Online Appendix: Assessing the Role of Workplace Heterogeneity in Recent Trends of the Gender Wage Gap

Benjamin Bruns
Humboldt University

April 15, 2018

A Data Set and Additional Summary Statistics

A.1 The LIAB Mover Model 9308

My analysis uses the LIAB Mover Model 9308, a linked employer-employee data set prepared and provided by the German Institute for Employment Research (IAB). This data set links three main data sources: 1. Integrated Employment Biographies (IEB) with worker (spell) data, 2. Establishment History Panel (EHP) with administrative data on establishments, and 3. IAB Establishment Panel (EP) with survey data on establishments. The three data sources can be matched via three identifiers: a unique person-identifier (PID) and two establishment identifiers (EID1, EID2) that are all stored in the IEB file. In the following, I label the identifier for the EHP data as EID1, and the identifier for the EP as EID2.

The mover model (MM) belongs to the longitudinal versions of the linked employer-employee data sets, including the complete employment biographies of all workers selected into the sample. It differs from earlier versions of the longitudinal models (LM1-3) as well as the current version (LM) mainly in terms of the sampling design (see also Heining et al., 2012, 2014). The construction starts with all workers who are employed at two different establishments from the EP as of June 30th in two different (not necessarily consecutive) years. In practice, one could imagine the sampling procedure in two steps. First, all workers with at least one spell with non-missing EID2 are selected (previously restricting spells to main employment and covering June 30th). From this sample, only those workers are retained, whose employment biographies contain at least two distinct EID2's. For these EID2's, it must be ensured that the establishment departed from has a valid interview in the year of departure, and that the establishment moved to has a valid interview in the year of arrival. After eliminating recalls (dupicate PID-EID2 observations), this gives a sample of unique PID-EID2 combinations. For each (unique) EID2 in this sample, up to 500 additional employment biographies who are also linked to that firm are then randomly sampled and added to the data set. If an establishment employs less than 500 different workers between 1993 and 2008, all workers are sampled. These 'additional' workers either do not move or move to an establishment not participating in the EP. The total number of observed individuals at an establishment finally consists of the randomly sampled individuals plus all movers.

The sampling design of the mover model implies that I do not observe all workers employed at an establishment at a given point in time if that establishment is associated with more than 500 workers over the 1993-2008 period. However, I know the "true" number of full time employees as of June 30th in any year from the corresponding administrative variable, hence I can evaluate the fraction of workers that I do observe at each establishment. This is done in Figure A.1 which illustrates the ratio of the *observed* to the *actual* number of full-time employees for three different samples used in my analysis: the largest connected set, the dual-connected, and the EP sample. The actual number of full-time employees comes from the EHP, refers to June 30th each year, and is aggregated from the universe of the IEB. The observed number comes from the mover model, and refers to all workers that can be linked with a firm in a year, subject to the sample restrictions. For establishments in the largest and the dual-connected sets, I observe some 20-30% of all workers (note that the largest connected set contains the dual-connected set and the EP sample), i.e., also in establishments that never participate in the EP. For the EP sample, the coverage ratio is much higher, but it is not 100% for establishments with less than 500 employees. This is partly due to the sample processing, where I lose some connections when I define unique person-year observations. Another reason is that, by design, high turnover firms might be relatively small but still employ more than 500 workers over time.

A.2 Data Processing and Initial Sample Restrictions

The worker-level information is stored in continuous spells (daily accuracy) extracted from the IEB which contains the universe of employment histories of workers subject to social security (roughly 80% of total employment). The sample selection of the mover model imposes that at least one spell of any worker contained in the sample covers June 30th of some year, but conditional on doing so, the complete employment biographies for the 1993-2008 period are added to the sample. Hence, further modifications are necessary. The IEB combines employment histories of workers subject to social security with process-generated data from the Federal Employment Agency (Heining et al., 2012). The employment histories are constructed from employer notifications that detail the start and end date of a job match. An employer has to file at least one notification for all workers employed at the end of a year, and each time the salary or employment status of a worker changes. This produces a data set with multiple consecutive and potentially parallel spells (if a worker holds multiple employment relationships) within a year. In order to arrive at a sample with unique worker-establishment-year observations, several modifications are necessary. In an initial sample selection, I exclude job-spells associated with marginal employment and apprenticeships, and trim the data on location (West Germany), age (20-60), and employment status (full-time). I also exclude all non-employment spells as well as parallel employment spells by adopting the definition of the 'main job' of the IAB.

Next, I combine multiple observations with the *same* employer during a year into a single worker-establishment-year observation. Following CHK, I retain the maximum occupation code if the occupation changes between multiple worker-establishment observations with a constant establishment identifier in the same year. I compute the average daily wage for a worker-establishment-year match by weighting each wage observation with the corresponding length in days. Afterwards, I define the main job within a year as the worker-establishment match that yields the highest total earnings in a year, and drop all other job matches of that worker in the course of the same year. I deflate wages to the year 2000 using an aggregate CPI, and drop all observations with an average (real) daily wage of less than 10 Euro after collapsing the data to unique person-establishment-year combinations.

From the establishment register data, I extract three-digit industry classifications, location (district level), and the number of full-time employees as well as the number of all employees. I measure firm size by the number of full-time employees (allowing for 0) which I average over the current and the preceding year. Moreover, I modify industry and region-variables so as to be constant at the establishment-level, assigning the modal value to all establishment-year observations. The following table provides an overview of the main variables used in the analysis.

A.3 Imputation of Wages and Education

Wages are top-coded at the social security maximum. Top-coding affects around 10-15% of male and around 5% of female observations each year. Following CHK, I impute an upper tail of the cross-sectional wage distributions based on a series of Tobit models. Specifically, I fit a series of 560 Tobits (14 years: 1995-2008; 5 education groups: no/missing, primary, vocational, some college, university; 4 age groups: 20-29, 30-39, 40-49, 50-60; 2 genders: male and female) of the real log daily wage on a linear term in age, an individual's wage in all other years, the fraction of censored wages in all other years, an indicator for large firm (>10 employees), a quadratic polynomial in firm size, and a dummy for singleton worker observations. Then I replace the censored wage by an uncensored prediction obtained from the estimated parameters and a random draw from the corresponding truncated normal distribution.

B Measures of Productivity

B.1 Definition of Productivity Measures

The primary measure of firm surplus used in this study is log value added, constructed from a combination of variables from the EP and administrative information from the EHP. Formally:

$$VA_{jt} = ln\left(\frac{Sales - Cost\ of\ inputs}{Firm\ size}\right)_{jt} \tag{1}$$

where *Sales* and *Cost of inputs* are derived from the EP, and *Firm size* is measured by the number of full-time employees taken from the administrative data and averaged over the past and current year. I exclude establishments in banking/insurance and public administration as their output is not measured in terms of sales. Other studies using EP data to construct productivity measures include Gürtzgen (2009, 2012), Frick (2002), and Wolf and Zwick (2002).

Cost of inputs are calculated by multiplying the reported cost share of inputs (in sales) by sales. Note that the cost share of inputs is missing in a number of cases where sales figures are available. Hence, in analyses based on sales, I can draw on a larger sample of firms compared to analyses based on value added. Most of my paper focuses on between-firm variation in log value added. I therefore average log value added per worker (and likewise log sales per worker) across firm-years, weighting by the number of person-years, $VA_j = \overline{VA_{jt}}$. A descriptive overview of the sample with mean log value added is provided in Table J.1, columns 5-6 and 11-12.

B.2 Descriptive Evidence of Productivity Measures

Figure J.2 shows histograms of *current* log value added per worker (panel A) and log sales per worker (panel B) in 1995 and 2008, weighting each firm equally. The distribution of log value added per worker

Table 1: Description of Variables and Data Sources

Variable	Data Source	Description
Education	IEB	The education level is imputed using the IP1-procedure of Fitzenberger et al. (2006). I combine the resulting 7 education levels (including missing) into 5 categories: missing, dropouts/primary, vocational without high school ('apprenticeship'), high school with/without vocational ('some college'), polytech/university. To compute the mean years of schooling shown in Tables 1 and J.1, I follow CHK and assign each category the expected years of schooling required to obtain the corresponding degree: missing - 10.5 years; dropouts/primary - 11 years; apprenticeship - 13 years; some college - 15 years; university - 18 years.
Occupation	IEB	I use a three-digit occupational classification standard of 1988. It contains around 330 occupations after dropping the following categories from the analysis: 555 (Disabled), 666 (Rehabilitants), 888 (Nursing staff), 924 (Employees by household cheque procedure), 971-999 (Family assistants, unpaid interns, workforce not further specified, employees in partial retirement, among others). In the decompositions by occupations, I aggregate this measure to 9 ISCO Major Groups, combining the groups 1 (elementary) an 2 (agriculture) as well as 8 (professionals) and 9 (senior officials and managers).
Industry	EHP	I use a three-digit industry classification standard of 1993. It distinguishes between 224 industries. I combine the following industries due to very small cell sizes: 14/15, 120/131/132, 182/183, 602/603. In the associated industry-decompositions, I aggregate the three-digit classification into 16 broader categories based on a definition used in the IAB-EP.
Full-time employees	EHP	The number of full-time employees is derived from the corresponding variable in the administrative records (az_ges_vs). I interpolate the average of the past and current year, and allow for zeros.
Sales p. w.	EP	This variable is reported retrospectively for the previous year. I forward impute last year's information to obtain current year values. Sales per worker is computed by dividing through the number of full-time employees, and deflated to the year 2000 using an aggregate GDP-deflator from the World Bank. Moreover, I divide sales by 1,000.
Share of inputs	EP	This variable is also reported retrospectively, and, hence, imputed forward. It contains the percentage share of inputs in sales. I trim the variable at 1% and 99% (not percentiles!) prior to calculating value added.
Union coverage	EP	Establishments are requested to state the type of coverage, distinguishing between a) industry-wide agreements, b) firm-wide agreements, and c) no coverage. The variable is consistently available since 1995. I impute a constant union coverage by calculating the modal value of reported coverage types, and assign the stricter regime in case of ties, i.e., industry-level over firm-level over no coverage.

Note: Table summarises the main variables used in the analysis.
Source: LIAB Mover Model 9308

is trimmed below the 1st and above the 99th percentile, but still reveals a sizeable dispersion: in 1995, it spans a range of about 387 thousand Euro, and in 2008 it covers some 483 thousand Euro. For log sales, the corresponding values are 1.06 million and 1.46 million Euro. These figures also illustrate the sizeable expansion of productivity dispersion over time, though I acknowledge that some of this expansion may be driven by changes in the sampling frame (note, however, that by standardising with the number of full-time employees, this concern should be reduced).

The expansion noted in Figure J.2 blends changes in the composition of firms over time (entry/exit) with time variation in productivity for continuing firms. Histograms for *average* productivity measures, i.e., which are constant at the firm level and correspond to the concept of productivity that I adopt in most of my analysis, are plotted in Figure J.3. A first observation is that the variation is smaller compared with the non-averaged case, though the overall spread is somewhat larger. It is reasonable to assume that part of the lower variation is achieved through the fact that averaging at the firm level cancels out some of the annual fluctuations that simply reflect noise. To the extent that the spikes and fatter tails in the non-averaged case (Figure J.2) are generated by fluctuations around the true productivity level of a firm, the distributions illustrated in Figure J.3 should give a better approximation of firm level productivity. A second observation is that also in the case of average productivity, which shuts of time variation in productivity for continuing firms, the distributions expand over time. Note that this is not a straightforward implication of Figure 2 (main text), which represents employment-weighted distributions. In that case, the changes over time might simply reflect a reallocation of workers across differently productive firms. In light of one of my main results in the paper, which is that this reallocation does not play a role in driving the growth in firm premium inequality between genders over time, it is reassuring that a similar implication follows indirectly from comparisons of the weighted and unweighted productivity distributions. Finally, note that when I distinguish between periods in the analysis, the averaging of firm productivity is conditional on whether I consider the 1995-2001 or the 2001-2008 time interval, which is consistent with allowing for unrestricted firm premiums in the AKM estimations. The histograms in contrast refer to the distribution when lumping all years together.

To understand how the productivity distribution expanded, I show in Figure J.4 the evolution of different percentiles and percentile spreads for mean log value added per worker, normalised to zero in 1995. The series now refer to the employment-weighted distributions to ensure comparability with Figure 2. The patterns reveal that the bulk of the overall expansion comes from a relative expansion of lower tail inequality, measured by the distance between the 10th and 50th percentile.

The IAB expanded the EP sample substantially in the year 2000, almost doubling the number of establishments in West Germany with a particular focus on small firms. It is thus reasonable to wonder whether the productivity trends shown above are merely an artefact of changes in the sampling frame. The trends shown in Figure 2 and Figure J.4 do not provide a strong case for this as all lines evolve rather smoothly. Of course, this is only indicative, but given theoretical as well as empirical support for a positive correlation between firm size and firm productivity (Burdett and Mortensen, 1998; Foster et al., 2008), one would suspect a disproportionate addition of small firms to imply a rather sharp drop of the lower percentiles of the distribution (although this might be alleviated to some degree by the fact that I average productivity at the firm level).

Since I can neither exclude the possibility that changes in the sample composition drive the patterns in Figure 2 and Figure J.4 (though, as noted above, dividing by the number of full-time employees should at least reduce this problem) or use appropriate weighting factors (see below), I resort to census data to derive a point of comparison for the growth in productivity dispersion illustrated above. While this data source has the obvious limitation of not referring to the firm level, it provides all the necessary information to compute measures of value added that are conceptually similar to the measure I use in my analysis (sales - cost of inputs). In particular, I obtained input-output tables of manufacturing industries for the four-digit level (~250 cells per year) and of all industries for the two-digit level (~70 cells per year). In each case, I deflate nominal values to the year 2010 based on a producer price index from the Federal Statistical Office, divide by the number of employees in a year (unfortunately, I cannot distinguish between full-time and part-time), take logs, and compute the employment-weighted variance. These computations lead to very similar trend increases of productivity dispersion over time. For example, between 1995 and 2007 (the last year with consistent industry codes), I find that the dispersion in manufacturing rises by about 6.4 log points and for all industries it increases by 4.4 log points. For comparison, my EP based variance of average productivity rises by 4.3 log points between 1995 and 2008. I plot the time series of the productivity variance in manufacturing in Figure 2.

Person-Year Weights vs. IAB Cross-Section Weights: Larger firms are overrepresented in the Establishment Panel, even after the expansion took place in 2000. To obtain results that are nationally representative one thus needs to correct for the particular sampling design. For the more widely used EP data sets, the cross-sectional models and the longitudinal models (Heining et al., 2014), the IAB provides appropriate weighting factors. These weights adjust (apart from non-response) for the disproportionate sampling within industry categories, firm size groups, and federal states, and aim to make the sample of firms in each year representative for the population of firms at the federal state level. Due to the specific sampling design of the mover model, which contains only establishments connected through worker mobility (see section A.1), the standard weights provided by the IAB are not useful. However, the IAB also does not provide adjusted weights.

Throughout my analysis, I therefore weight summary statistics and regressions with the number of person-years. This is a good approximation for firms with less than 500 employees, and tends to downweight larger firms in the sample, which in a sense might counterbalance the oversampling of these firms in the EP. A similar person-year weighting has been used in CHK, section VIII, in Goldschmidt and Schmieder (2017), section IV.C, and in Gürtzgen (2012), among others.

Alternative Normalisations of Productivity: The EP is currently the only data source that combines information on firm productivity with detailed worker level data. It is clear, however, that the survey nature comes along with considerable measurement error due to misreporting, roundings, and approximations on the part of respondents. To deal with this issue, I implement several cleaning steps. As regards the cost share of inputs, I drop values below 1 and above 99 percent before computing total value added. To convert total sales and value added into per capita terms, I divide by the number of full-time employees as reported in the EHP, which I previously averaged over the past and current year (more on this below). I then trim per capita measures of productivity below the 1st and above the 99th percentile, and calculate person-year weighted averages at the firm-level from the truncated distribution.

Normalising by full-time employees might produce an odd ranking of firms if some firms employ disproportionately more part-time workers than others, which in turn could affect the analysis of productivity-related gender differentials, e.g., if the share of part-time workers is particularly high in female-dominated firms. Put differently, productivity per worker might be higher in firms with a larger share of female workers because the share of females correlates negatively with the share of full-time workers. Comparing the number of full-time employees with the number of all employees (including part-time workers) in the gender-pooled sample, I find a correlation of 0.993. Moreover, a regression of the number of full-time employees on the total number of employees yields a coefficient of 0.89 with a standard error (clustered at the establishment level) of 0.023. Accordingly, switching to this variable basically generates a left-shift in the overall distribution of value added, suggesting that firms with more full-time workers employ (roughly) proportionately more part-time workers. Reassuringly, estimates for men and women yield coefficients of 0.894 and 0.852 with cluster-robust standard errors of 0.0248 and 0.0181. Hence, I cannot reject the null of no systematic upward shift in the productivity measure at female-dominated firms.[1]

A problem with the total number of employees according to the administrative data is that it is not measured consistently over time: from 1999 onwards, marginal employment became subject to social security contributions, generating a jump in the total number of employees driven by part-time employees. Indeed, including an interaction term for the post-1999 period in the above regression yields a coefficient of -0.038 with a standard error of 0.013. Thus, it is possible that based on this variable, the productivity ranking is confounded by different numbers of marginal employees showing up in the administrative data from 1999 onwards. Again, I do not observe significant gender differences in these estimates (male coefficient: -0.0379 (0.0124); female coefficient: -0.0398 (0.0184)).

To address this consistency issue, one could draw on survey information regarding the number and structure of employees, which is even more detailed and would enable a distinction between full-time and part-time workers. Informal experimentation, however, reveals some considerable differences between the number of employees stated by employers and the number according to the administrative data. This may be due to misreporting on the part of survey respondents, e.g., if they refer to the firm rather than the establishment. Of course, one may speculate that similar misreportings prevail in the variables used to measure productivity, but ultimately, I feel more comfortable using administrative information whenever possible.

Regardless of which data source is used, the inclusion of part-time workers requires an appropriate weighting to adjust for the fewer working hours (for year-round employees) and possibly fewer days or weeks worked in a year. While the latter could be approximated by the employment biographies (daily accuracy) if one takes them as representative for all workers at an establishment (note that I do not necessarily observe all workers, in particular, at large firms), the former, and indeed far more relevant, would require to weight part-timers with an *assumed* number of hours relative to full-time workers — like in most administrative data sources, I do not observe hours of work. This is facilitated by two categories of part-time work that can be identified in the data — small and large part-time — based on the maximum hours per week, but it remains a crude approximation of the labour input that appears more appropriate at higher levels of aggregation (e.g., industries or regions), but not at the firm level.

A final questions concerns the issue whether I should use the calculated number of full-time workers using only individuals observed in the data, or the administrative information. If I observed the universe of social security records, this choice would be inconsequential. However, as I only observe a subsample of individuals, the sampling design of the mover model tends to put a lower weight on very large firms (> 500 employees) as it only guarantees a direct linkage of up to 500 workers. As the very large firms are typically more productive, this would generate longer upper tails of per capital surplus distribution. In sum, I believe that a normalisation based on the actual number of full-time workers provides the most consistent measure of firm productivity.

C The Wage Bargaining Model

The log-linear wage equation with worker- and firm-fixed effects can be derived from a simple wage bargaining model. In the following paragraphs, I will provide a detailed derivation of the estimation equation. The exposition closely follows CCK.

In a given sample period p (1995-2001 or 2001-2008), the data contain N^* observations on $i =$

[1]The t-test for equal coefficients gives a test statistic of 1.368.

$1, \ldots, N$ individuals (workers) employed at $j = 1, \ldots, J$ establishments (firms). Following the notation of Abowd et al. (1999), the function $J(i,t)$ provides a mapping of worker i to firm j in year $t \in \{1, \ldots, T\}$. Note that T_i may vary across i. Worker i's gender is denoted by $G(i)$ with realisations on the set $g = \{M, F\}$ as shorthand for male and female. Assume that the logarithm of the real daily wage of worker i in year t and period p who is employed at firm j is determined by (I omit the p-subscript for clarity)

$$w_{it} = a_{it} + \gamma^{G(i)} S_{i,J(i,t),t} \tag{2}$$

where a_{it} denotes the alternative wage available to worker i in year t, and $S_{i,J(i,t),t}$ is the contemporaneous match surplus generated through the match between worker i and firm j in year t. The coefficient γ^g measures the average share of surplus appropriated by gender g. The match surplus has the following log-linear structure:

$$S_{i,J(i,t),t} = \overline{S}_{J(i,t)} + \phi_{J(i,t),t} + \eta_{i,J(i,t)} \tag{3}$$

The first term represents the average surplus level at firm j. It captures long-term firm-specific components such as brand recognition or monopolistic profits from patents. ϕ_{jt} captures firm-level transitory shocks to productivity such as short-term output demand shocks and other short-run factors with instantaneous impacts on surplus. $\eta_{i,j}$ is a time-invariant match-specific component that reflects productive complementarities between worker i and firm j. I assume that $\eta_{i,j}$ is mean zero for each firm and gender in each interval. The alternative wage, a_{it} can be decomposed as follows:

$$a_{it} = \alpha_i + X_{it}'\beta^{G(i)} + \epsilon_{it} \tag{4}$$

where α_i is a person-specific component that captures the idiosyncratic earnings potential which is assumed to be fully portable across employers, ϵ_{it} is a mean-zero worker-specific transitory earnings shock, and X_{it} is a covariate index that contains time-varying covariates that affect a worker's earnings potential irrespective of the employer that she works for. In the analysis, I include in X_{it} a set of education dummies fully interacted with year dummies, and a cubic polynomial in age (omitting the linear age term due to collinearity with year). β^g is a gender-specific vector of coefficients that contains the returns to observables. Under these assumptions, log wages can be expressed as follows:

$$w_{it} = \alpha_i + \psi^{G(i)}_{J(i,t)} + X_{it}'\beta^{G(i)} + r_{it} \tag{5}$$

where $r_{it} \equiv \gamma^{G(i)} \left(\phi_{J(i,t),t} + \eta_{i,J(i,t)}\right) + \epsilon_{it}$ is a composite error term, and $\psi^{G(i)}_{J(i,t)} \equiv \gamma^{G(i)}\overline{S}_{J(i,t)}$ is a gender-specific firm-wage premium paid to all workers of gender g conditional on working at firm j. This equation, augmented by a period-specific index, is the basis for the analysis in the main text.

D Exogenous Mobility and Additive Separability

While the data I use in this analysis are drawn from the same IEB source as in CHK, it is a smaller and non-random subsample of the universe of employment biographies used in their analysis. Thus, it deems worthwhile to investigate whether the basic mobility patterns are similar, and to probe the proximate validity of the exogenous mobility assumption along the way. In this section, I replicate several of the "tests" proposed by CHK and CCK, concluding that the basic assumption of exogenous mobility, as in their studies, seems to hold approximately. This finding may be useful for other researchers estimating worker-firm models since access to the universe of administrative records is often restrictive. The following discussion draws on arguments in CHK and CCK. Further useful discussions of the AKM assumptions can be found in CCHK.

D.1 Construction of Event Study

I begin with a sample of all male and female workers from the gender-specific largest connected sets in each interval. From this set, I select all firms that employ at least one male and one female worker over the analysis period. Then, I calculate for each worker the mean log wage of her coworkers in a given year and assign all jobs to quartiles of mean log coworker wages. Due to the sampling design, it is possible that the coworker wage is not based on all coworkers at an establishment, especially if firms are large.

Next, I identify event spells by following workers over a period of 4 years including 2 years at the old and 2 years at the new job, where the first year at the new job is coded as zero.[2] Hence, in the first period, the analysis is based on job spells beginning in 1997-2000, and in the second period it is based on job spells starting in 2003-2007. The associated transition profiles of men and women moving from the 1st and 4th quartile to any of the other quartiles are plotted in Figures D.1 and D.2.

D.2 Probing the Exogenous Mobility Assumption

The exogenous mobility assumption states that conditional on worker- and firm-characteristics, the probability of a match between worker i and firm j is independent of the time-varying error term for each gender in each period. If it is satisfied, estimation of worker- and firm-fixed effects via OLS yields unbiased parameters. To investigate possible violations of the exogenous mobility assumption, recall the definition of the composite error term of the AKM models specified in eq. (5):

$$r_{i,t} = \gamma^{G(i)} \left(\phi_{J(i,t),t} + \eta_{i,J(i,t)} \right) + \epsilon_{it} \tag{6}$$

Consider first the term which reflects a worker's share in firm-level transitory productivity shocks, $\gamma^g \phi_{j,t}$. Under the proposed wage model, this shock affects the wages of all workers employed at a given firm, scaled by the gender-specific elasticity of wages with respect to productivity. For example, workers might be more likely to leave firms that are hit by negative productivity shocks, and join firms that are hit by positive shocks. This suggests an average wage loss just before a move occurs, and larger wage growth of recent joiners. The transition profiles in Figures J.5 and J.6 do not show evidence of such trends.

Another possible violation arises if mobility is correlated with an individual's transitory wage shock, $\epsilon_{i,t}$. For example, workers that climb up the hierarchy in their old firm might use this momentum to enter firms that pay higher average wages. This rationale suggests a rise in mean wages prior the transition date. Inspection of the transition profiles again does not point to such trends.

Perhaps the most controversial implication of the exogenous mobility assumption is that job changes are independent of the match-specific permanent earnings component, $\eta_{i,J(i,t)}$ (Eeckhout and Kircher, 2011; Hagedorn et al., 2016). The transition profiles suggest that wage gains and losses are approximately symmetric, and that within-quartile mobility does not lead to notable wage gains. Neither observation is consistent with mobility being driven by match-specific wage components, since in that case downward

[2]Since job changes refer to changes in the main occupation which I define as the job that generates the maximum income within a year, I reduce the problem of long gaps in employment between two seemingly consecutive jobs. However, some job-to-job transitions may be interrupted by periods of unemployment or short-term employment at other employers. In general, the restriction to the main occupation introduces some ambiguity as to the exact timing of the transition.

and upward moves would be "more positive", leading to asymmetric transition profiles. To assess this more formally, Tables J.5 and J.6 offer a more detailed analysis of the wage profiles of men (panel A) and women (panel B). The entries in columns 1 and 2 contain the absolute and relative frequencies of moves between job quartiles. In total, the analysis in 2001-2008 is based on 262,374 job transitions of male workers and 116,863 job transitions of female workers. These transitions occur between 165,565 unique origin-destination firm combinations (not reported). Around 20-30% of all job transitions fall into each origin quartile with a minor exception being the 1st quartile among women from which 34.0% of all job transitions originate. Columns 3-6 show the mean wage profiles of movers from each gender for all origin-destination quartiles. The entries in column 7 report the 3-year log point change in wages calculated by subtracting column 6 from column 3 and multiplying by 100. In column 8, I perform a simple trend-adjustment for each origin quartile by subtracting off the 3-year raw wage change of workers who also change jobs but stay within the same job quartile. The entries in this column allow for an interpretation as the excess wage change that is associated with inter-quartile mobility beyond the return to mobility itself and conditional on heterogeneity between stayers and movers (since only movers are contained in the analysis). Conditional on this trend, it holds that wage changes associated with upward mobility are uniformly positive in sign, whereas changes associated with downward mobility are uniformly negative in sign and approximately symmetric. This holds for every origin-quartile, for every gender, and in each period.

Movers between different coworker quartiles may differ in terms of human capital. For example, young workers who just recently entered the labour market are more likely to be observed in lower coworker wage quartiles, and, depending on their education, are also more likely to transit to higher coworker quartiles later on. Such heterogeneity not only results in different initial wage levels within the same origin quartiles, but may also lead to seemingly asymmetric wage changes associated with opposite transitions. Similarly, general (economy-wide) wage gains and losses might affect the wage profiles over time. These possibilities do not pose a violation of the exogenous mobility assumption as long as the AKM models include appropriate covariates to account for such heterogeneity (X-covariates).

A cleaner picture of the wage profiles is presented in column 9, where I show regression-adjusted 3-year wage changes that account for observed heterogeneity in education and age/experience between movers and stayers. Following CCK, I regress, for each year, the 3-year wage change of stayers on dummies for the base year, education and a quadratic polynomial in age fully interacted with education dummies. I fit these models separately to the sample of stayers from each gender, and use the estimated parameters predict the associated wage change of movers with similar human capital characteristics. Finally, I subtract the predicted wage change of movers from the actual change to arrive at an excess wage growth associated with mobility between coworker quartiles. In column 10, I additionally report the corresponding two-way clustered standard errors (workers/establishments/combined) using the method of Kline (2014) which accounts for sampling errors in the regression adjustment. Overall, I find that these "cleaned" wage changes associated with *inter*-quartile mobility are sizeable and almost uniformly different from zero.

In Figure J.7, I plot mean adjusted 3-year wage changes associated with downward transitions against the corresponding upward transition for each gender. The negatively sloped 45-degree line indicates perfectly symmetric wage changes of opposing pairs of downward- and upward-mobility. Although the points do not exactly lie on a line with slope -1, they are clustered reasonably close around it, especially since this analysis is based on quartiles which leaves scope for considerable distributional effects. For example, wage changes of workers moving from 99th percentile in the 2nd quartile to the 1st percentile in the 3rd quartile are smaller (in absolute values) than wage changes of workers moving from the 99th percentile in the 3rd quartile to the 1st percentile in the 2nd quartile. If the probability of upward and downward moves depends on the location within each quartile, then this might lead to deviations from symmetry though this does not immediately imply a violation of the AKM assumption. I checked for such distributional differences by inspecting transition profiles for different percentiles (25th, 50th, 75th) within each quartile. The basic patterns were very similar.

Figures J.8 and J.9 show the regression-adjusted mean wage changes experienced by movers who move between different origin-destination quartiles relative to the wage change experienced by workers who also move but stay within the same quartile. The construction is identical to the calculations used to obtain the entries in column 8, only that I now use regression-adjusted wages. The results again confirm the proximate symmetry of gains and losses from moving between quartiles of coworker wages.

E Model Fit and Additive Separability

In this section, I provide additional model fit analyses of the AKM models described in Table 2 of the main text. The key issue here is whether the data show signs that the additive separability assumption is not satisfied. In an initial exercise, I compute the mean AKM residuals across the joint distribution of worker- and establishment effects. The results are provided in Figures J.11 and J.12 for the two periods and genders, respectively. In general, the mean residuals are small, reflecting the relatively good fit of the wage model to the overall wage distribution (see Table 2). The largest deviation is -0.0084 for the 1-1 interaction among men in the first period. The fact that the AKM models perform relatively poorly for the lowest ability workers at firms with the lowest premiums has also been documented in CHK, CCK, and Macis and Schivardi (2016). In a second exercise, I calculate the event study using quartiles of estimated establishment effects rather than coworker wages.I find very similar patterns as those shown in Figures J.5 and J.6, suggesting that the wage model captures the mobility process and associated wage changes reasonably well. The results are available upon request.

F Trends in Average Hours

Although my analysis focuses on full-time workers, it is still possible that variation in average working hours between men and women might overstate the gender gap. My conclusions regarding comparisons across periods are unaffected by this, unless male and female hours (of full-time workers) follow differential trends. However, an hours bias might confound the gender gap when considering each period separately. To probe into this issue, Figure J.14 illustrates the ratio of female to male usual weekly working hours calculated from OECD data. As shown by the top line (excluding part-time), the hours gap is very small (-0.9 to -1.3%), and virtually flat over the sample period. Of course, this is only suggestive since hours differences may be particularly pronounced in terms of overtime and unusual working hours (Goldin, 2014). For comparison, I also show the evolution of relative hours including part-time. In this case, the relative hours gap declines from -20% to -30% between 1995 and 2008. Even though I cannot fully rule out the presence of an hours bias for full-time workers, these figures seem to suggest that it is relatively small.

G Gender-Pooled Firm Premiums

The stagnation of the overall gender wage gap on the one hand, and the rise in gender inequality due to firm-specific wages on the other hand, suggests that the wage components attributable to $X\beta$ and α

contributed to a narrowing of gender wage differentials. To gauge the contribution of these components, I apply Gelbach's decomposition (Gelbach, 2016; Cardoso et al., 2016). A challenge in identifying the contributions of each component of the AKM model to the gender wage gap is that one cannot include a gender dummy in the estimations directly since it is absorbed by the person effect. Drawing on Cardoso et al. (2016), I provide a detailed derivation of how one can still identify the contributions to the gender gap based on this decomposition approach in section I.2.

The results of estimating Gelbach's decomposition are summarised in Table J.13, showing in even-numbered columns the estimates for the dual-connected set. Row 1 reports the unconditional gender gap, and rows 2-4 show the partition into human capital ($X\beta$), worker effects (α), and firm wage premiums (ψ). Comparing the estimates across columns, I find that the contribution of firm premiums rises from 5.4 log points (21.9%) to 7.3 log points (29.4%), yielding a cross-period growth of 1.8 log points. This is about equal to the 2.0 log points increase shown in Table 3, but only half the magnitude of the overall rise when introducing a bargaining channel.[3] One interpretation of this finding is that omitting the bargaining component would lead one to understate the impact of firm rent differentials on the gender gap (3.6 log points) by around 50%, suggesting that it is important to account for group-specific firm premiums in analyses of between-group inequality. Given a widening of firm rent differentials coupled with a stagnation of overall gender inequality, it is natural to wonder whether it is worker quality or human capital that counterbalances the trend in firm gender gaps, or both. Rows 3 and 4 show that both wage components would have compressed the gender gap, with human capital accounting for 1/3 and worker quality for 2/3 of the total compression.

H The Impact of Children

H.1 Identification of Childbirth Events

In this section, I describe how I identify the event of childbirth from female employment biographies. A key advantage of my data set is the availability of raw employment biographies with daily accuracy covering the years 1993-2008. Combined with information on the reason of leave (reported by employers as part of the notification procedure), restrictions on mother's age at birth, the duration of leave, and the gap between consecutive births, this enables me to identify work interruptions that are likely due to childbirth. In particular, I require that a valid childbirth event must satisfy four main criteria: i) the reason of leave indicates a temporary work interruption or maternity protection; ii) mothers are between 18 and 40 years of age; iii) the duration of leave is not below 14 weeks (the statutory maternity protection period); and iv) the gap between consecutive births is not below 40 weeks (the usual pregnancy duration). Using these restrictions, I identify about 550,000 births, of which 432,000 births (78.5%) refer to the first child.

I restrict the analysis to the birth of the first child, so that I have at most one event per year, allowing me to match the childbirth events at the person-year level with my worker sample. Doing so, however, comes along with two issues important for the interpretation. First, many women do not return to full-time employment after childbirth, while my analysis is based on full-time workers only. This might generate a selection bias, though it is not very clear in which direction this would show. For example, those returning immediately to full-time employment might be the needy who must earn money to provide their maintenance; but they could also be the privileged with high career ambitions. Second,

[3] The firm premium component in Table J.13 is a weighted average of the sorting components shown in Table 3 and Table J.7.

the event of childbirth may occur throughout a year while my analysis is based on unique person-year combinations by focusing on the main job. The main job, however, may refer to the pre-birth or the post-birth period, depending on whether the birth event occurs rather early or rather late in a year. In brief, this means that the immediate (t=0) impact of childbirth on wages might confound pre- and post-birth wages, and the medium run effect (2-3 years after childbirth) may be a more representative statistic for the wage impact of children. In fact, this observation could also be one reason why the impact of children declines gradually over the first couple of years after the event date, though I presume that this pattern is dominated by the fact many mothers only return to full-time employment several years after they gave birth.

Keeping these caveats in mind, I am able to match some 196,000 births to the dual-connected sample of full-time women in their main job between 1995 and 2008. Note that the sample of "actual mothers" in the analysis is somewhat larger, consisting of about 270 thousand women (see section 4.6). This is because I include in this group also those women who gave birth at ages 18 and 19 (my analysis is restricted to ages 20-60), or in years 1993 and 1994 (my analysis is based on 1995-2008), or women who return to full-time years after childbirth. In these cases, where I know that a women is a mother and also know the event year, but where I do not observe the corresponding person-year combination in my main analysis sample (full-time, age 20-60, 1995-2008, dual-connected), I still include them in the regression to estimate the long run effects. This is because I consider potential wage penalties due to work interruptions and temporary hours reductions as part of the long run effect of childbirth.

H.2 Identifying Placebo Mothers

I follow the approach described in Appendix A of Kleven et al. (2017) for identifying a control group of non-mothers, and assigning them a placebo childbirth date. However, because my sample differs in several dimensions from the data used in Kleven et al. (2017), I have to make some adjustments which I explain in the following.

One main problem is that the data do not provide direct information on the number of children, i.e., whether a women is a mother or not. As described above, I assume that the maximum childbearing age is 40, so in a first step, I select women who were born between 1953 and 1968, i.e., who are not yet 40 by 1993 and who will have turned 40 by 2008. For some of these women, I have observed a childbirth, and for others I have not. I call the last group "never mothers", though it is conceivable that they have given birth prior to 1993, which I cannot identify using the restricted employment biographies at hand. The remaining group of mothers, which does not turn 40 until 2008, might be giving birth after the terminal year of the sample (right-censoring). To include these women in my analysis, I estimate a linear probability model for the birth cohorts 1953-1968, with the dependent variable equal to one for a never mother. I include the following list of controls: quartiles of female firm premiums and quartiles of wages (constructed for each birth cohort), education dummies, and dummies for more than 300 districts. I assign each woman of the birth cohorts 1969-1989 a predicted probability of never giving birth, and, starting from the highest predicted probability, add as many women into the control group so that the fraction of never mothers in the birth cohorts 1969-1989 is equal to the fraction of never mothers in the birth cohorts 1953-1968.

Altogether, my analysis sample thus consists of three types of women: one group that comprises of "actual mother", and two groups considered as "never mothers". The latter two groups are relevant for the identification of placebo mothers. The first of them consists of mothers born between 1953 (1993-40) and 1968 (2008-40), i.e., whose childbearing age is not right-censored (they turn 40 before 2008). This group embeds women who are actual "never mothers", plus an unknown number of women that may have given birth before 1993, hence, for whom I am unable to observe the birth date. As I cannot separately identify the latter, I keep them in the control group, noting that this might bias the childbirth penalty towards zero (their wage is already reduced due to childbirth in the past). The second group indicated above consists of women who also do not give birth between 1993 and 2008, but who only turn 40 after 2008, i.e., whose childbearing age is right-censored (and left-censored for birth cohorts 1969-1974, i.e., women who are older than 18 before 1993). For this group, which in total spans the age cohorts 1969-1989, I use the procedure described above to identify a subgroup of women who are most unlikely to give birth later in life.

For the two groups of never mothers, I then assign placebo births as in Kleven et al. (2017), using a random draw of a log normal distribution of age at first birth with location and scale parameters obtained from the associated distribution of actual mothers.

I Technical Issues

I.1 Decompositions in Terms of Union Coverage

In this section, I derive equations (4) and (5), and show how one can estimate the contribution of gender differences in deunionisation to the male-female wage gaps. For consistency with eq. (2), I use notation in terms of conditional expectations instead of probability density functions. Moreover, I maintain the convention of using male employment and female returns to construct the counterfactuals.

Denote by ψ_p^g the firm premium of gender $g = \{M, F\}$ in period $p = \{1995, 2008\}$, and let U_p denote the distribution of union coverage in period p. In the following exposition, the period-subscript on the firm premiums (1995, 2008) is only added to clarify whether I refer to the first or second period. In the empirical implementation, I use the estimated parameters from the AKM models described in Table 2, i.e., based on 1995-2001 (here: 1995) and 2001-2008 (here: 2008).

The period p gender gap in firm premiums (FGG) can be represented as follows:

$$FGG_p = E\left[\psi_p^M | M_p, U_p\right] - E\left[\psi_p^F | F_p, U_p\right] \tag{7}$$

To simplify, I assume that U is binary (though, empirically, I distinguish between three types of union coverage: industry level, firm level, no coverage). Then, one can rewrite the eq. (7) as follows:

$$\begin{aligned}
&\underbrace{E\left[\psi_p^M | M_p, U_p = 0\right] - E\left[\psi_p^F | M_p, U_p = 0\right]}_{FGG_p^{nu}} \\
&+E\left[\psi_p^M | M_p, U_p = 0\right] + \underbrace{Pr(U_p = 1 | M_p)}_{s_{M,p}^u} \times \underbrace{\left\{E\left[\psi_p^M | M_p, U_p = 1\right] - E\left[\psi_p^M | M_p, U_p = 0\right]\right\}}_{UP_p^M} \\
&-\left(E\left[\psi_p^F | F_p, U_p = 0\right] + \underbrace{Pr(U_p = 1 | F_p)}_{s_{F,p}^u} \times \underbrace{\left\{E\left[\psi_p^F | F_p, U_p = 1\right] - E\left[\psi_p^F | F_p, U_p = 0\right]\right\}}_{UP_p^F}\right)
\end{aligned} \tag{8}$$

Adding and subtracting $s_{M,p}^u \times UP_p^F$ yields:

$$FGG_p = FGG_p^{nu} + \left(s_{M,p}^u - s_{F,p}^u\right) \times UP_p^F + s_{M,p}^u \times \left(UP_p^M - UP_p^F\right) \tag{9}$$

where $UP_p^M - UP_p^F = FGG_p^u - FGG_p^{nu}$, which corresponds to eq. (4) in the main text. Next, evaluate p in 1995 and 2008 and denote the difference between 2008 and 1995 by Δ:

$$\begin{aligned}\Delta FGG =& \Delta FGG^{nu} + \left(s^u_{M,2008} - s^u_{F,2008}\right) \times UP^F_{2008} - \left(s^u_{M,1995} - s^u_{F,1995}\right) \times UP^F_{1995} \\ & + s^u_{M,2008} \times \left(UP^M_{2008} - UP^F_{2008}\right) - s^u_{M,1995} \times \left(UP^M_{1995} - UP^F_{1995}\right)\end{aligned} \tag{10}$$

Adding and subtracting the terms $\left(s^u_{M,1995} - s^u_{F,1995}\right) \times UP^F_{2008}$ and $s^u_{M,2008} \times \left(UP^M_{1995} - UP^F_{1995}\right)$ gives

$$\begin{aligned}\Delta FGG =& \Delta FGG^{nu} + \left(s^u_{M,1995} - s^u_{F,1995}\right) \times \Delta UP^F + \left(UP^M_{1995} - UP^F_{1995}\right) \Delta s^u_M \\ & + UP^F_{2008} \Delta \left(s^u_M - s^u_F\right) + s^u_{M,2008} \Delta \left(UP^M - UP^F\right)\end{aligned} \tag{11}$$

where, as before, $UP^M - UP^F = FGG^u - FGG^{nu}$.

To derive the contribution of the gender bias, I begin by evaluating eq. (7) in 1995 and 2008, and subtracting the latter from the former. Rearranging yields:

$$\begin{aligned}\Delta FGG =& E\left[\psi^M_{2008} | M_{2008}, U_{2008}\right] - E\left[\psi^M_{1995} | M_{1995}, U_{1995}\right] \\ & - \left(E\left[\psi^F_{2008} | F_{2008}, U_{2008}\right] - E\left[\psi^F_{1995} | F_{1995}, U_{1995}\right]\right) \\ =& E\left[\psi^M_{2008} - \psi^M_{1995} | M_{1995}, U_{1995}\right] + E\left[\psi^M_{2008} | M_{2008}, U_{2008}\right] - E\left[\psi^M_{2008} | M_{1995}, U_{1995}\right] \\ & - \left(E\left[\psi^F_{2008} - \psi^F_{1995} | F_{1995}, U_{1995}\right] + E\left[\psi^F_{2008} | F_{2008}, U_{2008}\right] - E\left[\psi^F_{2008} | F_{1995}, U_{1995}\right]\right)\end{aligned} \tag{12}$$

Let $\Delta^1 = E\left[\psi^M_{2008} - \psi^M_{1995} | M_{1995}, U_{1995}\right] - E\left[\psi^F_{2008} - \psi^F_{1995} | F_{1995}, U_{1995}\right]$, then rearranging gives:

$$\begin{aligned}& \Delta^1 + E\left[\psi^M_{2008} | M_{2008}, U_{2008}\right] - E\left[\psi^F_{2008} | F_{2008}, U_{2008}\right] \\ & - \left(E\left[\psi^M_{2008} | M_{1995}, U_{1995}\right] - E\left[\psi^F_{2008} | F_{1995}, U_{1995}\right]\right)\end{aligned} \tag{13}$$

In the next step, I assign women the male unionisation rate in each period:

$$\begin{aligned}& \Delta^1 + E\left[\psi^M_{2008} - \psi^F_{2008} | M_{2008}, U_{2008}\right] + E\left[\psi^F_{2008} | M_{2008}, U_{2008}\right] - E\left[\psi^F_{2008} | F_{2008}, U_{2008}\right] \\ & - \left(E\left[\psi^M_{2008} - \psi^F_{2008} | M_{1995}, U_{1995}\right] + E\left[\psi^F_{2008} | M_{1995}, U_{1995}\right] - E\left[\psi^F_{2008} | F_{1995}, U_{1995}\right]\right)\end{aligned} \tag{14}$$

Denoting $\Delta^2 = E\left[\psi^M_{2008} - \psi^F_{2008} | M_{2008}, U_{2008}\right] - \left(E\left[\psi^M_{2008} - \psi^F_{2008} | M_{1995}, U_{1995}\right]\right)$, the final expression reads:

$$\begin{aligned}\Delta FGG =& \Delta^1 + \Delta^2 + \left(E\left[\psi^F_{2008} | M_{2008}, U_{2008}\right] - E\left[\psi^F_{2008} | F_{2008}, U_{2008}\right]\right) \\ & - \left(E\left[\psi^F_{2008} | M_{1995}, U_{1995}\right] - E\left[\psi^F_{2008} | F_{1995}, U_{1995}\right]\right)\end{aligned} \tag{15}$$

Overall, the decomposition consists of three parts. Δ^1 measures the contribution of the change in the firm gender gap ΔFGG that is due to changes in the firm-specific wage structure of men and women between 1995 and 2008, weighted by 1995 employment/unionisation rates. For example, if women experienced a relatively greater wage growth than men (unrelated to unionisation), then this would reduce the overall change in the gender wage gap by an amount of Δ^1. The second component, Δ^2, measures the contribution of shifts in male deunionisation over time, weighted by the wage structure difference of men and women in 2008.

The third part measures the contribution of gender differentials in deunionisation rates. The first term of this component, $E\left[\psi^F_{2008}|M_{2008}, U_{2008}\right] - E\left[\psi^F_{2008}|F_{2008}, U_{2008}\right]$, can be obtained by reweighting the female distribution with male characteristics in 2008. The empirical counterpart of this is shown in columns 4-6 of Table 7, with column 5 referring to the counterfactual distribution of women if they were given male union coverage rates. The second part, however, consists of one element that would require to reweight female firm premiums in 2008 with male unionisation rates in 1995. To avoid this, I approximate ψ^F_{2008} by ψ^F_{1995}. Since the distribution of female firm premiums in 1995 differs from that in 2008, this step implies that I potentially blend changes in female returns with changes in union coverage rates. Based on the results in section 4.3, I suspect that weighting by ψ^F_{2008} generates a larger firm premium gap, i.e., $E\left[\psi^F_{2008}|M_{1995}, U_{1995}\right] - E\left[\psi^F_{2008}|F_{1995}, U_{1995}\right] > E\left[\psi^F_{1995}|M_{1995}, U_{1995}\right] - E\left[\psi^F_{1995}|F_{1995}, U_{1995}\right]$. Hence, the approximation might overstate the contribution of the gender bias in deunionisation. However, for the sake of empirical tractability, I believe that this step is legitimate. The empirical counterpart of this second term is reported in columns 1-3 of Table 7.

I.2 Gelbach's Decomposition

In section G, I apply Gelbach's decomposition approach to gauge the separate contributions of worker, firm, and human capital components to the gender wage gap. In the following, I illustrate the working of this decomposition approach, and derive the final expressions. The exposition is inspired by Gelbach (2016) and Cardoso et al. (2016). As a point of departure, consider the following baseline (b) model:

$$w_{it} = X'_{it}\beta_b + W'_i\theta_b + r_{it} \tag{16}$$

where θ_b is a $k^c_b \times 1$ vector of returns to time-constant (c) worker characteristics (e.g., gender) stored in a $k^c_b \times 1$ vector W_i, and β_b is a $k^v_b \times 1$ vector of time-varying (v) worker observables (e.g., age interacted with education) contained in a $k^v_b \times 1$ vector $X_{i,t}$. For the exposition, I assume that $k^c_b = 1$ with gender, G_i, being the only time-constant variable of interest. However, the following statements generalise to $k^c_b > 1$. The assumptions so far, yield the following wage equation:

$$w_{i,t} = X'_{i,t}\beta_b + G_i\theta_b + r_{i,t} \tag{17}$$

In matrix notation, the stacked system can be written as:

$$w = \underset{(N^*\times k^v_b)}{X}\ \underset{(k^v_b\times 1)}{\beta_b} + \underset{(N^*\times 1)}{G}\ \underset{(1\times 1)}{\theta_b} + \underset{N^*\times 1}{r} \tag{18}$$

I incorporate G into X, and denote the resulting $N^* \times (k^v_b + 1)$ covariate index as $\tilde{X} = [X, G]$ and the associated $(k^v_b + 1) \times 1$ parameter vector as $\tilde{\beta}_b = [\beta_b, \theta_b]'$. Note that this model embeds the special case of $\tilde{X} = G$ when $k^v_b = 0$, which is the setting used in the main text. In this case, $\hat{\tilde{\beta}}_b = \hat{\theta}_b = (G'G)^{-1} G'w$ measures the unconditional gender gap, where $M_g = (G'G)^{-1} G'$ is a matrix that, if left-multiplied to a conformable vector of covariates, computes the mean difference between men and women. In general, if $k_b^v > 0$, eq. (18) can be reformulated as follows:

$$w = \underset{(N^*\times k_b)(k_b\times 1)}{\tilde{X} \quad \tilde{\beta}_b} + \underset{N^*\times 1}{r} \quad \text{where } \hat{\tilde{\beta}}_b = \underbrace{\left(\tilde{X}'\tilde{X}\right)^{-1}\tilde{X}'}_{M_{\tilde{x}}} w \tag{19}$$

Here, I use $k_b = k_b^v + k_b^c = k_b^v + 1$ to denote the total number of parameters in the baseline model. I specify the full model consistent with eq. (1), rewritten in matrix notation, but omit the gender-specific parameters, and instead run pooled regressions on

$$w = \underset{(N^*\times N)(N\times 1)}{D \quad \alpha} + \underset{(N^*\times J)(J\times 1)}{F \quad \psi} + \underset{(N^*\times k_f)(k_f\times 1)}{X \quad \beta_f} + \underset{(N^*\times 1)}{r} \tag{20}$$

where D is a $N^* \times N$ design matrix of indicators for every worker, F is a $N^* \times J$ design matrix of indicators for every establishment, and X is a $N^* \times k_f$ matrix of time-varying observables ($k_f = k_f^v$) that may (or may not) be included in eq. (18). This means, in particular, that the number of time-varying covariates may differ between the full and the baseline model. In this case, the full model contains a (weakly) larger number of covariates, i.e., $k_b^v \leq k_f^v$. β_f is an associated vector of pooled returns, and r is a mean zero error term. Note that in eq. (20), any time-invariant worker level variable — such as gender — is absorbed in the worker effect, and similarly, any time-invariant firm level variable — such as industry — is absorbed in the firm effect. Hence, one cannot include such covariates directly in the estimation.

Consider a case when $k_f^v = k_b^v > 0$ and $k_b^c = 1$, so that the baseline model and the full model both contain a positive number of time-varying observables, and the baseline model additionally includes a single time-invariant variable (gender). The 'twist' of Gelbach's decomposition is to interpret $\hat{\tilde{\beta}}_b$ in the baseline model as a biased estimator of the covariate returns on $\tilde{X}$. To illustrate, consider the fitted regression of eq. (20):

$$w = \underset{(N^*\times N)(N\times 1)}{D \quad \hat{\alpha}} + \underset{(N^*\times J)(J\times 1)}{F \quad \hat{\psi}} + \underset{(N^*\times k_f)(k_f\times 1)}{X \quad \hat{\beta}_f} + \underset{(N^*\times 1)}{\hat{r}} \tag{21}$$

Since eq. (20) does not contain time-invariant worker-level variables — as these parameters are absorbed in the person-fixed effects — I expand the X matrix of the full model by the time-invariant worker-level variables contained in eq. (18), and add a conformable number of zero rows to $\hat{\beta}_f$. Denote the corresponding components as $\tilde{X}$ and $\hat{\tilde{\beta}}_f$. Then, I left-multiply by $M_{\tilde{x}}$ to obtain:

$$\hat{\tilde{\beta}}_b = \underset{(k_b\times N^*)(N^*\times 1)}{M_{\tilde{x}} \quad D\hat{\alpha}} + \underset{(k_b\times N^*)(N^*\times 1)}{M_{\tilde{x}} \quad F\hat{\psi}} + \underset{(k_b\times N^*)(N^*\times 1)}{M_{\tilde{x}} \quad \tilde{X}\hat{\tilde{\beta}}_f} \tag{22}$$

where I use the orthogonality of $M_{\tilde{x}}$ with respect to r. Rearranging yields

$$\underset{k_b\times 1}{\hat{\tilde{\beta}}_b} - \underset{k_b\times 1}{\hat{\tilde{\beta}}_f} = \underset{k_b\times 1}{\hat{\delta}_{\alpha\tilde{X}}} + \underset{k_b\times 1}{\hat{\delta}_{\psi\tilde{X}}} \tag{23}$$

Each term on the right hand side of eq. (23) is a regression of the least-squares estimates of the full model on the covariate index $\tilde{X}$. Now, focusing on the row that contains the coefficient on G in each regression, one obtains:

$$\underset{1\times1}{\hat{\theta}_b} - \underset{1\times1}{\hat{\theta}_f} = \underset{1\times1}{\hat{\delta}_{\alpha G}} + \underset{1\times1}{\hat{\delta}_{\psi G}} \tag{24}$$

This means that the difference of the conditional gender wage gap between the baseline model and the full model can be decomposed into the sum of coefficients on the gender dummy obtained from two auxiliary regressions of the worker and firm components on $\tilde{X}$. Since $\hat{\theta}_f = 0$, this result yields an unambiguous partition of the conditional gender wage gap into worker and firm components.

In the main text, the baseline model only includes the gender dummy (and no time-varying variables) because I want to decompose the unconditional gender gap (rather than the covariate-adjusted gender gap). Therefore, consider now a baseline specification without time-varying variables, i.e., $k_b^v = 0$. Then a comparison of othe baseline and the full model additionally requires to run a regression of the fitted time-varying covariates of the full model, $X\hat{\beta}_f$, on the time-invariant component(s) contained in $\tilde{X}$. As G is the only time-invariant variable in the baseline specification, replacing $M_{\tilde{x}}$ in eq. (22) by M_g yields:

$$\hat{\theta}_b = \underset{(1\times N^*)(N^*\times1)}{M_g \quad D\hat{\alpha}} + \underset{(1\times N^*)(N^*\times1)}{M_g \quad F\hat{\psi}} + \underset{(1\times N^*)(N^*\times1)}{M_g \quad X\hat{\beta}_f} \tag{25}$$

In principle, the last term of eq. (25) specifies as many individual regressions as there are time-varying covariates in X. The total contribution of observable characteristics to the gender wage gap would then be the sum of these coefficients. However, as shown in Gelbach (2016), one can first create a heterogeneity term by summing over all covariates at the worker-year level, and then regress the resulting compound heterogeneity index on the gender dummy. This is the procedure that I implement in the main text. Using the notation from eq. (23), one obtains the final result:

$$\underset{1\times1}{\hat{\theta}_b} = \underset{1\times1}{\hat{\delta}_{\alpha G}} + \underset{1\times1}{\hat{\delta}_{\psi G}} + \underset{1\times1}{\hat{\delta}_{\beta G}} \tag{26}$$

References

Abowd, John M., Francis Kramarz, and David N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–333.

Burdett, Kenneth, and Dale T. Mortensen (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review 39*(2), 257–273.

Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline (2016). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal for Labor Economics* (Forthcoming), 51.

Card, David, Ana Rute Cardoso, and Patrick Kline (2016). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *The Quarterly Journal of Economics* (Forthcoming).

Card, David, Jörg Heining, and Patrick Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics 128*(3), 967–1015.

Cardoso, Ana Rute, Paulo Guimarães, and Pedro Portugal (2016). What drives the gender wage gap? A look at the role of firm and job-title heterogeneity. *Oxford Economic Papers 1*, 1–19.

Destatis (2016). Registrierte Arbeitslose, Arbeitslosenquote nach Geschlecht. (https://www.destatis.de/DE/ZahlenFakten/Indikatoren/LangeReihen/Arbeitsmarkt/lrarb002.html).

Eeckhout, Jan, and Philipp Kircher (2011). Identifying Sorting–In Theory. *Review of Economic Studies 78*(3), 872–906.

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 126*(3), 405–436.

Foster, Lucia, John Haltiwanger, and Chad Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review 98*(1), 394–425.

Frick, Bernd (2002). Mandated Codetermination, Volunary Profit Sharing and Firm Performance. *Unpublished Manuscript*.

Gelbach, Jonah B. (2016). When Do Covariates Matter? And Which Ones, and How Much? *Journal for Labor Economics* (Forthcoming).

Goldin, Claudia (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review 104*(4), 1091–1119.

Goldschmidt, Deborah, and Johannes F. Schmieder (2017). The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure. *Quarterly Journal of Economics* (Forthcoming).

Gürtzgen, Nicole (2009). Rent-sharing and collective bargaining coverage: evidence from linked employer-employee data. *The Scandinavian Journal of Economics 111*(2), 26.

Gürtzgen, Nicole (2012). Estimating the Wage Premium of Collective Wage Contracts - Evidence from Longitudinal Lined Employer-Employee Data. *ZEW Discussion Paper 12*(73), 43.

Hagedorn, Marcus, Tzuo-Hann Law, and Iourii Manovskii (2016). Identifying Equilibrium Models of Labor Market Sorting. *Econometrica* (Forthcoming).

Heining, Jörg, Peter Jacobebbinghaus, Theresa Scholz, and Stefan Seth (2012). Linked-Employer-Employee-Daten des IAB: LIAB-Mover-Modell 1993-2008 (LIAB MM 9308). *FDZ-Datenreport 01/2012*.

Heining, Jörg, Wolfram Klosterhuber, and Stefan Seth (2014). An Overview on the Linked Employer-Employee Data of the Institute for Employment Research (IAB). *Schmollers Jahrbuch 134*, 141–148.

Kleven, Henrik J, Camille Landais, and Søgaard, Jacob E. (2017). Children and gender inequality: Evidence from Denmark. *Working Paper*, 1–56.

Kline, Patrick (2014). A note on variance estimation for the Oaxaca estimator of average treatment effects. *Economics Letters 122*, 428–431.

Macis, Mario, and Fabiano Schivardi (2016). Exports and Wages: Rent Sharing, Workforce Composition or Returns to Skills? *Journal of Labor Economics* (Forthcoming).

Wolf, Elke, and Thomas Zwick (2002). Reassessing the impact of high performance workplaces. *ZEW Discussion Paper 7*, 1–36.

J Appendix Tables and Figures

Table J.1: Descriptive Statistics of Firms Grouped by Mean log Value Added per Worker

	1995-2001						2001-2008					
	Overall Worker Sample		Sample connected with EP-estab's		Sample with 1+ years of VA		Overall Worker Sample		Sample connected with EP-estab's		Sample with 1+ years of VA	
	Male	Female	Male	Female	Male	Female	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Observations												
Person-years	10,586,204	4,657,423	7,762,277	3,334,606	2,360,729	723,285	11,985,517	5,143,438	8,676,703	3,617,634	3,770,299	1,140,803
Percent of overall sample	*100.0*	*100.0*	*73.3*	*71.6*	*22.3*	*15.5*	*100.0*	*100.0*	*72.4*	*70.3*	*31.5*	*22.2*
Workers	1,858,422	973,038	1,552,007	764,642	517,656	180,425	1,884,818	972,199	1,560,252	750,394	707,799	247,884
Percent of overall sample	*100.0*	*100.0*	*83.5*	*78.6*	*27.9*	*18.5*	*100.0*	*100.0*	*82.8*	*77.2*	*37.6*	*25.5*
Establishments	389,253	257,083	18,247	17,397	4,731	4,485	382,775	259,279	18,777	17,936	6,930	6,578
Percent of overall sample	*100.0*	*100.0*	*4.7*	*6.8*	*1.2*	*1.7*	*100.0*	*100.0*	*4.9*	*6.9*	*1.8*	*2.5*
Number of job-spells per worker	1.60	1.47	1.12	1.09	1.03	1.03	1.61	1.49	1.10	1.08	1.04	1.03
Worker characteristics												
Mean age	39.1	37.4	40.2	38.5	40.1	38.4	41.5	39.9	42.3	40.8	41.8	40.5
Mean tenure	8.4	7.2	10.0	8.6	9.4	7.9	9.4	8.0	11.6	10.1	11.1	9.4
Share low education	0.14	0.18	0.15	0.20	0.17	0.28	0.12	0.15	0.13	0.16	0.14	0.21
Share medium education	0.67	0.63	0.67	0.63	0.69	0.58	0.65	0.60	0.66	0.61	0.68	0.59
Share high education	0.19	0.18	0.18	0.17	0.14	0.14	0.23	0.25	0.22	0.23	0.18	0.20
Wages												
Mean of log daily wage	4.55	4.30	4.59	4.35	4.56	4.28	4.57	4.32	4.62	4.38	4.59	4.31
Std. dev. of log daily wage	(0.373)	(0.395)	(0.337)	(0.352)	(0.358)	(0.367)	(0.419)	(0.449)	(0.372)	(0.400)	(0.369)	(0.423)
100 x Log gender wage gap	24.8		24.0		27.9		25.0		24.1		28.0	
Percent of overall sample	*100.0*		*96.8*		*112.5*		*100.0*		*96.4*		*112.0*	
Workplace characteristics												
Mean firm size (full-time)	1,403	729	1,788	888	862	654	1,343	654	1,745	829	2,038	1,012
Share of female coworkers	0.21	0.53	0.22	0.48	0.18	0.41	0.21	0.52	0.22	0.47	0.18	0.41
Share at all-male firms	0.10	–	0.006	–	0.01	–	0.09	–	0.005	–	0.01	–
Share at all-female firms	–	0.082	–	0.005	–	0.00	–	0.080	–	0.005	–	0.00
Mean log value added p.w.					4.36	4.32					4.43	4.37
SD of mean log value added p.w.					(0.60)	(0.64)					(0.64)	(0.68)

Note: Table shows person-year weighted summary statistics for 1995-2001 and 2001-2008. The overall worker sample is the baseline sample after basic sample processing described in the main text and the Appendix. The sample connected with EP-establishments contains all workers and establishments that participate in the IAB Establishment Panel. The sample with 1+ years of mean log value added contains all establishments that participate in the EP and for which I can calculate mean log value added from available information in at least one year. The sample with mean log sales per worker is not summarised separately. Value added and sales are measured in thousands of real year 2000 Euros, using an aggregate World Bank GDP-deflator. Per capita values are computed by dividing through the number of full-time employees, and to compute "mean" productivity measures, I average annual information at the firm level within a period across all firm-years, weighting each year by total employment (person-years).
Source: LIAB Mover Model 9308

Table J.2: Descriptive Statistics of Firms Grouped by Mean log Value Added per Worker

	1995-2001			2001-2008		
	Low VA	High VA	No VA	Low VA	High VA	No VA
	(1)	(2)	(3)	(4)	(5)	(6)
Mean firm size (full-time employees)	160	341	110	157	265	101
Mean share of female employment (in %)	37.8	30.1	42.1	34.2	32.2	41.5
Industry distribution (in %)						
Agriculture, hunting, forestry, fishing	2.1	0.7	0.6	2.2	1.0	0.5
Energy and mining	0.0	0.3	0.2	0.0	0.4	0.1
Production of rubber and plastic products	2.1	4.6	2.2	3.4	3.6	1.9
Chemical industry	0.5	3.0	1.4	0.4	2.4	1.3
Metal production and processing	10.8	16.8	6.1	4.5	13.7	5.6
Automotive, production of data processing	5.1	10.3	4.6	3.8	8.4	4.2
Consumer goods	10.8	13.6	8.5	11.3	11.1	6.5
Hospitality industry	9.2	2.7	4.2	9.0	3.1	4.4
Sales (retail/wholesale)	12.8	13.3	16.3	10.5	12.8	13.8
Maintenance, repair of motor vehicles	2.6	2.1	2.0	2.6	2.6	2.0
Building industry	4.1	5.9	3.7	3.8	4.7	2.6
Transport and communication	9.2	4.6	6.0	7.1	5.5	5.9
Credit and insurance intermediation	20.0	11.8	23.6	31.6	16.2	29.0
Public and personal services	4.6	3.8	4.7	2.6	4.4	4.9
Education, social, and health care	5.6	6.2	11.4	6.8	9.6	13.0
Public administration, social security	0.5	0.3	4.4	0.4	0.4	4.4
Firm Productivity						
Avg exc. log value added per worker (1,000)	–	1.17	–	–	1.26	–
Avg mean log value added per worker (1,000)	2.88	4.37	–	2.69	4.36	–
Mean log sales per worker (in 1,000)	4.51	5.13	5.15	4.34	5.14	5.12
Observations						
# of person-years	104,579	2,804,343	10,140,115	156,747	4,489,321	10,031,367
# of firm-years	1,230	20,981	420,826	1,834	34,011	479,099
# of firms	195	3,196	85,577	266	4,643	88,397

Note: Statistics are based on the dual-connected set. Each firm enters with weight 1. Columns 1-2 and 4-5 refer to firms for which value added is observed in the data. Firms are grouped into "low" and "high" according to the critical threshold value obtained from the non-linear system of equations (see main text). Entries in columns 3 and 6 show statistics for firms with no value added. Some of these establishments provide information on sales. Value added and sales are measured in thousands of real year 2000 Euros, using a World Bank GDP-deflator.
Source: LIAB Mover Model 9308

Table J.3: Summary Statistics of Samples Used in Stayer Analysis

	Men (1)	Women (2)
Worker characteristics		
3-year wage change	0.06	0.07
3- year lagged wage level	4.54	4.27
3-year lagged age	39.9	39.0
3-year lagged female employment share	0.17	0.36
Excess log value added per worker	1.29	1.27
Education shares (in percent)		
Low education (missing/primary)	17.9	29.0
Medium education (apprenticeship)	67.5	55.7
High education (some college/university)	14.6	15.4
Industry employment shares (in percent)		
Agriculture, hunting, forestry, fishing	0.3	0.2
Energy and mining	0.9	0.3
Production of rubber and plastic products	5.6	4.3
Chemical industry	6.1	7.6
Metal production and processing	32.1	19.6
Automotive, production of data processing	17.5	16.3
Consumer goods	12.2	15.8
Hospitality industry	0.2	0.6
Sales (retail/wholesale)	4.2	9.5
Maintenance, repair of motor vehicles	1.5	1.3
Building industry	6.3	2.7
Transport and communication	3.3	1.9
Credit and insurance intermediation	4.6	9.1
Public and personal services	4.0	3.0
Education, social, and health care	1.0	7.3
Public administration, social security	0.1	0.5
Observations		
Number of person-years	639,644	153,121
Number of establishments	2,371	2,371

Note: Table shows summary statistics of men and women who stay with their employer for at least three years between 1995 and 2008, and whose employer has non-missing excess log value added over the entire time frame. Value added is measured in thousands of real year 2000 Euros, using an aggregate World Bank GDP-deflator.
Source: LIAB Mover Model 9308

Table J.4: Summary of Estimation Results for Additional AKM Models

	Gender-Pooled				Period-Pooled			
	1995-2001		2001-2008		Men		Women	
	(1)		(2)		(3)		(4)	
Estimation results								
# of person-effects	2,744,864		2,769,838		2,070,621		1,128,590	
# of firm-effects	436,807		426,126		542,993		331,429	
Standard dev. of person-effects (a)	0.290		0.318		0.275		0.285	
Standard dev. of firm-effects (y)	0.160		0.199		0.181		0.222	
Standard dev. of covariates (Xb)	0.095		0.077		0.110		0.087	
Correlation of person- and firm-effects	0.101		0.135		0.123		0.043	
Correlation of male and female firm-effects					0.713			
Model fit								
Adjusted R2 of AKM model	0.911		0.919		0.901		0.858	
Root MSE of AKM model	0.103		0.112		0.117		0.144	
Fit of match effects model								
Adjusted R2 of match-effects model	0.930		0.935		0.929		0.889	
Root MSE of match-effects model	0.104		0.111		0.106		0.139	
Standard dev. of match-effects	0.055		0.059		0.069		0.077	
Percent of residual variance of AKM	*(28.7)*		*(27.4)*		*(34.8)*		*(28.4)*	
Variance decomposition								
Var(w)	0.152	*100*	0.191	*100*	0.158	*100*	0.174	*100*
Var(α)	0.084	*55*	0.101	*53*	0.076	*48*	0.081	*47*
Var(ψ)	0.026	*17*	0.040	*21*	0.033	*21*	0.049	*28*
Var($X\beta$)	0.009	*6*	0.006	*3*	0.012	*8*	0.008	*4*
2 × Cov(α, ψ)	0.009	*6*	0.017	*9*	0.012	*8*	0.005	*3*
2 × Cov(α, $X\beta$)	0.010	*7*	0.011	*6*	0.008	*5*	0.008	*4*
2 × Cov(ψ, $X\beta$)	0.003	*2*	0.004	*2*	0.003	*2*	0.002	*1*
Var(r)	0.011	*7*	0.013	*7*	0.014	*9*	0.021	*12*

Note: Table shows results from OLS estimations of AKM models specified in eq. (1) for gender-pooled samples (columns 1-2), and for period-pooled samples (columns 3-4). All models are estimated on the largest connected sets as defined in the main text. The correlation of male and female firm effects is calculated for the subset of dual-connected firms. Match effect models contain a dummy for each job-match. The correlations of fixed effects between the pooled (P) and the gender-specific (M, F) models, summarised in Table 2, are as follows. For the 1990s: Corr(ψ^M,ψ^P) = 0.9488, Corr(ψ^F,ψ^P) = 0.8139, Corr(α^M,α^P) = 0.9740, Corr(α^F,α^P) = 0.9362; for the 2000s: Corr(ψ^M,ψ^P) = 0.9646, Corr(ψ^F,ψ^P) = 0.8572, Corr(α^M,α^P) = 0.9820, Corr(α^F,α^P) = 0.9432
Source: LIAB Mover Model 9308

Table J.5: Mean Log Real Daily Wages of Movers Between Origin- and Destination Quartiles in 1995-2001

Origin-Destin. Quartile	Nb. of moves (1)	Perc. of moves (2)	Mean log daily wage of movers: 2 years prior (3)	1 year prior (4)	1 year post (5)	2 years post (6)	3-Year Wage Change (in percent): Raw (7)	Trend-adjusted (8)	Regr.-adjusted (9)	Stand. Error (10)
Panel A. Men										
1 to 1	32,722	53.9	4.230	4.232	4.253	4.272	4.2	0.0	-1.1	(0.3)
1 to 2	14,769	24.3	4.350	4.350	4.445	4.464	11.4	7.1	5.6	(1.3)
1 to 3	7,824	12.9	4.322	4.332	4.501	4.556	23.4	19.2	15.8	(0.6)
1 to 4	5,371	8.9	4.407	4.425	4.645	4.715	30.8	26.6	21.4	(0.7)
2 to 1	9,061	18.0	4.435	4.438	4.382	4.411	-2.4	-6.9	-8.1	(0.5)
2 to 2	20,766	41.3	4.492	4.497	4.520	4.537	4.5	0.0	-0.3	(0.5)
2 to 3	14,515	28.9	4.550	4.560	4.613	4.637	8.7	4.2	2.6	(0.5)
2 to 4	5,882	11.7	4.646	4.662	4.775	4.835	18.9	14.4	9.6	(0.5)
3 to 1	5,181	10.7	4.518	4.530	4.418	4.444	-7.4	-14.4	-13.5	(0.8)
3 to 2	7,794	16.1	4.584	4.600	4.603	4.629	4.6	-2.5	-1.7	(0.6)
3 to 3	22,200	45.9	4.638	4.647	4.675	4.708	7.0	0.0	1.2	(0.4)
3 to 4	13,223	27.3	4.723	4.739	4.824	4.875	15.2	8.2	6.7	(0.5)
4 to 1	3,429	4.3	4.644	4.661	4.510	4.546	-9.8	-18.9	-16.7	(1.2)
4 to 2	4,199	5.2	4.730	4.752	4.721	4.763	3.2	-5.8	-4.3	(0.5)
4 to 3	10,261	12.8	4.757	4.783	4.781	4.822	6.4	-2.6	-0.8	(0.5)
4 to 4	62,501	77.7	4.875	4.896	4.928	4.965	9.1	0.0	2.2	(0.3)
Panel B. Women										
1 to 1	23,575	60.5	4.086	4.095	4.123	4.138	5.2	0.0	-0.6	(0.4)
1 to 2	8,439	21.6	4.218	4.229	4.325	4.345	12.8	7.6	6.3	(0.5)
1 to 3	4,099	10.5	4.187	4.198	4.356	4.396	21.0	15.8	13.3	(0.8)
1 to 4	2,880	7.4	4.219	4.236	4.462	4.517	29.8	24.6	21.0	(0.8)
2 to 1	5,019	22.8	4.280	4.283	4.231	4.251	-2.9	-8.7	-9.2	(0.6)
2 to 2	8,957	40.7	4.358	4.372	4.402	4.417	5.9	0.0	-0.5	(0.3)
2 to 3	5,653	25.7	4.401	4.415	4.465	4.485	8.4	2.5	1.2	(0.4)
2 to 4	2,382	10.8	4.437	4.455	4.566	4.608	17.2	11.3	8.5	(0.6)
3 to 1	2,395	5.2	4.325	4.337	4.246	4.271	-5.4	-13.5	-12.2	(0.8)
3 to 2	32,268	70.3	4.388	4.406	4.434	4.454	6.7	-1.3	-0.3	(0.5)
3 to 3	6,969	15.2	4.438	4.455	4.495	4.519	8.0	0.0	1.1	(0.4)
3 to 4	4,271	9.3	4.529	4.553	4.636	4.673	14.4	6.4	5.9	(0.5)
4 to 1	1,606	6.1	4.422	4.442	4.299	4.343	-8.0	-18.8	-15.3	(1.0)
4 to 2	1,657	6.3	4.464	4.488	4.478	4.508	4.4	-6.4	-3.2	(0.7)
4 to 3	3,568	13.6	4.559	4.575	4.599	4.631	7.9	-3.0	0.1	(0.5)
4 to 4	19,423	74.0	4.663	4.696	4.747	4.771	10.8	0.0	3.4	(0.3)

Note: Sample contains all event spells defined as job transitions with 2+ preceding and 2+ succeeding observations at the old and new employer, respectively. Origin-destination quartiles are constructed from co-worker wages, focusing on firms in the dual-connected set. Firms that, in a given year, can be linked to a single worker only ("singleton-years") are excluded from the analysis. Columns 1 and 2 show the distribution of even spells by origin-destination quartile indicated in the row-heading. Entries in columns 3 to 6 are mean log real daily wages of movers around the transition date. Columns 7 to 9 report the (approximate) percent changes of wages for movers comparing two years prior and two years after the event. Trend-adjusted changes are constructed by subtracting the mean wage change of movers who move between the same origin-destination quartile (1-1, 2-2,...) during the period. Regression-adjusted wage changes are computed by fitting a linear regression separately to male and female stayers (model includes 5 education dummies fully interacted with a quadratic term in age, and a dummy for the initial year), and using the coefficients of that model to predict the wage change of movers. Column 9 reports twoway-clustered standard errors (workers/establishments/combined) of the predicted wage change using the method of Kline (2014). This accounts for sampling errors in the regression adjustment. The SAS/Stata routines for this analysis were provided by Patrick Kline.
Source: LIAB Mover Model 9308

Table J.6: Mean Log Real Daily Wages of Movers Between Origin- and Destination Quartiles in 2001-2008

			Mean log daily wage of movers				3-Year Wage Change (in percent)			
Origin-Destin. Quartile	Nb. of moves (1)	Perc. of moves (2)	2 years prior (3)	1 year prior (4)	1 year post (5)	2 years post (6)	Raw (7)	Trend-adjusted (8)	Regr.-adjusted (9)	Stand. Error (10)
Panel A. Men										
1 to 1	44,528	68.1	4.156	4.151	4.131	4.139	-1.6	0.0	-5.0	(0.3)
1 to 2	10,543	16.1	4.324	4.327	4.437	4.455	13.0	14.7	9.0	(0.7)
1 to 3	6,071	9.3	4.315	4.329	4.522	4.562	24.7	26.3	19.5	(0.9)
1 to 4	4,219	6.5	4.377	4.400	4.650	4.706	32.9	34.5	26.9	(1.0)
2 to 1	13,094	23.0	4.445	4.445	4.292	4.305	-14.0	-14.7	-17.2	(0.6)
2 to 2	26,030	45.8	4.548	4.544	4.555	4.555	0.7	0.0	-2.0	(0.3)
2 to 3	11,968	21.0	4.646	4.649	4.684	4.700	5.4	4.7	1.4	(0.4)
2 to 4	5,800	10.2	4.755	4.771	4.845	4.892	13.7	13.0	8.1	(0.7)
3 to 1	7,036	11.9	4.542	4.544	4.302	4.329	-21.4	-24.7	-25.4	(0.8)
3 to 2	10,369	17.5	4.674	4.688	4.652	4.671	-0.3	-3.6	-4.2	(0.4)
3 to 3	26,367	44.6	4.711	4.720	4.738	4.744	3.3	0.0	-0.3	(0.3)
3 to 4	15,395	26.0	4.824	4.833	4.884	4.925	10.1	6.8	5.1	(0.4)
4 to 1	4,202	5.2	4.692	4.707	4.425	4.463	-22.9	-28.8	-27.4	(1.3)
4 to 2	4,740	5.9	4.828	4.862	4.776	4.813	-1.5	-7.5	-6.3	(0.8)
4 to 3	11,823	14.6	4.863	4.881	4.851	4.886	2.3	-3.6	-2.2	(0.4)
4 to 4	60,189	74.3	4.988	5.004	5.006	5.047	5.9	0.0	1.5	(0.2)
Panel B. Women										
1 to 1	28,373	71.3	4.066	4.059	4.038	4.036	-2.9	0.0	-5.1	(0.3)
1 to 2	6,167	15.5	4.230	4.230	4.327	4.334	10.4	13.3	7.4	(0.8)
1 to 3	3,029	7.6	4.213	4.216	4.372	4.400	18.7	21.6	14.6	(0.9)
1 to 4	2,211	5.6	4.217	4.235	4.463	4.500	28.3	31.2	23.4	(1.2)
2 to 1	6,639	25.3	4.317	4.311	4.185	4.189	-12.8	-12.7	-15.4	(0.7)
2 to 2	11,641	44.3	4.426	4.421	4.434	4.424	-0.2	0.0	-2.5	(0.4)
2 to 3	5,543	21.1	4.483	4.481	4.524	4.525	4.2	4.4	0.8	(0.5)
2 to 4	2,455	9.3	4.539	4.547	4.623	4.650	11.2	11.3	6.7	(0.8)
3 to 1	3,587	14.9	4.376	4.370	4.174	4.189	-18.7	-20.3	-21.8	(0.9)
3 to 2	4,368	18.2	4.511	4.516	4.496	4.506	-0.5	-2.1	-3.9	(0.4)
3 to 3	10,917	45.4	4.554	4.561	4.580	4.570	1.6	0.0	-1.1	(0.4)
3 to 4	5,149	21.4	4.634	4.641	4.701	4.721	8.7	7.1	4.7	(0.5)
4 to 1	2,109	7.9	4.507	4.510	4.222	4.248	-25.9	-31.8	-29.3	(1.1)
4 to 2	1,981	7.4	4.626	4.638	4.577	4.593	-3.2	-9.2	-7.0	(0.7)
4 to 3	4,196	15.7	4.647	4.664	4.640	4.653	0.6	-5.3	-3.1	(0.5)
4 to 4	18,498	69.1	4.802	4.819	4.843	4.861	6.0	0.0	2.4	(0.3)

Note: See notes to Table J.5.
Source: LIAB Mover Model 9308

Table J.7: Decomposition of Changes in the Gender Gap of Firm Premiums Across Periods – Reversing the Reference Group

	1995-2001				2001-2008				Differences			
	Male	Female	Δ Male-Female		Male	Female	Δ Male-Female		Male	Female	Δ Male-Female	
			Change	Share			Change	Share			Change	Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. All workers												
Mean log real daily wages (unadjusted)			0.247	100.0			0.247	100.0			0.000	0.0
Firm-specific wage premium	0.103	0.075	0.028	11.3	0.176	0.112	0.064	25.9	0.073	0.037	0.036	14.6
Sorting (using female FE's)	0.137	0.075	0.062	25.1	0.189	0.112	0.077	31.1	0.052	0.037	0.015	6.0
Bargaining (using male employment)	0.103	0.137	-0.034	-13.6	0.176	0.189	-0.013	-5.3	0.073	0.053	0.021	8.3
Panel B. Skill Groups												
i. Missing / Primary												
Mean log real daily wages (unadjusted)			0.223	100.0			0.226	100.0			0.003	1.3
Firm-specific wage premium	0.093	0.068	0.024	10.9	0.151	0.095	0.056	24.9	0.058	0.027	0.032	12.9
Sorting (using female FE's)	0.119	0.068	0.051	22.9	0.156	0.095	0.061	26.9	0.036	0.027	0.010	4.0
Bargaining (using male employment)	0.093	0.119	-0.027	-12.0	0.151	0.156	-0.005	-2.0	0.058	0.036	0.022	9.0
ii. Apprenticeship												
Mean log real daily wages (unadjusted)			0.204	100.0			0.201	100.0			-0.002	-1.2
Firm-specific wage premium	0.103	0.074	0.028	13.9	0.173	0.109	0.063	31.5	0.070	0.035	0.035	14.2
Sorting (using female FE's)	0.137	0.074	0.063	30.7	0.186	0.109	0.076	38.0	0.049	0.035	0.014	5.6
Bargaining (using male employment)	0.103	0.137	-0.034	-16.8	0.173	0.186	-0.013	-6.5	0.070	0.049	0.021	8.6
iii. College / University												
Mean log real daily wages (unadjusted)			0.276	100.0			0.292	100.0			0.016	5.7
Firm-specific wage premium	0.117	0.096	0.021	7.6	0.195	0.136	0.059	20.3	0.079	0.040	0.038	15.5
Sorting (using female FE's)	0.161	0.096	0.066	23.8	0.220	0.136	0.085	29.0	0.059	0.040	0.019	7.6
Bargaining (using male employment)	0.117	0.161	-0.045	-16.3	0.195	0.220	-0.025	-8.7	0.079	0.059	0.019	7.9

Note: Sample contains all firms in dual-connected sets in each sample period. Columns 1-4 show results for the 1990s, columns 5-8 for the 2000s, and columns 9-12 calculate changes across periods. In each set of columns, the first two columns calculate the (counterfactual) levels of firm rents using eq. (2). The next column calculates the difference, and the last column reports percent shares of row 1 in each panel. Entries in column 12 refer to the percent change relative to the level in the 1990s. For example, 3.6/24.7 = 14.6%.
Source: LIAB Mover Model 9308

Table J.8: Decomposition of Changes in the Gender Gap of Firm Premiums Across Periods – Alternative Normalisations

	1995-2001				2001-2008				Differences			
	Male	Female	Δ Male-Female		Male	Female	Δ Male-Female		Male	Female	Δ Male-Female	
			Change	Share			Change	Share			Change	Share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Mean log value added per worker (Table 3)												
Mean log real daily wages (unadjusted)			0.247	100.0			0.247	100.0			0.000	0.0
Firm-specific wage premium	0.103	0.075	0.028	11.3	0.176	0.112	0.064	25.9	0.073	0.037	0.036	14.6
Sorting (using female FE's)	0.118	0.075	0.042	17.2	0.175	0.112	0.063	25.4	0.057	0.037	0.020	8.3
Bargaining (using male employment)	0.103	0.117	-0.014	-5.8	0.176	0.175	0.001	0.3	0.073	0.058	0.015	6.1
Panel B: Mean log sales per worker (bottom 6%)												
Mean log real daily wages (unadjusted)			0.247	100.0			0.247	100.0			0.000	0.0
Firm-specific wage premium	0.130	0.086	0.044	17.9	0.261	0.196	0.065	26.3	0.131	0.110	0.021	8.4
Sorting (using female FE's)	0.128	0.086	0.042	17.2	0.259	0.196	0.063	25.4	0.131	0.110	0.020	8.3
Bargaining (using male employment)	0.130	0.129	0.001	0.6	0.261	0.259	0.002	0.8	0.131	0.130	0.001	0.2
Panel C: Hospitality Industry												
Mean log real daily wages (unadjusted)			0.247	100.0			0.247	100.0			0.000	0.0
Firm-specific wage premium	0.289	0.244	0.045	18.2	0.300	0.249	0.051	20.6	0.011	0.005	0.006	2.4
Sorting (using female FE's)	0.287	0.244	0.043	17.2	0.312	0.249	0.063	25.4	0.025	0.005	0.020	8.2
Bargaining (using male employment)	0.289	0.287	0.002	1.0	0.300	0.312	-0.012	-4.8	0.011	0.025	-0.014	-5.8
Panel D: Business Service Providers												
Mean log real daily wages (unadjusted)			0.247	100.0			0.247	100.0			0.000	0.0
Firm-specific wage premium	0.185	0.162	0.023	9.1	0.269	0.199	0.070	28.4	0.084	0.037	0.047	19.2
Sorting (using female FE's)	0.205	0.162	0.042	17.2	0.262	0.199	0.063	25.4	0.057	0.037	0.020	8.2
Bargaining (using male employment)	0.185	0.205	-0.020	-8.0	0.269	0.261	0.007	3.0	0.084	0.057	0.027	11.0

Note: See notes to Table 3 and main text for descriptions of the alternative normalisation schemes.
Source: LIAB Mover Model 9308

Table J.9: Decomposition of the Gender Gap in Firm Premiums by Occupations in the 1990s and 2000s

	1995-2001				2001-2008			
	Overall gender wage gap	Gender gap in firm premiums	Decomposition: Sorting Effect	Decomposition: Bargaining Effect	Overall gender wage gap	Gender gap in firm premiums	Decomposition: Sorting Effect	Decomposition: Bargaining Effect
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All workers	0.247	0.028	0.042	-0.014	0.247	0.064	0.063	0.001
		(11.4)	(17.2)	(-5.8)		(25.8)	(25.4)	(0.3)
Occupation: ISCO88 major groups								
Elementary/agricultural occupations	0.288	0.032	0.054	-0.022	0.279	0.059	0.067	-0.008
		(11.1)	(18.8)	(-7.7)		(21.0)	(24.0)	(-3.0)
Plant/machine operator	0.275	-0.005	0.004	-0.009	0.261	0.014	0.012	0.002
		(-1.7)	(1.5)	(-3.2)		(5.3)	(4.6)	(0.7)
Craft and related workers	0.367	0.067	0.069	-0.002	0.364	0.078	0.063	0.015
		(18.4)	(18.9)	(-0.5)		(21.4)	(17.4)	(4.1)
Service/sales workers	0.224	-0.016	0.010	-0.026	0.232	0.033	0.029	0.004
		(-6.9)	(4.6)	(-11.5)		(14.3)	(12.4)	(1.9)
Clerks	0.221	-0.005	0.022	-0.023	0.235	0.024	0.031	-0.007
		(-2.3)	(10.0)	(-10.2)		(10.3)	(13.3)	(-3.0)
Technicians/Associate Professionals	0.310	0.026	0.043	-0.017	0.326	0.070	0.069	0.001
		(8.4)	(13.8)	(-5.4)		(21.5)	(21.2)	(0.3)
Professionals/Senior officials/Managers	0.361	0.043	0.068	-0.025	0.382	0.103	0.112	-0.009
		(12.0)	(18.8)	(-6.8)		(27.0)	(29.3)	(-2.3)

Note: Sample contains all firms in the dual-connected set in 1995-2001. Occupation groups indicated in the row heading are aggregates from three digit occupational codes into ISCO88 major groups (categories 1 and 2 as well as 8 and 9 are each merged into a single group). Column 1 shows the overall gender gap conditional on the group indicated in the row heading. Column 2 shows the difference in normalised gender-specific firm effects. Entries in columns 3-4 and 7-8 refer to the sorting and bargaining components, where sorting effects are evaluated using female firm premiums and bargaining effects using male employment. Values in parentheses are percentage shares of the overall gender gap in column 1.
Source: LIAB Mover Model 9308

Table J.10: Decomposition of the Gender Gap in Firm Premiums by Industries in the 1990s and 2000s

	1995-2001				2001-2008			
	Overall gender wage gap	GG in firm rents	Decomposition		Overall gender wage gap	GG in firm rents	Decomposition	
			Sorting Effect	Bargaining Effect			Sorting Effect	Bargaining Effect
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All workers	0.247	0.028	0.042	-0.014	0.247	0.064	0.063	0.001
		(11.4)	(17.2)	(-5.8)		(25.8)	(25.4)	(0.3)
Industry Groups (aggregates of 3-digit codes)								
Agriculture, hunting, forestry, fishing	0.258	0.074	0.081	-0.007	0.273	0.053	0.113	-0.059
		(28.9)	(31.6)	(-2.6)		(19.6)	(41.2)	(-21.6)
Energy and mining	0.070	0.033	0.016	0.017	0.038	0.002	0.000	0.002
		(47.6)	(22.9)	(24.7)		(6.6)	(1.1)	(5.5)
Production of rubber and plastic products	0.271	0.017	0.028	-0.011	0.266	0.049	0.041	0.007
		(6.3)	(10.3)	(-3.9)		(18.3)	(15.5)	(2.8)
Chemical industry	0.188	-0.010	0.014	-0.024	0.164	0.006	0.018	-0.012
		(-5.4)	(7.3)	(-12.6)		(3.8)	(11.1)	(-7.3)
Metal production and processing	0.248	0.003	-0.004	0.007	0.239	0.011	0.007	0.004
		(1.2)	(-1.5)	(2.6)		(4.6)	(3.1)	(1.5)
Automotive, production of data processing	0.308	0.009	0.046	-0.038	0.286	0.037	0.029	0.008
		(2.8)	(15.0)	(-12.2)		(12.9)	(10.0)	(2.9)
Consumer goods	0.314	0.057	0.054	0.004	0.295	0.080	0.054	0.026
		(18.3)	(17.0)	(1.3)		(27.0)	(18.2)	(8.8)
Hospitality industry	0.219	-0.005	0.011	-0.015	0.221	0.024	0.011	0.013
		(-2.2)	(4.8)	(-7.0)		(11.0)	(5.2)	(5.9)
Sales (retail/wholesale)	0.284	0.023	0.012	0.011	0.271	0.043	0.017	0.026
		(8.1)	(4.1)	(4.0)		(15.8)	(6.2)	(9.6)
Maintenance, repair of motor vehicles	0.243	0.041	0.032	0.008	0.238	0.039	0.012	0.027
		(16.8)	(13.3)	(3.5)		(16.4)	(5.1)	(11.3)
Building industry	0.240	0.101	0.013	0.088	0.235	0.104	0.010	0.094
		(42.2)	(5.4)	(36.8)		(44.3)	(4.3)	(40.1)
Transport and communication	0.181	0.011	0.013	-0.002	0.165	0.022	0.035	-0.012
		(6.1)	(7.4)	(-1.3)		(13.6)	(21.1)	(-7.6)
Credit and insurance intermediation	0.222	-0.028	-0.006	-0.022	0.180	-0.023	-0.011	-0.012
		(-12.4)	(-2.6)	(-9.8)		(-12.7)	(-6.1)	(-6.6)
Public and personal services	0.243	0.020	0.037	-0.017	0.242	0.045	0.051	-0.006
		(8.3)	(15.4)	(-7.2)		(18.5)	(20.9)	(-2.4)
Education, social, and health care	0.217	-0.068	-0.005	-0.063	0.240	-0.031	0.005	-0.026
		(-31.3)	(-2.2)	(-29.1)		(-12.8)	(2.1)	(-10.7)
Public administration, social security	0.143	-0.070	-0.003	-0.068	0.128	-0.042	0.004	-0.046
		(-49.4)	(-2.0)	(-47.4)		(-32.7)	(3.2)	(-35.9)

Note: Sample contains all firms in the dual-connected set in 1995-2001. Industries are aggregates of 3-digit industry classifications. Columns 1 and 5 show the overall gender gap conditional on the group indicated in the row heading. Columns 2 and 6 show the gender gap in normalised gender-specific firm effects. Entries in columns 3 and 7 show the sorting effects using female firm effects. Entries in columns 4 and 8 show the corresponding bargaining effects using male employment. Values in parentheses represent percentage shares of the corresponding component relative to the overall gender gap in column 1.
Source: LIAB Mover Model 9308

Table J.11: Summary of AKM Models for Low/High Union Coverage Industries and Decomposition of Firm Gender Wage Gaps

	Below Median Union Coverage				Above Median Union Coverage			
	1995-2001		2001-2008		1995-2001		2001-2008	
	Men	Women	Men	Women	Men	Women	Men	Women
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Summary of largest connected sets and AKM models								
i. Observations								
Overall	6,656,101	3,131,508	7,856,050	3,506,962	4,022,936	1,663,767	4,404,665	1,757,217
Largest connected set	6,470,272	2,896,927	7,637,726	3,229,079	3,900,287	1,522,286	4,248,275	1,596,072
Percentage share	*97.2*	*92.5*	*97.2*	*92.1*	*97.0*	*91.5*	*96.4*	*90.8*
ii. Descriptive Statistics								
Mean log daily wage	4.557	4.287	4.577	4.305	4.632	4.415	4.686	4.471
Mean age	38.9	37.3	41.3	39.9	40.0	38.1	42.5	40.6
iii. Estimation results								
# of person-effects	1,226,555	631,779	1,285,048	634,053	749,920	330,926	727,677	316,690
# of firm-effects	182,922	123,995	191,893	126,274	34,058	21,572	33,495	21,586
Standard dev. of person-effects (a)	0.283	0.303	0.315	0.340	0.270	0.276	0.290	0.317
Standard dev. of firm-effects (y)	0.164	0.208	0.208	0.245	0.143	0.170	0.172	0.209
Standard dev. of covariates (Xb)	0.106	0.078	0.083	0.068	0.105	0.078	0.086	0.070
Correlation of person- and firm-effects	0.039	-0.086	0.098	-0.050	-0.162	-0.237	-0.155	-0.287
Panel B. Baseline decompositions for dual-connected set								
Threshold identifying zero surplus firms	3.15		3.15		3.10		3.25	
Gender gap in log daily wages	0.270		0.272		0.217		0.216	
	(100.0)		(100.0)		(100.0)		(100.0)	
Gender gap in firm-specific wage premiums	0.025		0.115		0.069		-0.018	
	(9.2)		(42.3)		(31.7)		(-8.4)	
Sorting (using female FE's)	0.030		0.055		0.054		0.067	
	(11.1)		(20.3)		(25.0)		(31.3)	
Bargaining (using male employment)	-0.005		0.060		-0.012		-0.086	
	(-2.0)		(22.0)		(-5.5)		(-39.7)	

Note: Panel A summarises AKM models estimated for firms/workers in the largest connected sets of industries with below (columns 1-4) or above (columns 5-8) median union coverage as of 2010. The grouping is based on information on union coverage from Statistical Yearbooks. Panel B shows results of the associated decomposition of firm premiums as specified in eq. (2), where I previously normalised the AKM firm effects based on the identified threshold values indicated in the first row using the CCK procedure.
Source: LIAB Mover Model 9308; Federal Statistical Office

Table J.12: Decomposition of Rent-Sharing Elasticities

	Model 1		Model 2		Model 3	
	Overall	Fraction	Overall	Fraction	Overall	Fraction
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. 1995-2001						
Log Real Daily Wage	0.156	100.0	0.131	100.0	0.097	100.0
	(0.010)		(0.009)		(0.007)	
Person Effects	0.066	42.2	0.057	43.1	0.042	43.2
	(0.005)		(0.005)		(0.004)	
Establishment Effects	0.089	57.0	0.074	56.0	0.054	55.9
	(0.006)		(0.005)		(0.005)	
Covariates	0.001	0.8	0.001	0.8	0.001	0.9
	(0.000)		(0.000)		(0.000)	
# of observations			2,674,024			
# of workers			603,961			
# of establishments			4,115			
Panel B. 2001-2008						
Log Real Daily Wage	0.195	100.0	0.147	100.0	0.105	100.0
	(0.014)		(0.010)		(0.008)	
Person Effects	0.066	33.7	0.060	40.8	0.045	43.5
	(0.006)		(0.005)		(0.005)	
Establishment Effects	0.128	65.6	0.086	58.5	0.059	56.1
	(0.011)		(0.006)		(0.005)	
Covariates	0.001	0.7	0.001	0.7	0.000	0.4
	(0.000)		(0.000)		(0.000)	
# of observations			4,134,895			
# of workers			812,767			
# of establishments			6,027			

Note: Table shows decomposition of firm-level rent-sharing elasticities into worker, establishment, and covariate components by reporting the corresponding coefficients on mean log value added from separate estimations. Models are estimated on the pooled sample. Only establishments with 1+ year of mean log value added are included, and the distribution of mean log value added is trimmed above 5. All regressions include a cubic polynomial in experience and a set of year dummies fully interacted with 5 education groups. Model 2 additionally includes dummies for 16 major industries and 11 federal states. Model 3 uses detailed three-digit industry codes (up to 254 dummies; some cells empty). Standard errors are clustered at the establishment level.
Source: LIAB Mover Model 9308

Table J.13: Gelbach's Decomposition of the Overall Gender Wage Gap Based on Pooled AKM Models

	1995-2001		2001-2008		Differences	
	All	Dual-connected	All	Dual-connected	All	Dual-connected
	(1)	(2)	(3)	(4)	(5)	(6)
Unadjusted gender wage gap	0.239	0.247	0.240	0.247	0.001	0.000
	(100.0)	(100.0)	(100.0)	(100.0)		
Share of component						
Firm-effects (ψ)	0.055	0.054	0.072	0.073	0.017	0.018
	(23.1)	(21.9)	(30.2)	(29.4)	(7.2)	(7.5)
Person-effects (α)	0.163	0.170	0.153	0.158	-0.010	-0.011
	(68.1)	(68.7)	(63.6)	(64.2)	(-4.3)	(-4.6)
Covariate index ($X\beta$)	0.021	0.023	0.015	0.016	-0.006	-0.007
	(8.8)	(9.3)	(6.2)	(6.4)	(-2.6)	(-2.9)
Observations						
# of person-year	14,964,310	13,049,037	16,805,826	14,677,435		
# of workers	2,744,864	2,526,929	2,769,838	2,544,888		
# of establishments	436,807	88,968	426,126	93,306		

Note: Sample in columns 1, 3, and 5 contains all workers in the largest connected set of the pooled model. Sample in remaining columns is restricted to firms in the dual-connected set, using the parameters of the pooled model. Entries are estimated using Gelbach's exact decomposition. The contribution of the covariate index ($X\beta$) is obtained by summing over all covariates for each worker, and regressing the compound heterogeneity index on the gender dummy (see Gelbach, 2016). Entries in parentheses are the percentage share of the raw gender gap attributable to the component indicated in the row heading. For ease of interpretation, I reversed the sign on the female dummy.
Source: LIAB Mover Model 9308

Figure J.1: Ratio of Observed to Actual Number of Full-time Workers (Male + Female) for Different Analysis Samples

(a) 1995-2001

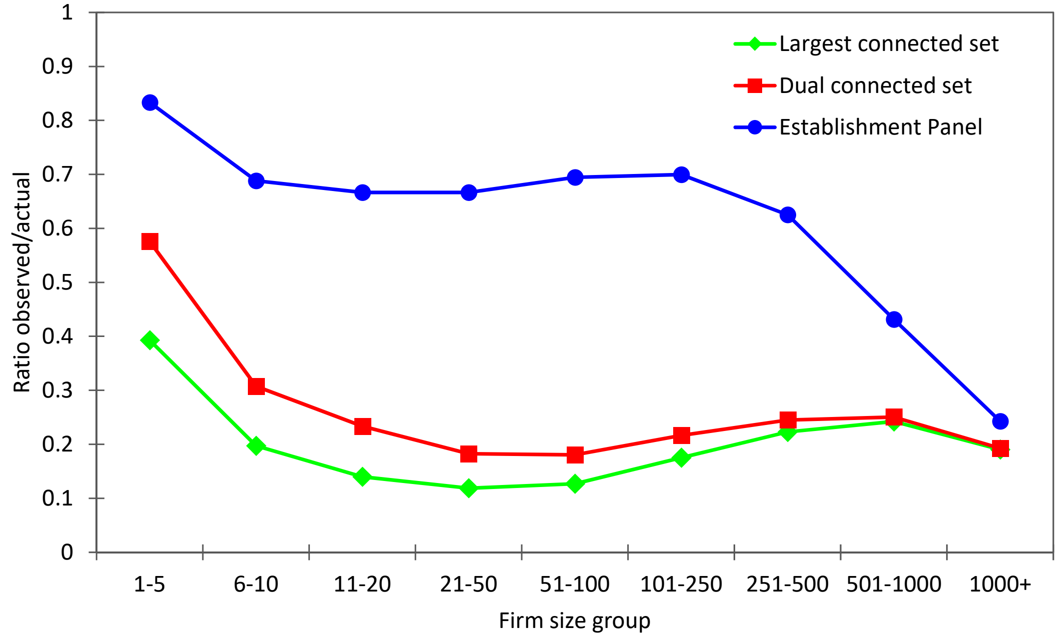

(b) 2001-2008

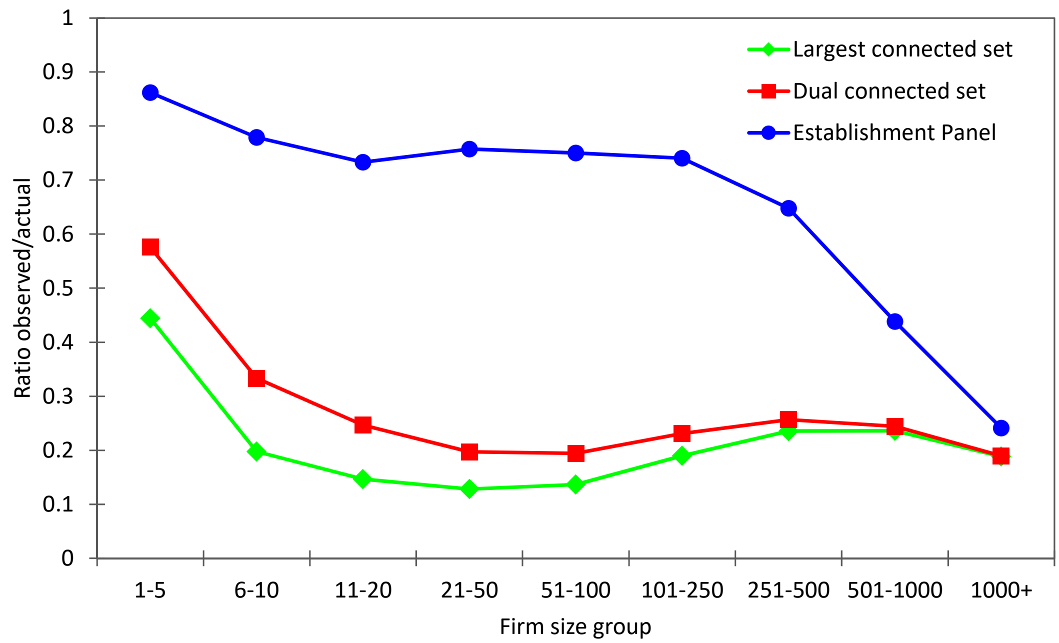

Note: Figure plots the shares of observed over actual number of full-time workers for three different analysis samples in two periods. The observed number of workers refers to the individuals that can be linked to an establishment via the establishment identifier. The actual number of workers refers to the figure recovered from the universe of employment biographies. The analysis samples are described in the main text. The unit of observation is the firm-year, and all employment refers to full-time workers.
Source: LIAB Mover Model 9308

Figure J.2: Histograms of Log Productivity per Worker in 1995 and 2008

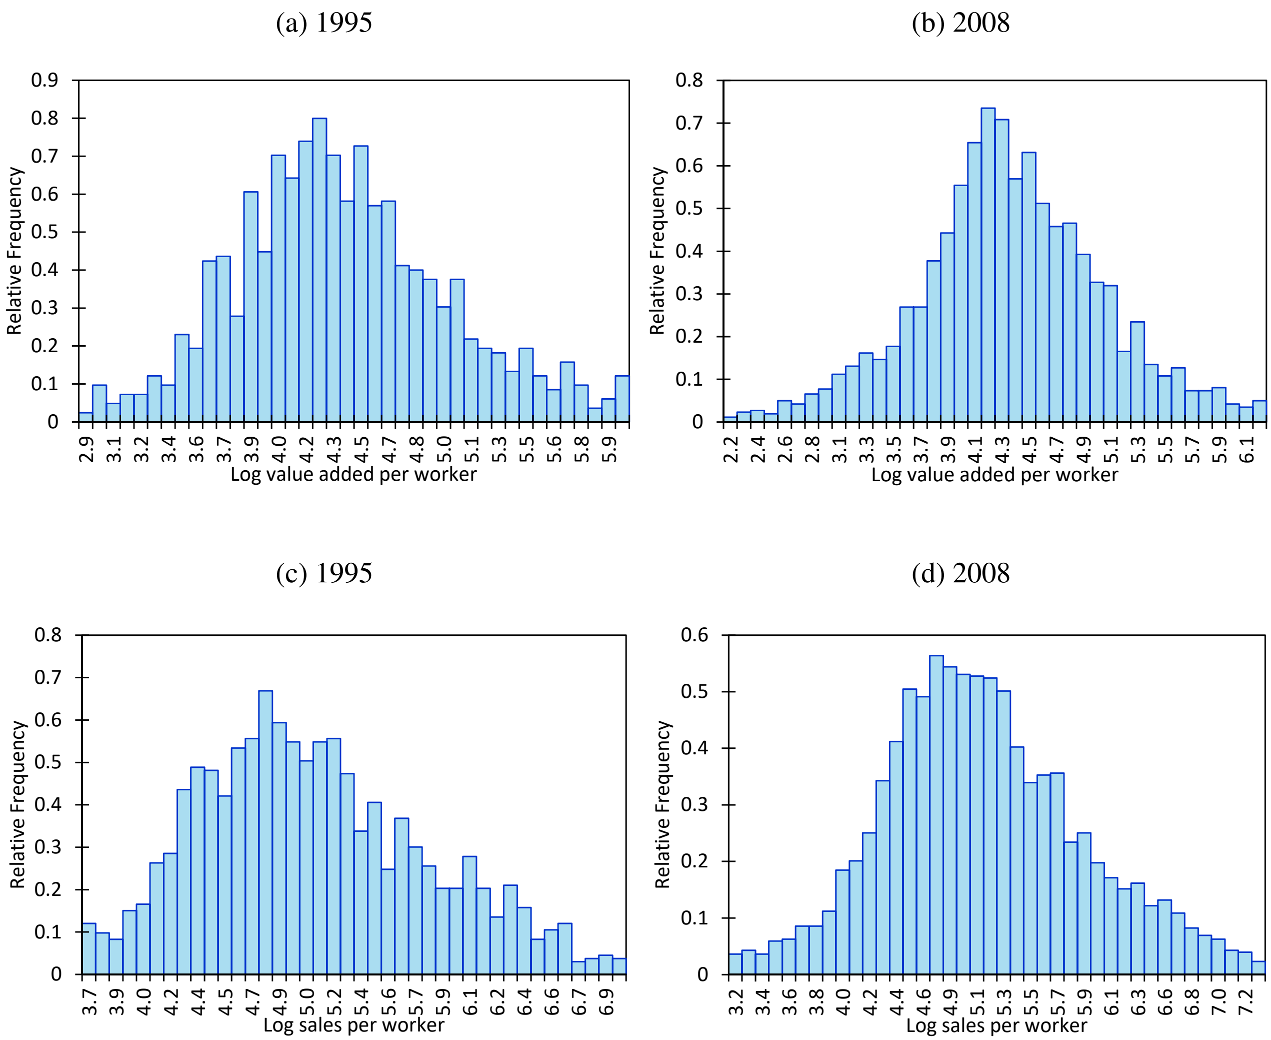

Note: Figure shows histograms of log value added per worker (panels A and B) and log sales per worker (panels C and D) in 1995 and 2008. For construction of these variables, see main text and Appendix. Histograms weight each firm equally.
Source: LIAB Mover Model 9308

Figure J.3: Histograms of Mean Log Productivity per Worker in 1995 and 2008

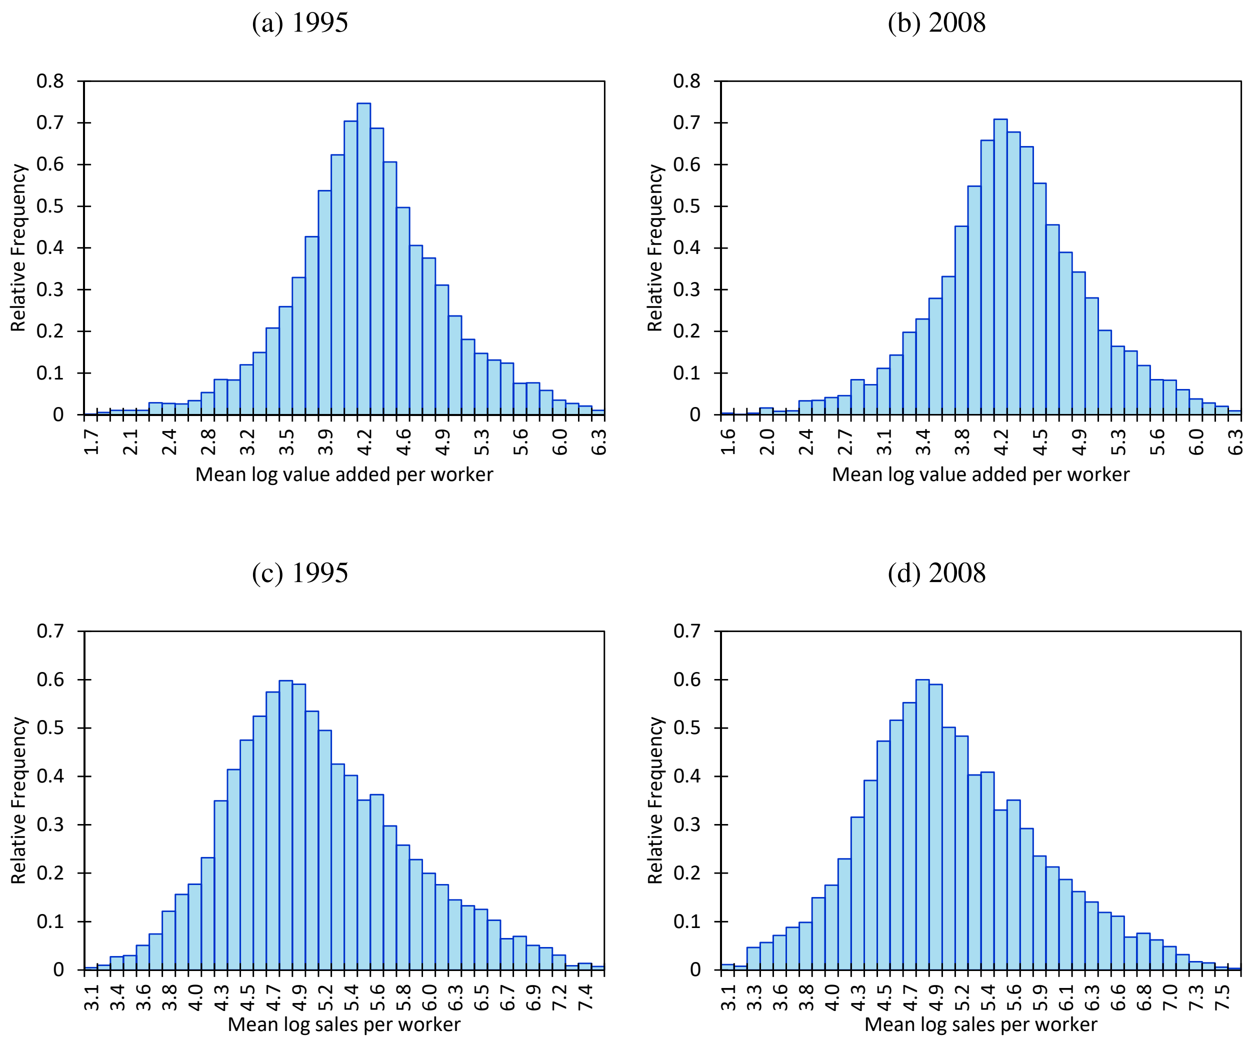

Note: Figure shows histograms of mean log value added per worker (panels A and B) and mean log sales per worker (panels C and D) in 1995 and 2008. For construction of these variables, see main text and Appendix. Histograms weight each firm equally.
Source: LIAB Mover Model 9308

Figure J.4: Percentile Growth of Mean Log Value Added Per Worker between 1995 and 2008

(a) Evolution of Selected Percentiles of Mean Log Value Added per Worker

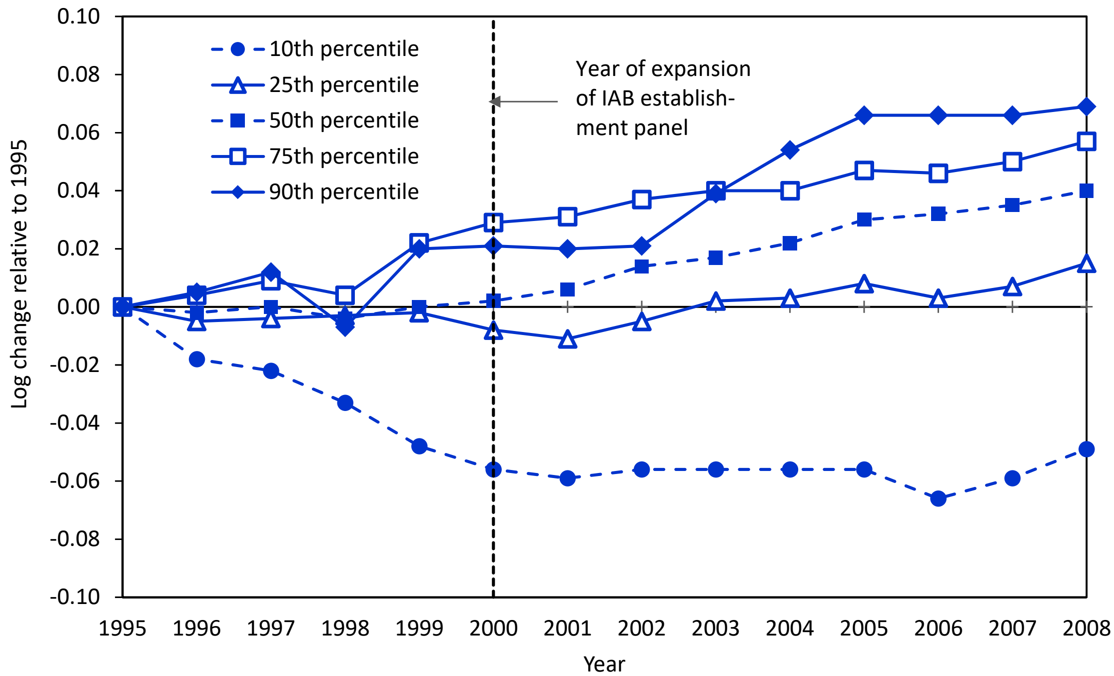

(b) Evolution of Selected Percentile Gaps of Mean Log Value Added per Worker

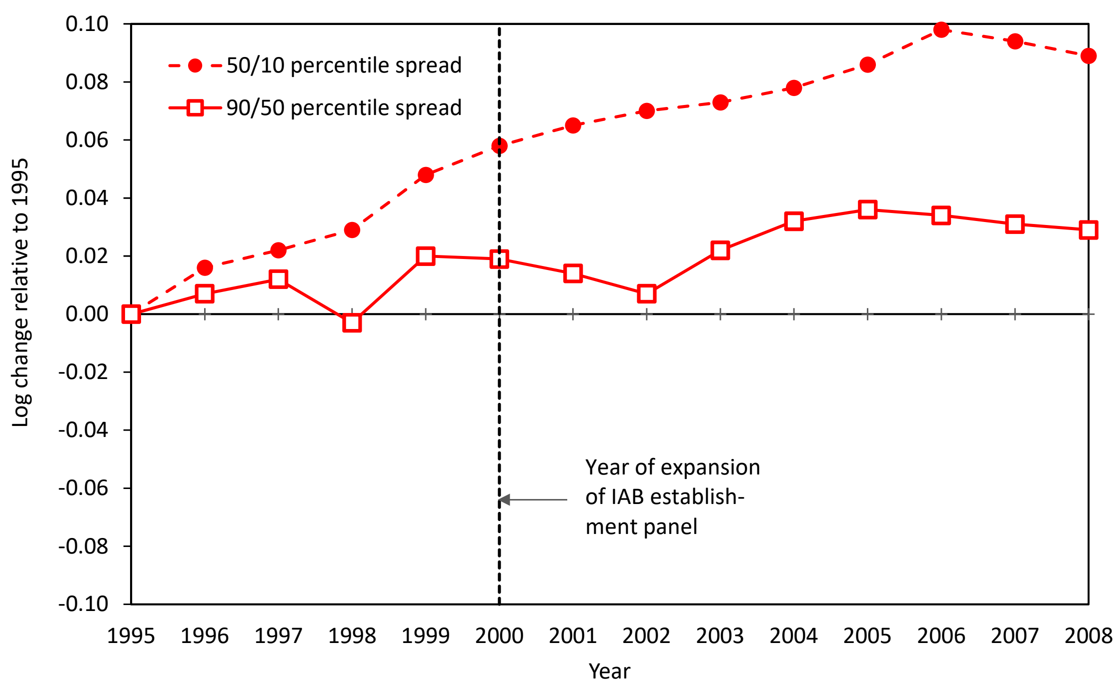

Note: Figure plots percentiles (panel A) and percentile spreads (panel B) of mean log value added over time. The distribution of productivity measures is based on establishments with at least one year of value added/sales (EP sample), and averaged across all years a firm is observed. All lines are weighted by person-years and scaled to 1995=0.
Source: LIAB Mover Model 9308

Figure J.5: Mean Wage Profiles of Male Movers Between Origin- and Destination Jobs Classified by Quartiles of Co-Worker Wages

(a) 1995-2001

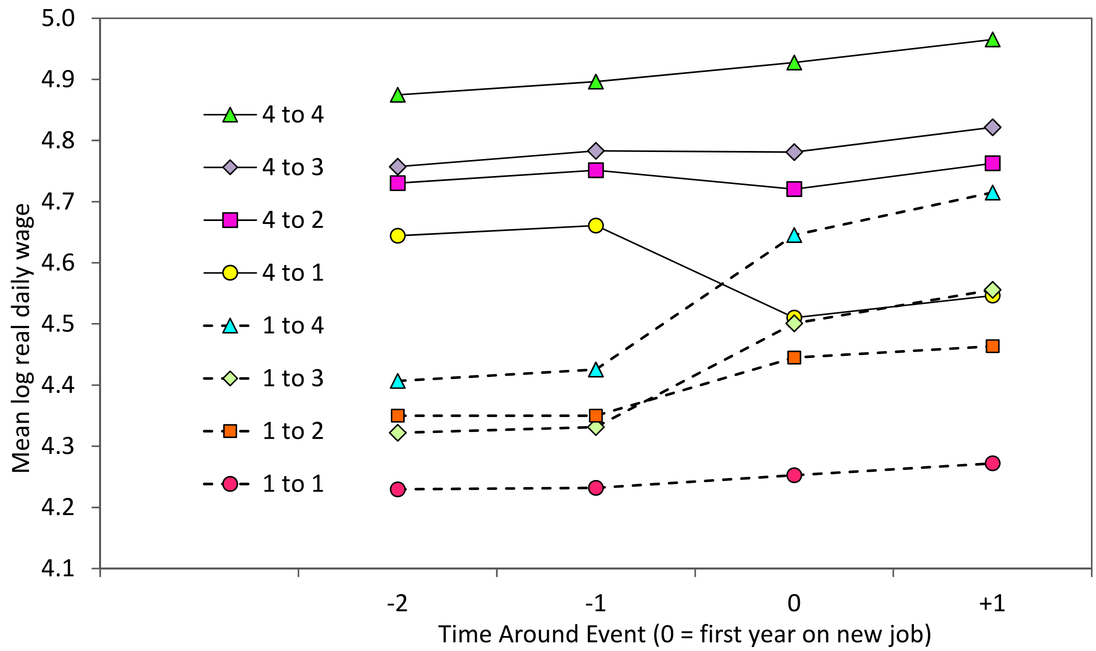

(b) 2001-2008

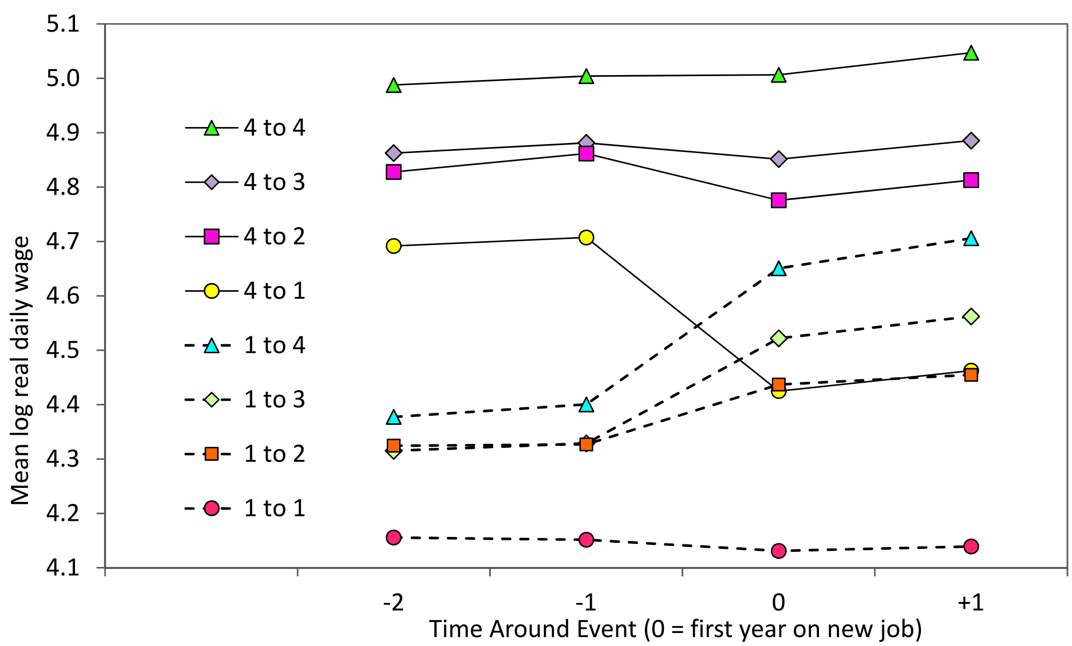

Note: Figure shows 4-year mean wage profiles of workers who move between jobs in the dual-connected set. Time 0 is the first year on the new job. Jobs are grouped into quartiles based on coworker wages in a given year. For a description of the event study, see Appendix.
Source: LIAB Mover Model 9308

Figure J.6: Mean Wage Profiles of Female Movers Between Origin- and Destination Jobs Classified by Quartiles of Co-Worker Wages

(a) 1995-2001

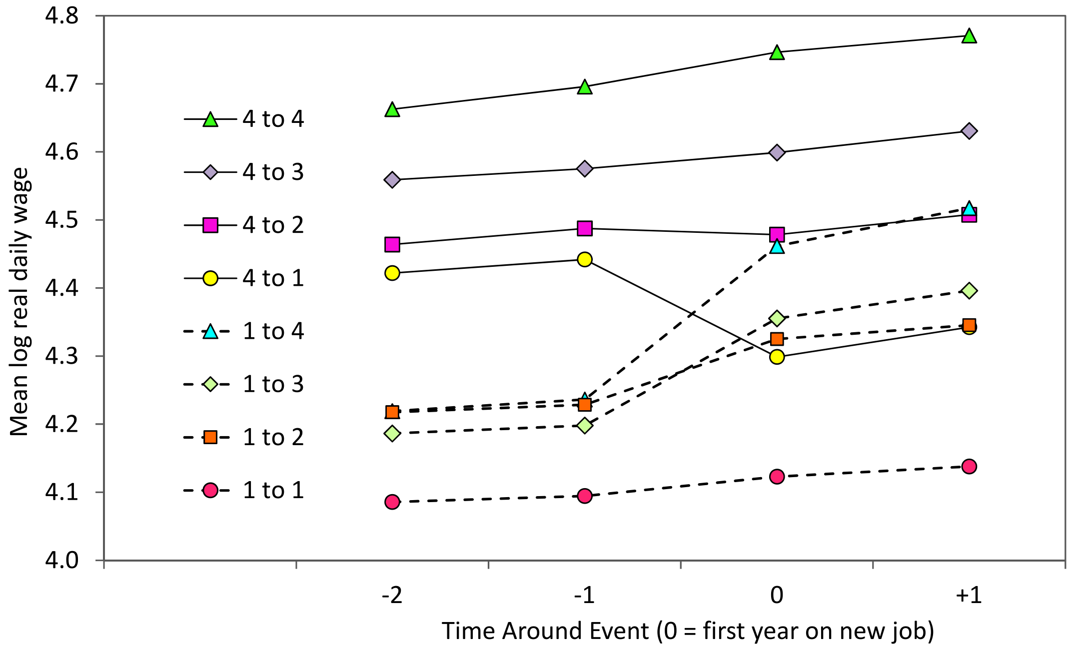

(b) 2001-2008

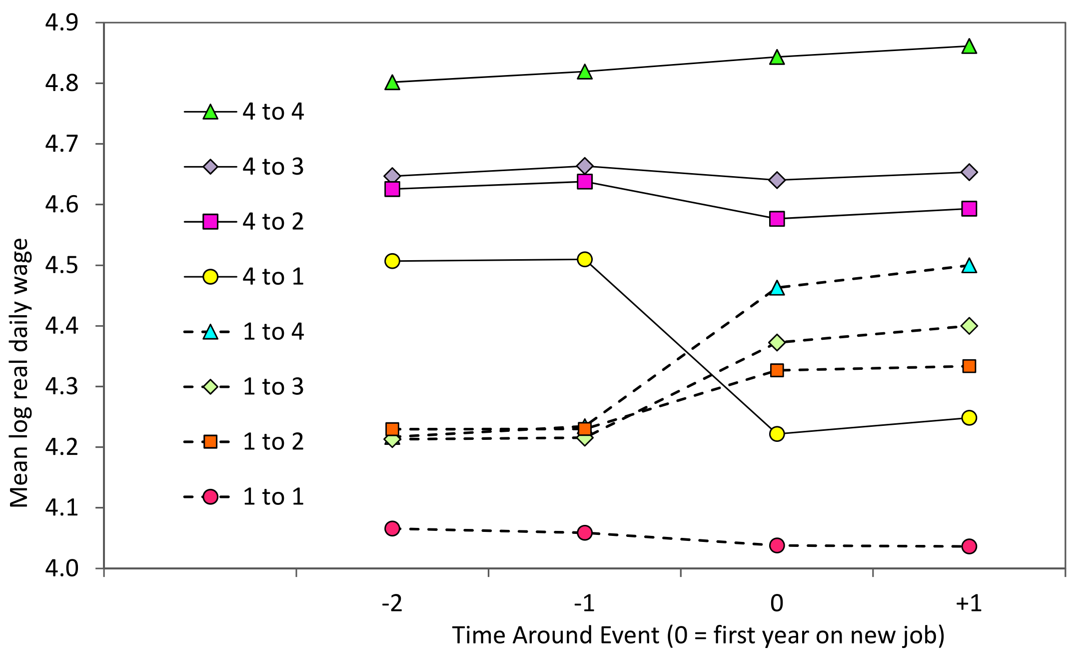

Note: See notes to Figure J.5.
Source: LIAB Mover Model 9308

Figure J.7: Regression-Adjusted Mean Wage Changes of Male and Female Movers Grouped into Co-Worker Wage Quartiles

(a) 1995-2001

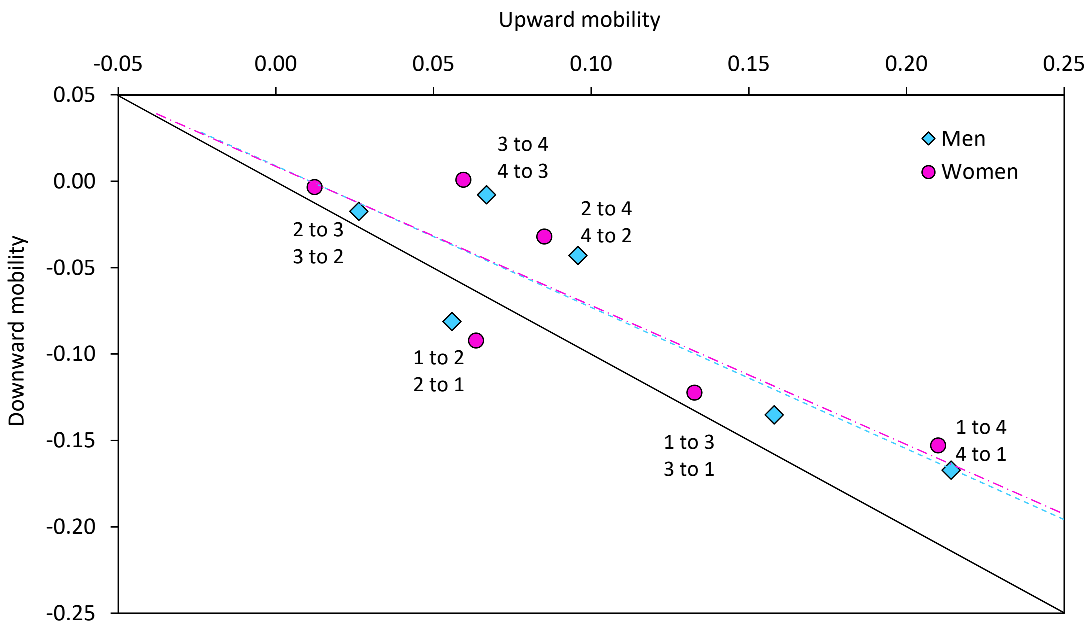

(b) 2001-2008

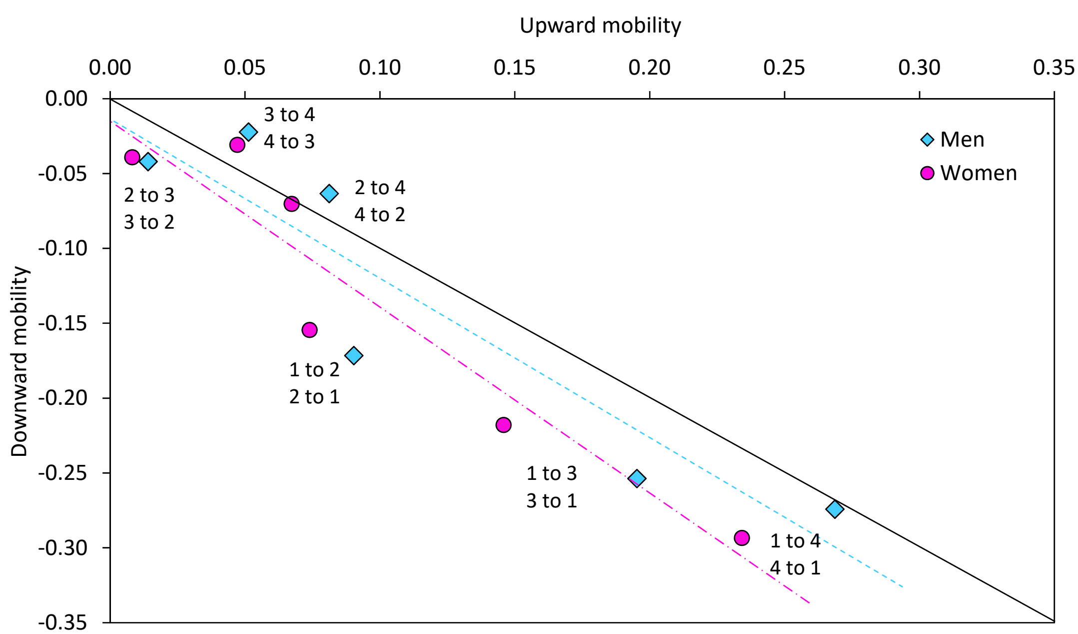

Note: Figure plots regression-adjusted average 3-year wage changes associated with downward and upward transitions for each gender. Points on the 45-degree line indicate perfect symmetry of upward and downward changes. The sample contains all firms in the dual-connected set.
Source: LIAB Mover Model 9308

Figure J.8: Regression- and Trend-Adjusted Mean Wage Changes of Men by Origin-Destination Quartiles

(a) 1995-2001

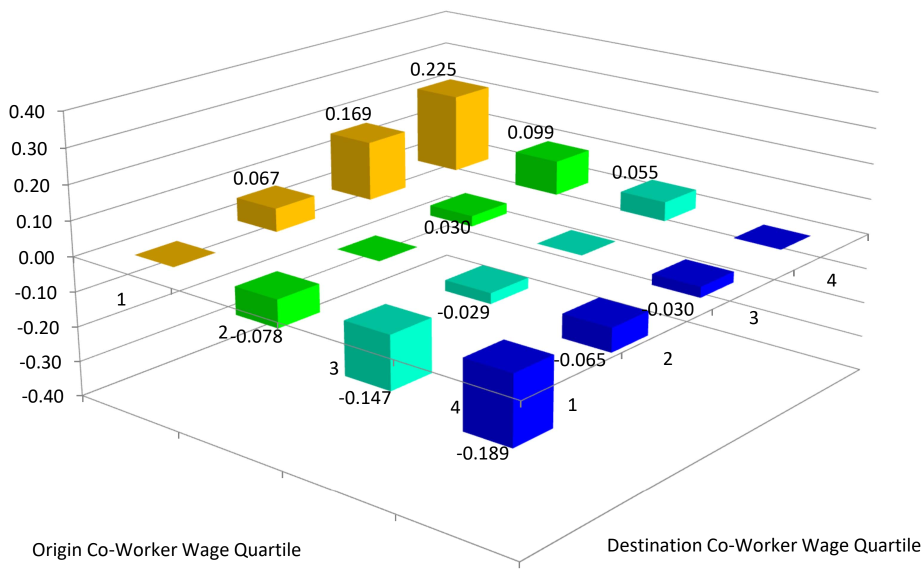

(b) 2001-2008

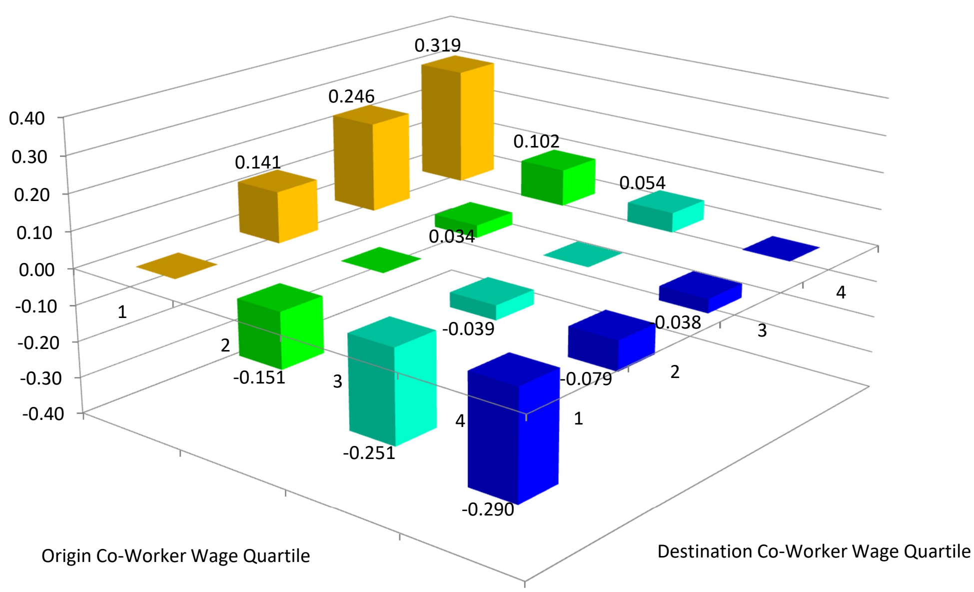

Note: Figure plots regression-adjusted 3-year wage changes of male movers between origin- and destination quartiles. In addition to a regression-adjustment, all values are trend-adjusted by subtracting the average 3-year wage change of workers who also move but stay within the same quartile.
Source: LIAB Mover Model 9308

Figure J.9: Regression- and Trend-Adjusted Mean Wage Changes of Women by Origin-Destination Quartiles

(a) 1995-2001

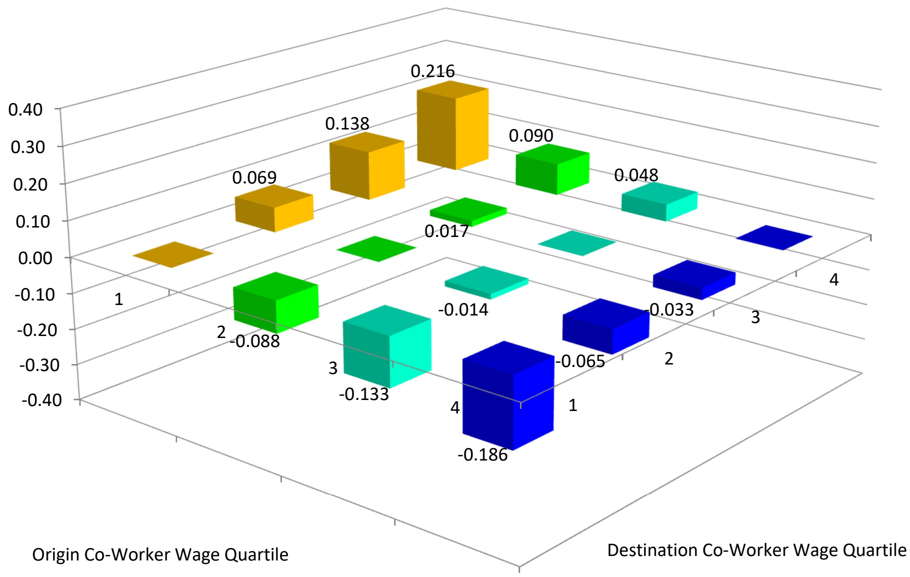

(b) 2001-2008

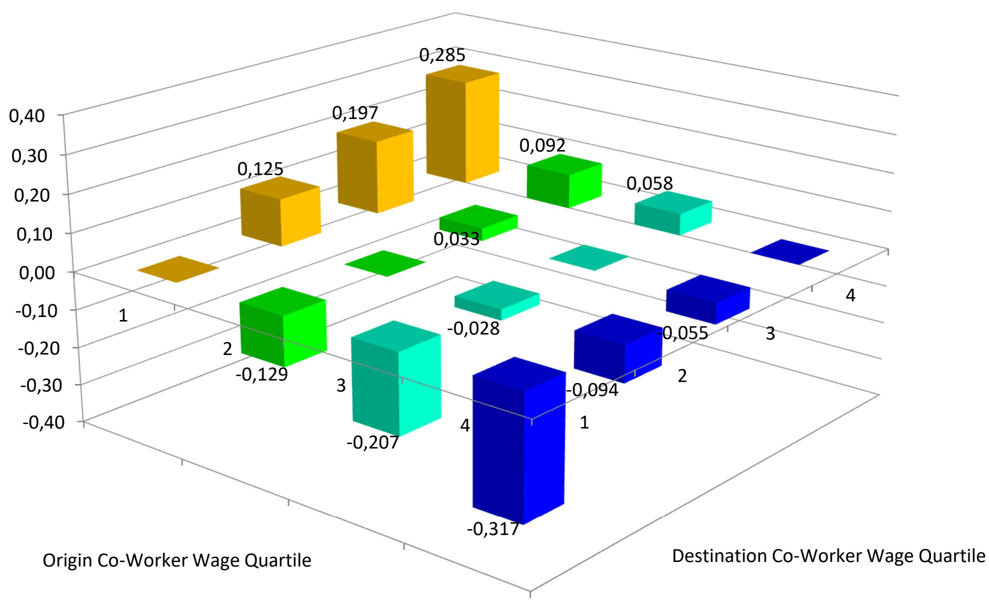

Note: Figure plots regression-adjusted 3-year wage changes of female movers between origin- and destination quartiles. In addition to a regression-adjustment, all values are trend-adjusted by subtracting the average 3-year wage change of workers who also move but stay within the same quartile.
Source: LIAB Mover Model 9308

Figure J.10: Raw and Regression-Adjusted Wage Changes of Female and Male Workers Moving Between Same Origin and Destination Quartiles

(a) 1995-2001

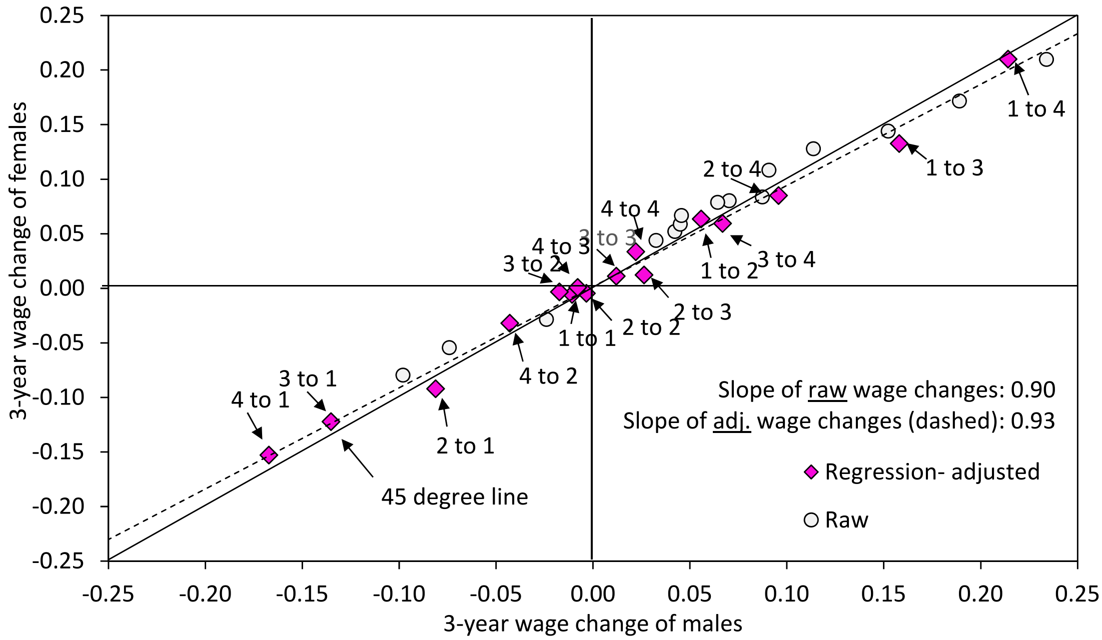

(b) 2001-2008

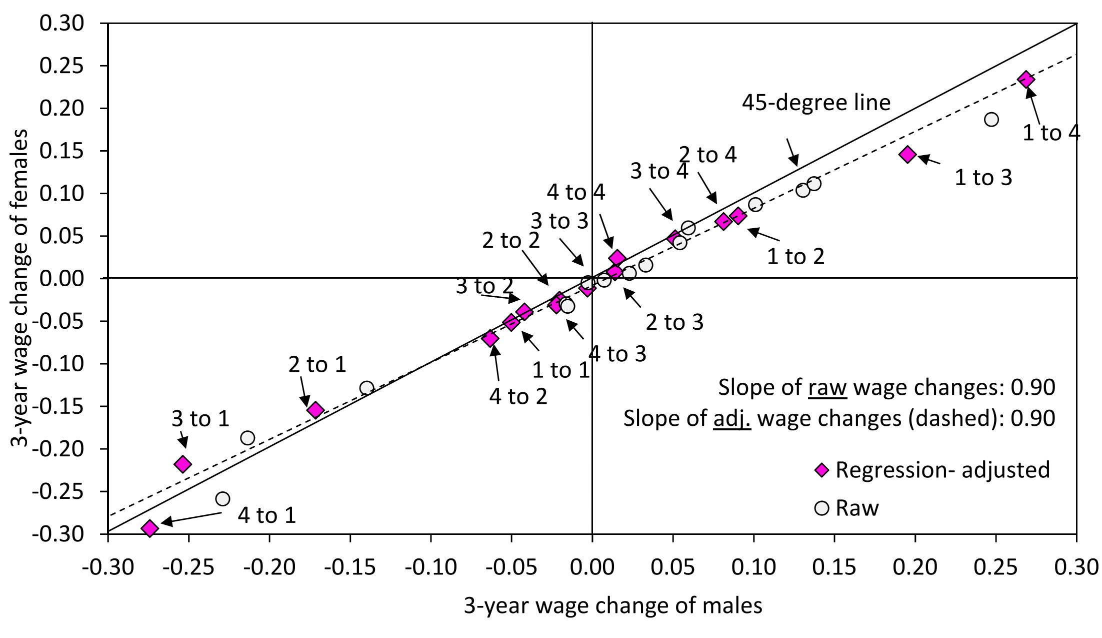

Note: Figure plots raw and regression-adjusted 3-year female mean wage changes associated with job transitions between co-worker wage quartiles against corresponding male wage changes. Results based on the dual-connected set. The dashed line is a linear regression of female changes on male changes. The solid line is a 45-degree line. One outlier of raw wage changes not shown.
Source: LIAB Mover Model 9308

Figure J.11: Mean AKM Residuals Across Deciles of Person and Firm Effects: Men

(a) 1995-2001

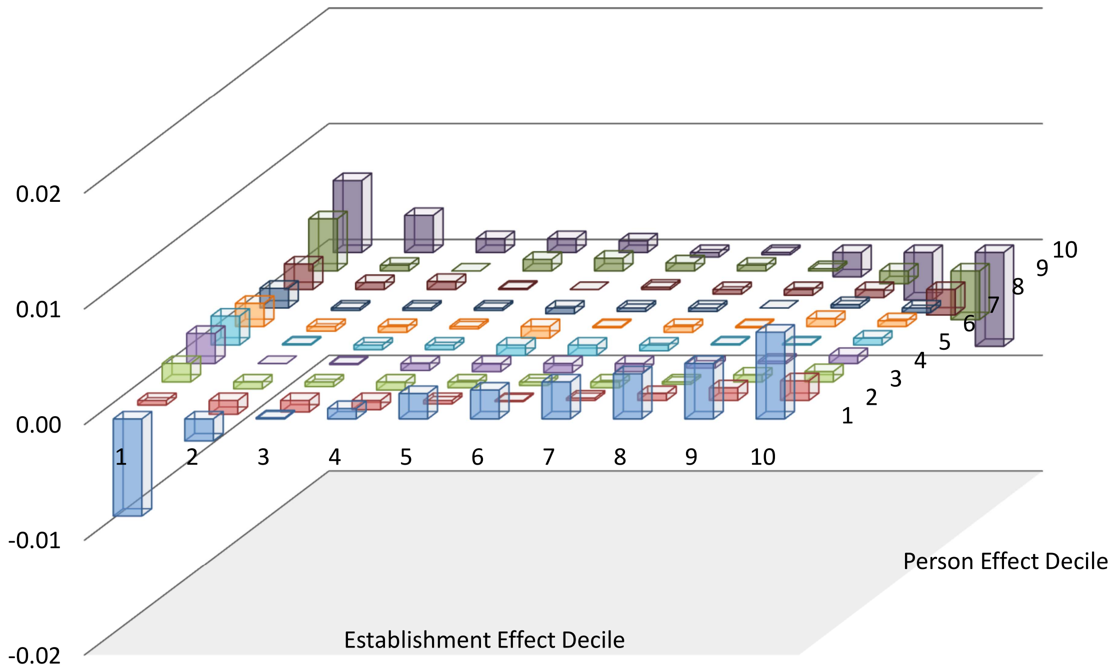

(b) 2001-2008

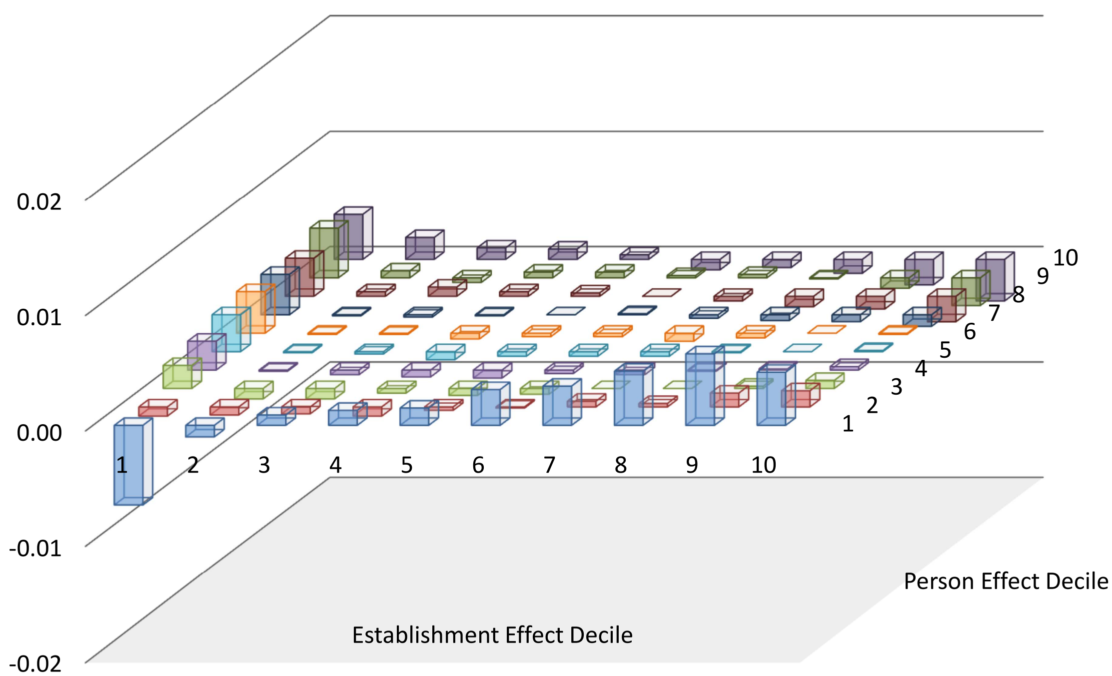

Note: Figure shows mean AKM residuals across 100 cells of person and firm effect interactions (10 deciles of person effects interacted with 10 deciles of firm effects). The corresponding AKM models are summarised in Table 2.
Source: LIAB Mover Model 9308

Figure J.12: Mean AKM Residuals Across Deciles of Person and Firm Effects: Women

(a) 1995-2001

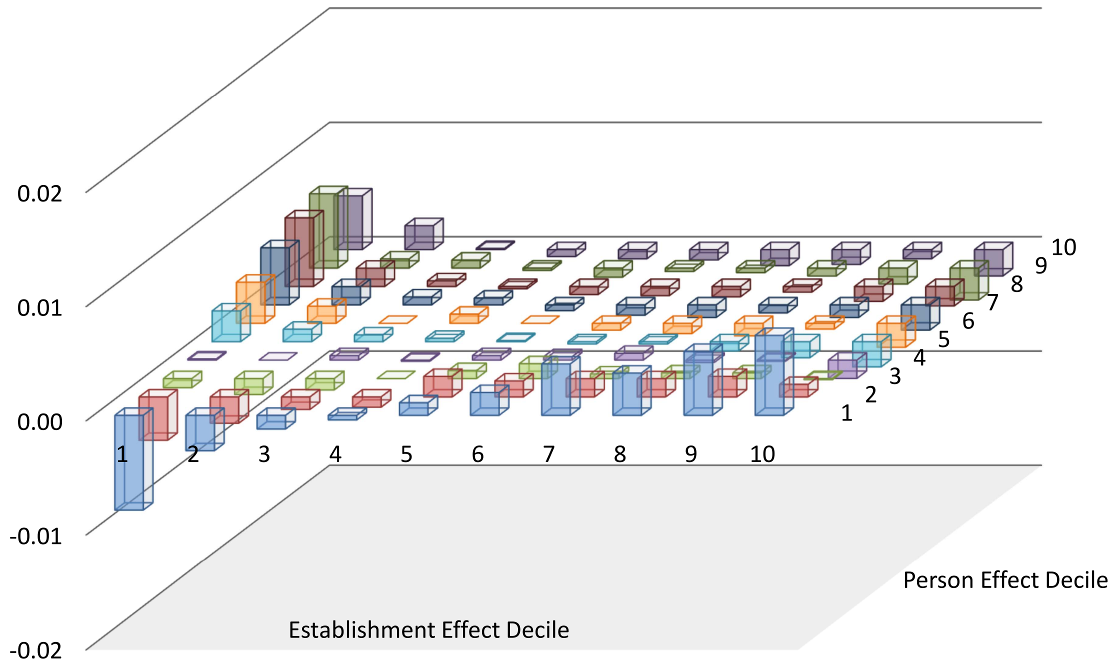

(b) 2001-2008

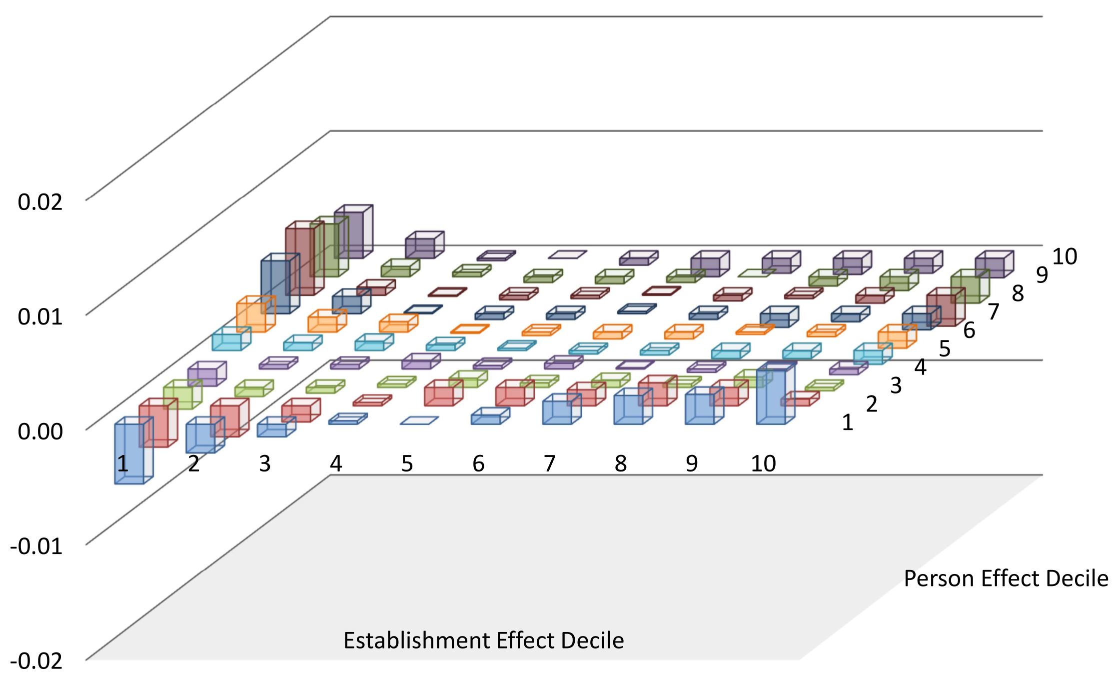

Note: Figure shows mean AKM residuals across 100 cells of person and firm effect interactions (10 deciles of person effects interacted with 10 deciles of firm effects). The corresponding AKM models are summarised in Table 2.
Source: LIAB Mover Model 9308

Figure J.13: Estimated Male and Female Firm Wage Premiums Across 16 Major Industry Groups in 1995-2001 and 2001-2008

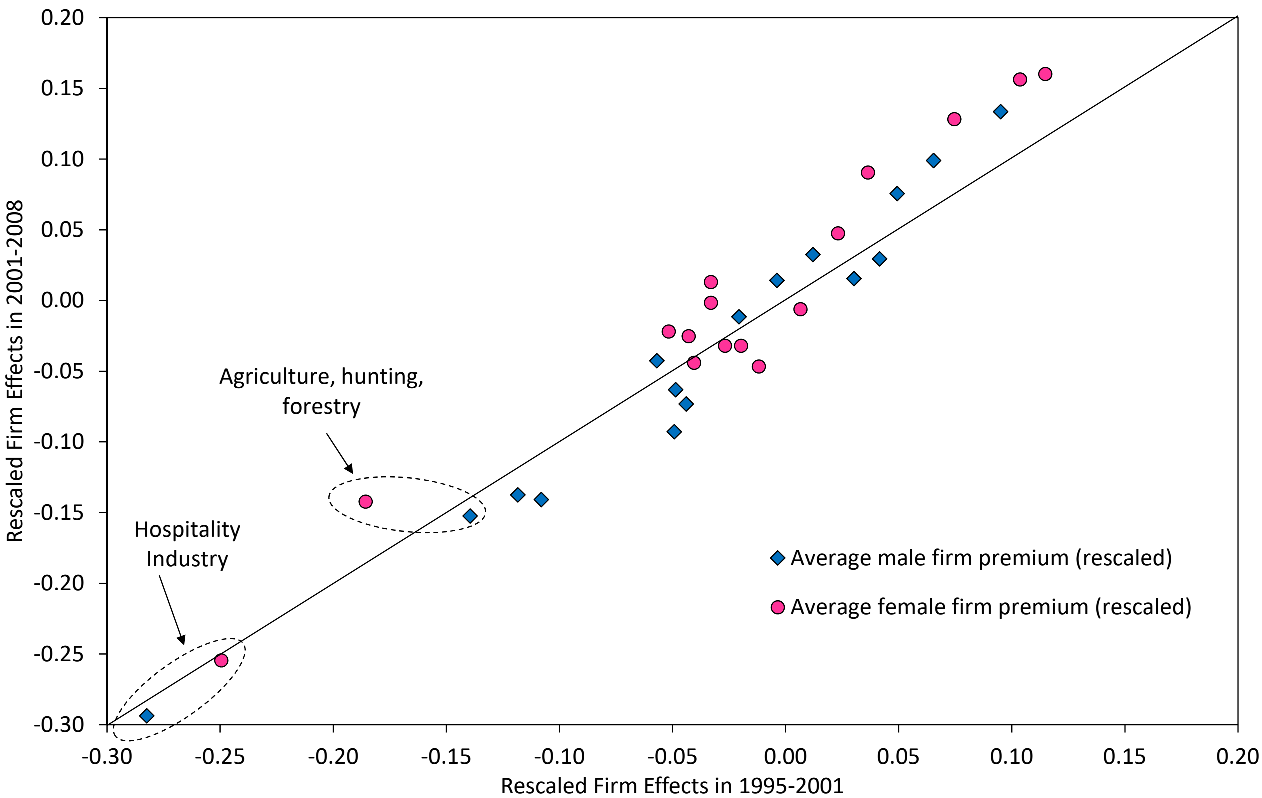

Note: Figure plots average estimated firm effects of men and women within 16 industry groups in 2001-2008 against the corresponding values in 1995-2001. Results based on dual-connected sample in each period. For the purpose of comparability, the gender-specific firm effects are rescaled so that they sum to zero for each gender in each period.
Source: LIAB Mover Model 9308

Figure J.14: Average Usual Weekly Working Hours of Full-Time Men and Women Between 1995 and 2008

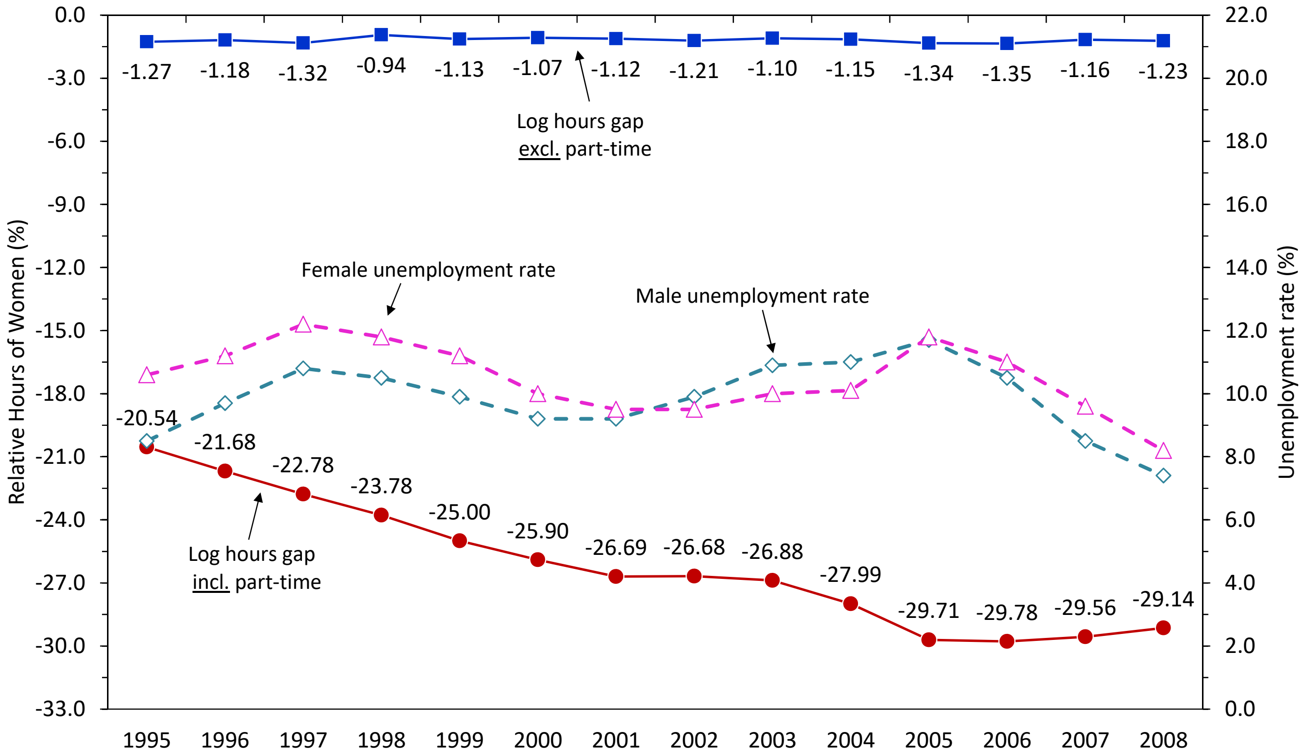

Note: Plot shows the evolution of relative weekly working hours of women. Calculations are based on the OECD concept of usual weekly working hours of the working age population between 15 and 64. Data on unemployment rates refers to the total civilian labour force. Gender-specific unemployment rates refer to the same age groups.
Source: LIAB Mover Model 9308; OECD.StatExtract data base, accessed on 17th April 2016; Federal Statistical Office

Figure J.15: Age Profiles of the Men/Mother and Men/Non-mother Wage and Firm Rent Differentials

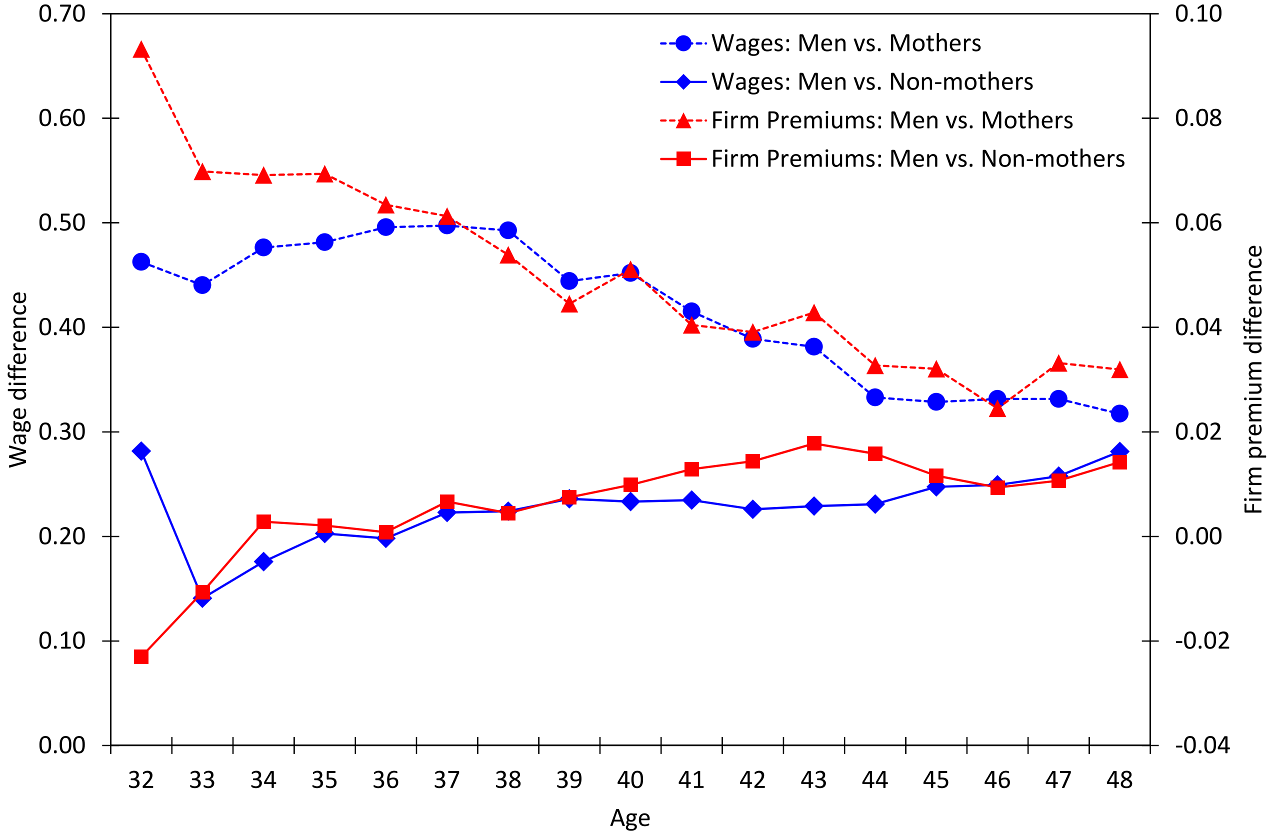

Note: Figure shows the mean wage/firm premium gap between men and mothers (dashed) and men and non-mothers (solid) at ages 32 to 48. At each age, I select mothers who gave birth 10 years earlier. Ages 30/31 and 49/50 are omitted due to very small samples of mothers.
Source: LIAB Mover Model 9308