Labor Market Sorting in Germany*

Benjamin Lochner[1] and Bastian Schulz[2]

[1] *University of Erlangen-Nuremberg and IAB*

[2] *Ifo Institute at the University of Munich*

February 15, 2016

Abstract

This paper analyzes how workers are allocated to jobs in Germany. Our main contribution is to show how labor market sorting has evolved over time across worker types and how this developments are related to wages and wage inequality. We show direct empirical evidence that wages are not necessarily increasing in the productivity of the firm a worker is matched with. This feature of the data is predicted by theoretical models of labor market sorting (Shimer and Smith, 2000; Eeckhout and Kircher, 2011), but is at odds with two-way fixed-effect models (Abowd et al., 1999), the most common empirical tool to analyze wage dispersion. We use a structural search model and the identification procedure proposed by Hagedorn et al. (2014) to empirically identify worker and firm rankings as well as the bivariate density of matches in Germany. We compute rank correlations and find that low-type workers have become increasingly sorted into low-type firms in our period of observation (1998-2008), especially out of unemployment. Sorting of high-type workers into high-type firms has, if anything, slightly decreased during our period of observation.

Keywords: Sorting, Wage Dispersion, Unobserved Heterogeneity, Optimal Rank Aggregation, Firm Performance, Job Mobility, Matched Employer-Employee Data

JEL Classifications: J31, J62, J64, L25

*We thank Michele Battisti, Bernd Fitzenberger, Timo Hener, Christian Holzner, Patrick Kline, Johannes Koenen, Christian Merkl, Helmut Rainer, Johannes Schmieder, and Heiko Stüber as well as numerous seminar participants in Munich, Nuremberg, and Regensburg for comments, suggestions, and inspiring discussions. 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). The usual disclaimer applies. Contact: benjamin.lochner@fau.de; schulz.b@ifo.de

1 Introduction

Wage inequality has been increasing in recent decades, making it a topic of high interest for both policymakers and academics. Germany is no exception: Dustmann et al. (2009) find 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 the now widely available matched employer-employee data, Abowd et al. (1999) (henceforth AKM) were the first to demonstrate how to quantify the respective contributions of both observed and unobserved worker and firm heterogeneity to overall wage dispersion in a Mincerian setting. This was a big step forward in the study of wage inequality because unobserved heterogeneity typically explains most of the observed wage dispersion. Card et al. (2013) (henceforth CHK) apply the AKM methodology to the universe of German social security records. Thanks to the richness of these data, CHK are able to decompose the increase of wage dispersion in Germany into the components due to of rising worker heterogeneity, rising dispersion in the wage premiums paid by different employers, and due to an increase in sorting based on observables, that is, the sorting of different education and occupation groups to certain firms.[2]

Our starting point is the fairly recent discovery that the AKM two-way fixed-effects methodology is potentially problematic. As pointed out by Eeckhout and Kircher (2011) and Lopes de Melo (2013), the assumption that log wages are additively separable into worker and firm fixed effects leaves, by construction, no role for match-specific effects and excludes sorting based on *unobservables*. The AKM model is therefore at odds with the theoretical literature on labor market sorting. Certain features of wage data, in turn, are hard to explain without a match-specific component of wage dispersion, for example the fact that firms pay different wage premiums across their employees. In sorting models, this observation is readily explained by a production complementarity between worker and firm types (determining match-specific output) and search frictions (rationalizing the existence of suboptimal matches). CHK make the point that additive separability and, hence, the abstraction from match-specific effects is a defendable assumption because deviations in terms of wage residuals appear to be small for most (but not all) combinations of firm and worker types. In this paper, we focus on sorting based on unobservables (which we call simply sorting) and abandon the assumption of additive separability. We

[1]Remarkably, the yearly increase of wage dispersion in Germany had an order of magnitude comparable to the United States and this trend can be traced back well into the 1980s or even 1970s. Dustmann et al. (2009) find that the gap between the 85th and the 50th percentile of the German wage distribution has increased by about 0.6 log points per year between 1975 and 2004, comparable to Autor et al. (2006), who report that the gap between the 90th and the 50th percentile in the United States has increased at a rate of roughly 1 log point per year during the same period.

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

analyze German matched employer-employee data[3] through the lens of a sorting model featuring heterogeneity on both sides of the market and search frictions. Our main analytical tool is the structural search model proposed by Hagedorn et al. (2014) (henceforth HLM). This model allows us to estimate economy-wide rankings of both workers and firms and to identify the empirical bivariate density of matches between workers and firms without imposing additive separability of worker and firm heterogeneity.

Our main contribution is to quantitatively investigate how labor market sorting has evolved over time and across worker types. First, we confirm that matching patterns are positive assortative in the German labor market. The overall rank correlation coefficient is 0.25. Second, sorting increased significantly over time. It rose particularly for matches formed by workers out of unemployment. Third, we find that this trend is primarily driven by low-skill workers who are increasingly sorted into low-productivity firms. Interestingly, these workers do not necessarily incur wage losses from being matched with lower firm types. We are the first to provide direct empirical evidence that wages are not monotonically increasing in the type of firm with which a worker is matched. In fact, low-type workers receive higher wages while working at low-type firms and have to accept lower wages if hired by more productive firms. More sorting thus reduced the wage inequality among low-type workers. This non-monotonicity is a key feature of theoretical sorting models like Shimer and Smith (2000) and Eeckhout and Kircher (2011), but it is at odds with the AKM two-way fixed-effects approach, because in the reduced-form model wages mechanically increase in the firm effect. Sorting of high-type workers into high-type firms has, if anything, slightly decreased during our period of observation. Fourth, we find that patterns of sorting based on unobservables are most pronounced for worker types that have the largest deviations in terms of residuals from additive separability in CHK.[4] These are particularly the very low and very high types of workers. For large parts of the data in between, however, we confirm that match-specific effects do not appear to be very important, consistent with CHK. Wages are falling in the firm type only for the lowest worker types. This leads us to conclude that sorting patterns are highly asymmetric across different worker and firm types. For a large part of the type space, the two-way fixed-effects model appears to be a valid approximation. Nevertheless, it is important to emphasize that the reduced-form model performs worst for low-type workers. The outcomes for this group of employees are the ones with which labor market policy is typically most concerned.

Some recent structural work has explored possibilities to quantify the efficiency loss incurred by a non-optimal allocation of workers to jobs in frictional labor markets, see among others Gautier and Teulings (2015), Lise et al. (2015), and Lamadon et al. (2013).

[3]German 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).

[4]See Figure VI on p. 996 and the discussion on pp. 989-991 in Card et al. (2013).

These papers develop complex structural search models in the spirit of Shimer and Smith (2000) and estimate them. Bagger and Lentz (2015) use Danish data and estimate a structural model in which search intensity varies across workers, thereby generating sorting of more intensely searching high-type workers. Bonhomme et al. (2015) propose an empirical model that allows for unrestricted interaction effects between firm and worker unobservables. Our analysis is inspired by Hagedorn et al. (2014), which also belongs to this group of papers. HLM show how match-specific output and the sign and strength of sorting are non-parametrically identified with the standard matched employer-employee data that is available for many countries. The key challenge for the empirical study of structural sorting models is to identify credible global rankings of workers and firms. While studies in the AKM tradition implicitly rank workers and firms by their fixed effects,[5] HLM solve a Kemeny-Young rank aggregation problem to merge intra-firm wage rankings into a global ranking of workers. This procedure makes use of the fact that matched employer-employee data allows comparing wages across firms: co-workers at one firm move from firm to firm over time and form a network of comparable wage observations. A computational algorithm running on this network can effectively maximize the likelihood of the correct global ranking, as proven by Kenyon-Mathieu and Schudy (2007). Furthermore, HLM show that once workers are globally ranked, it is possible to rank firms based on the respective value of a job vacancy within them. This quantity is identified from the difference of observed wages and estimated reservation wages of similarly ranked workers under mild technical assumptions about wage bargaining. Having ranked both workers and firms, HLM then solve the wage equation for the output of every possible match to identify the production function. In the Shimer and Smith (2000) class of models, the wage equation can be simply inverted. Using their methodology on German data, HLM report a rank correlation of 0.71, which, compared to other empirical studies of sorting, appears to be very high. While the big advantage of the elegant identification procedure in HLM is that it does not require much additional information on the firm side, the question arises whether this high correlation might be driven to some extent by the fact that both workers and firms are effectively ranked based on the same information, namely, wages.

Our paper's second contribution is to test the HLM identification procedure with alternative firm rankings. We use detailed firm information from the IAB Establishment Panel, which can be linked to our primary data set, to construct measures of firm performance. Ranking firms based on a rather crude measure—profits per worker—works surprisingly well. This idea is supported by Bartolucci et al. (2015), who use very detailed firm-level balance-sheet data to study sorting in the Italian context. They argue that profits are an appropriate way of ranking firms because all firms have the same objective of maximizing them. Furthermore, firms are typically matched to a large number of

[5] Interpreting their correlation as a measure for the sign and strength of sorting.

workers, so match-specific noise tends to be integrated out of firm-level profits. We check the validity of our simple but transparent profit measure by running additional panel data regressions with fixed effects and a variety of controls on our firm data. Ranking firms by their estimated fixed effect does not significantly change our results and, thus, we use the profit-based firm ranking as our baseline in this paper. A paper closely related to ours is Kantenga and Law (2014). They use the full HLM methodology to make a contribution complementary to ours. Using the model's wage equation, they structurally decompose increasing overall wage variance into the contributions of worker heterogeneity, firm heterogeneity, and sorting. They reject the assumption of additive separability statistically by following pairs of workers that jointly change employers. They find that changes in production complementarities can explain a large part of the increase of wage dispersion in Germany, what is in line with our disaggregated approach that studies the changing allocation of different worker types across firms directly.

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. Sections 5 presents our results: correlations of our estimated worker and firm rankings, the dynamics of sorting, and how the distribution of different worker types across firms has changed over time and how this relates to wages and wage inequality. 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"[6] and restrict our analysis to the years 1998-2008 for computational reasons discussed below. 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. This section provides an overview of our data preparation procedures; are more detailed explication 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 various plausibility checks by the social security institutions followed by sanctions for non- and misreporting. In our raw data, we have nominal gross daily wages, which we deflate using the consumer price index from German national accounts

[6]File: `LIAB_MM_9308`. See Heining et al. (2012) for a detailed description of the data set and Appendix A for more details on sample selection and descriptive statistics.

with base year 2005. The education variable in German social security data suffers from missing values and inconsistencies, essentially because misreporting an employee's highest degree has no negative consequences. Therefore, we impute missing and inconsistent observations using the methodology proposed in Fitzenberger et al. (2006). We cannot impute missing values for about 2% of the data so we drop the remaining spells for which education information is missing. A limitation of the wage data is that the German social security system tracks earning only up to a certain threshold, the contribution assessment ceiling ("Beitragsbemessungsgrenze"). 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.[7]

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. It would be very interesting to analyze the labor market sorting of the excluded subgroups, but we put a higher priority on using a dataset that, on the one hand, is homogenous and, on the other, comparable to data used in previous studies of wage dispersion in Germany. To minimize working-time effects in our wage data, we also exclude part-time and marginal employment.[8] A job spell can last from one day to one year, due to the reporting rules of the German social security system, which 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. On average, each worker in our sample has 6.4 job spells between 1998 and 2008 (with a standard deviation of 5.0) and a job spell lasts 289 days (with a standard deviation of 114). We drop workers with more than 150 job spells (less than 0.01%). We sometimes observe identical spells except for the ending date and the corresponding wage, which could be due to multiple reports by an employer when there has been a change in contract ending dates. Longer spells are typically associated with a higher wage, possibly due to Christmas bonuses or other one-off salary supplements. We always keep the spell with the highest wage in such cases.

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 whom they belong. We use the terms "firm", "establishment", and "employer" interchangeably throughout the paper. The standard deviation of log wages in the final sample is 0.455.

[7]The average censoring rate is 13.6% across years. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold.

[8]To 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").

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.

This value is in line with measures of wage dispersion reported by CHK for our period of observation, indicating that our data preparation procedure generated a comparable sample.[9]

2.2 Residual Wages

Since our analysis focuses on *unobservable* determinants of wage dispersion, we use residual wages as the basis for our quantitative exercise. Residual wages are wages net of the effects of observable characteristics. To compute them, we use an empirical model of wage dispersion close to CHK and regress log wages on person fixed effects, an unrestricted set of year dummies, and quadratic and cubic terms in age fully interacted with educational attainment:

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

$\ln w_{it}$ denotes the log real daily wage of worker i in year t, x'_{it} includes time-varying observable characteristics, α_i is a worker fixed effect, and r_{it} is the error term. As expected, the explanatory power of this AKM regression model is very high. The adjusted R^2 from estimating Equation (1) is 81%, slightly below the around 90% reported by CHK. This difference is easily explained by the lack of establishment fixed effects in our specification, which is, however, inconsequential for our purposes.[10] 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 effects alone explain 74% of log wage variance; the residual absorbs another 17%. To compute the residual

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

[10] In contrast to CHK, who have access to the universe of German social security records, we cannot reliably estimate establishment fixed effects in our sample because we do not observe the whole workforce of the establishments in our data. This is not a problem for our analysis, however. The ranking procedure based on residual wages, following Hagedorn et al. (2014), relies on pairwise comparisons of wage observations of workers who are employed at the same firm (see Section 4.1). The unobserved firm effect therefore affects the wage of both workers by exactly the same amount. The bilateral ranking of a pair of workers is not affected by the firm effect.

wages, we simply subtract the share of the wage predicted by observable characteristics from the observed wage for every individual:

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

The log residual wage has a standard deviation of 0.433 (variance 0.187) and is only slightly less dispersed than observed wages in our sample due to the small role that observable characteristics play for wage dispersion. The correlation between observed log wages and log residual wages is very high (0.98). To address any concern about the unaccounted for influence of occupation on wages, we augment Equation (1) by controlling for 32 categories of occupations.[11] We interact occupations with education and year dummies. Reassuringly, including occupations does not add much explanatory power to the wage regression. The adjusted R^2 stays virtually the same (81%) and the portion of variance explained by observable characteristics increases only slightly from 3.7% in the baseline CHK regression to 6.7% when including occupation controls. The correlation between baseline residual wages and residual wages net of occupational effects is very high, above 0.99. Thus, controlling for occupation does not meaningfully alter our worker ranking based on residual wages. We are confident that residual wages, $\tilde{w}_{it}$, are an appropriate input for our analysis.[12]

2.3 Firm Data

We use the IAB Establishment Panel to calculate measures of firm performance which will be employed and described in more detail in Section 4.2. The IAB Establishment Panel draws a stratified random sample of establishments from the register data.[13] We restrict our sample to establishments that employ at least 10 workers on average during our period of interest. The data contains no direct information about the capital stock on which an establishment operates. To circumvent this shortcoming, we use a modified perpetual inventory approach as 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 of firms' net investments in the different categories of capital goods. We exclude non-competitive sectors of the economy

[11] In accordance with Bundesagentur für Arbeit (1988) the occupational classification available to us consists of a 3-digit code comprising about 330 values. We use 32 2-digit values ("Berufsabschnitte") in order to being able to estimate a large set of interaction terms.

[12] CHK provide a comprehensive analysis of the explanatory content of additional controls in their setting. In line with our finding, they report that adding occupational (and industry) dummies does not significantly improve upon the explanatory power of AKM-type models of wage dispersion.

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

and firms that do not report revenues as their primary measure of output.[14] Finally, we merge our firm sample with the data on employment spells. This leaves us with 4,754 establishments for which we observe 3,712,194 employment spells of 628,462 workers.

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 restrict our attention to the first employment spells of a new match. 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 one and exclude job recalls. Our final sample consists of 407,484 matches, 168,592 from out of unemployment allocations and 238,892 from job to job switches.

3 Sorting Theory

This paper analyzes labor market sorting in Germany over time. To support our analysis, this section briefly summarizes the theory of labor market sorting that guides our thinking about the allocation of workers to jobs and the determination of wages.

Conceptually, sorting requires explicit heterogeneity of agents in a two-sided market. Since this is a paper about the German labor market, we call our heterogeneous types *workers* and *firms*. Workers and firms carry identifiers i and j, respectively. Assume for now that an unambiguous ranking of worker and firm types exists and is observable. Let the rank of worker i be denoted $x(i)$ while the rank of firm j is $y(j)$. For simplicity, heterogeneity across workers and firms is assumed to be one-dimensional.[15] Let the one-dimensionally-varying characteristic be called *productivity* for firms and *skill* for workers. The starting point for the theory of labor market sorting is the neoclassical optimal assignment model proposed by Becker (1973) in the context of the marriage market. In the absence of search frictions, every worker finds his optimal firm assignment instantaneously and a match is formed. This leads to a Walrasian first-best allocation of workers to

[14] Broadly speaking, this applies to health care, education, non-profit organizations, and the public sector (non-competitive sectors). The most important groups of establishments that does not report revenues are banks.

[15] For two recent studies of multidimensional sorting see Lindenlaub (2014) and Lise and Postel-Vinay (2015).

firms in the labor market equilibrium. The specific nature of the assignment depends on the production structure of the economy. One possible assignment pattern is positive assortative matching (PAM). In theory, a complementarity between worker and firm type, modeled as a feature of the production function, gives rise to PAM.[16] To maximize output, this complementarity requires that the most productive firm employs the most skilled worker, the second most productive firm employs the second most skilled worker, and so forth. The worker with the lowest level of skill is optimally matched with the least productive firm. This allocation implies a perfect positive correlation of worker and firm ranks: $\rho = \text{corr}(x(i), y(j)) = 1$. ρ is Spearman's rank correlation coefficient, a natural measure for sorting that we use in the following. A negative rank correlation, $\rho < 0$, implies negative assortative matching (NAM), an allocation pattern in which highly productive firms would tend to employ low-skilled workers to maximize output. In turn, $\rho = 0$ indicates the complete absence of sorting patterns in allocating workers to jobs.

A major step forward in theory of labor market sorting was incorporating search frictions in the Beckerian model. Shimer and Smith (2000) show the existence and properties of equilibrium in a two-sided assignment model with frictions. Search frictions imply that the first-best allocation cannot be realized with a continuum of types because it is costly to wait for the optimal partner. Therefore, a subset of firm and worker types in the vicinity of the optimal allocation becomes acceptable. These subsets, the matching sets, are determined by the option value of a match.[17] Figure 1 depicts two equilibrium features of a simple Shimer and Smith (2000) economy. This illustrative economy is populated by a stationary population of 1,000 discrete, heterogeneous, and rankable worker and firm types, each having an equal mass normalized to 1. The production function exhibits a complementarity and induces positive assortative matching. Panel (a) of Figure 1 shows the space of all worker and firm-type combinations. The 45-degree line represents the first best allocation, unreachable in the presence of search frictions. The blue boundaries are the cutoffs of the matching sets for all worker and firm types. The black line is the matching set of a type 500 worker. We see that this worker type is willing, more or less, to match with all firm types above type 180 and below type 820. This is a random search model, so unmatched agents can meet types with which they are not willing to form a match. The acceptability of potential matches is always mutual. Within the matching sets, the surplus of matches is positive. Thus, according to the model, matches are only formed within these boundaries. The predicted match density is zero in the upper-left and the lower-right corners of the type space, indicating the absence of matches between high-type workers and a low-type firms and vice versa. We show later that this simple

[16]Technically speaking, the production function takes worker and firm types as arguments and is supermodular, that is, the cross-derivatives are strictly positive.

[17]Matches are acceptable once their option value is higher than the value of continued search.

Figure 1: A Simple Model of Sorting

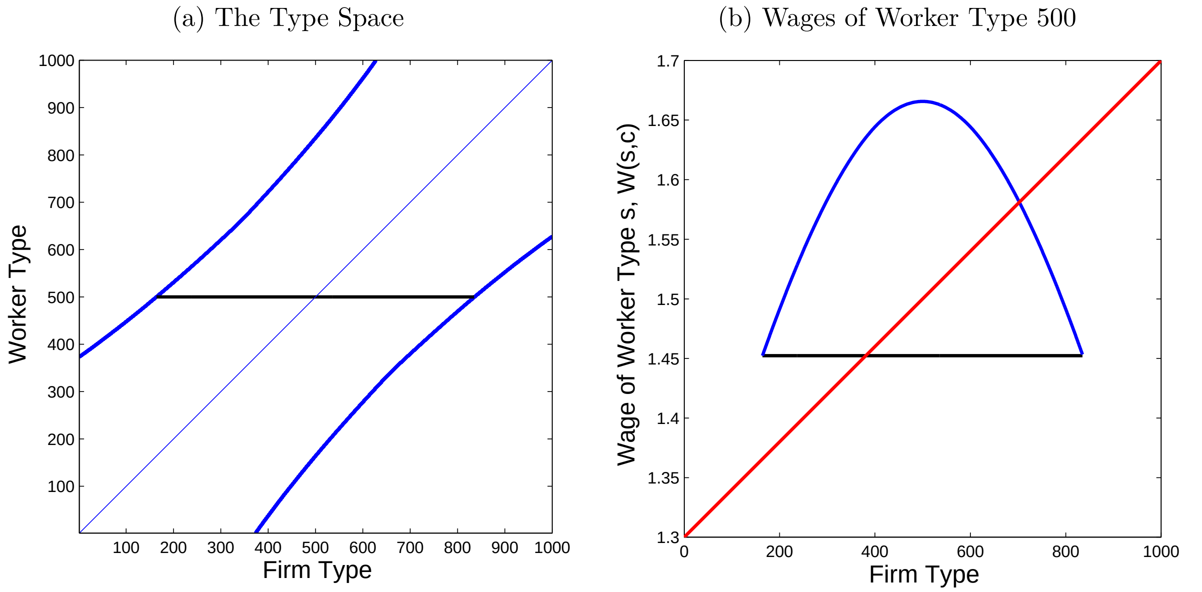

Note: The blue curves depict the equilibrium of a simple Shimer and Smith (2000) economy. The black line represents the equilibrium matching set of a worker of type 500. The model is a simplified version of the structural search model presented in Hagedorn et al. (2014). It is solved numerically using standard parameter values. For the derivation and solution procedure see 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 fixed-effects models of wage dispersion.

prediction is (partly) confirmed by the data. Panel (b) shows how the wage of a type 500 worker varies across firm types. In the model, wages are determined by splitting the surplus according to the Nash bargaining solution. Note that due to the production complementarity in this economy, the wage of this worker type is maximized in a match with a type 500 firm. A match with a firm type below 500 results in lower output since the production complementarity is not fully exploited. On the other hand, in a match with a higher firm type the worker needs to compensate this firm for not waiting to hire a better worker and accepts a lower wage. This non-monotonicity lies at the heart of the critique of the AKM two-way fixed-effects approach. In a fixed-effect model, the wage mechanically increases in the firm effect. The reduced-form model implies a monotonic association of wages and firm type, which might look like the stylized red line in Panel (b) of Figure 1. The higher the firm fixed effect, interpreted as the firm's type, the higher is the wage of every worker the firm employs, a simple wage premium. The theoretical model's predictions regarding wages are clearly at odds with the empirical model of wage dispersion in the spirit of AKM. We show below that, in fact, the theoretical prediction of non-monotonous wages across firm types is met by the data for certain (but not all) worker types.

The outlined simple theoretical model is similar to the model used by HLM to develop their constructive proof of identification. Empirically, the key challenge remains to establish plausible global rankings of both workers and firms, which are simply assumed to exist in theoretical models of sorting. The next section presents our identification strategy, which builds on HLM to a large degree.

4 Identifying Sorting

Our approach to the identification of labor market sorting in this paper combines the worker ranking procedure proposed by HLM with an independent firm ranking. The main contribution of HLM is to show how to rank workers and firms empirically to identify labor market sorting using commonly available matched employer-employee data. A distinct feature of their approach is that it works on wage data alone. We want to test whether the construction of a firm ranking independent of wage data changes the empirical rank correlations and the findings regarding the strength of sorting.[18] Using German data to evaluate how well their non-parametric identification of the structural search model works, HLM find a very high rank correlation of 0.71 for the years 1993-2007. This value is much higher than the correlation of estimated worker and firm fixed effects reported by CHK using the AKM methodology on German data. For roughly the same period, they report correlations between 0.17 and 0.25. This large difference indicates a huge role of the match-specific effects that CHK abstract from. The HLM results also lie above correlations reported by other authors using alternative identification procedures. Bagger and Lentz (2015) use Danish data and estimate a structural model in which search intensity varies across workers, thereby generating sorting of more intensely searching high-type workers. They report a correlation of only 0.11. Bonhomme et al. (2015) propose an empirical model that allows for unrestricted interaction effects between firm and worker unobservables. They use Swedish data and find a correlation of 0.46. Bartolucci et al. (2015), in turn, rely on high-quality Italian firm data (balance sheets). They develop a refined reduced-form approach to identify sorting and find a correlation of worker and firm types of 0.52. The increasing number of studies that report positive correlations suggests that patterns of positive sorting are indeed an important feature of labor market allocations.[19]

The correlations we report in the following using the HLM worker ranking and our alternative firm ranking are lower than 0.71. With rank correlations in the range 0.25-0.30, we are closer to CHK who use the AKM methodology and the universe of the German employment registers. We do not conclude from this, however, that match-specific effects play no role. Quite the contrary, we observe that the tendency of workers to sort into jobs is most pronounced for those worker types that show the largest deviations from the reduced-form CHK-AKM model, namely, the highest and lowest types of workers.[20] In

[18]HLM rank firms based on the firm-specific value of vacant jobs which can be identified in the context of their model from the difference of wages paid to different worker types and the estimated reservation wages of these worker types.

[19]It remains unclear, however, to what extent the large differences between the reported positive correlations are driven by alternative identification strategies, cross-country differences in sorting, or by cyclical fluctuations of the degree of labor market sorting. Looking more closely into these differences is an exciting topic for future research.

[20]In fixed-effects models, the estimated fixed effects of both workers and firms are implicitly interpreted as their cardinally rankable type. CHK show residuals for type combinations in Figure VI on p. 996.

Section 5 we analyze the nature of sorting patterns in Germany for different worker types and over time. Before, we briefly describe our ranking procedures.

4.1 Ranking of Workers

To construct a global ranking of workers, we rely on the procedure proposed by HLM. They suggest a computational algorithm to merge intra-firm wage rankings into a global ranking of workers by solving a Kemeny-Young rank aggregation problem.[21] The procedure makes use of the fact that matched employer-employee data allows wage comparisons across firms: 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 discussed 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.[22] Using a simple Bayesian approach and 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:

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

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 overall 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, thus, the more plausible is the ranking implied by the wage observations at this firm, resulting in a higher value of $c(i,j)$. To justify using the overall wage variance, the assumption that all variation in wages for a specific job stems from measurement error only is necessary.[23] The probability $c(i,j)$ is defined for pairs of workers employed at the same firm only. In case a pair of workers is

[21]Rank 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.

[22]Hagedorn et al. (2014) prove that, in the context of their model, the reservation wage, the maximum accepted wage, and the adjusted average wage of a worker are monotonically increasing in the unobserved type. Importantly, average wages, sometimes used to rank workers according to their wages in empirical applications, are not monotonically increasing in the type because they do not factor in the values of workers' interjacent unemployment spells.

[23]For details of the derivation, see Appendix III.1 in Hagedorn et al. (2014).

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	0.47

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.

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:

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

$\Pi(i,j)$ ($\Pi(j,i)$) 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.[24] Importantly, we do not have to observe all workers of a given establishment because the pairwise comparison of residual wages of two workers at the same firm is not affected by a potential wage premium. Both workers receive it.[25] We particularly 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.[26] It is worth noting, however, that the formation of coworker pairs increases the size of our data set enormously, beyond the size we can handle with the computational tools available to us (>2.15 billion observations). We have to drop the first five years of our data and therefore restrict our analysis to the years 1998-2008. We arrive at a final ranking which gives an estimate of the unobserved type, $\hat{x}(i)$, of every individual worker in our data. We group workers into 100 bins of equal size for convenience. 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 final worker ranking $\hat{x}(i)$, we compare it to other wage statistics commonly used to rank

[24]This is due to the fact that residual wages are deflated and net of time effects.

[25]This 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.

[26]See Appendix A for sampling details.

Table 3: Properties of Worker Bins

	$\ln w_{it}$	$\ln \tilde{w}_{it}$	age	education
Overall Variance	0.175	0.157	96.63	1.85
Between bins	0.114 (65%)	0.103 (66%)	4.64 (5%)	0.72 (34%)
Within bins	0.061 (35%)	0.054 (34%)	91.99 (95%)	1.23 (66%)

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".

workers in the empirical literature, mean wages and person fixed effects (AKM), as well as rankings based on the observable characteristics age and education. Table 2 shows the correlations of the rankings. Notably, the rank aggregation procedure produces a ranking which is quite different from rankings based on alternative wage statistics or observable characteristics. 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. In turn, the final ranking is only weakly correlated with age and education. This is reasonable since our ranking is based on residual wages, which are net of the effect of observables.

Regarding the binning of workers, Table 3 shows a decomposition of observed variances of workers' wages, residual wages, age, and education in the shares explained within our bins and between them. In line with the fact that our ranking is based on wages, the share of variance explained between the bins is roughly two-thirds both for log wages ($\ln w_{it}$) and log residual wages ($\ln \tilde{w}_{it}$). Hence, our bins are internally homogenous in terms of wage variation. As one would expect, the bins are much less homogenous 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 in our sample (20-60) in every bin. A high-type worker is not necessarily of a high age 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 A.1. Despite the low correlation and the high within variance, one can still observe that the bin means of age and education increase monotonically in the estimated worker type.

4.2 Ranking of Firms

The baseline results in this paper make use of a very simple and transparent firm ranking based on 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. They 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. The main argument in favor of profits is that all firms share a similar objective: maximizing profits. In addition, firms tend to be matched to a large number of workers over time and match-specific noise should be largely integrated out of firm-level profits. Conversely, workers are typically matched only to a small number of employers throughout their career, creating noisy wage histories. Using the IAB Establishment Panel, we can compute economic profits per worker of firm j in period t, π_{jt}, by simply subtracting the reported costs from firms' revenues:

$$\pi_{jt} = \frac{\Pi_{jt}}{N_{jt}} = R_{jt} - C_{jt} - W_{jt} - K_{jt}. \tag{5}$$

Π_{jt} denotes aggregate profits, N_{jt} the size of the workforce, R_{jt} revenues, C_{jt} input costs, and W_{jt} the wage bill. We do not observe K_{jt}, the capital costs of firm j in year t, directly. We approximate the capital stock of each individual establishment using a perpetual inventory method as outlined in section 2.3 and assume a yearly interest rate of 7.1% to compute capital costs.[27] We rank firms based on the average of profits per worker over time, $\bar{\pi}_j$, and combine them in 30 bins of equal size, denoting the estimated rank of firms $\hat{y}(j)$.

Profits per worker are undoubtedly a simplistic measure of firm performance, despite the data and arguments presented by Bartolucci et al. (2015). To show that our results do not hinge on the profit ranking, we also estimate a rich empirical model of firm performance with a large number of controls on our firm data to ensure the robustness of our findings. A natural empirical tool to rank firms is, of course, a fixed effect. We seek to identify the impact of unobservable firm characteristics on firm performance independently of wages, as opposed to the AKM model of wage dispersion discussed before. Conceptually, identifying a firm fixed effect is closely related to stochastic frontier analysis using panel data. The estimated fixed effect is inversely related to an individual firm's distance to the production frontier as a measure of relative performance.[28] Given inputs, capital and labor, the most productive firm in the economy is located on the production frontier and has the highest fixed effect. Less productive firms generate less output relative to their inputs. They have a certain distance to the frontier due to technical inefficiency and, accordingly, have a lower fixed effect and also a lower rank.

We estimate a translog production function with firm specific effects. We use log value added as our measure of output, $\ln y_{jt}$. x'_{jt} contains time-varying explanatory variables. In the translog setting, these include capital (measured by the capital stock), labor (measured by the size of the workforce), capital and labor squared as well as their

[27]This number is backed by Evers et al. (2015), who analyze investment decisions of German firms during our period of observation. Empirically, the time variation of the interest rate relevant for firms' investment decisions is minimal and we assume a constant interest rate for simplicity.

[28]Stochastic frontier analysis (SFA) has been developed in the context of cross-sectional data, see Aigner et al. (1977). SFA has been extended to the panel data context by, among others approaches, Schmidt and Sickles (1984) and Cornwell et al. (1990).

Table 4: Properties of Firm Ranking

	$\ln \bar{y}_j$	$\ln \bar{y}_j / \bar{N}_j$	$\hat{\phi}_j$	$\bar{N}_j$
Correlation with $\hat{y}(j)$	0.52	0.70	0.75	0.15

Note: The table shows correlations of our firm ranking ($\hat{y}(j)$) with other statistics that could be used to rank firms: log mean value added ($\bar{y}_j$), log mean value added per worker ($\ln \bar{y}_j / \bar{N}_j$), and estimated firm fixed effects (extracted from running regression 6, $\hat{\phi}_j$), as well as the average size of a firm's workforce, $\bar{N}_j$.

interaction. z'_{jt} includes dummies for 32 sectors of the economy which we include as additional controls.[29] Time effects are captured by ω_t and v_{jt} is the residual.

$$\ln y_{jt} = \phi_j + \omega_t + x'_{jt}\beta + z'_{jt}\gamma + v_{jt}. \tag{6}$$

After running this regression, we rank firms based on their estimated fixed effect, $\hat{\phi}_j$. We use the fixed effects ranking as a robustness check and present all our main results using the alternative ranking in Appendix C. All our main conclusions regarding sorting in the German labor market are unaffected by the choice of the ranking, indicating that our simple profit ranking does indeed a good job in summarizing firm heterogeneity. Table 4 underlines this by showing correlations of the profit-based ranking with other possible firm rankings based on log value added, log value added per worker, the estimated fixed effect, and the size of the firms' workforce. These correlations underline that the profit ranking works well. It is highly correlated with the estimated firm fixed effect, the basis of our robustness ranking and with mean log value added per worker. The correlation with mean log value added alone is somewhat lower. This seems reasonable since the profit ranking is also computed based on profits per worker. The correlation with the average size of a firms workforce, $\bar{N}_j$, is very low, indicating that our firm ranking is very different from a simple firm ranking based on the size of the workforce. Just like in case of the worker ranking, we also decompose the variance of some key variables in our data into the shares explained between and within our firm bins to show in which dimension the bins are internally homogenous and in which they are not. Table 5 shows the decomposition. As one would expect, the bins are internally homogenous in terms of average profits per worker. The profit ranking bins are much less homogenous in terms of log value added, however, the majority of the variance of log value added is, reassuringly, between the bins. Furthermore, very much like in the case of workers, our bins are not directly related to observable firm characteristics, in this case the size of the workforce and the sector a firm operates in. The respective shares of variance within in the bins are 89% and 91% respectively, indicating that we have both badly performing large firms and

[29] We 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 5: Properties of Firm Bins

	$\bar{\pi}_j$	$\ln \bar{y}_j$	#workers	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 profits ($\bar{\pi}_j$), log value added ($\ln \bar{y}_j$), the size of the firms' workforce ($\bar{N}_j$), and of the sectors the firms operate in into the respective shares explained within and between the worker 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".

very well performing small firms in our sample. Our firm ranking has only a weak link to the size of the workforce of an establishment. The same is true for the industry structure, since we find in every bin firms from almost every sector. The plotted distributions of firm size and sectors are shown in Appendix B.2, Figure A.2.

5 Labor Market Sorting in Germany

Having ranked both workers and firms, we can now analyze the empirical bivariate distribution of matches across all possible combinations of worker and firm types in the German labor market. We use 100 worker and 30 firm bins based on the estimated types with a high degree of internal homogeneity as detailed in the preceding section. To compute the density of matches, we use the definition of a match described in Section 2.4. A match is always the first new person-year observation (employment spell) between a worker and an establishment in our data. We exclude recall. This definition is well-suited for the analysis in this paper because our primary interest is the allocation of workers to jobs. We focus on the first wage observation which arguably represents the wage offered to a worker when he decides to accept a job. The abstraction from the length of an employment spell and tenure related wage increases within one firm makes it easier to compare the allocation of workers to jobs out of unemployment from the allocation of workers who switch jobs between firms, a distinction we will make use of in the following.[30] We refer to the two different origins of a match as "match types". Note that our definition of matches leads to a rather conservative measurement of sorting. The rank correlations we compute in the following would be significantly higher if we included all employment spells we observe in the calculation and not just the first match, particularly if matches with a small degree of mismatch tend to last longer.

[30] In the light of the theory of search on the job, the comparison of the allocation of workers matched out of unemployment and job-to-job switchers is by no means perfect, e.g. because theory suggests that wages are determined very differently for the two groups. The abstraction from the spell length ensures, however, that certain match types are not over-represented (under-represented) in our sample simply because they last longer (shorter) on average, leading to more (less) person-firm-year observations.

Figure 2: Empirical Bivariate Density of Matches in Germany (1998-2008)

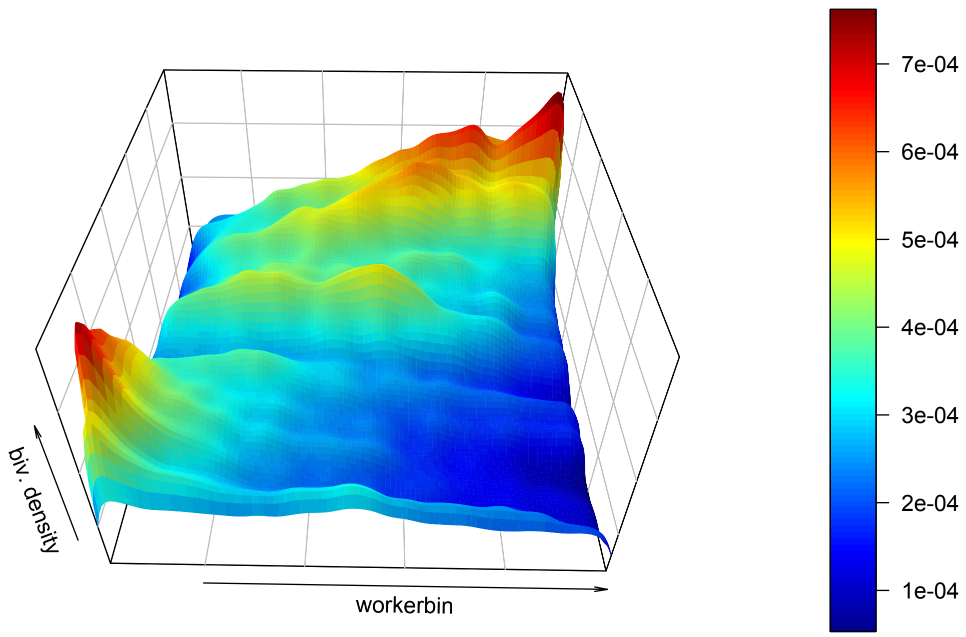

After we correlate our binned worker and firm rankings, we find that the overall rank correlation coefficient (Spearman's ρ) for all matches is significantly positive with a value of 0.250. This indicates that the matching process in the German labor market features positive sorting of workers to jobs, albeit not to a very high degree. Sorting is more pronounced for matches out of unemployment with a rank correlation of 0.256. In the sample of job-to-job switchers, the rank correlation is 0.223. The degree of sorting we find for all matches is somewhat higher than the values reported by CHK, especially in the first half of our sample. CHK report a correlation of the estimated person and establishment effects of 0.169 for the years 1996-2002. We find a rank correlation of 0.248 for 1998-2002. For the second half of our sample, the correlation found by CHK is 0.249 (2002-2009); ours is roughly similar at 0.251 (2003-2008).[31] Figure 2 plots the bivariate density of matches to illustrate in more detail the sorting patterns found in our data. Interestingly, the propensity to sort is by no means homogenous across worker and firm types. There is a distinct tendency of high-skilled workers to be matched with high-productivity firms and, at the same time, low-skilled workers tend to work at low-productivity firms, which is illustrated by the spikes in the upper-right and lower-left corners of the plot. Conversely, labor market sorting is not very pronounced in the center of the type space. For workers and firms in the middle of the ranking, the gains from being optimally matched appear to be small. Figure 3 presents a direct comparison of the theoretical model and our empirical findings. Panel (a) shows the theoretically optimal allocation along the 45-degree diagonal

[31]See 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.

Figure 3: Sorting: Theory and Empirical Evidence

(a) Sorting in Theory (b) Contour Plot of Empirical Match Density

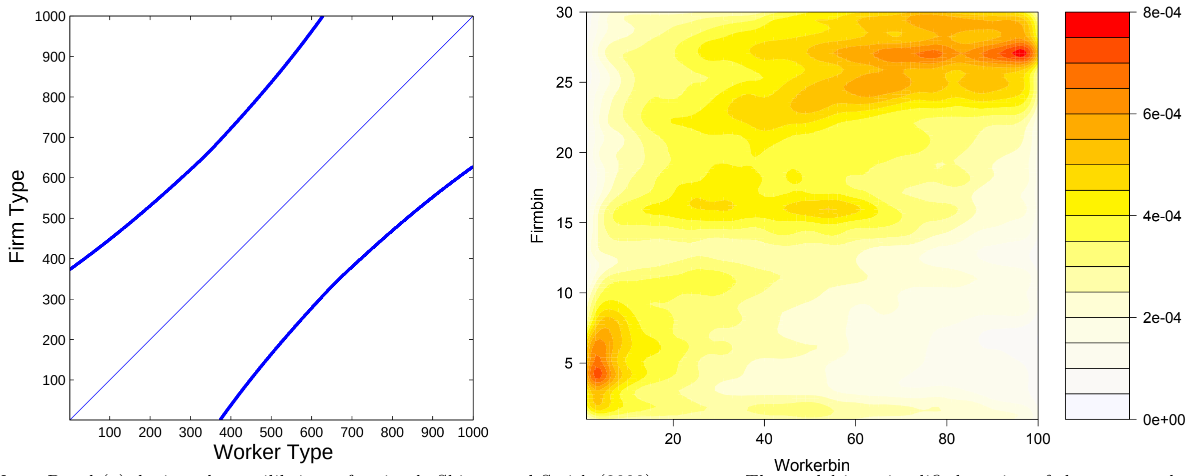

Note: Panel (a) depicts the equilibrium of a simple Shimer and Smith (2000) economy. The model is a simplified version of the structural search model presented in Hagedorn et al. (2014). It is solved numerically using standard parameter values. For the derivation and solution procedure see Schulz (2015). 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 fixed-effects models of wage dispersion.

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 in a labor market with frictions no matches are formed by type-combinations outside the matching sets, that is, there should be no matches of high-type firms and low-type workers or vice versa. Empirically, we clearly see in Panel (b) that the majority of matches is located in the vicinity of the diagonal, in line with theory. Moreover, the density of matches in the lower-right corner is very low, indicating that high-skilled workers almost never work at low-productivity firms. The prediction regarding the opposite corner, however, is not confirmed by our data. We see a moderate density of matches between low-type workers and high-type firms. The symmetry of the simple theoretical plot is thus clearly rejected by the data.[32]

5.1 Changes of Sorting over Time

The large number of observations in our data allows us to go beyond simply analyzing aggregate rank correlations and glean some insight into the dynamics of sorting patterns in our data. Are these a time-invariant, structural feature of the German labor market? Or did they change over time, possibly in conjunction with institutional reforms or cyclical dynamics of the labor market? As a first step towards this goal, we run our analysis separately by year to understand how sorting changed over time. Figure 4 plots rank correlation coefficients by year; the blue line shows the correlation computed for all matches. We further differentiate between the two types of matches: those made with

[32]In the model, symmetry hinges on the assumption of equal bargaining powers of workers and firms in the Nash bargaining solution. By relaxing this assumption, it is easily possible to align the theoretical model with the apparent asymmetry of matching patterns in the data.

Figure 4: Rank correlations by match type over time (1998-2008)

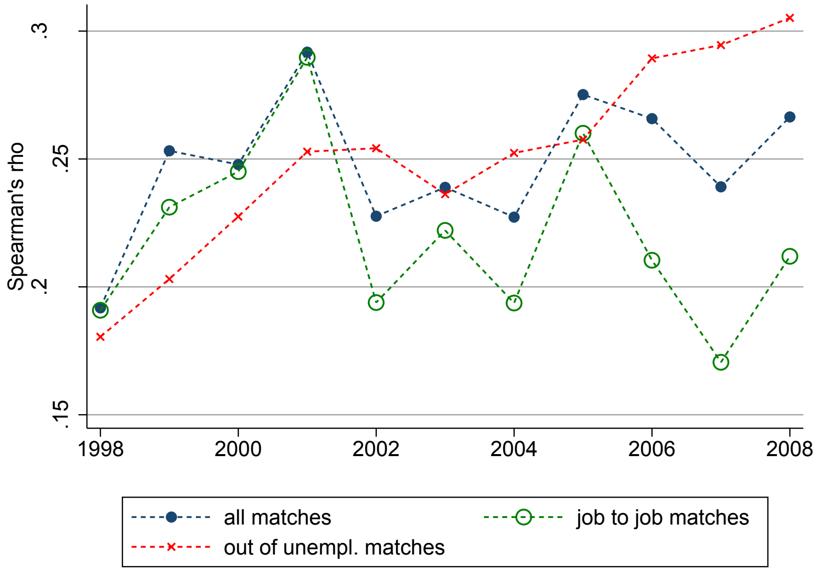

workers out of unemployment and those made with workers who switch jobs, meaning they changed their employer without an intervening unemployment spell. About 41% of the matches in our sample are matches out of unemployment, the remaining 59% are matches of job switchers. Overall, the extent of positive sorting in the German labor market increased during our period of observation. For all matches, the rank correlation rose from 0.192 in 1998 to 0.266 in 2008 (blue line). The year-to-year variation of the correlation coefficient for all matches is high, however. There was a steep increase, up to a maximum of 0.292, of the correlation between 1998 and 2001. This was followed by a severe drop almost back to the initial level. After 2002, the overall rank correlation increased again, albeit at a slower rate and again with small corrections in 2006 and 2007. Conditioning the rank correlation on the different match types sheds some light on the time variation. For matches out of unemployment, a surprisingly clear tendency is visible (red line): the rank correlation increased almost uniformly from 0.181 in 1998 to 0.310 in 2008. In turn, most of the variance 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 or, if anything, has slightly decreased.[33]

Our analysis of the changing rank correlations 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 higher for job-to-job switch-

[33]A table with all rank correlations by year and match type can be found in Table A.2 in Appendix B.

ers, but starting in 2002, the rank correlation becomes higher for new matches out of unemployment (except in 2005) due to the steady increase of the rank correlation for those matches. This is remarkable since theoretical models of labor markets with on the job search typically suggest that sorting patterns should be more pronounced for new matches resulting from search on the job. For instance, in 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. Hence, our observation of lower sorting on the job as compared to sorting out of unemployment is puzzling from a theoretical perspective.

5.2 Distributional Dynamics

We have established that labor market sorting in Germany has increased between 1998 and 2008 in terms of the aggregate rank correlation. This increase is primarily associated with a higher sorting propensity out of unemployment. In this section, we analyze this observation in more detail. To understand which worker-firm combinations contributed to the increase of sorting, we track the changing distribution of certain worker types across firms over time, both for matches out of unemployment and for job-to-job switches. This allows us to show for which worker types the distribution across firms shifted towards the theoretically predicted optimal allocation, explaining the observed increase of the overall rank correlation. Since we have seen that sorting patterns in the German labor market are far from being homogenous across worker and firm types, it is not surprising that we also observe worker types which became less prone to sort. In the following, we condition on a number of worker bins and, figuratively, slice through the empirical bivariate density of matches depicted in Figure 2. We compare these "density slices" for different time intervals to see where the distribution across firms changed and where it remained constant. Figure 5 shows estimated density functions for the worker bins 1, 10, 50, 90, and 100.[34] We compare the first half of our sample, 1998-2002 (red line), to the second half, 2003-2008 (black line). Notably, the distributions changed significantly over time for both low-type (bins 1 and 10) and high-type (bins 90 and 100) workers. For medium-type workers, however, the density functions largely lie on top of each other and can thus not 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. At the same time, the density of matches with firms in bins 9 to 18 has decreased significantly. Apparently, the distribution of the workers in bin

[34] Histograms of the raw match data can be found in Appendix B, Figures A.3 and A.4

Figure 5: Estimated Density Functions: 1998-2002 (red) vs. 2003-2008 (black)

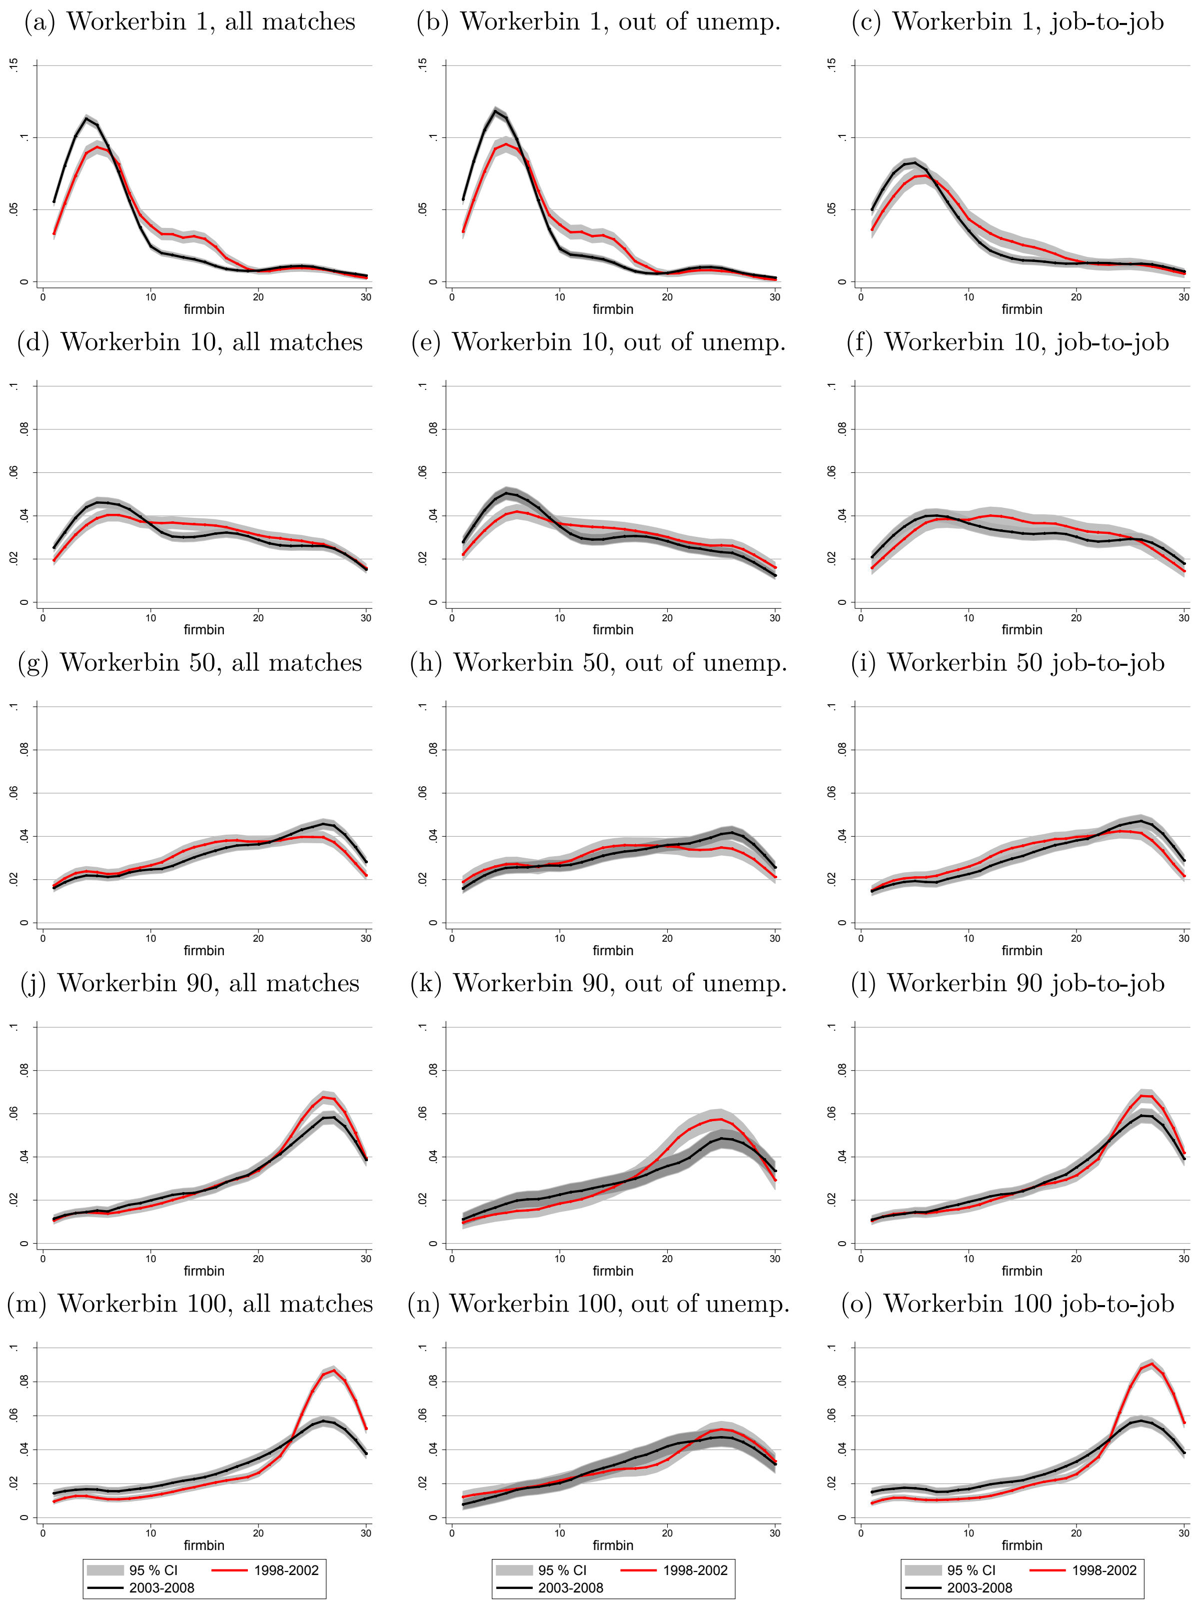

Note: Estimated univariate kernel densities of matches conditional on worker bins, time, and match type. The kernel is estimated using an Epanechnikov kernel function. The bandwidth is calculated by Silverman's rule of thumb. Pointwise confidence intervals are calculated using a quantile of the standard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

1 across firms shifted to the left during our period of observation, see Panel 5a. Recall that this kind of shift must lead to an increase of our measure of sorting because workers of the lowest type moved towards their theoretically predicted "optimal" counterparts, the lowest firm type. 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. A further decomposition of the distribution of bin 1 workers across firms into out of unemployment and job-to-job matches reveals that the described shift is present for all matches of bin 1 workers but more pronounced for matches which were 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 4. Due to space constraints, we cannot plot estimated densities for all 100 worker bins. The described distributional shift is present for all workers up to bin 10, see Panels 5d and 5e. Panel 5f confirms that the shift is less pronounced for job-to-job matches. Above worker bin 10, the distributions are largely flat and do not change significantly over time. Medium worker types are much more equally distributed across firms and, apparently, sorting patterns for them have not changed. If anything, we observe a slight significant increase of bin 50 workers with high-type firms. For the highest worker ranks in bins 90 and 100, the observed distributional changes are very different. Overall, the distribution has shifted away from the optimal allocation at the highest firm types, see Panel 5m, even though its mode remains well above firm bin 25 also in the second half of our sample. Overall, sorting has become much less pronounced for high-type workers in our sample. Interestingly, this development is driven primarily by job-to-job switches. While the distribution of matches with high-type workers out of unemployment has not changed significantly over time (Panel 5n), the density of new job-to-job matches with high-type firms has decreased by about one third (Panel 5o). We observe the described reduction of the sorting propensity of high-type workers for all workers above, roughly, bin 85 in our sample.

Since we use estimated kernel density functions in Figure 5 to analyze the changing distributions of workers across firms, we show additionally that our findings do not hinge on the density estimation procedure. They are present in raw data as well. In the context of our two samples, 1998-2002 and 2003-2008, we plot the difference of the two empirical densities conditional on worker bin, time interval, and match category in Figure 6. In case of a positive (negative) difference for a certain firm bin, the empirical density of matches has shifted upwards (downwards) for a given worker-firm combination. We show density difference plots for the same worker bins as before and confirm the observed distributional shifts towards more sorting out of unemployment for low-type workers and less sorting from job-to-job switches of high-type workers.

Figure 6: Densitiy Differences: 2003-2008 − 1998-2002

(a) Workerbin 1, all matches

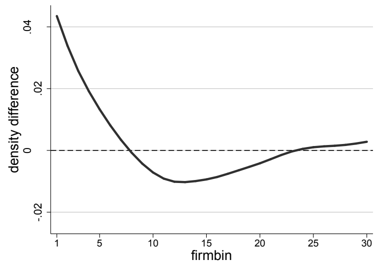

(b) Workerbin 1, out of unemp.

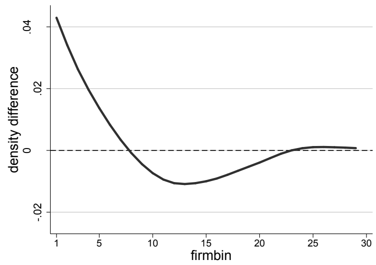

(c) Workerbin 1, job-to-job

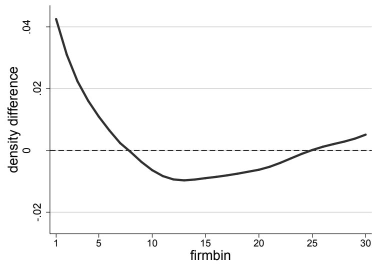

(d) Workerbin 10, all matches

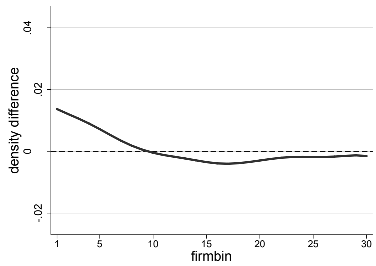

(e) Workerbin 10, out of unemp.

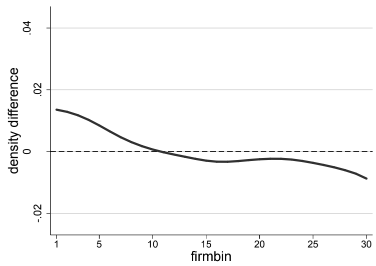

(f) Workerbin 10, job-to-job

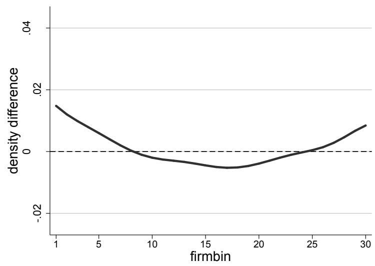

(g) Workerbin 50, all matches

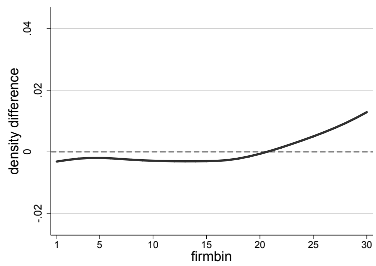

(h) Workerbin 50, out of unemp.

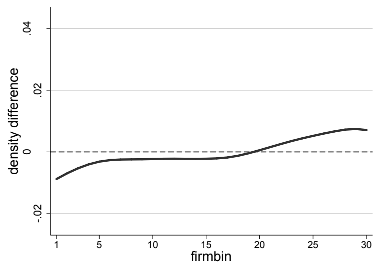

(i) Workerbin 50 job-to-job

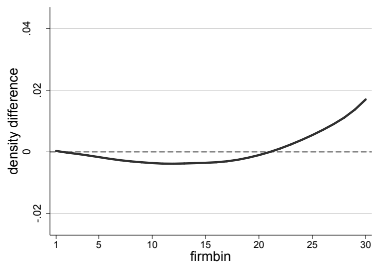

(j) Workerbin 90, all matches

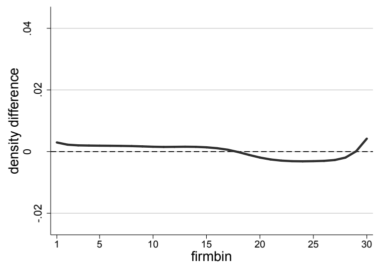

(k) Workerbin 90, out of unemp.

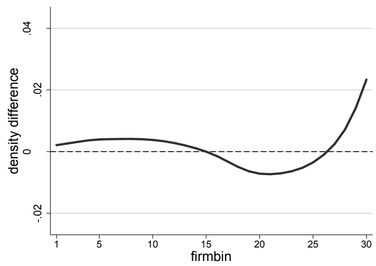

(l) Workerbin 90 job-to-job

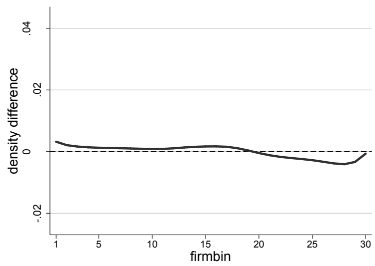

(m) Workerbin 100, all matches

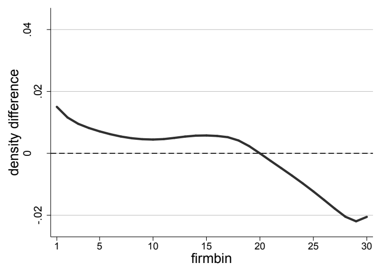

(n) Workerbin 100, out of unemp.

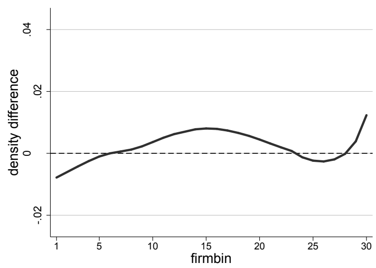

(o) Workerbin 100 job-to-job

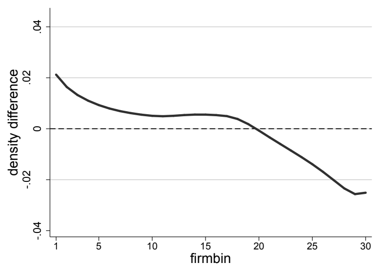

Note: Density differences are calculated by subtracting the raw match frequencies by bin. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

5.3 Wage (Non-)Monotonicity and Inequality

Our analysis of labor market sorting is motivated by the theoretical prediction of non-monotonic wages of a given worker type across firm types. This is at odds with the assumptions of the AKM model. In theory, a given worker's wage can be negatively related to the type of the firm he is matched with. The wage of every worker is maximized at the type-specific optimal allocation, arguably due to a production complementarity in case of positive sorting. The wage falls in other matches apart from the optimal allocation what gives rise to the non-monotonic wage pattern. In the simple theoretical model we considered, the optimal match of a worker is simply a firm with a similar position in the ranking. In the final part of the paper, we check if we can indeed observe these theoretically predicted wage patterns in our data. We then relate the wage patterns we observe to the changes of sorting in Germany documented before. Additionally, we look at the changes of wage inequality within and across our worker bins in the light of our findings.

We have shown that the distribution of low-type workers across firms has shifted towards the theoretically predicted optimal allocation while the distribution of high-type workers has shifted away from it. These contrary trends add up to an overall increase of sorting in Germany, which is most pronounced for new matches out of unemployment. It turns out that the observed higher tendency of low-type workers to sort into low-type firms is indeed associated with a wage pattern across firms that corresponds to the theory of sorting. We use our rankings to show evidence that low-type workers earn the highest wages at low-type firms. Figure 7 illustrates this finding. Clearly, the mean wage of our lowest worker type (worker bin 1) is decreasing in the type of the firm. To the best of our knowledge, we are the first to provide direct empirical evidence for a negative relation of a worker's wage and a performance measure of the firms he is matched with. It is noteworthy, however, that this negative relation is visible in our data for the lowest worker types only. The AKM methodology assumes monotonically increasing wages in the type of the firm a worker is matched with. This pattern is visible for medium-type workers in bins 25 and 50 in our data. For workers in bins 75 and 100, the AKM assumption is met for almost all firm types. We see deviations from monotonicity when high worker types are matched with firms of the lowest type, possibly because high-type workers happen to be matched with low-type firms as managers or executives. This might enable them to extract higher wages even at non-performing firms.[35] We also observe a severe drop of wages for matches with the highest firm types in bin 100, which we take to be an artifact of our firm ranking procedure in the absence of a better explanation.[36]

[35]Unfortunately, we cannot test this possible explanation because our data do not allow us to identify managers and executives.

[36]By construction, the within-bin variance ist maximized in the highest and lowest worker and firm bins.

Figure 7: Estimated Mean Wages across Worker Types

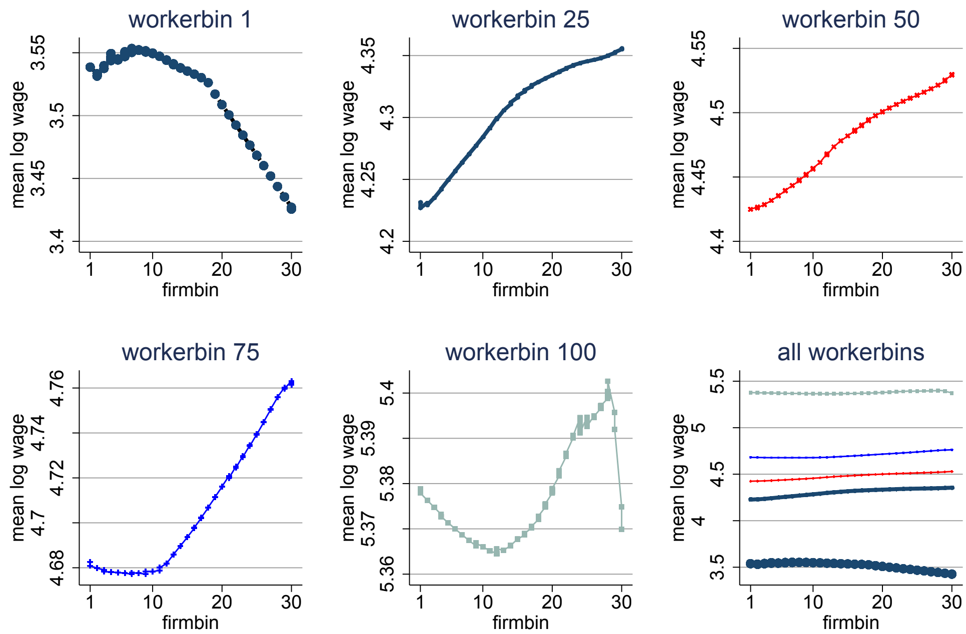

Note: Wages are means of real log cell-wages for the given worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function.

The analysis of estimated mean wages per worker bin shows that the AKM assumption of monotonically increasing wages in the firm type is clearly violated for low-type workers. Arguably, match-specific effects due to possible production complementarities are strong for this type of workers and should not be excluded by assumption. Our finding corresponds to the large negative wage residuals that CHK report for low type workers. On the other hand, the AKM assumption appears to describe wages well for a large share of medium and high-type workers in our sample. For high-type workers, the monotonicity assumption and the theoretically predicted wage pattern overlap. Their wage is maximized in matches with high-type firms both in the reduced-form fixed effects model and in theoretical models of sorting. Our results are in line with CHK in the sense that deviations from the two-way fixed effect model may very well be small on average. We have shown, however, that deviations are huge for low types of workers while the model predictions are hard to disentangle for high-type workers. Depending on the question at hand, the AKM model does a good job in explaining the link between unobservable characteristics of workers and firms and the wages paid. From the perspective of labor market policy however, which is commonly most concerned with labor market outcomes of low-type workers, taking into account the documented negative relation between these worker's wages and the type of the firm they are matched with, which cannot be captured with the reduced-form model, appears to be crucial.

At last, we take a look at wage inequality with the tools we developed and plot differences of the log wage variances over time for the same worker bins as before, see Figure 8.

Figure 8: Differences of Log Wage Variances over Time and across Worker Types

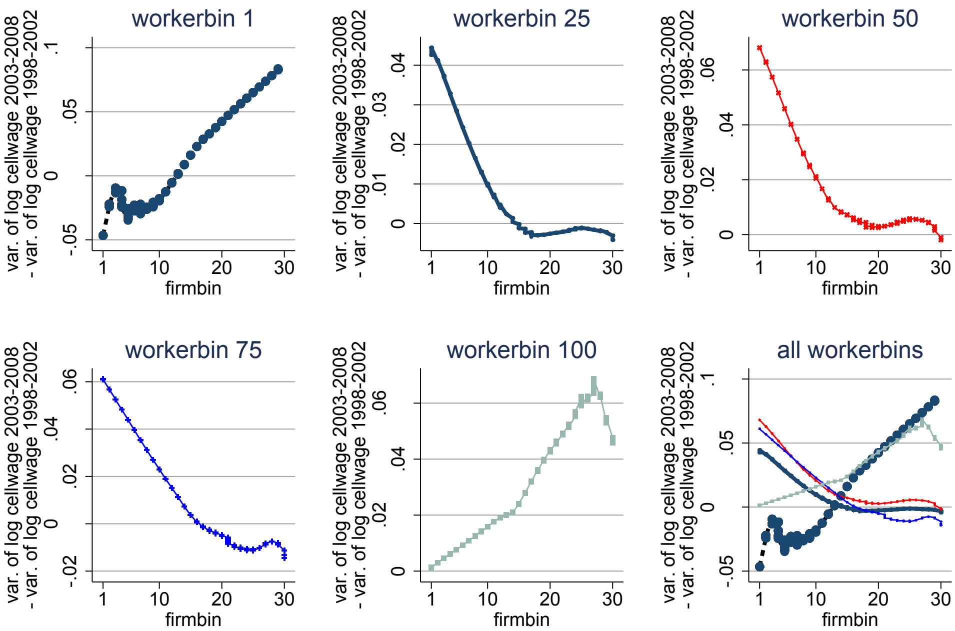

Note: Variances are calculated from log real daily wages of 30x100 bin-cells. The Figure shows differences between the variance in 2003-2008 and the variance in 1998-2002 for certain worker bins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and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

We clearly see negative differences of log wage variances for low worker types (here bin 1) and firms in the bottom half of the firm ranking. The negative differences indicate that the wage variance within these worker bins has decreased over time. Interestingly, we also observe the strongest increase of sorting for these matches. The distributional shift towards the optimal allocation of low-type workers hence appears to have reduced wage inequality for low type workers who are matched with low type firms. As low type workers became increasingly matched with lower firm types over time in our sample, the higher wages earned in those matches compressed the wage distribution of low type workers in these matches. For worker bins 25, 50, and 75 the plotted variance differences look like the mirror image of worker bin 1. The positive differences show that the variances of log wages within these worker bins increased in matches with low type firms and did not change or decreased slightly in matches with the upper half of the firm ranking. In worker bin 100, wages have become more unequal in matches with all firm types. The largest variance differences are found around firm bin 25 indicating that the increasing wage variance for these matches could be linked to the downward shift of the density of matches in this area, recall Figure 5o. Apparently, the severe decrease of job-to-job sorting has made wages more unequal for high-type workers matched with high-type firms.

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 are related to the dynamics of overall wage inequality in Germany. Therefore,

Figure 9: Wage Dynamics

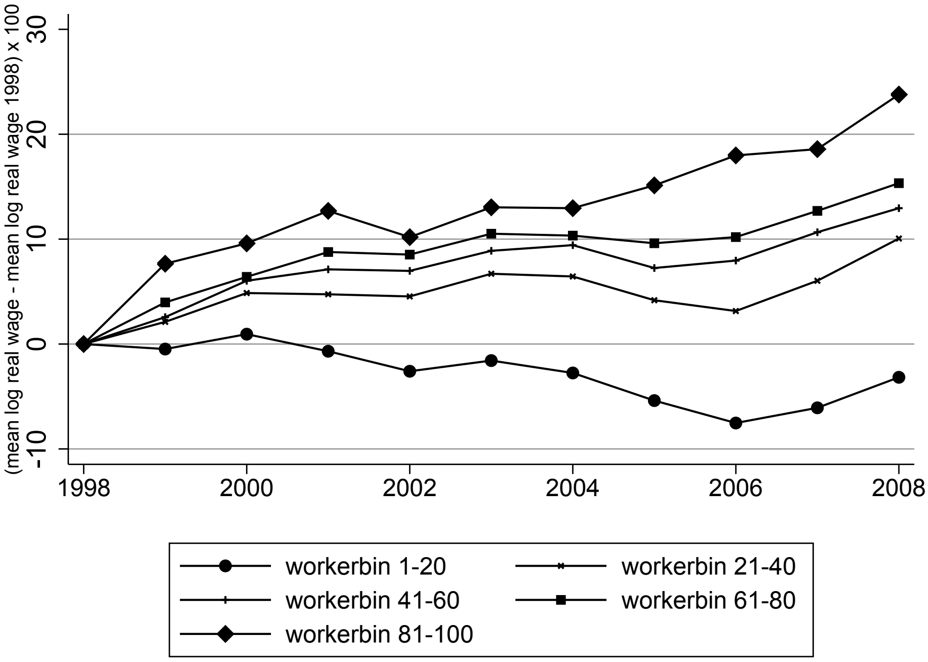

Note: The Figure shows deviations of mean log real daily wages for certain worker bins from the same value in 1998 (multiplied by 100).

in Figure 9 we plot how the mean wages for our worker bins evolved over time. The plot shows, in line with CHK and Dustmann et al. (2009), that overall wage inequality has increased over time in Germany during our period of observation. While mean log real wages have increased for all worker bins above 20 almost uniformly over time, we observe wage decreases for the 20 lowest worker bins. We conclude that the increase of sorting for low types had no significant effect on overall wage inequality, despite the higher homogeneity of wages within the lowest worker bins. Apparently, the shift towards higher wages for these workers was not pronounced enough to revert the trends in overall wage inequality or the wage increases were too small. Conversely, the increasing within-bin wage inequality that we observe for high-type workers is in line with the wage development of worker bins 81-100 in Figure 9.

6 Conclusion

The empirical study of labor market sorting in Germany documented in this paper 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 ranked workers who show an increased propensity to match with their theoretically predicted optimal counterparts. At the same time, we find that sorting has decreased for high-type workers who change from one employer to another. This decrease does not revert the overall trend of increasing sorting, however. Second, we find that the distributional shift of low-type workers matching more often with low-type firms is associated with higher wages indicating that wages guide the sorting of workers to firms supporting the non-monotonicity prediction in theory. This finding is at odds with reduced-form models of wages dispersion in the spirit of AKM that assume that more productive firms pay the same wage premium to all worker types. Third, we find that positive sorting is in line with previous results showing rising overall wage inequality. However, we also show that the increased sorting is connected to a decreasing wage dispersion over time among low-type workers in low-type firms.

In this paper, we did not attempt to make causal claims linking recent changes of labor market policy in Germany to the observed changes of the distribution of workers across jobs, it would appear far-fetched to assert that the changes occurred in isolation. Taking our data at face value, however, we do not see evidence for a discontinuity in the sorting patterns which could have been triggered by, say, 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—increasing wage dispersion and a severe reduction of aggregate unemployment—can be traced back well into the 1990s and are not a direct result of the Hartz reforms. Analyzing the relation between labor market policy and sorting in more detail remains a fascinating topic for further research.

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.

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.

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 (2015) "A Distributional Framework for Matched Employer Employee Data."

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.

Gautier, Pieter A. and Coen N. Teulings (2015) "Sorting and the Output Loss Due to Search Frictions," *Journal of the European Economic Association*, Vol. 13.

Hagedorn, Marcus, Iourii Manovskii, and Tzuo Hann Law (2014) "Identifying Equilibrium Models of Labor Market Sorting," Revision.

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, Nürnberg.

Kantenga, Kory and Tzuo Hann Law (2014) "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.

Lamadon, Thibaut, Jeremy Lise, Costas Meghir, and Jean-Marc Robin (2013) “Matching, Sorting, and Wages,” unpublished.

Lindenlaub, Ilse (2014) “Sorting Multidimensional Types: Theory and Application,” unpublished.

Lise, Jeremy, Costas Meghir, and Jean-Marc Robin (2015) “Matching, Sorting and Wages,” *Review of Economic Dynamics*.

Lise, Jeremy and Fabien Postel-Vinay (2015) “Multidimensional Skills, Sorting, and Human Capital Accumulation,” unpublished.

Lopes de Melo, Rafael (2013) “Firm Wage Differentials and Labor Market Sorting: Reconciling Theory and Evidence,” discussion paper, Universtiy of Chicago.

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.

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 who employ at least one employee who is employed at least at two different establishments of the IAB Establishment Panel. Additionally, up to 500 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 earning are right censored at the contribution assessment ceiling ('Beitragsbemessungsgrenze'). We use the pension insurance of workers and employees. This earning 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 observation by comparing wages with the upper earning limit. We define a wage observation as censored whenever the reported wage is higher than 99% of the censoring threshold. On average about 13 percent among all wage

Table A.1: Summary Statistics of 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.

observation 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 A.2: Rank correlations for different time intervals and match types

	Rank correlations (ρ)		
	all matches	out of unemp.	job-to-job
1998-2008	0.250	0.256	0.223
1998-2002	0.248	0.227	0.239
2003-2008	0.251	0.274	0.209
1998-1999	0.225	0.192	0.215
2000-2001	0.270	0.239	0.267
2002-2003	0.233	0.245	0.207
2004-2005	0.251	0.255	0.226
2006-2008	0.256	0.295	0.197
1998	0.192	0.181	0.191
1999	0.253	0.203	0.231
2000	0.248	0.228	0.245
2001	0.292	0.253	0.290
2002	0.228	0.254	0.194
2003	0.239	0.236	0.222
2004	0.227	0.252	0.194
2005	0.275	0.258	0.260
2006	0.266	0.289	0.210
2007	0.239	0.295	0.171
2008	0.266	0.305	0.212

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.

B.2 Distribution of Observables Across Worker and Firm Bins

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

(a) Age

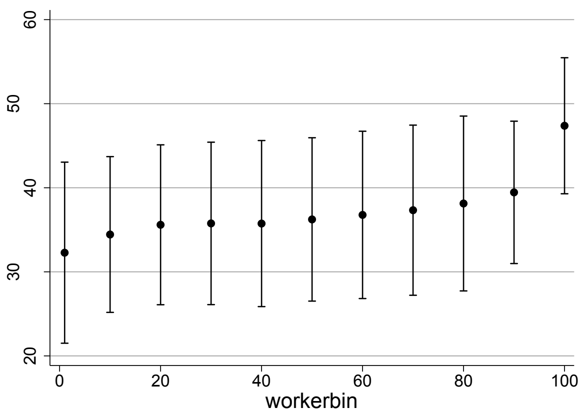

(b) Education

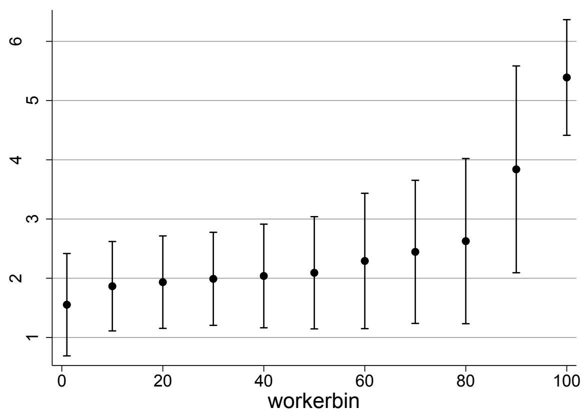

Note: The age of individual workers in our sample ranges from 20 to 60. There are 6 education categories: 1 = "no degree", 2 = "vocational training", 3 = "high school", 4 = "high school and vocational training", 5 = "technical college", 6 = "university".

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

(a) Number of Workers

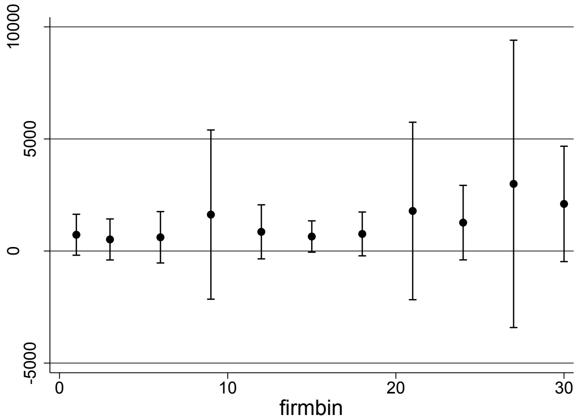

(b) Sector

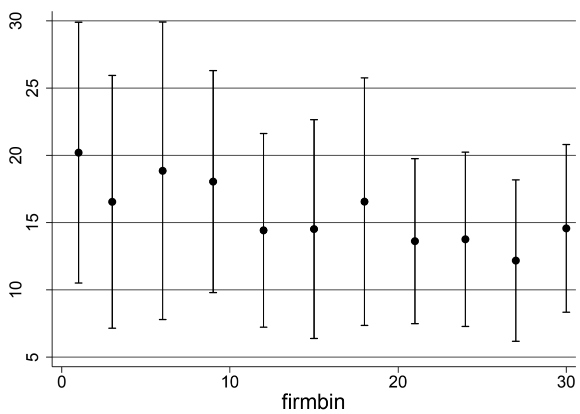

Note: Panel (a) shows means and standard deviations of the average workforce size of firms within bins. Panel (b) shows means and standard deviations of the firms' sectoral classification within the bins. We use the WZ93/WZ03 classification of industries available in the IAB Establishment Panel, which 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".

B.3 Histograms and Kernel Density Estimates

Figure A.3: Histograms of Raw Data and Estimated Density Functions: 1998-2002

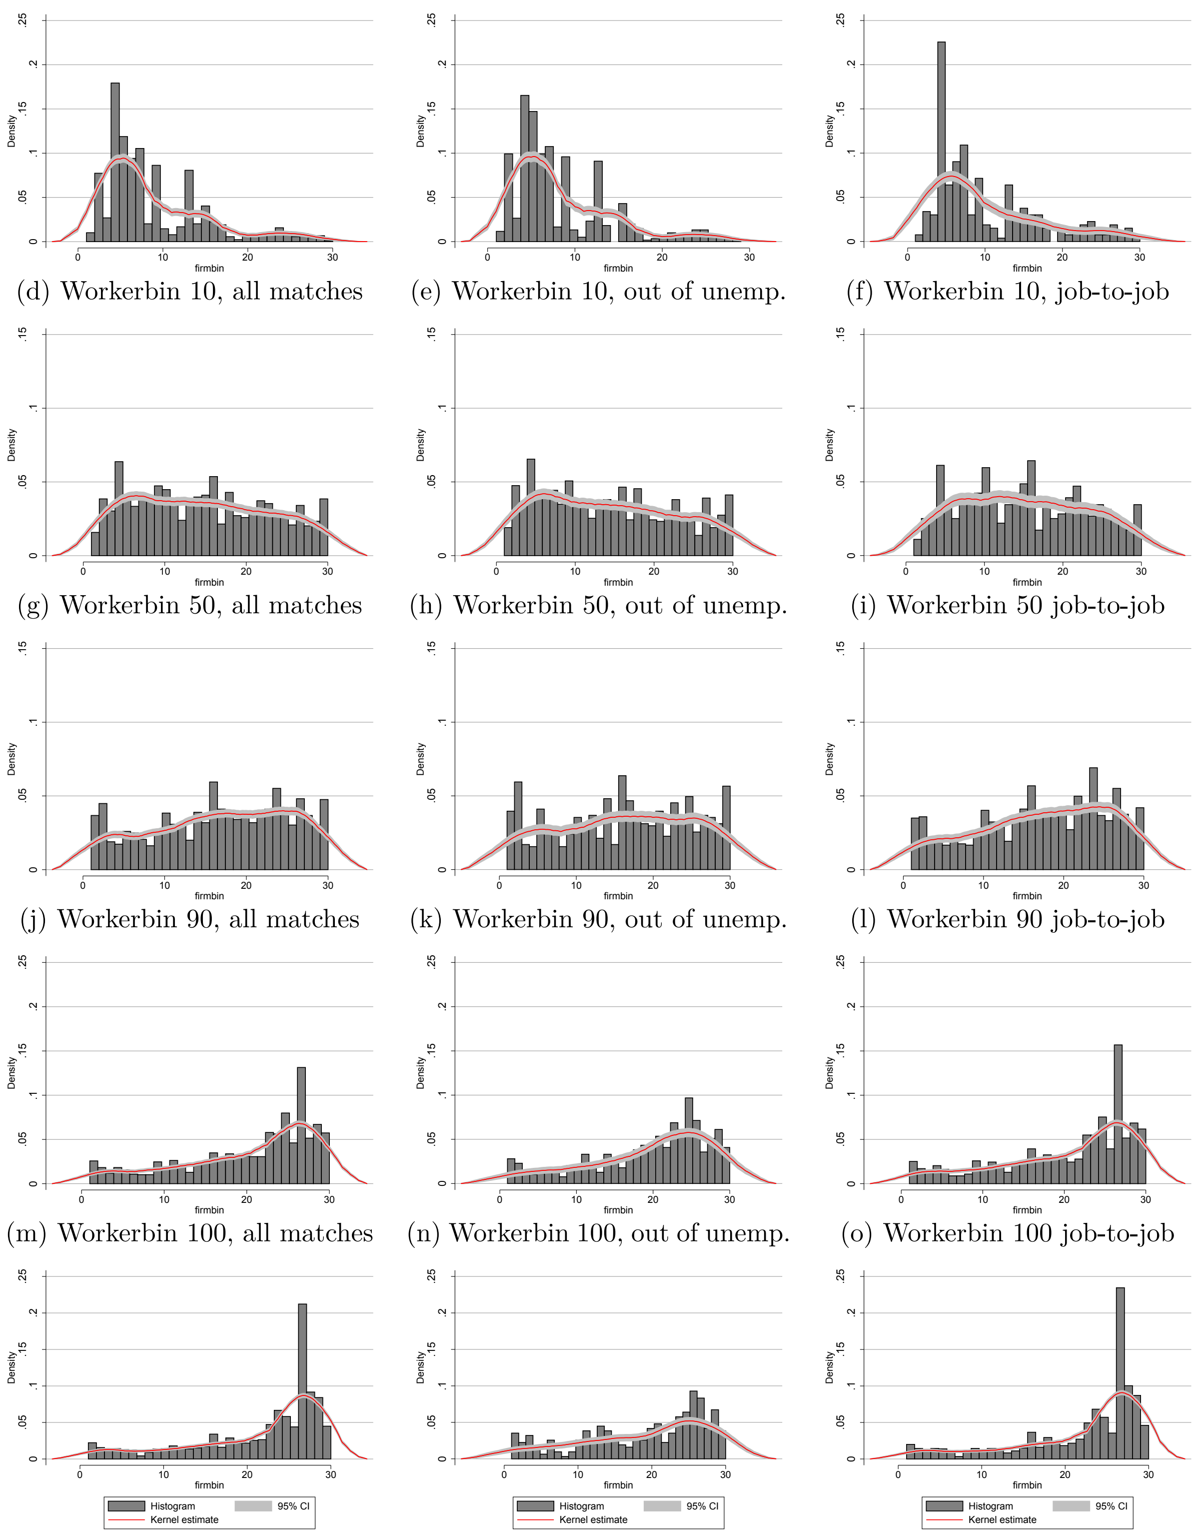

(a) Workerbin 1, all matches (b) Workerbin 1, out of unemp. (c) Workerbin 1, job-to-job

(d) Workerbin 10, all matches (e) Workerbin 10, out of unemp. (f) Workerbin 10, job-to-job

(g) Workerbin 50, all matches (h) Workerbin 50, out of unemp. (i) Workerbin 50 job-to-job

(j) Workerbin 90, all matches (k) Workerbin 90, out of unemp. (l) Workerbin 90 job-to-job

(m) Workerbin 100, all matches (n) Workerbin 100, out of unemp. (o) Workerbin 100 job-to-job

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.

Figure A.4: Histograms of Raw Data and Estimated Density Functions: 2003-2008

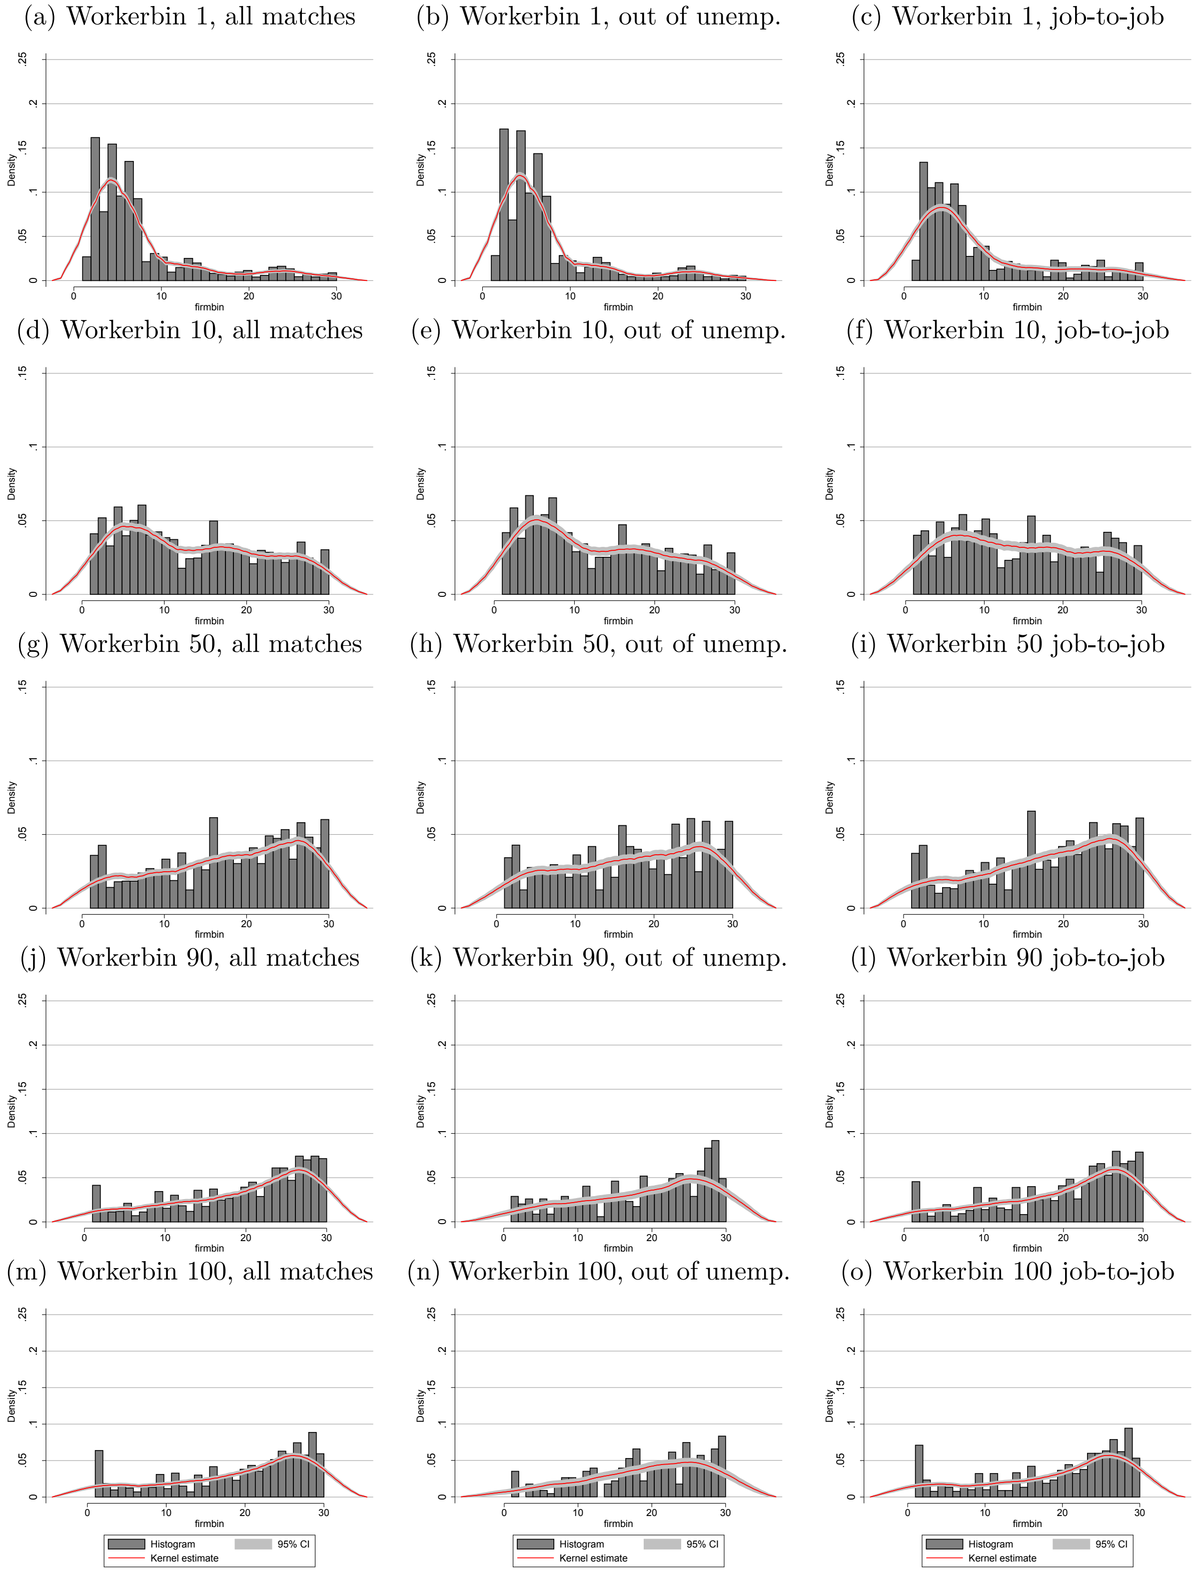

(a) Workerbin 1, all matches (b) Workerbin 1, out of unemp. (c) Workerbin 1, job-to-job

(d) Workerbin 10, all matches (e) Workerbin 10, out of unemp. (f) Workerbin 10, job-to-job

(g) Workerbin 50, all matches (h) Workerbin 50, out of unemp. (i) Workerbin 50 job-to-job

(j) Workerbin 90, all matches (k) Workerbin 90, out of unemp. (l) Workerbin 90 job-to-job

(m) Workerbin 100, all matches (n) Workerbin 100, out of unemp. (o) Workerbin 100 job-to-job

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.

C Robustness Checks

C.1 Rank Correlations using Alternative Firm Ranking

Table A.3: Rank correlations for different time intervals and match types

	Rank correlations (ρ)		
	all matches	out of unemp.	job-to-job
1998-2008	0.233	0.256	0.194
1998-2002	0.211	0.212	0.191
2003-2008	0.248	0.282	0.197
1998-1999	0.209	0.186	0.192
2000-2001	0.203	0.216	0.184
2002-2003	0.232	0.242	0.201
2004-2005	0.244	0.260	0.209
2006-2008	0.253	0.304	0.186
1998	0.191	0.160	0.197
1999	0.222	0.210	0.181
2000	0.230	0.200	0.233
2001	0.173	0.235	0.133
2002	0.232	0.249	0.196
2003	0.232	0.236	0.207
2004	0.218	0.253	0.174
2005	0.271	0.267	0.245
2006	0.268	0.301	0.205
2007	0.239	0.303	0.164
2008	0.254	0.308	0.187

Notes: Rank correlations using alternative firm ranking as detailed in section 4.2. 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.

C.2 Main Plots using Alternative Firm Ranking

Figure A.5: Empirical Bivariate Density of Matches in Germany (1998-2008)

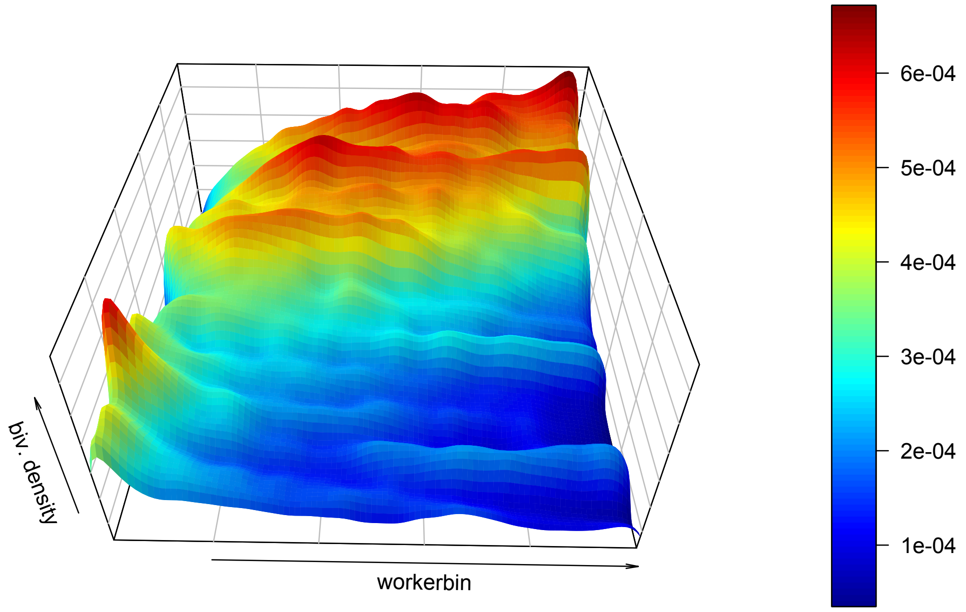

Note: Alternative firm ranking as detailed in section 4.2.

Figure A.6: Contour Plot of Empirical Bivariate Match Density

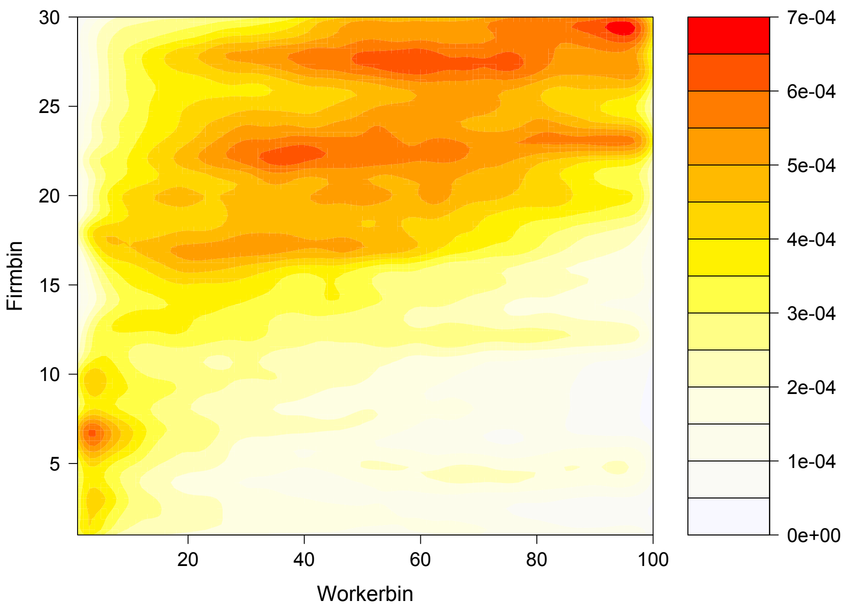

Note: Alternative firm ranking as detailed in section 4.2.

Figure A.7: Rank correlations by match type over time (1998-2008)

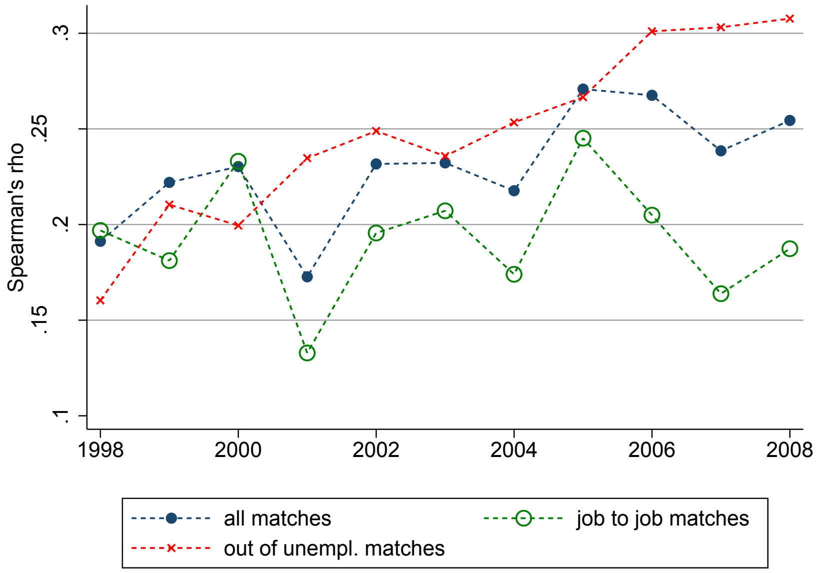

Note: Alternative firm ranking as detailed in section 4.2.

Figure A.8: Estimated Mean Wages across Worker Types

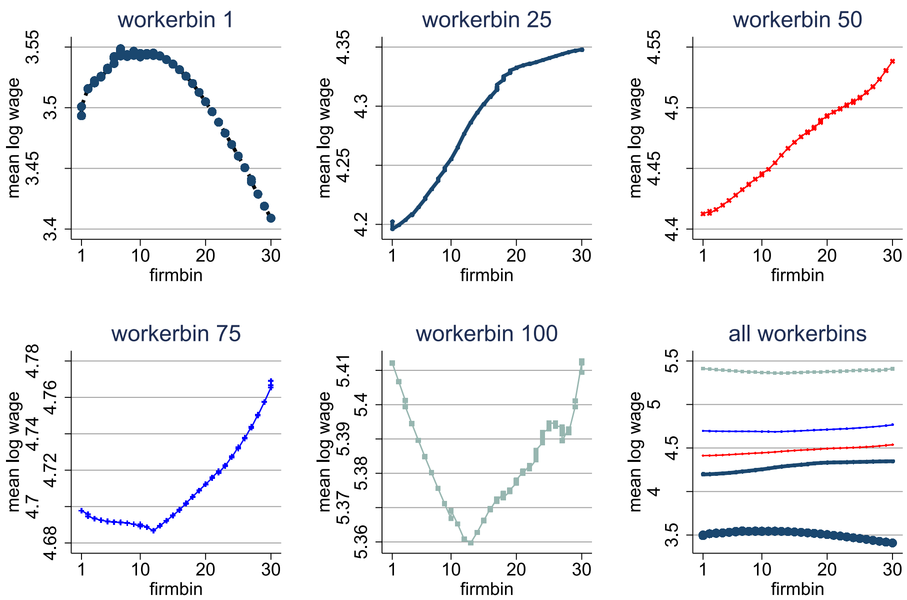

Note: Wages are means of real log cell-wages for the given workerbins. Line plots are the result of a locally weighted regression with running-line least squares smoothing and a tri-cube weighting function. Alternative firm ranking as detailed in section 4.2.