type firms as indicated by the two spikes in the lower-left corner of the plot. Most of the matches in our data are located in the upper half of the plot, representing matches with firms in the upper half of the firm ranking. Dispersion is higher in this part of the

Figure 3: Sorting: Theory and Empirical Evidence



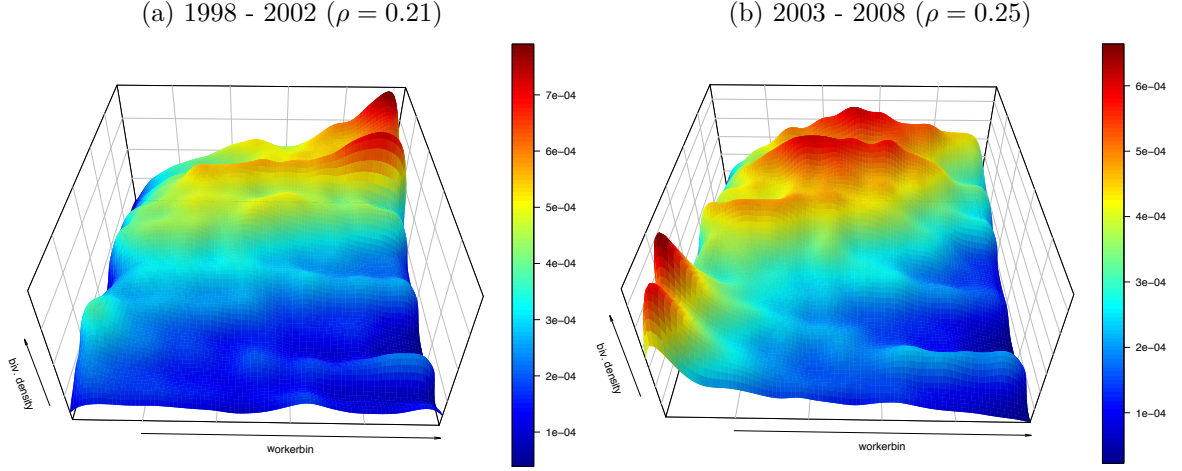
Note: In both Panels, worker and firm types are normalized into the unit interval to facilitate comparison. Panel (a) depicts the equilibrium of a simple Shimer and Smith (2000) economy. The model is a simplified version of the structural search model presented in Hagedorn et al. (2016). It is solved numerically using standard parameter values. For the derivation and solution procedure see Schulz (2015). Panel (b) shows a contour plot of the empirical match density depicted in Figure 2.

distribution, many low-type workers are employed by medium and high firm types. The maximum match density, however, is observed for matches between the highest worker and firm types. Correlating our binned worker and firm rankings for the full sample, we find that the rank correlation coefficient (Spearman's ρ) is significantly positive with a value of 0.24. This indicates that the matching process in the German labor market features positive sorting of workers to jobs, albeit not to a very high degree. Sorting is more pronounced for matches out of unemployment with a rank correlation of 0.26. In the sample of job-to-job switchers, the rank correlation is 0.20. The degree of sorting we find for all matches is somewhat higher than the values reported by CHK for the first half of our sample. CHK report a correlation of the estimated person and establishment effects of 0.17 for the years 1996-2002. We find a rank correlation of 0.21 for 1998-2002. For the second half of our sample, the correlation found by CHK is 0.25 (2002-2009); ours is virtually similar at 0.25 (2003-2008).⁵¹ HLM do not report rank correlations for sub samples, only an overall correlation of 0.76 (1993-2007).

Figure 3 presents a direct comparison of the theoretical model and the empirical density of matches. Panel (a) shows the theoretically optimal allocation along the 45-degree diagonal and the matching sets around it; Panel (b) is a contour plot of our bivariate match density in Figure 2. Recall that theory predicts that no matches are formed by type combinations outside the matching sets, that is, no matches of high-type firms and low-type workers or vice versa. Empirically, we clearly see in Panel (b) that

⁵¹See Table III in Card et al. (2013), “correlation of person/establ. Effects”, p. 994. Note that CHK compute the correlation for all person-firm-year observations in their sample, not just for new matches. All rank correlations reported here are rounded to two decimal points.

Figure 4: Empirical Bivariate Match Density by Sub-Period



Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 (#worker types \times #firm types). Ranks of workers and firms are time invariant, the plots are simply conditioned on the year of the match.

a large number of matches is located in the vicinity of the diagonal, in line with theory. This explains why we find a positive rank correlation. If one would connect the points with the highest match density for every worker bin, however, the result would not be a straight line from the lower-left to the upper-right corner. Rather, we would get a concave curve above the diagonal, suggesting that the empirical allocation of workers to jobs is different from the optimal allocation in the simple model. The density of matches in the lower-right corner is very low, indicating that high-skilled workers almost never work at low-productivity firms, in line with the theoretical matching sets. The prediction regarding the opposite corner, however, is not met by the data. The density is stretched out towards the upper-left corner, indicating that matches between low-type workers and high-type firms are common. The symmetry of the simple theoretical plot is thus rejected by the data, along with the diagonal optimal allocation.⁵²


5.1 Sorting over Time

Interestingly, the depicted sorting patterns in Germany changed considerably over time. We show this in Panels (a) and (b) of Figure 4. Splitting our sample into two distinct time periods reveals that sorting between high-type workers and high-type firms was very pronounced in the first half of our sample (Panel (a), 1998-2002).⁵³ In the second half

⁵²In the model, symmetry hinges on the assumption of equal bargaining powers of workers and firms. By relaxing this assumption, it is easily possible to align the theoretical model with the apparent asymmetry of matching patterns in the data. HLM show how the bargaining power can be measured in the data.

⁵³The time periods are 1998-2002 and 2003-2008. We interpret our worker and firm ranks as time-constant and do not rerun the ranking procedures on the subsamples. We simply condition on the year

Figure 5: Rank correlations by match type over time (1998-2008)



(Panel (b), 2003-2008), the allocation in the upper half of the plot became much more dispersed, the spike disappeared. The opposite is true at the bottom of the distribution. There was almost no match density for low-type workers and low-type firms before 2003⁵⁴, but in the second sub period we observe a huge increase of new matches between low worker and firm types. These contrary trends in the lower and upper half of the match density indicate that the overall increase of the rank correlation from 0.21 to 0.25 across the two time periods is the combined effect of two opposed developments: sorting strongly increased for low-type workers, their allocation moved closer to the theoretically predicted optimal allocation on the diagonal, increasing the overall rank correlation. The allocation of high-type workers, however, became more dispersed and moved away from the diagonal, reducing the overall rank correlation. Note that our sample split corresponds to the announcement and implementation of a big labor market reform program in Germany, the Hartz reforms.⁵⁵ Simply conditioning the match density on years suggests that the reforms potentially had a large effect on the allocation of workers to jobs in the labor market. We will investigate this hypothesis in more detail throughout the remainder of this paper.

Figure 5 shows a different decomposition of sorting patterns over time. The blue

when computing the match density.

⁵⁴The small hump is not present in the conditional densities before the year 2000.

⁵⁵The so-called Hartz reforms consist of four “Acts for Modern Services on the Labor Market” (‘Gesetze für moderne Dienstleistungen am Arbeitsmarkt’), which came into effect on January 1, 2003 (Hartz I and II), on January 1, 2004 (Hartz III), and on January 1, 2005 (Hartz IV). Hartz reforms I – III were mainly concerned with active labor market policies (ALMPs). Hartz IV significantly reformed the unemployment benefit system.

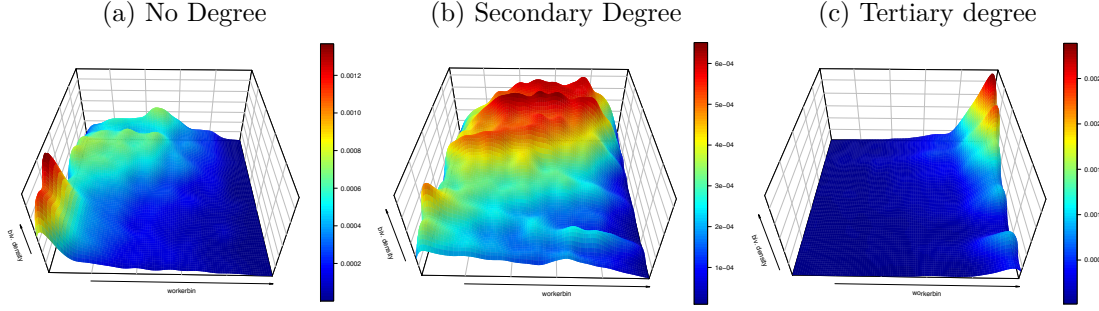
line plots rank correlation coefficients for all matches. We then differentiate between the two types of matches: new matches of workers out of unemployment and matches from workers who switch jobs, meaning they changed their employer without an intervening unemployment spell. About 41% of the matches in our sample are formed out of unemployment, the remaining 59% are matches of job switchers. This relation is stable over time. Overall, the extent of positive sorting in the German labor market increased during our period of observation. For all matches, the rank correlation rose from 0.20 in 1998 to 0.26 in 2008 (blue line). We see no uniform trend and considerable year-to-year variation of the correlation coefficient. Specifically, the correlation fell by about 24% from 2000 to 2001. This might be related to the recession in the early 2000s, which is suggestive of a cyclical component in sorting patterns. Conditioning the rank correlation on the different match types sheds some more light on the time variation. For matches out of unemployment, a surprisingly clear trend emerges (red line): the rank correlation increased almost uniformly from 0.17 in 1998 to 0.31 in 2008. In turn, most of the year-to-year variation we observe for the overall rank correlation appears to be driven by job-to-job switchers. The green line exhibits considerable fluctuation and there is no clear time trend for this group of matches. The observed degree of sorting for new matches resulting from search on the job did not change significantly during our period of observation. If anything, it has slightly decreased from a correlation of 0.20 in 1998 to 0.19 in 2008.⁵⁶

Our analysis of the changing rank correlation over time suggests that the relative importance of the two match types for the overall degree sorting has reversed. Specifically, in the beginning of our sample, the rank correlation was slightly higher for job-to-job switchers, but starting in 2001, the rank correlation for new matches out of unemployment dominates. This is remarkable since theoretical models of on the job search typically suggest that sorting patterns should be more pronounced for new matches resulting from search on the job.⁵⁷ Recall that we use only wage information from employment spells formed out of unemployment to rank workers using the HLM procedure. This implies that we cannot rank workers who switch jobs throughout our sample but are never observed as being unemployed. We therefore suspect that the observed rank correlation for job-to-job switchers might be somewhat depressed in our results. This is, however, inconsequential for the remaining analysis because we largely focus on low-type workers and match formation out of unemployment.

⁵⁶A table with all rank correlations and numbers of observations by year and match type can be found in Table B.1 in Appendix B.

⁵⁷In a sequential bargaining model à la Postel-Vinay and Robin (2002) and Cahuc et al. (2006), workers gradually move toward their optimal employer (in terms of wages) and observed sorting is expected to increase. Unemployed workers, however, have an incentive to accept all job offers as long as the option value of the job lies above their reservation value, including unemployment benefits. The observation of a lower degree of sorting from search on the job as compared to sorting out of unemployment is somewhat puzzling from a theoretical perspective.

Figure 6: Sorting by Education



Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 ($\#$ worker types \times $\#$ firm types). Panel (a) contains workers in category 1 of our education variable (no degree), Panel (b) contains workers in categories 2, 3, and 4 (vocational training only, high school only, high school and vocational training), Panel (c) contains workers in categories 5 and 6 (technical college, university).

5.2 Sorting by Education

To better understand empirical sorting patterns, it is also instructive to condition on education. For the U.S., Lise et al. (2016) similarly find that matches are not uniformly distributed across the type space. By splitting their sample into college graduates and workers with a high school degree or less, they show that sorting patterns are very pronounced for highly educated workers but less so for the low education group. For these workers, their results are only suggestive of a mild production complementarity, which cannot be distinguished from linearity. The German picture looks different. Panel (c) of Figure 6 shows, consistent with the findings for the U.S., a high concentration of workers with tertiary degrees at the highest firm types. These are the workers at the top of our ranking and they are strongly sorted. The allocation of workers with secondary degrees in Panel (b) is much more dispersed. We find these workers in all bins besides the highest ones. They contribute to most of the match density in the upper half of the firm ranking. However, we also observe a spike for matches of the lowest worker types in this education category with low-type firms, consistent with our finding that sorting of low-type workers into low-type firms is very pronounced in Germany. Panel (a) underlines this picture. We find workers with no educational degree exclusively in the bottom half of our worker ranking and they are concentrated in matches with low-type firms to a high degree. In contrast to what Lise et al. (2016) find for the U.S., the sorting of low-type workers is very pronounced in Germany and it has increased over time (recall Figure 4).

5.3 Distributional Dynamics

We have established that labor market sorting in Germany has increased between 1998 and 2008 in terms of aggregate rank correlations. This increase is primarily associated with a higher sorting propensity out of unemployment, while the degree of sorting of

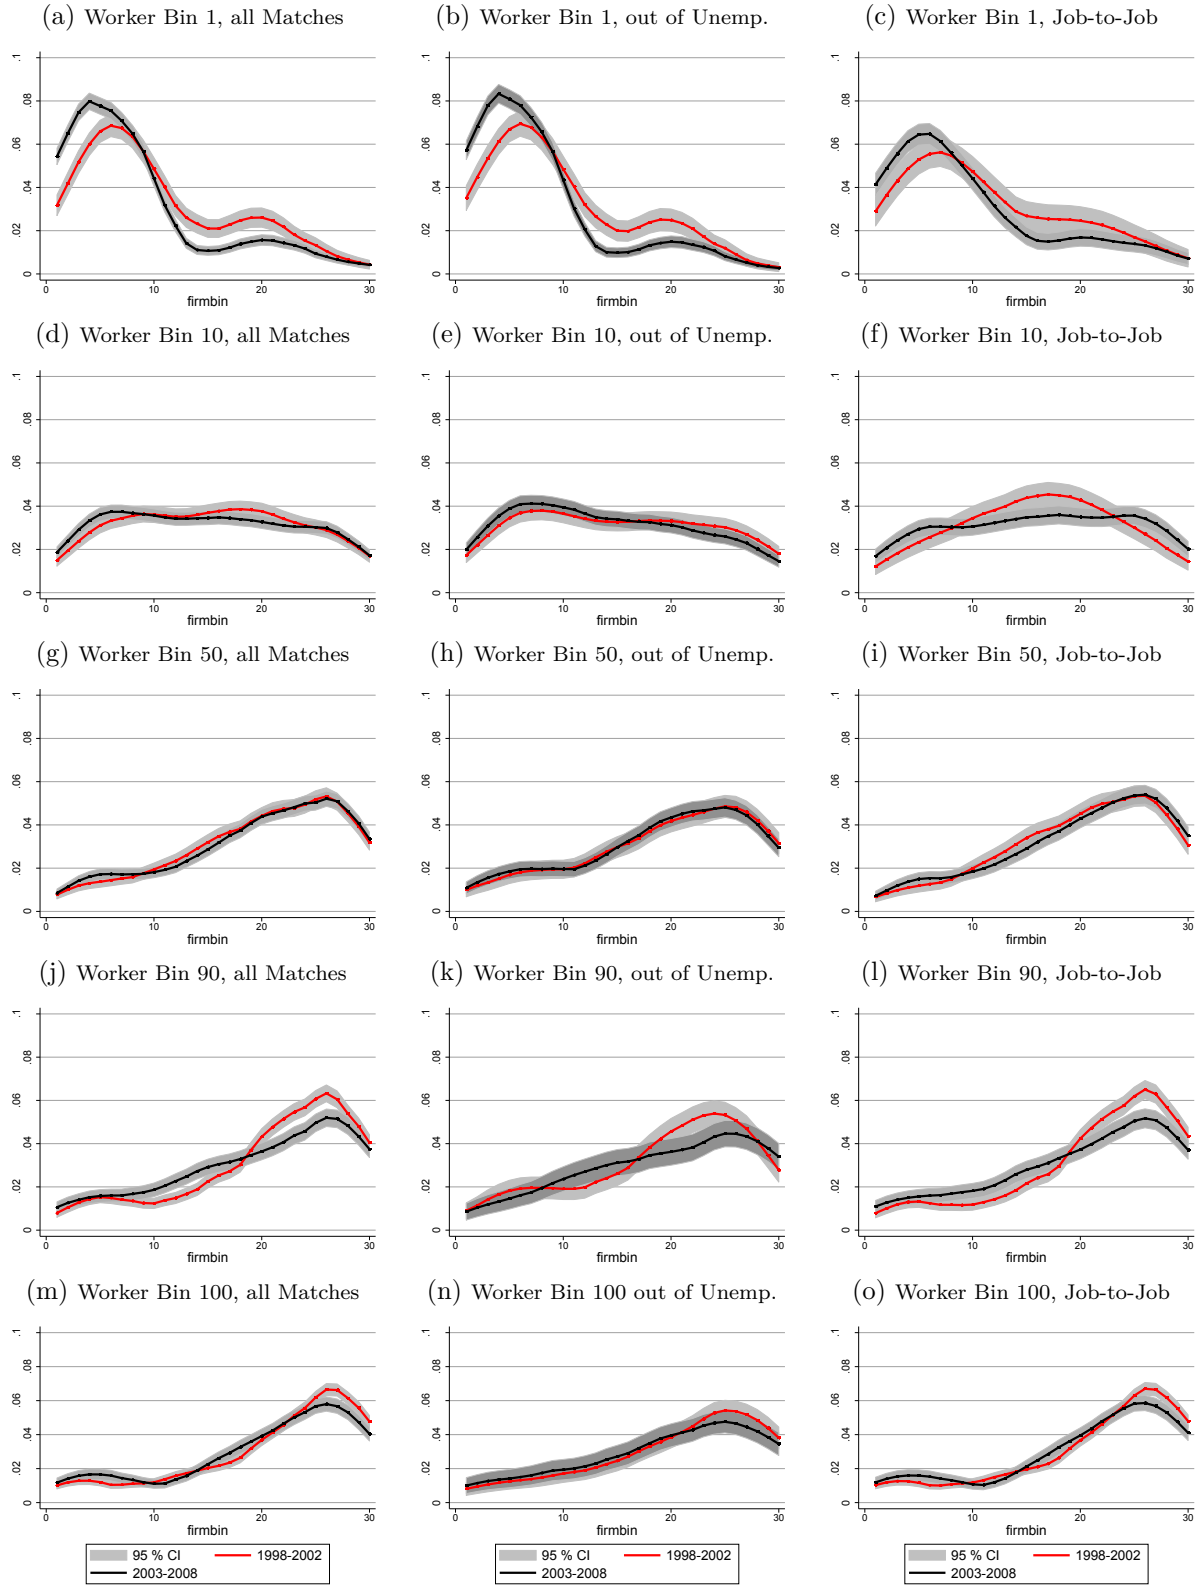
matches resulting from job switches is roughly constant over time in our data. To understand which worker-firm combinations contributed to the respective trends, we track the changing univariate distributions of worker types across firms over time, both for matches out of unemployment and for job-to-job switches. This allows us to show precisely for which worker types the distribution across firms shifted towards the theoretically predicted optimal allocation, where it shifted away from it, and where no significant change occurred. Figuratively, we slice through the empirical bivariate density of matches depicted in Figure 2 and compare the “density slices” for different time intervals to see where the distribution across firms changed and where it remained constant. Figure 7 shows the estimated univariate density functions for the worker bins 1, 10, 50, 90, and 100.⁵⁸ We compare the first half of our sample, 1998-2002 (red line), to the second half, 2003-2008 (black line). Notably, the distribution changed significantly primarily for low-type workers, see bin 1 in Panel 7a. For workers of medium and high-types, the density functions largely lie on top of each other and cannot be distinguished statistically.

For the lowest type of workers in bin 1, we observe an increase of the density of matches with firms of the lowest types. The density of matches with firms below, roughly, bin 8 has increased, significantly so below bin 6.⁵⁹ At the same time, the density of matches with firms in bins 12 to 22 has decreased significantly. Apparently, the distribution of workers in bin 1 across firms shifts to the left during our period of observation, see Panel 7a. Low-type workers become more concentrated in low-type firms, leading to increased sorting because workers of the lowest type moved towards their theoretically predicted optimal allocation. Having in mind the bivariate density in Figure 2, the spike of the bivariate density in the lower-left corner moved closer to the 45-degree line and became more pronounced. Distinguishing between matches out of unemployment and job-to-job switches reveals that the described shift is significant only for matches formed after an unemployment spell of the worker. This is in line with our observation that the aggregate rank correlation increased strongly for matches out of unemployment, recall Figure 5. Due to space constraints, we cannot plot estimated densities for all 100 worker bins. A significant distributional shift is present for all workers up to bin 8. The shifts are consistently more pronounced for matches out of unemployment than for job-to-job switches. In worker bin 10, we start observing distributions which do not change significantly over time, see Panels 7d through 7i. Bin 10 workers are distributed almost uniformly across firms. As we move up the worker ranking, the estimated density becomes more concentrated at the top of the firm ranking. We observe a small but

⁵⁸Histograms of the raw match data can be found in Appendix B, Figure B.3.

⁵⁹We deliberately choose to be conservative by determining statistical significance based on the overlap of confidence intervals. It is always true that with non-overlapping confidence intervals two statistics are significantly different. The converse, however, does not necessarily hold. As long as the difference of two statistics is significantly different from zero, a possible overlap of the confidence intervals does not imply insignificance.

Figure 7: Estimated Density Functions: 1998-2002 (red) vs. 2003-2008 (black)



Note: Estimated univariate kernel densities of matches conditional on worker bins, time, and match type. The kernel is estimated using an Epanechnikov kernel function. The bandwidth is calculated by Silverman's rule of thumb. Pointwise confidence intervals are calculated using a quantile of the standard normal distribution.

significant decrease of the match density at high-type firms starting around worker bin 90, which is mainly driven by job-to-job switches, consistent with the slightly shrinking rank correlation that we observe for this match type. For the highest worker type in bin 100, this distributional shift is still visible but only marginally significant. Overall, the distribution of high-type workers shifts away from high-type firms. Dispersion increases, in line with the development depicted for the whole type space in Figure 4, and sorting becomes less pronounced. The mode of the distribution of high-type workers, however, is well above firm bin 25 in both halves of our sample. This is already true in bin 50, see Panels 7g through 7o.


We use estimated kernel density functions in Figure 7 to analyze the changing distributions of workers across firms with confidence intervals. Our findings do not hinge on the density estimation procedure, they are present in raw data as well. To show this in the context of our two sub periods, 1998-2002 and 2003-2008, we plot the difference of the two empirical densities conditional on worker bin, time interval, and match type in Figure B.4. In case of a positive (negative) difference for a certain firm bin, the empirical density of matches has shifted upwards (downwards) for this worker-firm combination. The density differences confirm the observed distributional shifts towards more sorting at the bottom of the worker ranking, small changes in the middle, and slightly less sorting at the top.

5.4 (Non-)Monotonicity of Wages in Firm Types

Wages are the main determinant of how workers select themselves into jobs.⁶⁰ Hence, they are the key to understanding the distributional shifts and the increase of labor market sorting we observe in Germany. Our analysis is motivated by the theoretical prediction of non-monotonic wage patterns of worker types across firm types. With sorting, wages are maximized in the type-specific optimal allocation, arguably due to a production complementarity. Apart from this point, wages fall in both directions, what is at odds with the monotonicity assumption of the AKM model. Moving to a better firm does not necessarily lead to a higher wage because the worker needs to compensate the firm for not waiting for a better match. Our semi-structural identification method does not restrict the interaction of worker and firm heterogeneity. Nevertheless, we estimate rank correlations of the same order of magnitude (around 0.24) as the benchmark estimates provided by CHK, using the more restrictive AKM method. Figure 8 illustrates why this is the case. Simply plotting the mean of observed log wages for all combinations of worker

⁶⁰Sorkin (2015) uses revealed preference information contained in worker mobility patterns to construct a firm ranking and divide wage dispersion into the respective contributions of rents (or wages) and compensating differentials. Sorkin finds that compensating differentials explain up to 15% of the wage variance in the U.S. While this is a novel and interesting finding, the rents/wage-based explanation of worker selection into jobs that we and the papers we build upon focus on clearly dominates quantitatively.

Figure 8: Mean Wages for all Worker-Firm Type Combinations (1998-2008)

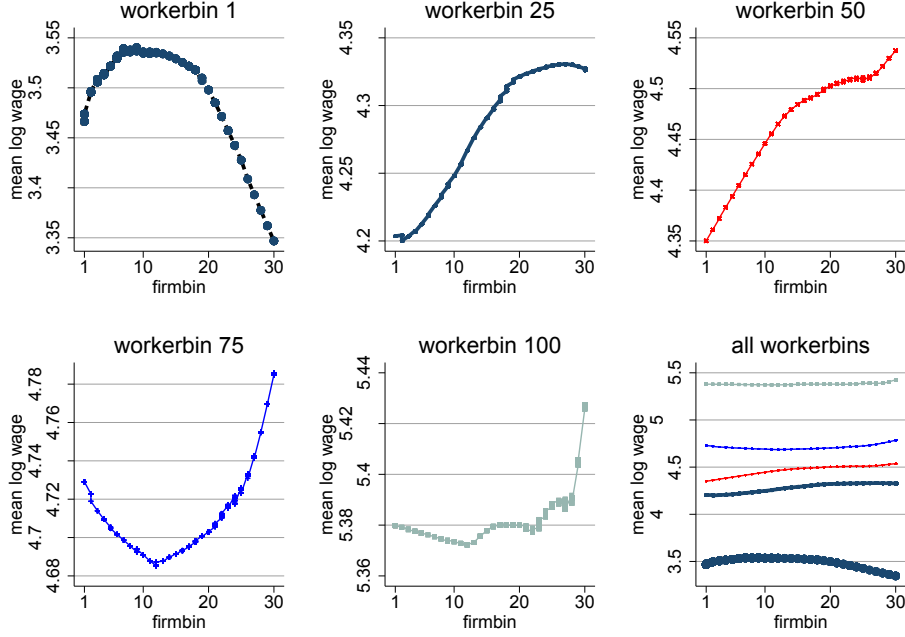


Note: The Figure shows the mean of the log real daily wage for all combinations of worker and firm types on a grid with dimensions 100×30 (#worker types \times #firm types).

and firm types according to our rankings reveals that linearity in logs, overall, is not a bad assumption to empirically analyze wage dispersion. The pattern of wages across the type space is not very different from a simple linear plane. This explains why the AKM framework produces a good fit. It does a good job in decomposing wage dispersion in the respective contributions of worker and firm heterogeneity, especially when the researcher has access to a data set as big and detailed as the universe of German social security records used by CHK.⁶¹ Consistent with CHK, we find that worker heterogeneity is the dominant source of wage dispersion. Wages increase strongly in the direction of the worker type in Figure 8. In the firm dimension, however, they hardly visibly increase. It almost appears as if, given the worker type, it does not matter much for the wage at what type of firm a worker is employed. As we will show in the following, this conclusion, which would call the importance of labor market sorting into question, is premature. We find that the large overall wage dispersion in a joint analysis of all workers and firms masks non-monotonic wage patterns for specific worker types, particularly at the bottom of the ranking. This is consistent with the fact that deviations from monotonicity are

⁶¹Recall that it is imperative for applying the AKM model to observe all workers at every establishment in order to consistently estimate firm-fixed effects. Most samples of matched employer-employee data do not meet this requirement and suffer from the limited mobility bias emphasized by Andrews et al. (2008, 2012). Our semi-structural method is well-suited also for smaller samples because it is not necessary to jointly estimate worker and firm-fixed effects.

Figure 9: Mean Wages across Worker Types



Note: Wages are means of real log cell-wages for the given worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

most pronounced for worker types with large residuals in an estimated AKM model. CHK themselves state that those residuals might indicate systematic departures from the monotonicity assumption and demand an in-depth analysis.⁶² We attempt to provide such an analysis in the remainder of this paper.

Figure 9 plots wage patterns across firms for single worker types. The observed higher tendency of low-type workers to sort into low-type firms indeed appears to be associated with a non-monotonic wage pattern across firms, in line with theoretical sorting models. The mean log wage of our lowest worker type (worker bin 1) increases at first, reaches its maximum in matches with firms around bin 10 and decreases thereafter in the type of the firm. This wage pattern suggests selection of low-type workers into low-type firms, simply because these are the wage-maximizing matches for them. This explains the increase of sorting at the bottom of the bivariate match distribution. To the best of our knowledge, we are the first to provide direct empirical evidence for a negative relation between a worker's wages and a performance measure of the firms he is matched with. Note, however, that this negative relation is visible in our data only for the lowest worker types. Consistent with the good overall fit of the AKM model, we find that wages are monotonically increasing in the firm type for the majority of workers in our sample. In bin 25, one can observe a small non-monotonicity at the top where wages level off and

⁶²A plot of the residuals and the related discussion can be found on p. 996 in CHK.

slightly decrease at the highest firm types. For medium-type workers in bin 50, wages are monotonically increasing everywhere. For high-type workers, the fixed effect model and the theoretical sorting model are observationally equivalent. Both frameworks predict the highest wages in matches with high-type firms. For workers in bins 75 and 100, monotonicity is only partly met. For the upper two thirds of firm types, wages increase as expected. In the bottom third, where high worker types are matched with firms of low types, we see stark deviations from monotonicity. Wages decrease at first, reach a minimum around firm bin 12 and increase only thereafter. One possible explanation for this pattern is that these workers are managers or executives at these low-type firms. Their wages might be determined in a way very different from the simple wage equations in CHK-AKM and HLM. For instance, managers could be able to extract rents from the firm in form of higher pay, even (or especially) at badly performing firms.⁶³ The analysis of mean wages per worker bin across firm types confirms that the AKM assumption of monotonically increasing wages is met for a large part of our data, as we already suspected from Figure 8. For low-type workers, it is clearly violated. Match-specific effects appear to be a more important determinant of wages for low-type workers as compared to the majority of higher-ranked workers. In Section 5.6, we conjecture that increased sorting at the bottom and the non-monotonicity of wages are related to the tendency of firms to increasingly outsource workers in certain occupations.

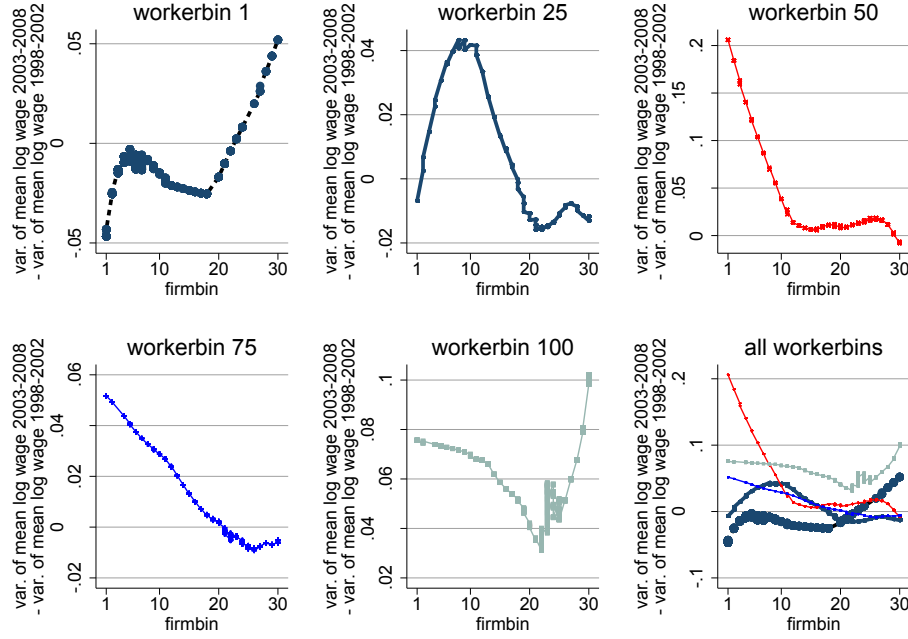
Our results are in line with CHK in the sense that deviations from the two-way fixed effect model are small on average. We have shown, however, that deviations are large and systematic for low-type workers. Depending on the question at hand, the AKM model can do a good job in explaining the link between unobservable characteristics of medium and high-type workers, firms, and the wages paid. However, the model leads to wrong conclusions regarding low-type workers because the non-monotonicity of their wages is at odds with AKM. From a policy perspective, it appears crucial to take the negative relation between low-type workers' wages and the type of firm they are matched with into account, for instance to optimally design unemployment insurance.

5.5 Wage Inequality

Overall, wage dispersion in Germany has strongly increased during our period of observation. This is well-known, e.g. from Dustmann et al. (2009). CHK report that the variance

⁶³Indeed, rent extraction by managers might be a cause of bad firm performance. We see identifying different compensation schemes for different worker types within the firm as a fascinating avenue for further research, related to recent advances in our understanding of CEO compensation, see Gabaix and Landier (2008). Unfortunately, we cannot distinguish normal workers from managers and executives in German data. We find, however, that dropping a fraction of workers at the very top of the wage distribution within every firm reduces the prevalence of higher wages of high-type workers at low-type firms as compared to medium-type firms.

Figure 10: Differences of the Variance of Log Wages over Time and across Worker Types



Note: Variances are calculated from log real daily wages of 30x100 bin-cells. The Figure shows differences between the variance in 2003-2008 and the variance in 1998-2002 for certain worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

of log real wages has increased from about 0.2 in 1995 to 0.34 in 2009.⁶⁴ This is an increase of 70%. The ranking tools we use in this paper allow us to analyze how wages change for single worker types and how these changes contribute to the overall trend of increasing wage dispersion. To this end, we plot the difference of the log wage variances over two time periods for the same worker bins as before (see Figure 10). We observe negative differences of log wage variances for low worker types (here bin 1) and firms in the bottom two thirds of the firm ranking. These negative differences indicate that the wage variance within these worker bins has decreased over time. This effect is most pronounced for workers in bin 1 who are employed at firms of type 1. For low-type workers, the variance has decreased stronger in the parts of the firm ranking most affected by the documented distributional shifts. In other words, increased sorting of low-type workers decreased the wage dispersion in these worker bins and, hence, compressed the wage distribution at the bottom. For worker bins 25, 50, and 75, the wage variance increased most for matches with low-type firms, particularly so for workers in bin 50. At the top of the firm ranking, wage dispersion did not change much or decreased even slightly. The wage distribution in bin 100 has become more unequal across all firm types. The lowest variance differences are found with firm types between bins 20 and 25, indicating that the slight downward

⁶⁴These variances include imputed wages beyond the censoring threshold. For more information about the imputation procedure and wage variances, see our Section 2 and Table I on p. 975 in CHK.

Figure 11: Wage Dynamics



Note: The Figure shows deviations of mean log real daily wages for single and grouped worker bins relative to 1998 (multiplied by 100).

shift of the match density we documented for these worker-firm combinations might have somewhat dampened wage growth (recall Figure 7m).

Since the observed changes of sorting and the related wage patterns have led to a reduction of within-bin wage inequality for low-type workers and an increase of within-bin wage inequality for high-type workers, it is worth asking the question how these developments relate to the overall increase of wage dispersion in Germany. Therefore, in Figure 11 we plot how the mean wages for our worker bins evolved over time. Panel (a) confirms, in line with CHK and Dustmann et al. (2009), that overall wage inequality in Germany has increased over time during our period of observation. For the highest worker types in bins 81-100, real log daily wages grew by more than 20%. For medium worker types in bins 21-80, wages increased by around 10% to 15%. On the other hand, real log daily wages for the 20 lowest worker bins decreased by about 6% relative to 1998.

Dustmann et al. (2014) point out that a large part of the recent favorable development of the German labor market⁶⁵, can be traced back to decreasing real wages at the lower end of the wage distribution. They argue that this wage moderation was largely a result of a decentralization of the wage-setting process.⁶⁶ Zooming into the lowest worker bins as depicted in Panel (b) of Figure 11 reveals a novel fact relative to Dustmann et al. (2014). Using our rankings, we observe that wages for low-type workers have not been uniformly decreasing over time. In fact, the wages of workers in bin 1 increased by

⁶⁵The German aggregate unemployment rate decreased from 11.7% in July 2005 to 6% in July 2016 according to data from the Federal Statistical Office.

⁶⁶Traditionally, wages in Germany are set by employer associations, trade unions, and works councils, typically at the industry level. The share of workers covered by this kind of industry-wide agreement decreased sharply because firms started opting-out increasingly from the mid 1990s, primarily to make the wage setting process more flexible.

about 10%. The wages of workers in the other bins depicted here (3, 5, 7, 9) largely evolved between -5% and 5% without a clear time trend.⁶⁷ Hence, not all low worker types contributed to increasing wage dispersion in Germany. The decentralization of the wage-setting process emphasized by Dustmann et al. (2014) supposedly affected low-type workers in the manufacturing sector in particular because collective bargaining is traditionally strong in this sector. This wage moderation, or negative real wage growth, quantitatively dominates in worker bins 1-20. Low-type workers in, for example, the service sector might have very different wage dynamics. In line with our finding of non-monotonic wage patterns for workers of the lowest type and their increased propensity to sort into their wage-maximizing jobs at the bottom of the firm ranking, workers in bins 10 and below exhibit wage dynamics which are not in line with the wage moderation hypothesis, see Panel (b). However, the increased sorting of low-type workers into their wage-maximizing jobs, which is in line with wage growth for low-type workers and their reduced within-bin wage dispersion, was apparently not strong enough to revert the overall trend of increasing wage dispersion in Germany.

5.6 Sorting and Domestic Outsourcing

Finally, we present an explanation for the increased sorting of low-type workers into low-type firms and their non-monotonic wage patterns. We observe that a large share of the new matches of low-type workers and low-type firms is concentrated in a specific sector of the economy, the so-called business service firms. These firms are at the center of a trend towards domestic labor service outsourcing in Germany.⁶⁸ They provide services that many firms traditionally organized internally, for instance cleaning, logistics, security, and food services. Temporary work agencies are another type of business service firms. They allow firms to hire workers for tasks directly related to production, for instance assembly line workers in manufacturing, without directly employing them. By outsourcing workers, firms can exclude them from high wage premia within the firm and reduce their overall wage bill. This point is made by Goldschmidt and Schmieder (2015), who provide a thorough analysis of domestic outsourcing in Germany using an AKM model and the universe of social security records. They show that the wages of displaced workers drop by around 10 log points due to foregone wage premia.⁶⁹ Note that in an AKM model, the establishment-fixed effect is the only channel capable of explaining a wage change

⁶⁷We selected bins 1, 3, 5, 7, and 9 for visibility reasons only. The same plot for bins 2, 4, 6, 8, and 10 can be found in Figure B.5 with the patterns of wage growth being very similar.

⁶⁸Outsourcing is commonly understood as the process of relocating tasks beyond the boundary of the firm. The term domestic implies that this happens within a country rather than internationally (offshoring).

⁶⁹By identifying on-site outsourcing events, where the worker and the location of the job stay the same but the employer changes, Goldschmidt and Schmieder (2015) can analyze the causal effect of outsourcing on wages. These are, however, only a small share of all outsourcing events.

when a worker switches employers. The estimated worker-fixed effect does not change by definition and the additive separability assumption excludes additional interactions of unobserved heterogeneity.

We re-evaluate the interplay between domestic outsourcing and observed wages using our more flexible worker and firm rankings. This makes a difference as compared to an AKM model because our framework allows a worker’s wage to increase when he switches jobs to a firm with a lower rank. A positive match-specific component of the wage may outweigh a negative change of the firm component and lead to an overall increase. As we demonstrated in Section 5.4, these patterns can be observed empirically. Wages do indeed rise on average when low-type workers move down the firm ranking. We find that this wage pattern and the increased sorting of low-type workers are primarily driven by business service firms. In other words, displaced workers do not necessarily suffer a wage loss and the reason lies in possible production complementarities as highlighted by sorting theory. For example, imagine a manufacturing plant with a small number of directly employed cleaning personnel. Since cleaning is not at the core of its business, the establishment does not have an incentive to invest in increasing the productivity of these workers. Once the cleaners are displaced, they might be employed at a highly specialized cleaning service provider. This business service firm has a large incentive to invest into the productivity of these workers because it directly affects its output. This is an example of a complementarity between the worker and the firm type, which determines wages in theoretical sorting models.

To show that domestic outsourcing is a quantitatively important phenomenon, we first introduce a distinction between four broad sectors in the economy: business service firms, manufacturing, consumer services, and other services. Business service firms are at the heart of the outsourcing hypothesis and include temporary work agencies, security services providers, cleaning companies, and other supplementary services specifically offered to firms. Manufacturing comprises all companies engaged in the industrial production of goods. Consumer services include trading firms and the retail sector. Insurance companies, leasing providers, and other firms supplying more complex services to firms are included in the other services category.⁷⁰ For simplicity, we condense our rankings and group all worker bins into four and all firm bins into three broad groups.

In our full sample, more than 20% of all new matches in the lowest 10 firm bins involve a business service firm. Most of these firms in our sample are temporary work agencies. The majority of them has been founded after 1998, so these are young firms. For medium and high firm bins, the percentages of new matches with business service firms are much smaller, only 4.6% (firm bins 11-20) and 2.5% (firm bins 21-30), respectively.

⁷⁰To be precise, we allocate companies to the four broad sectors using the WZ93/WZ03 classification of industries available in the IAB Establishment Panel. The respective WZ codes are 1-500 for manufacturing, 501-700 for consumer services, 701-744 for other services, and 745-748 for business service firms.

Table 6: Percentage Point Differences of Industry-Shares in New Matches by Worker Bins and Firm Bins (1998-2002 vs. 2003-2008)

	(a) Business Services			(b) Manufacturing		
Firm Bins	1-10	11-20	21-30	1-10	11-20	21-30
Worker Bins 1-25	10.82	1.24	0.08	-1.04	1.98	2.53
Worker Bins 26-50	2.28	0.50	0.23	-1.90	1.33	3.77
Worker Bins 51-75	1.06	0.09	0.31	-3.76	0.61	3.03
Worker Bins 76-100	0.32	-0.08	0.15	-6.51	-4.25	-8.42

	(c) Consumer Services			(d) Other Services		
Firm Bins	1-10	11-20	21-30	1-10	11-20	21-30
Worker Bins 1-25	1.50	1.02	0.88	0.20	0.68	0.06
Worker Bins 26-50	-1.01	-0.28	-0.09	-0.25	-0.05	0.28
Worker Bins 51-75	-0.65	-0.26	-0.15	-0.18	-0.90	-0.16
Worker Bins 76-100	-0.94	-0.52	-1.04	0.07	-1.12	-1.46

Note: WZ Codes 1-500 = “Manufacturing”, WZ Codes 501-700 = “Consumer Services”, WZ Codes 701-744 = “Other Services”, WZ Codes 745-748 = “Business Services”. All the percentage point differences in this Table sum to zero by construction.

Additionally conditioning on the worker type, leads to a share of new matches with business service firms in excess of 36% for worker bins 1-25. Hence, more than a third of all matches between low-type workers and low-type firms in our sample is likely a result of domestic outsourcing activities. The German labor market reforms⁷¹ play an important part in the observed increase of domestic outsourcing, particularly due to the deregulation of temporary work agencies. The Hartz I reform package, which came into effect on January 1, 2003, greatly liberalized temporary employment and subcontracted labor.⁷² Accordingly, we observe that of all matches which low-type workers (bins 1-25) formed with low-type firms (bins 1-10) in the business service sector, about 77% were formed after 2002.

Table 6 provides more details about the changing distribution of matches across our

⁷¹The so-called Hartz reforms consist of four “Acts for Modern Services on the Labor Market” (‘Gesetze für moderne Dienstleistungen am Arbeitsmarkt’), which came into effect on January 1, 2003 (Hartz I and II), on January 1, 2004 (Hartz III), and on January 1, 2005 (Hartz IV).

⁷²Before the reform, it was prohibited to repeatedly hire a worker on temporary contracts, to rehire a temporary worker within 3 months, to synchronize the length of the contract with an agency and the length of the assignment to a firm, and to assign a worker to a firm for more than 24 months. All these rules were abolished as a part of the Hartz I reform.

four sectors. The listed numbers are percentage point differences of the shares of new matches in every sector between the two sub periods used before, 1998-2002 and 2003-2008, for different combinations of worker and firm types. For instance, we observe that the share of new matches of low-type workers (bins 1-25) and low-type firms (bins 1-10) in the business service sector has increased by 10.82 percentage points. It rose from 8.90% (1998-2002) to 19.72% (2003-2008). This is by far the largest change we observe in our sample. Similarly, but to a lesser extent, new matches of low-type workers with medium-type firms and of medium-type workers with low-type firms has increased by 1.24 and 2.28 percentage points, respectively. Conversely, manufacturing firms have reduced their direct hiring, in line with the outsourcing hypothesis. Particularly, low-type manufacturing firms have negative percentage point differences with all worker types. Moreover, manufacturing firms particularly reduced the hiring of high-type workers between the two sub periods. For consumer and other services, the percentage point differences are small. Low-type consumer service firms increased their share of matches with low-type workers by 1.5 percentage points, the biggest change for these two sectors. It rose from 10.11% to 11.61%, so the share of new matches in this cell is substantial, so it contributes to observed sorting patterns, but it did not increase as much as in the business service sector. We conclude that most of the dynamics of sorting we document in this paper are driven by domestic outsourcing of manufacturing firms to business service firms. Manufacturing firms reduce their direct hiring of workers and, apparently, increasingly rely on business service firms to satisfy their labor demand. Interpreting this trend in relation to low-type workers' non-monotonic wage profiles leads to the conclusion that this development and the liberalization of subcontracted labor did not necessarily lead to wage losses for low-type workers as the AKM model would suggest.

6 Conclusion

The main task we conduct in this paper is to reconcile empirical models of wage dispersion in the spirit of AKM with recent structural work emphasizing the importance of production complementarities and match-specific effects for wage determination and the sorting of heterogeneous workers into heterogeneous firms. In the German case, the difference of estimated rank correlations using either an AKM model (CHK) or a structural equilibrium search model (HLM) is huge. We start out from the structural identification procedure proposed by HLM and test its main identifying assumption, wage bargaining, which allows them to rank both workers and firms based on wages. We use additional firm data to construct an efficiency-based alternative firm ranking which is independent of wage information. Correlating the HLM worker ranking with the independent firm ranking delivers rank correlation which are slightly higher but broadly in line with CHK. While this confirms that AKM models generate a good fit, we document important de-

viations from wage monotonicity and a tendency towards more labor market sorting in the portion of the type space where the AKM model produces the largest residuals.

Moreover, this paper provides a detailed empirical analysis of labor market sorting in Germany which reveals a number of novel empirical facts: first, we find evidence for an increasing degree of positive sorting in the German labor market throughout a period of profound institutional change. Sorting has increased particularly for low-type workers out of unemployment, who show an increased propensity to match with low-type firms, their theoretically predicted optimal match. Second, we present direct empirical evidence that the wages of low-type workers decline in the type of the firm they are matched with. This wage pattern drives the distributional shift of low-type workers towards low-type firms. Higher wages in these matches guide the sorting of low-type workers, supporting the non-monotonicity prediction of sorting theory. Such wage patterns are at odds with the AKM fixed effect model of wage dispersion, which assume that firms pay the same wage premium to all their workers. Third, many of the new matches that contributed to the increased sorting of low-type workers involve business service firms. Increased domestic outsourcing, which has been liberalized as part of the German labor market reforms, is an important driving force behind the observed shift of the distribution of low-type workers across firms. In contrast to an AKM analysis of this trend, we conclude that domestic outsourcing does not necessarily imply wage losses for low-type workers. Fourth, our results are in line with previous studies documenting the rising overall wage inequality in Germany. We see deviations from this trend for low-type workers using our ranking method. Their wages grew by about 10% over the eleven years of our data (1998-2008), in line with the wage non-monotonicity and increased sorting. However, this wage growth for low-type workers could not overturn the overall trend of increasing wage dispersion.

We did not attempt to make causal claims linking recent changes of labor market policy in Germany to the changing allocation of workers to jobs. It would appear far-fetched to assert, however, that the changes occurred in isolation. Taking our data at face value, we do not see a clear-cut discontinuity in the sorting patterns which could have been triggered by, for example, the Hartz IV reform in Germany, which reduced unemployment benefits for long-term unemployed workers. Rather, our analysis supports the conclusion of Dustmann et al. (2014), who find that the recent trends in the German labor market—the severe reduction of aggregate unemployment as well as increasing wage dispersion—can be traced back well into the 1990s and are not a direct result of the Hartz reforms. Rather, institutional changes which take effect over prolonged time periods—for instance the decline of collective bargaining at the industry level or the trend towards domestic outsourcing—have been determinative for the development of the wage distribution and labor market sorting in Germany.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis (1999) “High Wage Workers and High Wage Firms,” *Econometrica*, Vol. 67, pp. 251–333.
- Aigner, Dennis, C.A.Knox Lovell, and Peter Schmidt (1977) “Formulation and Estimation of Stochastic Frontier Production Function Models,” *Journal of Econometrics*, Vol. 6, pp. 21 – 37.
- 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 (Statistics in Society)*, Vol. 171, pp. pp. 673–697.
- Andrews, M.J., L. Gill, T. Schank, and R. Upward (2012) “High Wage Workers Match with High Wage Firms: Clear Evidence of the Effects of Limited Mobility Bias,” *Economics Letters*, Vol. 117, pp. 824 – 827.
- Atakan, Alp E. (2006) “Assortative Matching with Explicit Search Costs,” *Econometrica*, Vol. 74, pp. 667–680.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2006) “The Polarization of the U.S. Labor Market,” *American Economic Review*, Vol. 96, pp. 189–194.
- Bagger, Jesper and Rasmus Lentz (2015) “An Empirical Model of Wage Dispersion with Sorting,” unpublished.
- Bartolucci, Cristian, Francesco Devicienti, and Ignacio Monzón (2015) “Identifying Sorting in Practice,” unpublished.
- Bauer, Anja and Benjamin Lochner (2016) “History Dependence in Wages and the Role of Match Quality: Evidence from Germany,” unpublished.
- Beaudry, Paul and John DiNardo (1991) “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data,” *Journal of Political Economy*, Vol. 99, pp. 665–688.
- Becker, Gary S. (1973) “A Theory of Marriage: Part I,” *Journal of Political Economy*, Vol. 81, pp. 813 – 846.
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2016) “A Distributional Framework for Matched Employer Employee Data,” unpublished.
- Bundesagentur für Arbeit (1988) *Klassifizierung der Berufe - Systematisches und Alphabetisches Verzeichnis der Berufsbenennungen*, Bundesagentur für Arbeit, Nürnberg.

- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin (2006) “Wage Bargaining with On-the-Job Search: Theory and Evidence,” *Econometrica*, Vol. 74, pp. 323–364.
- Card, David, Jörg Heining, and Patrick Kline (2013) “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *Quarterly Journal of Economics*, Vol. 128, pp. 967–1015.
- Cornwell, Christopher, Peter Schmidt, and Robin C. Sickles (1990) “Production frontiers with cross-sectional and time-series variation in efficiency levels,” *Journal of Econometrics*, Vol. 46, pp. 185 – 200.
- Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener (2014) “From Sick Man of Europe to Economic Superstar: Germany’s Resurgent Economy,” *Journal of Economic Perspectives*, Vol. 28, pp. 167–88.
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009) “Revisiting the German Wage Structure,” *Quarterly Journal of Economics*, Vol. 124, pp. 843–881.
- Eeckhout, Jan and Philipp Kircher (2011) “Identifying Sorting-In Theory,” *Review of Economic Studies*, Vol. 78, pp. 872–906.
- Evers, Lisa, Christoph Spengel, and Julia Braun (2015) “Fiscal Investment Climate and the Cost of Capital in Germany and the EU,” ZEW policy brief, ZEW Mannheim.
- Fischer, Gabriele, Florian Janik, Dana Müller, and Alexandra Schmucker (2009) “The IAB Establishment Panel - Things Users Should Know,” *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 129, pp. 133–148.
- Fitzenberger, Bernd, Aderonke Osikominu, and Robert Völter (2006) “Imputation Rules to Improve the Education Variable in the IAB Employment Subsample,” *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 126, pp. 405–436.
- Gabaix, Xavier and Augustin Landier (2008) “Why has CEO Pay Increased So Much?” *The Quarterly Journal of Economics*, Vol. 123, pp. 49–100.
- Gartner, Hermann and Christian Holzner (2015) “Wage Posting as a Positive Selection Device: Theory and Empirical Evidence,” CESifo Working Paper 5494, CESifo.
- Gautier, Pieter A. and Coen N. Teulings (2006) “How large are search frictions?” *Journal of the European Economic Association*, Vol. 4, pp. 1193–1225.
- (2015) “Sorting and the Output Loss Due to Search Frictions,” *Journal of the European Economic Association*, Vol. 13.

- Goldschmidt, Deborah and Johannes F. Schmieder (2015) “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure,” Working Paper 21366, National Bureau of Economic Research.
- Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii (2016) “Identifying Equilibrium Models of Labor Market Sorting,” *Econometrica*, forthcoming.
- Hagedorn, Marcus and Iourii Manovskii (2013) “Job Selection and Wages over the Business Cycle,” *The American Economic Review*, Vol. 103, pp. 771–803.
- Heining, Jörg, Peter Jacobebbinghaus, Theresa Scholz, and Stefan Seth (2012) “Linked-Employer-Employee Data from the IAB: LIAB-Mover-Model 1993-2008 (LIAB MM 9308),” FDZ-Datenreport 01/2012, The Research Data Centre (FDZ) at the Institute for Employment Research (IAB), Nürnberg.
- Kantenga, Kory and Tzuo Hann Law (2016) “Sorting and Wage Inequality,” unpublished.
- Kemeny, John G. and J. Laurie Snell (1962) *Mathematical Models in the Social Sciences*: Blaisdell, pp.145.
- Kenyon-Mathieu, Claire and Warren Schudy (2007) “How to Rank with Few Errors,” in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC ’07, pp. 95–103, New York, NY, USA.
- Kölling, Arnd (2000) “The IAB-Establishment Panel,” *Schmollers Jahrbuch: Journal of Applied Social Science Studies / Zeitschrift für Wirtschafts- und Sozialwissenschaften*, Vol. 120, pp. 291–300.
- Lindenlaub, Ilse and Fabien Postel-Vinay (2016) “Multidimensional Sorting Under Random Search,” unpublished.
- Lise, Jeremy, Costas Meghir, and Jean-Marc Robin (2016) “Matching, Sorting and Wages,” *Review of Economic Dynamics*, Vol. 19, pp. 63–87.
- Lise, Jeremy and Fabien Postel-Vinay (2016) “Multidimensional Skills, Sorting, and Human Capital Accumulation,” unpublished.
- Lopes de Melo, Rafael (2015) “Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence,” unpublished.
- Müller, Steffen (2008) “Capital Stock Approximation Using Firm Level Panel Data,” *Jahrbücher für Nationalökonomie und Statistik*, Vol. 224, pp. 357–371.
- Postel-Vinay, Fabien and Jean-Marc Robin (2002) “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, Vol. 70, pp. 2295–2350.

- Schmidt, Peter and Robin C. Sickles (1984) “Production Frontiers and Panel Data,” *Journal of Business & Economic Statistics*, Vol. 2, pp. 367–374.
- Schulz, Bastian (2015) “Wage Rigidity and Labor Market Dynamics with Sorting,” Ifo Working Paper 199, Ifo Institute at the University of Munich.
- Shimer, Robert and Lones Smith (2000) “Assortative Matching and Search,” *Econometrica*, Vol. 68, pp. 343 – 369.
- Sorkin, Isaac (2015) “Ranking Firms Using Revealed Preference,” unpublished.

A Details of Data Preparation

Our analysis is based on German matched employer-employee data. We use the “LIAB Mover Model” (file: LIAB_MM_9308). This section serves to detail the various data preparation and imputation procedures we apply. The LIAB Mover Model is based upon the IAB Establishment Panel. First, establishments are selected which employ at least one employee who is employed at least at two different establishments of the IAB Establishment Panel. Up to 500 additional employees per establishment are chosen randomly. The sampling procedure includes a robustness check regarding the number of employees in a certain establishment, i.e. whenever the information in the IAB Establishment Panel survey data deviates by more than 50% from the information in the register data the establishment is excluded.

A.1 Education Imputation

The employee education information is reported by employers after every year and whenever a job ends. Its quality may suffer because employers do not face consequences for non- and misreporting. However, the existence of a reporting rule allows for correction. It prescribes that only the highest educational degree of an employee needs to be reported. Therefore the individual educational attainment should not decline over consecutive job spells. The imputation procedure (IP1) as suggested by Fitzenberger et al. (2006) builds upon this reporting rule by assuming that there is any over-reporting in the data.

The original education variable distinguishes the following four different educational degrees: high school, vocational training, technical college and university. By imputing following the IP1 procedure we extrapolate both back and forwards and do some additional adjustments using individual information on age and occupational status. As a result we get six education categories which can be ranked in increasing order. However, we still observe missing entries of about 2 percent of the initial data after imputation. We drop these observations because we simply cannot make any statement about their true educational background.

A.2 Wage Imputation

In the LIAB data earnings are right censored at the contribution assessment ceiling (‘Beitragsbemessungsgrenze’). We use the pension insurance of workers and employees. This earnings limit is given by the statutory pension fund and is adjusted annually due to changes in earnings. First we deflate daily wages by using the CPI with base year 2005. Then we identify censored wage observations by comparing wages with the contribution assessment ceiling. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold. On average about 13% among all wage

Table A.1: Summary Statistics of the Wage Distribution (1998-2008)

	Mean	Std. Dev.	Min	Max
Censored	4.582	0.393	2.411	5.153
Imputed	4.618	0.455	2.411	7.132

Note: Summary statistics of the distribution of daily real log wages.

Table A.2: Additional Variance-Covariance Matrices

(a) Without Top-Coded Wages					(b) Including Occupational Controls				
	$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	\hat{r}_{it}		$\ln w_{it}$	$x'_{it}\hat{\gamma}$	$\hat{\alpha}_i$	\hat{r}_{it}
$\ln w_{it}$	0.126				$\ln w_{it}$	0.207			
$x'_{it}\hat{\gamma}$	0.006	0.005			$x'_{it}\hat{\gamma}$	0.031	0.016		
$\hat{\alpha}_i$	0.091	0.002	0.089		$\hat{\alpha}_i$	0.143	0.014	0.128	
\hat{r}_{it}	0.029	0.000	0.000	0.029	\hat{r}_{it}	0.034	0.000	0.000	0.034

Notes: Variance-Covariance matrix of regression model 1 without imputation of the censored part of the wage distribution. Top-coded wages are dropped. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual (\hat{r}_{it}). Rounded to three decimal places.

Notes: Variance-Covariance matrix of regression model 1 with 32 additional occupational controls, interacted with education and time effects. The variance of log wages ($\ln w_{it}$) is decomposed into the variance of observable characteristics ($x'_{it}\hat{\gamma}$), the person-fixed effect ($\hat{\alpha}_i$), and the residual (\hat{r}_{it}). Rounded to three decimal places.

observations are censored according to our definition. Then we follow Dustmann et al. (2009) and fit a series of Tobit regression on age-education-year-combinations to impute the right tail of the wage distribution. In all of these regressions we control for eight five-year age-categories, six education categories and all possible interactions between these two categories. The imputation methodology assumes that the error term in the Tobit regression is normally distributed and each education and each age category can have different variances. Hence for each year, we impute censored wages as the sum of the predicted wage and a random component which is obtained from the standard error of the forecast. This component is separately drawn from a normal distribution with mean zero and a different variance for each education and age category. Table A.1 shows moments of the imputed wage distributions compared to the censored wage distribution.

B Further Results

B.1 Rank Correlations


Table B.1: Rank correlations for different time intervals by type of match

	Spearman's ρ (Number of Observations)		
	all matches	out of unemp.	job-to-job
1998-2008	0.235 (183,156)	0.257 (75,831)	0.195 (107,325)
1998-2002	0.213 (83,348)	0.218 (32,399)	0.190 (50,949)
2003-2008	0.250 (99,808)	0.281 (43,432)	0.200 (56,376)
1998-1999	0.211 (28,493)	0.196 (11,062)	0.189 (17,431)
2000-2001	0.206 (40,698)	0.217 (15,408)	0.188 (25,290)
2002-2003	0.229 (28,860)	0.239 (12,051)	0.197 (16,809)
2004-2005	0.245 (33,913)	0.257 (14,925)	0.211 (18,988)
2006-2008	0.259 (51,192)	0.306 (22,385)	0.191 (28,807)
1998	0.198 (13,924)	0.171 (5,739)	0.204 (8,185)
1999	0.221 (14,569)	0.220 (5,323)	0.170 (9,246)
2000	0.233 (21,741)	0.198 (8,589)	0.238 (13,152)
2001	0.176 (18,957)	0.240 (6,819)	0.135 (12,138)
2002	0.233 (14,157)	0.256 (5,929)	0.188 (8,228)
2003	0.225 (14,703)	0.221 (6,122)	0.206 (8,581)
2004	0.215 (17,442)	0.244 (7,622)	0.177 (9,820)
2005	0.276 (16,471)	0.270 (7,303)	0.247 (9,168)
2006	0.271 (18,177)	0.301 (8,202)	0.214 (9,975)
2007	0.250 (18,187)	0.310 (7,857)	0.169 (10,330)
2008	0.260 (14,828)	0.306 (7,857)	0.190 (8,502)

Notes: We test the null hypothesis that worker and firm bins are statistically independent. All rank correlation coefficients are different from 0 at 1% level of significance. Rounded to 3 decimal places. Numbers of observations (new matches according to our definition) are reported in brackets.

B.2 Distribution of Observables Across Worker and Firm Bins

Figure B.1: Observable Characteristics of Workers by Bin (Full Sample, 1998-2008)



Note: Panels (a) and (b) show the means \pm one standard deviation of workers' age and education across worker bins. The age of individual workers in our sample ranges from 20 to 60. There are 6 education categories: 1 = "no degree", 2 = "vocational training", 3 = "high school", 4 = "high school and vocational training", 5 = "technical college", 6 = "university".

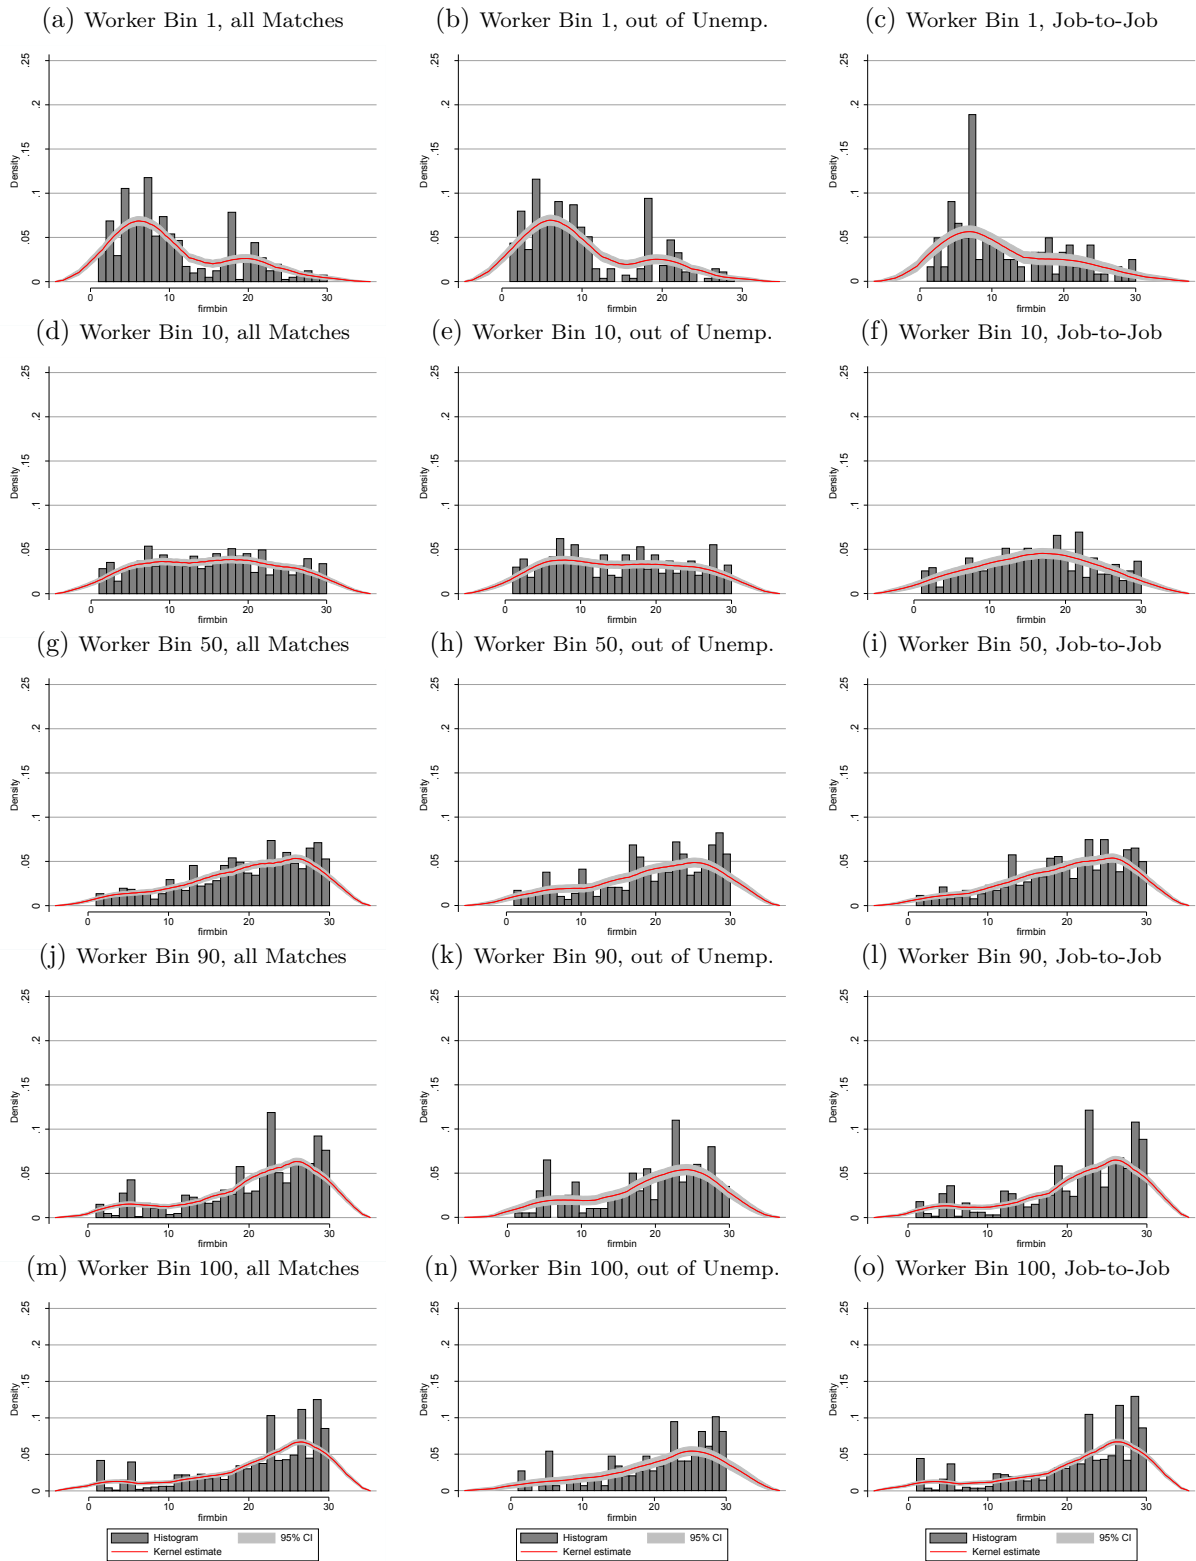
Figure B.2: Observable Characteristics of Firms by Bin (Full Sample, 1998-2008)



Note: Panel (a) shows means \pm one standard deviation of the average workforce size of firms within the bins. Panel (b) shows means \pm one standard deviation of the firms' sectoral classification within the bins. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = "Agriculture & Mining", 3-18 = "Manufacturing", 19-20 = "Construction", 21-23 = "Retail Trade", 24-32 = "Service Sector". Panel (c) shows means \pm one standard deviation of the variable indicating the application of a collective bargaining agreement within the bins: 1 = "sectorwide bargaining", 2 = "firmwide bargaining", 3 = "no collective bargaining". Panel (d) shows means \pm one standard deviation of the variable indicating the presence of formal employee representation within the bins: 1 = "employee representation exists", 2 = "no employee representation".

B.3 Histograms and Kernel Density Estimates

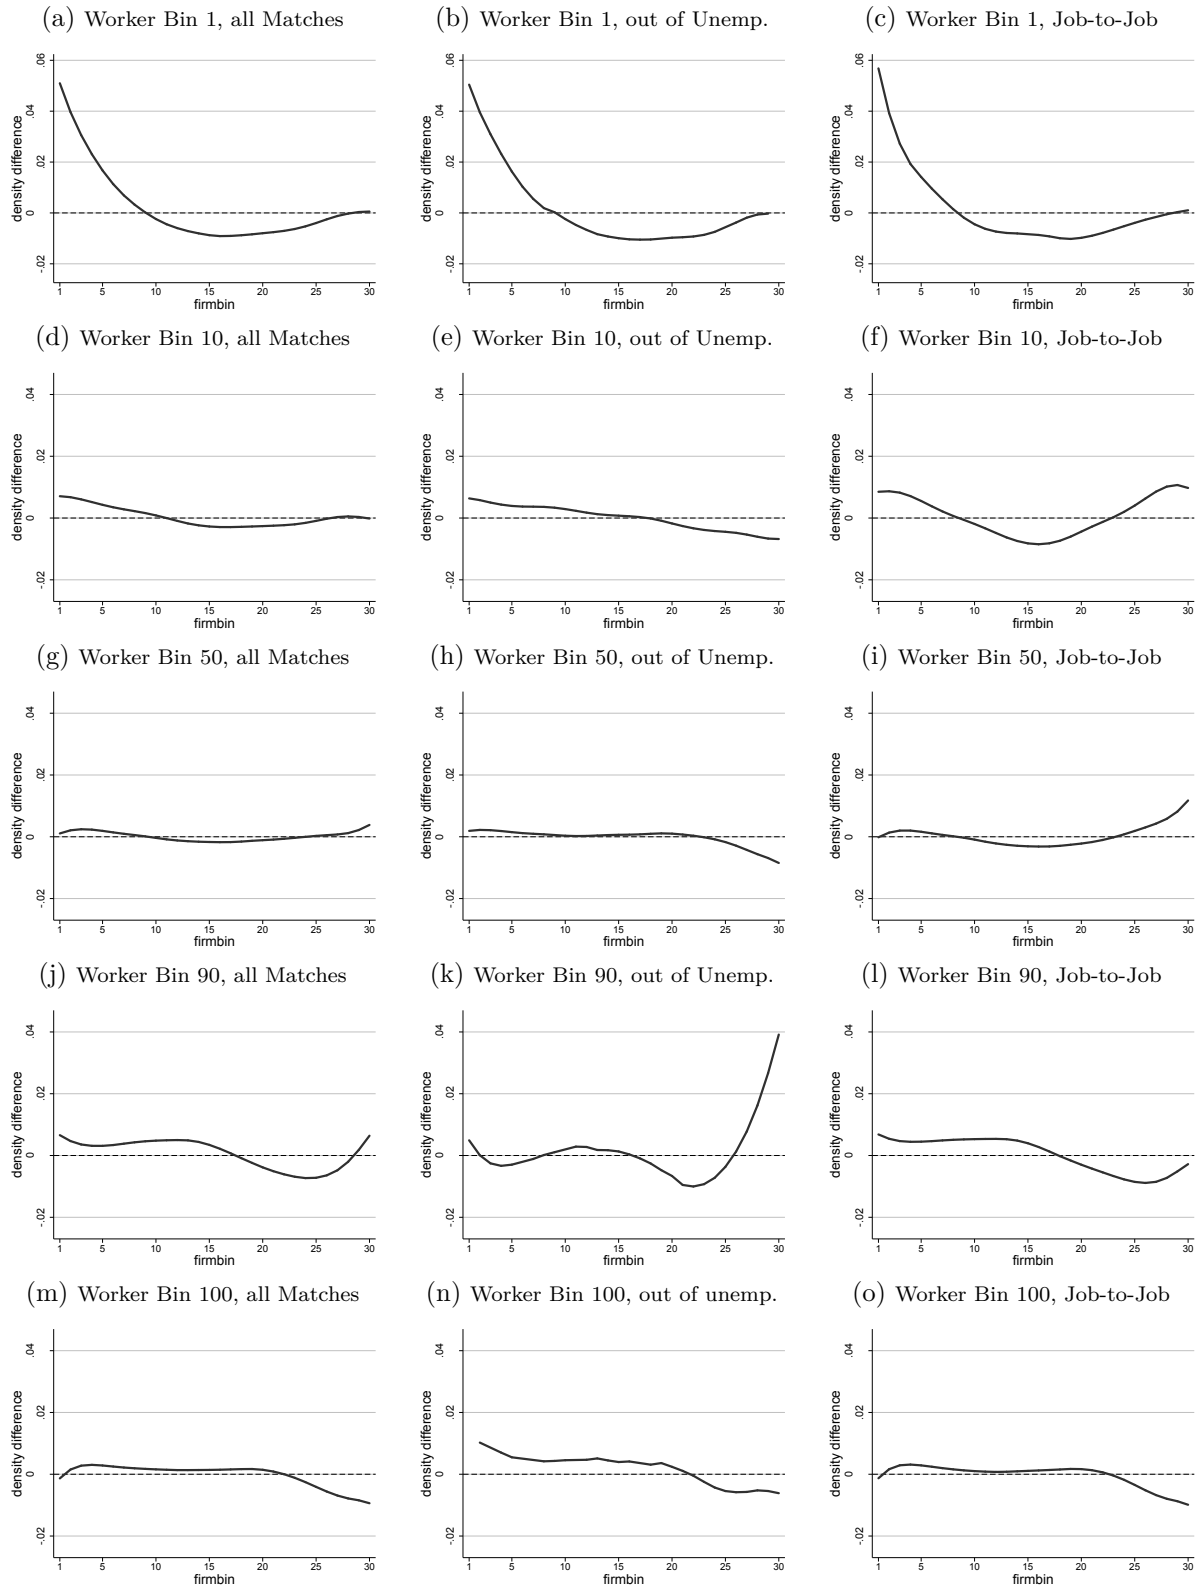
Figure B.3: Histograms of Raw Data and Estimated Density Functions



Note: Histograms of raw data and estimated univariate kernel densities of matches conditional on worker bins, time, and match type. Kernel: Epanechnikov. The bandwidth is calculated by Silverman's rule of thumb. Pointwise confidence intervals are calculated using a quantile of the standard normal distribution.

B.4 Density Differences

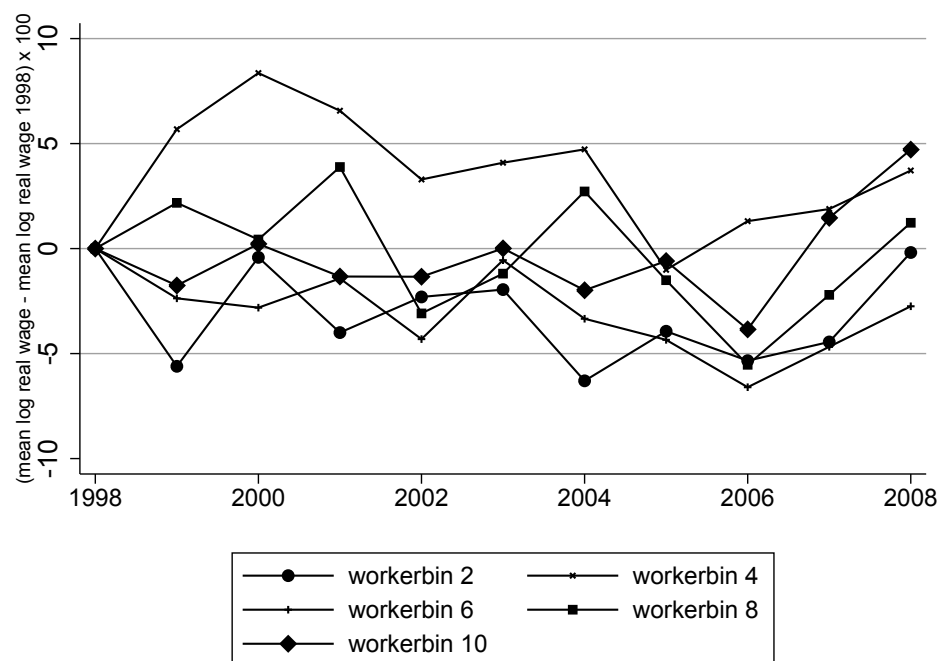
Figure B.4: Density Differences: 2003-2008 – 1998-2002



Note: Density differences are calculated by subtracting the raw match frequencies by bin. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

B.5 Wage Dynamics

Figure B.5: Wage Dynamics in Worker Bins 2, 4, 6, 8, and 10



Note: The Figure shows deviations of mean log real daily wages for single worker bins relative to 1998 (multiplied by 100).

C Robustness

C.1 Alternative Firm Ranking based on Profits per Worker

The construct a firm ranking based on average profits per worker, we build on Bartolucci et al. (2015) who use very detailed firm data (balance sheets) to study labor market sorting in the Italian region of Veneto. Using the IAB Establishment Panel, we can compute economic profits per worker of firm k in period t , π_{kt} , by simply subtracting the reported costs from firms' revenues: $\pi_{kt} = \frac{\Pi_{kt}}{N_{kt}} = R_{kt} - C_{kt} - W_{kt} - K_{kt}$. Π_{kt} denotes aggregate profits, N_{kt} the size of the workforce, R_{kt} revenues, C_{kt} input costs, and W_{kt} the wage bill. K_{kt} represents the capital cost of firm k in year t , which we compute by multiplying the capital stock with a yearly interest rate of 7.1%.⁷³ We do not observe the capital stock directly and approximate it for each establishment-year observation using a perpetual inventory method as outlined in section 2.3. We then rank firms based on the average of profits per worker over time, $\bar{\pi}_k$. We drop the most extreme outliers and group firms into 30 bins of equal size, denoting the estimated rank of all firms within one group $\hat{y}(k)$. Table C.1 shows correlations of the profit based firm ranking with other firm-level statistics: log value added, log value added per worker, the log of the firms' workforce size (all in firm-level means), as well as the estimated firm-fixed effect. Table C.2 decomposes the variance of some key firm variables in our data into the shares explained between and within the firm bins based on the profit ranking.

Table C.1: Properties of Firm Ranking based on Profits

	$\ln \bar{v}_k$	$\ln \frac{\bar{v}_k}{\bar{N}_k}$	$\hat{\phi}_k$	$\ln \bar{N}_k$
Correlation with $\hat{y}(k)$	0.52	0.70	0.75	0.15

Note: The table shows correlations of our alternative firm ranking based on profits per worker with other statistics that could be used to rank firms: log mean value added (\bar{v}_k), log mean value added per worker ($\ln \bar{v}_k / \bar{N}_k$), and estimated firm-fixed effects (extracted from running regression 5, $\hat{\phi}_k$), the log of the average size of a firm's workforce, \bar{N}_k .

Table C.2: Properties of Firm Bins based on Profits

	$\bar{\pi}_k$	$\ln \bar{v}_k / \bar{N}_k$	\bar{N}_k	Sector
Overall Variance	4.78E+09	0.873	1.92E+07	65.64
Between bins	4.06E+09 (85%)	0.477 (55%)	2.02E+06 (11%)	5.99 (9%)
Within bins	7.27E+08 (15%)	0.396 (45%)	1.72E+07 (89%)	59.65 (91%)

Note: The table decomposes the overall variance of average profits per worker ($\bar{\pi}_k$), log value added per worker ($\ln \bar{v}_k / \bar{N}_k$), the size of the firms' workforce (\bar{N}_k), and of the sectors the firms operate in into the respective shares explained within and between the firm bins based on the alternative profit ranking. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which is compatible to the common international classifications of industries, NACE and ISIC. We use 32 industries, roughly classified as follows: 1-2 = "Agriculture & Mining", 3-18 = "Manufacturing", 19-20 = "Construction", 21-23 = "Retail Trade", 24-32 = "Service Sector".

⁷³This number is taken from Evers et al. (2015).

C.2 Rank Correlations using the Profit Ranking


Table C.3: Rank correlations for different time intervals by type of match

	Spearman's ρ (Number of Observations)		
	all matches	out of unemp.	job-to-job
1998-2008	0.254 (167,764)	0.253 (69,398)	0.228 (98,366)
1998-2002	0.251 (77,892)	0.228 (30,236)	0.242 (47,656)
2003-2008	0.255 (89,872)	0.270 (39,162)	0.215 (50,710)
1998-1999	0.227 (26,743)	0.196 (10,386)	0.215 (16,357)
2000-2001	0.272 (37,949)	0.238 (14,303)	0.271 (23,646)
2002-2003	0.234 (26,694)	0.241 (11,174)	0.211 (15,520)
2004-2005	0.251 (31,003)	0.251 (13,675)	0.228 (17,328)
2006-2008	0.264 (45,375)	0.293 (19,860)	0.209 (25,515)
1998	0.195 (13,070)	0.177 (5,411)	0.200 (7,659)
1999	0.254 (13,673)	0.214 (4,975)	0.225 (8,698)
2000	0.250 (20,171)	0.225 (7,906)	0.249 (12,265)
2001	0.295 (17,778)	0.253 (6,397)	0.294 (11,381)
2002	0.236 (13,200)	0.259 (5,547)	0.203 (7,653)
2003	0.232 (13,494)	0.222 (5,627)	0.220 (7,867)
2004	0.228 (15,974)	0.248 (6,983)	0.198 (8,991)
2005	0.275 (15,029)	0.253 (6,692)	0.259 (8,337)
2006	0.269 (16,549)	0.281 (7,467)	0.226 (9,082)
2007	0.251 (15,935)	0.299 (6,916)	0.185 (9,019)
2008	0.275 (12,891)	0.302 (5,477)	0.218 (7,414)

Notes: Rank correlations using alternative firm ranking based on average profits per worker as detailed in Section C.1. We test the null hypothesis that worker and firm bins are statistically independent. All rank correlation coefficients are different from 0 at 1% level of significance. Rounded to 3 decimal places. Numbers of observations (new matches according to our definition) are reported in brackets.

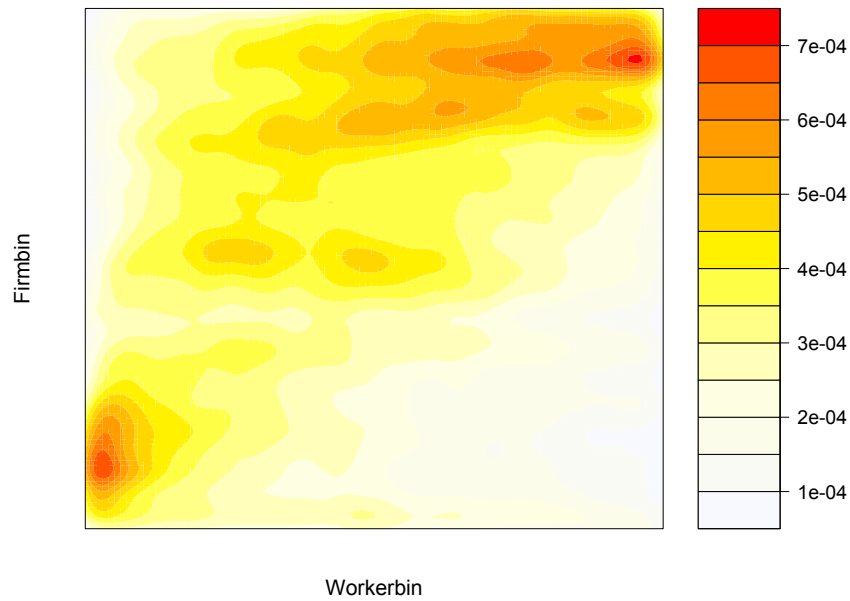
C.3 Main Plots using the Profit Ranking

Figure C.1: Empirical Bivariate Density of Matches in Germany (1998-2008)



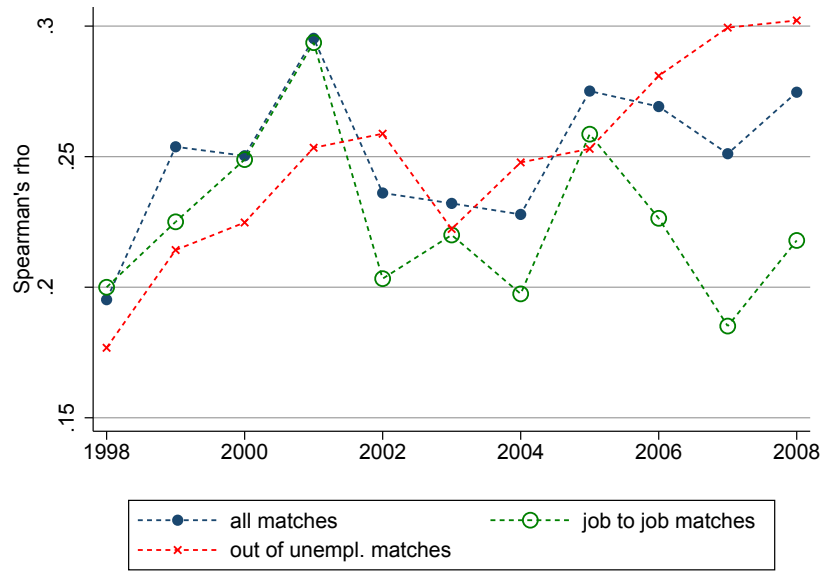
Note: Two-dimensional kernel density estimation with an axis-aligned bivariate normal kernel, evaluated on a grid with dimensions 100×30 ($\#$ worker types \times $\#$ firm types). Alternative firm ranking based on average profits per worker as detailed in Section C.1.

Figure C.2: Contour Plot of Empirical Bivariate Match Density




Note: Contour plot of Figure C.1. Alternative firm ranking based on average profits per worker as detailed in Section C.1.

Figure C.3: Rank correlations by match type over time (1998-2008)



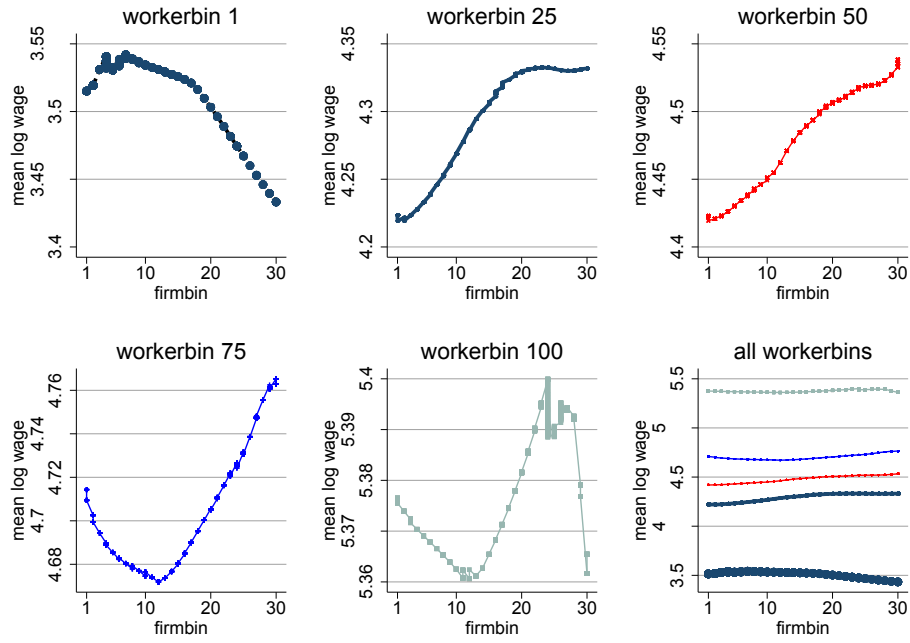
Note: Alternative firm ranking based on average profits per worker as detailed in Section C.1.

Figure C.4: Estimated Mean Wages



Note: Alternative firm ranking based on average profits per worker as detailed in Section C.1. The Figure shows the mean of the log real daily wage for all combinations of worker and firm types on a grid with dimensions 100×30 ($\#$ worker types \times $\#$ firm types).

Figure C.5: Estimated Mean Wages across Worker Types



Note: Wages are means of real log cell-wages for the given workerbins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function. Alternative firm ranking based on average profits per worker as detailed in Section C.1.