The Gender Gap in Earnings Between and Within Firms: Evidence from Linked Employer-Employee Data*

Alessandra Casarico† Salvatore Lattanzio‡

September 12, 2018

Abstract

This paper investigates the role of firms in determining the gender gap in earnings on average and at different quantiles of the earnings distribution. Using a linked employer-employee dataset for Italy, we show the existence of firm-specific premia which differ across gender and explain on average 30% of the gender pay gap in the period 1995-2015. We decompose differences in premia into sorting of women across firms and bargaining power within firms. We find that sorting explains a larger fraction of the gender pay gap than bargaining on average and at the bottom of the distribution, whereas bargaining dominates at the top. Moreover, bargaining has increased in importance over time. We propose sorting as the outcome of gender differences in mobility rates and provide evidence that women have a lower probability of moving towards firms with higher premia. Finally, we exploit a natural experiment on the gender composition of corporate boards and find no evidence that an increase in female presence in boards reduces differences in bargaining power between men and women.

Keywords: Bargaining, Sorting, Linked Employer-Employee Data, Mobility gap, Gender quotas
JEL codes: J16, J31, J71

*We would like to thank Kai Liu, Hamish Low, and audiences at the University of Cambridge, the Italian Social Security Institute (INPS) and the ECINEQ Conference in New York for comments and suggestions. Data access was provided through the VisitINPS Scholars B programme, whose support we gratefully acknowledge. We thank Angela Ciavarella (CONSOB) for providing data on the board renewals of listed companies and Luca Brugnara for excellent research assistance. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not reflect the views of the *Istituto Nazionale di Previdenza Sociale*.

†Università Bocconi, CESifo and Dondena. Email: alessandra.casarico@unibocconi.it
‡University of Cambridge. Email: sl828@cam.ac.uk

1 Introduction

The gender wage gap has decreased remarkably starting from the 1960s but its decline has stalled. The median gender wage gap in OECD countries was 13.9% in 2016 against a value above 30% in 1975, but only 1.7 percentage points below its value in 2005, with large cross-country differences.[1]

A large literature documents the extent of gender wage gaps and their evolution over time,[2] and offers explanations for their presence. Demand-side factors, such as taste or statistical discrimination, and supply-side factors, such as productivity differences due to human capital accumulation and work effort of women relative to men, are among the explanations surveyed in Altonji and Blank (1999). Recent explanations of the persistent gap in pay focus on the role of social norms and differences in psychological traits (Bertrand, 2011, Azmat and Petrongolo, 2014), and how these affect labour market outcomes of men and women. Clearly, such outcomes will depend not only on the characteristics of the workers, but also on those of the firms which employ them.[3] First, firm-related wage differences can show up through labour market segmentation of women into firms with lower pay rates (Groshen, 1991; Bayard et al., 2003; Ludsteck, 2014; Cardoso et al., 2016). Card et al. (2016) using Portuguese data show that women are less likely to work at firms that pay higher premia to either gender, with sorting effects being most important for workers with no university degree.[4] Second, women may show lower bargaining power compared to men working at the same firm: women may negotiate less aggressively compared to men (Bowles et al., 2007, 2005; Babcock et al., 2006; Rozada and Yeyati, 2018) and this can result in gender pay gaps. Last, men and women may face different standards of promotion, even when wages tend to be equal within the same occupations (Petersen and Morgan, 1995; Blau, 2012): in this case, the wage growth of women and men differs substantially as career advances (Manning and Swaffield, 2008), because the proportion of women at the top of the pay hierarchy is lower than the proportion of men (Del Bono and Vuri, 2011).[5]

In this paper we investigate the role of firms in determining the gender gap in earnings on average, and at different percentiles of the earnings distribution. To explain the impact of firm pay policy on the gender earnings gap, we divide the role of sorting across firms and bargaining within firms, and study their impact across the distribution of earnings and over time. We then investigate what drives sorting and bargaining. We hypothesise that sorting is the outcome of gender gaps in mobility across firms. We propose a novel definition of gender mobility gap, which takes into account where the

[1]Source: OECD (2018), LFS - Decile ratios of gross earnings, and OECD Family Database (2017).

[2]For cross-country evidence see, for example, Blau and Kahn (2003), Gregory (2009), Ponthieux and Meurs (2015), Olivetti and Petrongolo (2016); for a focus on the US, Blau and Kahn (1997, 2000, 2006).

[3]The literature shows that there are large earnings differentials across firms and that the change in the variance of earnings between different firms explains a significant part of the trend in earnings inequality (see Barth et al., 2016, and Song et al., 2017, for evidence on the US and Card et al., 2013, for Germany).

[4]Reasons for why women sort into firms with lower pay include preferences for jobs and/or firms that allow more flexibility and a better work-life balance. For instance, there is evidence that the presence of women is lower in firms more open to trade and more subject to competitive pressure, where work flexibility is harder to achieve (Black and Brainerd, 2004; Bøler et al., 2018; Heyman et al., 2013). Changes in the sorting of men and women across high- and low-pay establishments also add to the increase in the gender pay gap over the life-cycle, as shown by Barth et al. (2017).

[5]Fortin et al. (2017) show that vertical segregation as proxied by the different shares of men and women in the bottom 90%, the next 9%, the next 0.9% and the top 0.1% account for more than half of the gender pay gap in Sweden, Canada and the UK. Gobillon et al. (2015) calculate that the gender difference in the probability of getting a job increases along the wage ladder, with women displaying a much lower probability of obtaining a high-pay job.

origin and destination firms are located in the firm fixed effect distribution, and we study whether risk aversion or limited geographical mobility can contribute to explain the male-female gap in the probability of moving towards more generous firms. As to bargaining, we investigate which factors influence the willingness of women to bargain over their pay or over their career path. In firms where there is more gender balance at the top of the hierarchy, women may have fairer chances to close the pay gap with men (Flabbi et al., 2016). To obtain exogenous variation in the gender composition of boards, we exploit a recent Italian law which prescribes gender quotas in corporate boards of listed companies and study whether it had any impact on gender gaps in bargaining power, as a mediating factor for gender gaps in pay.

Our analysis is based on a large linked employer-employee dataset that records the work and pay history of the universe of Italian workers in the non-agricultural private sector between 1995 and 2015. The dataset is provided by the Italian Social Security Administration (INPS, *Istituto Nazionale di Previdenza Sociale*) via the "VisitINPS" program. The dataset contains more than 22 million workers employed by approximately 1.6 million firms.

We proceed as follows. We set up an AKM model (Abowd et al., 1999) in which earnings are related to observable individual time-varying characteristics, to worker fixed effects and to firm fixed effects. The firm fixed effects are the key to measure the link between firms and the gender earnings gap. In our model firm effects are proportional to firms' rents, and represent a premium paid by firms to their workforce.[6] We first assess whether firms pay different premia to men and women and we find that this is indeed the case and that men receive larger premia. We then measure the firm contribution to the gender earnings gap, i.e. the impact of differences in firm premia on differences between male and female earnings. Following Card et al. (2016), we decompose the firm contribution into sorting and bargaining power.

We find that differences in firm-specific premia account for approximately 30% of the Italian gender pay gap at the mean. This contribution is mainly explained by sorting, which accounts for two thirds of the firm contribution and one fifth of the overall gender gap in earnings. Women have lower bargaining power than men, which determines one third of the firm contribution and slightly less than one tenth of the mean gender gap. The dominant role of sorting compared to bargaining is persistent across age and cohorts and it is more evident for older women. For managers, though, bargaining is the main factor driving the firm contribution to the gender pay gap. The importance of bargaining for high-pay jobs is confirmed when we perform the decomposition analysis within percentiles of the pay distribution and focus on the top, where bargaining dominates sorting: even when women work for high-pay firms, their earnings are lower than those of men because of their worse bargaining power. We also find that the importance of bargaining has increased over time. When we estimate firm components and their decomposition into sorting and bargaining in four overlapping time intervals between 1995 and 2015, we find that, while the contribution of firms to the gender pay gap is practically unchanged over time, its decomposition between sorting and bargaining has varied considerably, with the former decreasing in importance and the latter sharply increasing. This may reflect the spreading of decentralised wage setting in the Italian labour market.

[6] Card et al. (2018) surveys the literature on the elasticity of wages to firms' rents.

Turning to the mechanisms behind our results, and starting from sorting, we find that a gender mobility gap is present and persistent, with women – especially high ability ones – displaying a lower likelihood of moving to better paying firms, compared to men with similar characteristics. There is some evidence that gender differences in risk aversion are at play, providing an explanation for differences in mobility. As to bargaining, we show that the law 120/2011, which requires gender balance in the composition of corporate boards of listed companies, did not modify the bargaining power of female versus male workers, as measured by the change in the rent-sharing coefficient in firms affected by the reform.

This paper adds to the existing evidence on how firms' pay setting policies may impact on the gender pay gap along different dimensions. It studies the link between firms and the gender pay gap over a long period of time, which spans from 1995 to 2015, covering more than 20 years. This is a unique point of strength of our dataset and, to the best of our knowledge, it is the first time that the gender gap in earnings within and between firms is analysed over such a long horizon, allowing us to understand how the relative importance of sorting and bargaining has changed over time.[7] Moreover, it provides evidence on how firm-related factors impact at different quantiles of the earnings distribution. This complements the literature documenting the gender pay gap across the distribution – with larger gaps at the top providing evidence of a *glass ceiling* (Albrecht et al., 2003, 2015; Arulampalam et al., 2007)–, and the smaller decrease over time at the top (Blau and Kahn, 2016 and Goldin, 2014).[8] It also provides an exploration into the mechanisms which may drive sorting and bargaining, introducing a novel definition of gender mobility gaps and studying the impact of the gender composition of boards on bargaining.[9]

The remainder of the paper is organised as follows: section 2 describes the dataset and provides evidence on the gender gap in earnings in Italy; section 3 explains the methodology used to measure and decompose the firm contribution to the gender pay gap; section 4 presents the results on the decomposition on average, across the distribution of earnings and over time; section 5 investigates firm-related mobility; section 6 discusses the impact of the gender quota law on the relative bargaining power of female employees; section 7 concludes.

2 Data and Descriptive Statistics

The analysis is based on data provided by the Italian Social Security Administration (INPS, *Istituto Nazionale di Previdenza Sociale*) that record the work and pay history of the universe of employees in the private non-agricultural sector. The main source of information for these data is the form that employers have to fill in to pay pension contributions to their employees. We focus on the period

[7] Bruns (2018) analyses the period 1995-2008, comparing changes in the impact of firms on the gender wage gap in Germany before and after 2001. We incorporate also more recent years and cover the aftermath of the Great Recession.

[8] Recent evidence on the relative absence of women at the top of the US earnings distribution is provided also by Guvenen et al. (2014) and Piketty et al. (2018). For 8 countries, including Italy, Atkinson et al. (2018) document that female presence has increased less at the very top of the income distribution compared to other percentiles.

[9] While there is literature on firms' and workers' performance of exogenous changes in the composition of corporate boards (e.g. Bertrand et al., 2018), the focus on bargaining is new, as far as we know

1995-2015.[10] The data provide information about the characteristics of the jobs held by workers in the sample and about some of their personal characteristics. In particular, we have information on gross annual earnings,[11] the number of days and weeks worked in a given year, the type of contract (whether full-time or part-time), the region of work, the position held at the firm (apprentice, blue-collar, white-collar, middle-manager from 1996, and executive), the gender and the year of birth. We also know the first year of work, which allows us to build a measure of labour market experience. For each worker in the dataset we have a firm identifier we can match with information coming from the firm side of the dataset. In a separate record, INPS provides information on location, industry,[12] and date of opening and closure of all firms in the data. Furthermore, we link firms to balance sheet information, coming from the AIDA-Bureau Van Dijk dataset. This database collects balance sheet information for all the companies that are obliged to file their accounts within the Italian Business Register (mainly, limited liability companies and joint stock companies). Specifically, we use information on sales and value added, which are publicly available online from 2005 onwards, and upon request for previous years.

2.1 Descriptive Statistics

We build a panel dataset that comprises one observation per worker per year. Since some individuals are observed more than once within a year, we select the observation corresponding to the main job, that is, the contract associated with the highest number of weeks worked. In case two or more observations are characterised by the same number of weeks, we keep the observation with the highest weekly earnings. In addition, we keep only workers who have been employed for at least 4 weeks during the year.[13] We further restrict our analysis to workers with age between 19 and 65, and with at least two years of labour market experience.

From the firms' side, we drop single gender firms, that is, those firms that employ all individuals of the same sex for the entire period under analysis. This means that our final sample covers firms that have employed at least two workers of different genders.[14]

Table 1 reports descriptive statistics. We first look at columns (1) and (2). We have 129 million person-year observations for the male sample and 80 million person-year observations for the female sample. The number of male workers is 13.3 million and that of female workers is 9.1 million. Firms are 1.6 million. Mean age is slightly higher for men than for women, and so is the average job tenure.[15] The average real weekly earnings – the measure of pay we focus on – are larger for men,

[10]Even though digitalised records for workers' histories are available since 1983, we focus on the period 1995-2015 for a number of reasons. First, before 1995 information on firms is less accurate (especially sectoral codes, which are fundamental for our purposes, as we will explain later). Second, the computational burden of our estimation procedure is particularly high: 21 years should represent a significant portion of the evolution of the Italian labour market. Third, in July 1993, there was a major reform of the system of collective bargaining in Italy, which restructured the links between sector and firm level bargaining. We therefore choose to start our analysis one year and half after this reform in order to capture all the relevant changes that it brought about.

[11]Besides the full net annual earnings, this includes all kinds of pecuniary compensation, grossed up with labour income taxes and social security contributions on the employee.

[12]Industry is classified according to NACE rev. 2 sectoral codes (whose Italian counterpart is ATECO 2007).

[13]If, after these restrictions, some individuals are still observed more than once within a single year, we retain only one observation. Doing so, we drop 91,511 observations, around 0.04% of total.

[14]Overall, after data cleaning we drop 126,491,382 observations in total, approximately 38% of the original population.

[15]Job tenure is a left-censored variable. Thus, true average job tenure may be higher.

Table 1: Summary statistics

	(1)	(2)	(3)	(4)
	All		Dual connected	
	Male	Female	Male	Female
Age	39.59	38.17	39.79	38.34
Tenure	5.17	5.00	5.25	5.02
Experience	19.35	17.33	19.53	17.50
Adjusted weeks	43.62	37.42	44.14	37.85
Weekly earnings	561.34	439.29	583.68	448.12
Number of workers per firm	8.33	5.34	10.39	6.67
Share blue-collar	63.54	44.31	61.19	44.52
Share white-collar	28.33	50.43	30.30	50.46
Share executive	1.72	0.36	1.92	0.40
Share middle manager	3.91	1.94	4.43	2.14
Share apprentice	2.50	2.95	2.16	2.48
Share part-time	6.14	31.18	5.69	29.95
Observations	129,048,272	79,620,898	112,721,072	70,341,016
Number of workers	13,330,473	9,060,341	12,248,104	8,315,143
Number of firms	1,618,072	1,618,072	1,205,878	1,205,878

Notes. Columns (1) and (2) report summary statistics for male and female workers in the entire sample. Columns (3) and (4) report summary statistics for the sample used in the analysis in section 3.3. It contains the firms that are present in both largest connected sets of the male and female samples (see sections 3.1 and 3.2 for details). *Tenure* is computed as the number of years the worker is with the same firm. *Weekly earnings* are expressed in 2010 real prices. *Adjusted weeks* is the number of weeks worked in a year, standardised to account for part-time work (see text for details). The number of workers per firm is computed as the average of the yearly male and female workforce at each firm.

with a 22% gender gap. The average number of male workers in a firm is 8 and of female workers is 5, both reflecting the small average firm size of Italian firms. The share of blue-collar workers is higher for males (64% versus 44%), whereas that of white-collar workers is higher for females (50% versus 28%). The percentage of executives and middle managers is higher for male workers (1.8% and 3.9%) than for female workers (0.4% and 1.9%). The share of apprentices is higher for women. Around 6% of male workers has a part-time job, with the figure for women being 5 times larger. We keep part-time workers in the analysis, since the number of weeks worked is standardised in the data to make them comparable to those of full-time workers. In particular, for full-time workers we have the number of weeks worked over the year; for part-time workers we have the number of full-time equivalent weeks.[16]

2.2 Evidence on the Gender Earnings Gap in Italy

Figure 1 reports the evolution of the average gap in log real weekly earnings between men and women over the period 1995-2015. Earnings are expressed in 2010 real prices. Overall, the raw gap has decreased over time, though at a lower pace between 1995 and 1999 and between 2005 and 2008.

[16] This measure is computed by multiplying the number of actual weeks worked by the ratio between the number of hours worked in a month and the number of contractual hours for the full-time equivalent position. In this way, weekly earnings of full-time and part-time workers are comparable.

Figure 1: Gender pay gap over the period 1995-2015

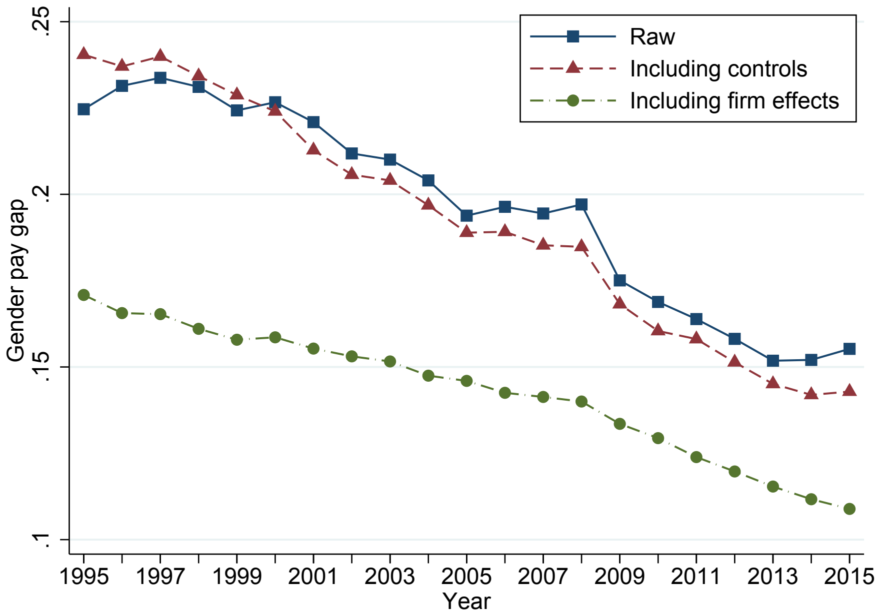

Notes. The figure plots coefficients of a dummy for male workers from log wage regressions, run for each year in three different specifications: without controls ("Raw"); controlling for observable characteristics of workers, i.e. cubic polynomials in age, experience and tenure, a dummy for full-time contract, the number of weeks worked, occupation and province of work fixed effects ("Including controls"); controlling for observable characteristics and, additionally, for firm fixed effects ("Including firm effects").

The raw average gender pay gap was approximately 22.5 log points in 1995 and 15.5 log points in 2015.

We ask how far the gender pay gap is related to firm-specific factors. A first evidence to address this question is provided in Table 2, where we report coefficients from log wage regressions. The first column of the Table is the unadjusted gender gap in log weekly earnings, which indicates that female earnings are 19.2 log points lower than men's over the period considered. Column (2) controls for a set of observable individual characteristics (cubic polynomials in age, tenure and experience, the number of weeks worked in each year, a dummy for full-time workers) and a full set of year, occupation and province dummies. The inclusion of these controls leaves practically unchanged the main coefficient of interest on the male dummy. Column (3) includes firm fixed effects. Their inclusion reduces the coefficient on the male dummy by 4.8 log points with respect to the specification in column (2) and by 5 log points with respect to that in column (1). This provides evidence that women tend to sort into firms that pay lower earnings on average. Controlling for firm heterogeneity across individuals and over time reduces the gender pay gap significantly. It is important to stress that we are not controlling here for non-random assignment of workers into firms via individual fixed effects. In addition, we are considering firm effects that do not vary by gender, assuming away within-firm differences in the ability of men and women to bargain over their pay. Hence, we can account only for the part of the gender pay gap explained by sorting of women into low-pay firms. Later in the paper we explicitly allow for firm effects to vary by gender and we control for non-random sorting of workers into firms via the inclusion of individual fixed effects.

Firm characteristics are relevant determinants of the gender pay gap over the entire period of

Table 2: Log wage regressions

	(1)	(2)	(3)
Male	0.192***	0.190***	0.142***
	(0.003)	(0.002)	(0.001)
Covariates	No	Yes	Yes
Year effects	No	Yes	Yes
Province effects	No	Yes	Yes
Firm effects	No	No	Yes
R-squared	0.040	0.509	0.705
Observations	208,669,170	208,669,170	208,669,170

Notes. Covariates include cubic polynomials in age, experience and tenure (linear term in age excluded), number of adjusted weeks worked in a year, a dummy for full-time workers, occupation dummies (blue-collar, white-collar, executive and middle manager; excluded category: apprentice). Robust standard error, clustered at firm level, in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

analysis: in Figure 1, besides the raw gender earnings gap, we plot the coefficients of the male dummy from regressions that control for individual observable characteristics (as in column (2) of Table 2) and from regressions that include firm fixed effects, in addition to individual observables (as in column (3) of Table 2). Figure 1 confirms that firm time-invariant characteristics represent an important determinant of the gender gap in earnings: the coefficient of the male dummy is lower in magnitude in each year when we control for firm effects.

The influence of the firm may vary across the distribution of earnings. In Figure D.1 in the Appendix we plot the gender pay gap across quantiles of the earnings distribution for 2015. Each dot represents the coefficient on a male dummy from a quantile regression that includes basic observable covariates (solid line) and that additionally controls for firm fixed effects (dashed line).[17] The figure shows the presence of a strong *glass ceiling* effect, meaning that the gap increases at the top of the distribution of earnings. At percentile 99.9, the gap between male and female weekly earnings is approximately 30% against a value slightly above 12% at the median. When firm effects are included, the gender gap in earnings decreases, especially in the bottom and middle portions of the distribution. The impact of firms at the very top is considerably smaller, highlighting that a large part of gender earning inequality for high earners originates within rather than between firms.

Up to now we have established that firms play a role in determining weekly earnings and the gender pay gap. In particular, we know that the coefficient on the male dummy declines when we include firm fixed effects, but we do not know how large the share of the gender pay gap explained by firms is. In the following section, we aim to measure this share and investigate whether women sort into firms that offer lower pay or whether, at the same firm, women are not able to negotiate the same contractual conditions as men.

[17]Following Canay (2011), we estimate fixed effects quantile regressions in two steps. In the first step, we run a simple regression at the mean including observable characteristics and firm effects. In the second step, we take the residual of earnings from firm effects and estimate a canonical conditional quantile regression.

3 Empirical strategy

We follow Card et al. (2016) and their novel decomposition method to estimate the share of the gender gap in earnings explained by firm-level pay setting strategies. In this section, we describe the details of such decomposition and the regression model used to retrieve the quantities of interest.

3.1 Two-way fixed effects model

We estimate log wage regressions separately by gender with the inclusion of both individual and firm effects to recover gender-specific firm fixed effects. In other terms, we estimate a two-way fixed effects model *à la* Abowd et al. (1999):

$$w_{it} = \theta_i + \psi_j^g + X_{it}'\beta^g + \varepsilon_{it} \tag{1}$$

where w_{it} is the natural logarithm of real weekly earnings, for worker i at time t, with $i \in \{1, ..., N\}$, and $t \in \{1, ..., T\}$, θ_i are the individual fixed effects, ψ_j^g are the gender-specific firm fixed effects in firm $j \in \{1, ..., J\}$ for gender $g \in \{M, F\}$, $X_{it}'\beta^g$ are the time-varying observable determinants of earnings multiplied by gender-specific coefficients and ε_{it} represents the residual unexplained component.

We interpret firm effects as quantities capturing the extent of gender-specific rent-sharing at each firm. Specifically, firm fixed effects are related to firms' rents as follows

$$\psi_j^g = \gamma^g \bar{S}_j \tag{2}$$

where $\bar{S}_j$ is the actual average surplus at firm j over the period of analysis and γ^g is the gender specific share associated to this measure of surplus. In other terms, firm effects capture the firm-level pay setting strategies, which we allow to vary by gender.[18]

3.2 Normalisation of Firm Effects

Male and female firm effects are not comparable. Fixed effects are identified only with respect to a reference firm or group of firms. Since male and female fixed effects are estimated separately, to compare their levels we need to normalise them with respect to a common criterion. Ideally, given equation (1), firm effects should be zero when firms do not share rents with their workers. Thus, we normalise firm effects with respect to the average firm effect in the accommodation and food industry, which is usually identified in the literature as a low-surplus sector (Card et al., 2016 and Coudin et al., 2018). The normalisation procedure entails rewriting our estimated firm effects as:

$$\psi_j^g = \widehat{\psi}_j^g - \mathbb{E}\left(\widehat{\psi}_j^g \mid \text{Accommodation and food}\right) \tag{3}$$

where ψ_j^g are the normalised firm effects, which are consistent with equation (2), $\widehat{\psi}_j^g$ are the estimated firm effects from model (1), and the conditioning event means that we are computing the average firm

[18] In Appendix A, we provide the modelling framework behind equation (1).

effect in the accommodation and food sector.

We report an alternative normalisation procedure in Appendix C, where we empirically identify the set of firms that pay zero rents to their workers. According to (2), firm effects are linearly related to firms' surplus. Figure C.1 plots female and male firm premia averaged across percentile bins of log value added per worker (our proxy for firm's average surplus $\bar{S}_j$) and shows a positive relationship. It also shows that firms with low value added per worker share very little rents with their workers. We identify the first decile of the value added distribution as the set of "low surplus firms". Thus, in the alternative normalisation procedure, we subtract from the estimated firm effects the average firm effects for firms in the first decile of the distribution of average log value added per worker. Results do not change and we leave to Appendix C a more thorough discussion of this alternative normalisation procedure.

3.3 Decomposition

Decomposition at the mean Once we obtain the normalised firm effects ψ_j^g, we evaluate the impact of firms on the gender pay gap by measuring the fraction of the gender pay gap that is explained by gender differences in firm pay policies. Following Card et al. (2016), we decompose the difference in firm premia into sorting and bargaining implementing the Oaxaca-Blinder decomposition (Oaxaca, 1973, Blinder, 1973) as follows:

$$\begin{aligned}\mathbb{E}\left[\psi_j^M \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right] =& \mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=M\right] \\ &+\mathbb{E}\left[\psi_j^F \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right] \qquad (4)\\ =&\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=F\right] \\ &+\mathbb{E}\left[\psi_j^M \mid g=M\right]-\mathbb{E}\left[\psi_j^M \mid g=F\right] \qquad (5)\end{aligned}$$

The left hand side of equation (4) takes the difference between the mean male firm premium across men, $\mathbb{E}\left[\psi_j^M \mid g=M\right]$, and the mean female firm premium across women, $\mathbb{E}\left[\psi_j^F \mid g=F\right]$. This difference captures the "firm contribution" to the gender pay gap. The two quantities are computed taking the average of the normalised firm effects across men and women in the double connected sample. So, $\mathbb{E}\left[\psi_j^M \mid g=M\right]$ is the male premium averaged across male observations, whereas $\mathbb{E}\left[\psi_j^F \mid g=F\right]$ is the female premium averaged across female observations. The conditioning event $\{g=M\}$ or $\{g=F\}$ indicates the set we are averaging in.

This difference can be decomposed in sorting and bargaining in two ways. In equation (4), the first term on the right hand side, $\mathbb{E}\left[\psi_j^M-\psi_j^F \mid g=M\right]$, represents the difference in firm premia between men and women, averaged across men. That is, it detects differences in firm premia, fixing the distribution of male jobs. This is a measure of the bargaining channel. It tells by how much the gender pay gap would change if women were given the same firm effects as men, weighted by the male distribution of jobs. The second block, $\mathbb{E}\left[\psi_j^F \mid g=M\right]-\mathbb{E}\left[\psi_j^F \mid g=F\right]$, represents the difference between the average female firm premia evaluated across men and the average female firm premia across women. This difference tells by how much the gender pay gap would change if women were employed in the same firms as men, weighted by the female firm effect.

Similarly, equation (5) splits the firm components into bargaining, evaluated using the female rather than the male distribution, and sorting, evaluated using male rather than female premia.[19]

Decomposition across the earnings distribution We know that lower and higher quantiles show a wider gender pay gap (see Figure D.1). Hence we investigate the impact of firm components on the gender pay gap at various quantiles of the distribution of earnings for a given year in our data. Specifically, we select groups in both the male and female samples corresponding to different percentiles of the male and female earnings distribution. For each gender-specific percentile group, we compute the mean male and female firm effects and then perform the Oaxaca-Blinder decompositions shown in equation (4) and (5). In other terms, for each gender-specific percentile group p_k^g, $k = 1, ..., 100$, we compute:

$$\begin{aligned}
&\mathbb{E}\left[\psi_j^M \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in p_k^F\right] \\
&= \mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = M, i \in p_k^M\right] + \mathbb{E}\left[\psi_j^F \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in p_k^F\right] \qquad (6)\\
&= \mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = F, i \in p_k^F\right] + \mathbb{E}\left[\psi_j^M \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^M \mid g = F, i \in p_k^F\right]. \qquad (7)
\end{aligned}$$

In both equations (6) and (7), the first term on the right hand side is the bargaining effect, whereas the difference between the second and the third term is the sorting effect. We then average sorting and bargaining as resulting from the two alternative decompositions of equations (6) and (7). That is, we compute:

$$\begin{aligned}
\text{Sorting} &= \frac{1}{2} \sum_{x \in \{F,M\}} \left\{\mathbb{E}\left[\psi_j^x \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^x \mid g = F, i \in p_k^F\right]\right\}, \\
\text{Bargaining} &= \frac{1}{2} \sum_{x \in \{F,M\}} \mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = x, i \in p_k^x\right].
\end{aligned}$$

4 Results

4.1 Estimation of two-way models

We estimate (1) separately for the largest connected groups of female and male workers.[20] We include as controls cubic polynomials in age,[21] tenure and experience, occupation dummies (blue-collar,

[19]The first block of equation (5), $\mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = F\right]$, evaluates the average difference in premia fixing the female distribution of jobs. A positive difference signals a different bargaining power within firm. The second block of equation (5), $\mathbb{E}\left[\psi_j^M \mid g = M\right] - \mathbb{E}\left[\psi_j^M \mid g = F\right]$, evaluates the difference in average male premia across male and female distribution of jobs. A positive difference signals the under-representation of women in high-pay firms.

[20]Abowd et al. (2002) show that identification of equation (1) is achieved within connected groups of firms and workers. Connected groups contain all the individuals that have ever been employed at one of the firms in the group and all the firms that have ever hired one of the workers in the group. Thus, two groups are not connected if one person of the second group has never been employed by a firm of the first group and a firm in the first group has never employed a person of the second group (or viceversa). Since fixed effects are identified up to a normalising constant, different connected groups give fixed effects estimates that are not comparable across each other. Thus, we perform the analysis on the largest connected group.

[21]We normalise the age profile to be flat at age 40 and we exclude the linear term in age to avoid potential collinearity with experience and year effects. See Card et al. (2018).

Table 3: Summary statistics for largest connected sets and AKM estimation

Largest connected sets		
	Male	**Female**
Number of p-y obs.	127,908,136	77,622,344
% of entire data	99.12%	97.49%
Number of workers	13,123,321	8,735,880
% of entire data	98.45%	96.42%
Number of firms	1,456,374	1,369,594
% of entire data	90.01%	84.60%
AKM estimation		
F-stat	60.180	23.020
Adjusted R-squared	0.871	0.741
RMSE	0.164	0.197
Mean log weekly earnings	6.189	5.997
St. dev. earnings	0.486	0.415
St. dev. worker effects	0.661	0.568
St. dev. firm effects	0.209	0.195
St. dev. xb	0.709	0.564
St. dev. residual	0.164	0.197
Corr(we,fe)	-0.043	0.000
Corr(we,xb)	-0.868	-0.879
Corr(fe,xb)	0.173	0.081

Notes. The Table reports summary statistics from the largest connected sets used for the estimation of the AKM two-way models. See text for details.

white-collar, executive, middle manager and apprentice) and a full set of year dummies. The upper part of Table 3 reports sample sizes of the largest connected sets in both the female and male samples. We retain 99.1% and 97.5% of the total person-year observations in the male and female samples, respectively. Men are 98.5% and women are 96.4% of those in the original data. Coverage of firms is 90% and 84.6% in the male and female samples, respectively, compared to the original population.

The lower part of Table 3 reports statistics about the fit of the model in equation (1) to our data for both samples of female and male workers and it shows that the fit is quite good and all the parameters are jointly significant.[22] Worker and firm effects display negative or no correlation (-0.04 and 0, respectively, in the female and male sample). In practice, this implies that the Italian labour market is characterised, if anything, by negative assortative matching. This result is consistent with Flabbi et al. (2016).

Finally, it is important to stress that the validity of the two-way fixed effects model in equation (1) relies upon the assumption of conditional random mobility of workers. We test this assumption in Appendix B. Overall, we conclude that it holds for both the female and the male sample.

[22] The standard deviation of the estimated worker effects is in both samples three times higher than the standard deviation of the firm effects. Thus, if we were to decompose the variance of earnings in its primary determinants, a greater part of such decomposition would be explained by individual, rather than firm variability. Moreover, the coefficient of correlation between female and male firm effects is positive. This means that firms paying high wages to men tend to do the same also with women.

4.2 Firm Contribution to the Gender Gap in Earnings, Sorting and Bargaining

4.2.1 Average Decomposition

We build a double connected set of workers and firms, by selecting the firms that appear in both largest connected sets of male and female samples. The structure of this set of workers and firms allows us to compare female and male firm effects and to measure counterfactual moments of the distribution of both female and male premia. Columns (3) and (4) of Table 1 report summary statistics for men and women in the double connected set. The number of person-year observations drops to approximately 113 million for males and 70 million for females, with 12.2 million male individuals and 8.3 million female individuals, employed by 1.2 million firms. Age, tenure and the distribution of occupations across genders is roughly comparable to the original dataset. Weekly earnings slightly increase for both men and women, as well as the number of workers per firm.

We normalise firm effects as detailed in section 3.2 and decompose the difference in firm pay premia as in equations (4) and (5). Results are in Table 4. Column (1) of the Table shows the overall firm contribution to the gender gap in earnings and its decomposition. In the double connected sample, the mean raw gender pay gap is 21.3 log points, compared to 19.2 in the overall sample. We can explain 30.4% of this gap as coming from the difference in premia recognised to men and women, since the gap in firm effects is approximately 6.5 log points. This contribution is mainly determined by sorting, irrespective of whether one uses the decomposition framework of equation (4) or (5). In both scenarios, sorting accounts for more than 20% of the overall gender pay gap, while bargaining accounts for a smaller share (between 7.6% and 9.8%). This result is similar to the one found by Card et al. (2016) for Portugal and by Coudin et al. (2018) for France. Thus, sorting is the main factor behind the different premia men and women receive.

In Columns (2) through (6) we report decompositions for subsamples defined by different occupations. The gender pay gap is small for apprentices (4.1 log points) – for whom salaries are usually low, irrespective of gender – and for middle managers (12.3 log points), whereas it is higher for blue-collar workers (22.7), white-collar workers (27.1) and executives (23.4). 39.4% of the gender pay gap for blue-collar workers (column 3) can be explained by firm components, mainly due to sorting of women into low-pay firms (roughly 31% of the gender pay gap). A similar result holds for white-collar workers (column 4): the gap in firm effects accounts for 26% of the gender gap in earnings, mainly due to sorting (18-21%) rather than bargaining (5-8%). Since the large majority of workers in our data is either classified as blue- or white-collar (roughly 91% of men and 95% of women), it comes as no surprise that, on average in the entire sample, we find that sorting is the main factor driving firm-related gender inequality. For apprentices and middle managers (columns 2 and 5), 49% and 19.5% of the gender pay gap, respectively, can be explained by differences in pay premia. This difference is mainly due to a lack of bargaining power of women compared to men: this channel accounts for at least 33% of the gender pay gap for apprentices and at least 22% for middle managers. Interestingly, sorting plays a negative role for the latter category of workers, meaning that men in this specific occupation are employed at low-pay firms compared to women. As to executives (column 6),

Table 4: Gender pay gap, firm effects, sorting and bargaining

	(1)	(2)	(3)	(4)	(5)	(6)
	Total	Appr.	Blue collar	White collar	Middle man.	Exec.
Gender pay gap	0.213	0.041	0.227	0.271	0.123	0.234
Male firm effects across males	0.113	0.035	0.074	0.167	0.275	0.222
Female firm effects across females	0.049	0.014	-0.015	0.097	0.251	0.165
Firm effects gap	0.065	0.020	0.089	0.070	0.024	0.058
% of gender pay gap	*30.4%*	*49.0%*	*39.4%*	*25.9%*	*19.5%*	*24.6%*
Decomposition:						
Sorting						
Using male coefficients	0.049	0.007	0.071	0.057	-0.004	0.047
% of gender pay gap	*22.8%*	*16.6%*	*31.1%*	*20.9%*	*-3.1%*	*20.3%*
Using female coefficients	0.044	0.003	0.070	0.049	-0.009	0.026
% of gender pay gap	*20.6%*	*7.9%*	*30.7%*	*18.2%*	*-7.2%*	*11.2%*
Bargaining						
Using male distribution	0.021	0.017	0.020	0.021	0.033	0.031
% of gender pay gap	*9.8%*	*41.1%*	*8.7%*	*7.7%*	*26.7%*	*13.5%*
Using female distribution	0.016	0.013	0.019	0.013	0.028	0.010
% of gender pay gap	*7.6%*	*32.5%*	*8.3%*	*5.0%*	*22.6%*	*4.3%*
Observations	183.1	4.2	100.3	69.7	6.5	2.4

Notes. See text for details on the decomposition method. Columns (2) to (6) report results for subsamples defined by occupation categories: apprentice, blue-collar, white-collar, middle manager and executive. The number of observations is expressed in millions.

the gap in firm effects accounts for almost a quarter of the gender pay gap. The relative importance of sorting and bargaining depends on the type of decomposition chosen.

Figure 2 shows the evolution of firm fixed effects, sorting and bargaining by age and cohort, with the upper (lower) panel reporting the analysis by age (cohort). In general, for younger workers (and more recent cohorts) firm-related inequality is lower relative to older workers. Looking at Panel (a), the gap in firm effects starts early in the career of male and female workers. It is around 5 log points when workers are 19 years old. The growth of firm effects is similar for men and women until age 27-29, when female firm effects stall and male firm effects keep rising until 57, when for both genders they start to decline. In Panel (b) we decompose this difference into sorting and bargaining. Both effects are similar in magnitude until age 25. After this age, bargaining remains roughly constant until age 55, when it starts to decline and becomes zero just before retirement age. On the contrary, sorting increases in importance, reaching a peak at age 60 and declining in the years before retirement. The results suggest that mobility may play a role in explaining the widening gap in firm effects between male and female workers, a channel that we explicitly investigate later in the paper. Note also that the start of the divergence between the two effects shown in Panel (a) roughly corresponds to the age of the first child. Panels (c) and (d) show that for younger cohorts, i.e. those born after 1976, the gap in firm pay premia is relatively small in magnitude and the impact of sorting and bargaining on this gap is roughly equal. On the other hand, for older cohorts the gap in firm effects rises up to

Figure 2: Firm effects, sorting and bargaining by age and cohort

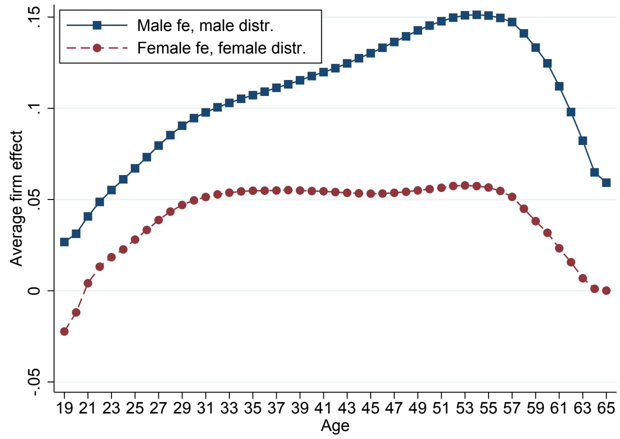

(a) Firm effects by age

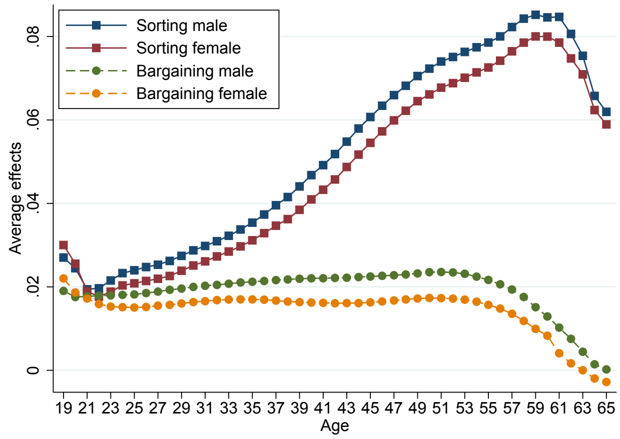

(b) Sorting and bargaining by age

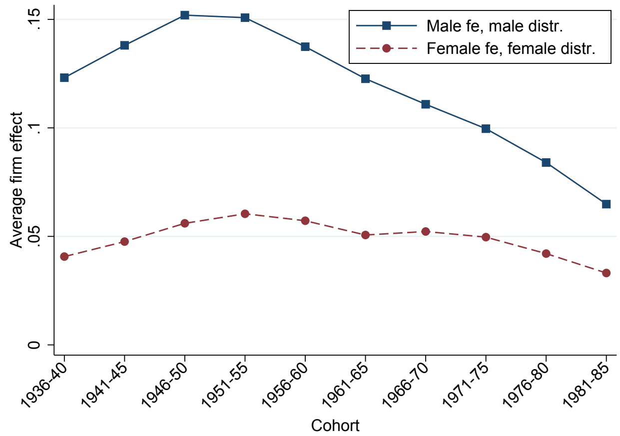

(c) Firm effects by cohort

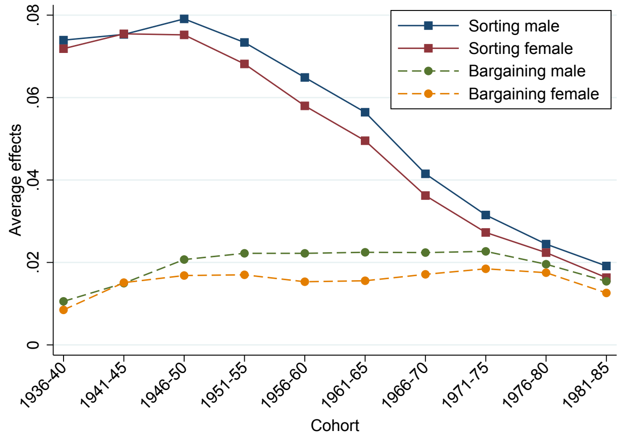

(d) Sorting and bargaining by cohort

Notes. "Male fe, male distr." is the mean male firm effect across males, $\mathbb{E}\left[\psi_j^M \mid g = M\right]$; "Female fe, female distr." is the mean female firm effect across females, $\mathbb{E}\left[\psi_j^F \mid g = F\right]$. "Sorting female" and "Sorting male" are the sorting channels evaluated by fixing female (eq. 4) and, respectively, male (eq. 5) firm effects. "Bargaining male" and "Bargaining female" are the bargaining channels evaluated by fixing male (eq. 4) and, respectively, female (eq. 5) distributions of jobs.

a maximum of about 10 log points for the cohort 1946-50, with sorting being the dominant driver relative to bargaining power. We report in Appendix Table D.1 details of the decompositions by age and cohort, showing the magnitude of the two effects and their contribution to the gender pay gap for aggregate age groups (19-29, 30-39, 40-49, 50-65) and cohorts (1940-49, 1950-59, 1960-69, 1970-79). The share of the gender pay gap explained by firm effects remains roughly constant between age 30 and 65, whereas it is higher for younger workers.

Figure 3 shows the sectorial decomposition of the gender pay gap (Panel A), the gap in firm effects along with the estimated bargaining and sorting effects (Panel B). The sectors are coded according to Ateco 2007 sectoral codes, which is the Italian version of the sectoral codes defined by the European Union (Nace rev. 2). We report results based on the decomposition in equation (4), thus evaluating

Figure 3: Decomposition by sector

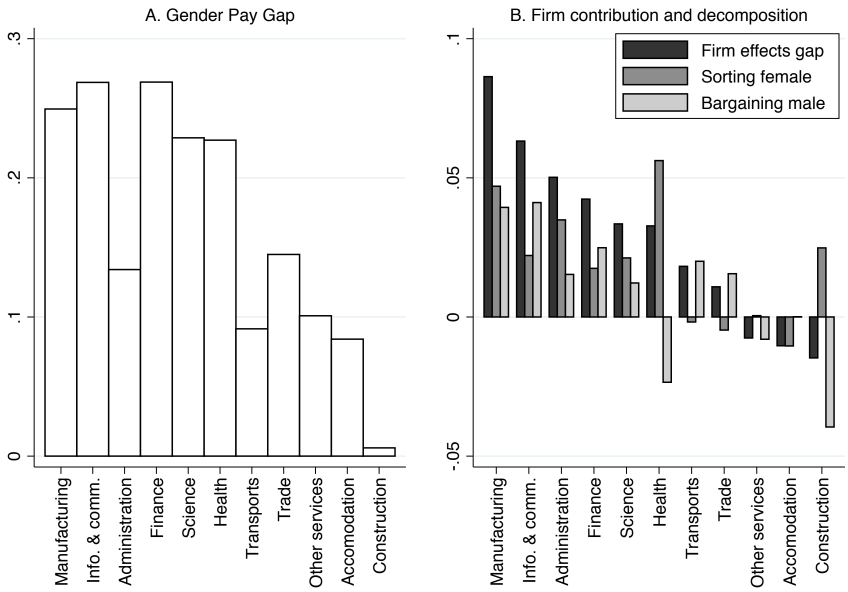

Notes. Sectors are defined according to Ateco 2007 sectoral codes and ordered according to the gap in firm effects (highest to smallest). We exclude sectors that employ less than 1% of the total workforce in all years considered. Overall, the sectors reported in the figure represent 95% of the total person-year observations between 1995 and 2015. See notes to Figure 2 for the definition of sorting and bargaining.

bargaining across male jobs and sorting across female firm effects.[23] Overall, the gender gap in earnings is the highest in the ICT and Finance sectors. This is in line with evidence for other countries (Denk, 2015). Firm effects increase the gender pay gap in all sectors, except construction, accommodation and food, and in the residual category "other services". Sorting is the main driver behind the firm contribution to the gender pay gap in manufacturing, construction, science, administration and health.[24] In finance, on the contrary, bargaining power explains a larger share of the firm effects gap relative to sorting.

In summary, in the Italian labour market, a combination of where women work and what bargaining power they have determines the opening of the gender pay gap, with the former being the more relevant driving force.

4.2.2 Decomposition Across the Earnings Distribution

As a first evidence on the magnitude of firm effects along the earnings distribution, we plot in Figure 4, the within-percentile mean male firm effect across the male distribution and the within-percentile mean female firm effect across the female distribution for 2015. The relationship is positive and monotonic for both men and women, suggesting that firm effects are a more important component of earnings for high-wage workers, irrespective of gender. The gender gap in firm effects is basically zero or negative in the very first percentile groups, but it starts to widen in the middle part of the

[23] Using the other decomposition does not change our conclusions. Results available upon request.

[24] In particular, it is likely that the results on manufacturing and trade sectors are behind the dominance of sorting in the overall dataset, being these two sectors those employing more than 50% of all person-year observations in the data.

Figure 4: Firm effects along the wage distribution

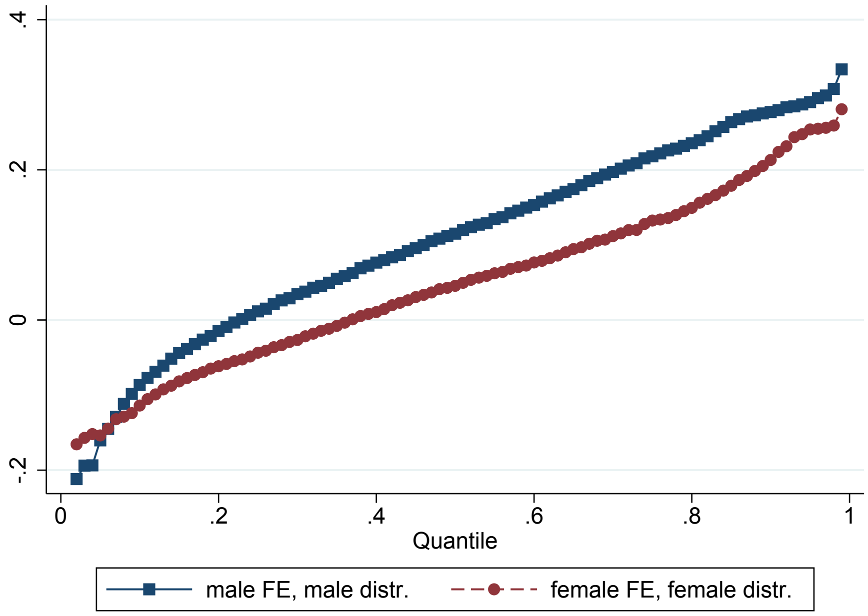

distribution. At higher percentiles the gap closes, especially in the last 10 percentile groups. The closing of the gap at the top of the pay hierarchy can be due to an increased presence of high-pay female workers in high-pay firms (thus, a better sorting) or to a higher bargaining power within firm of women relative to men.

We investigate which effect prevails, by looking at Figure 5, that shows the decomposition of the difference in firm effects for 2015. The figure shows the share of the gender pay gap at each percentile group explained by firm components. The contribution of firm effects is fairly stable in the central part of the wage distribution, and smaller at lower and higher quantiles. As to the determinants of this contribution, sorting is more relevant at the bottom and middle of the distribution. Its importance declines as we move along the distribution. On the contrary, bargaining is the most relevant factor after the 85th percentile. This result is consistent with what we found previously on occupations and sectors. Indeed, workers with higher earnings are likely to work in top positions at firms. For these workers, we find that bargaining is the main driving force behind the firm contribution to the gender gap in earnings, as we did for middle managers and executives.[25] On the other hand, for low-pay workers sorting is the main driver of the firm contribution to the gender pay gap, as we found for blue-collars.[26] High-pay workers are also likely to work in sectors with higher earnings on average. Coherently, we have shown that bargaining explains a larger portion of the gender pay gap in Finance and ICT.

In Appendix Figure D.2 we report the same decomposition for 1995, 2000, 2005 and 2010. The results are very similar.

[25]For the latter, only using the decomposition of equation (4).

[26]For apprentices, a category of workers which is likely in the bottom of the earning distribution, we find that bargaining is the main driver of the gap in firm effects. However, we do not see any spike in the bargaining at the bottom of the distribution. This is due to the small sample size of apprentices relative to blue-collar workers (4 million vs 100 million person-year observations in the overall sample). Any effect for apprentices is thus likely to be "masked" by what happens to blue-collar workers.

Figure 5: Impact of firm components on the gender pay gap along the wage distribution

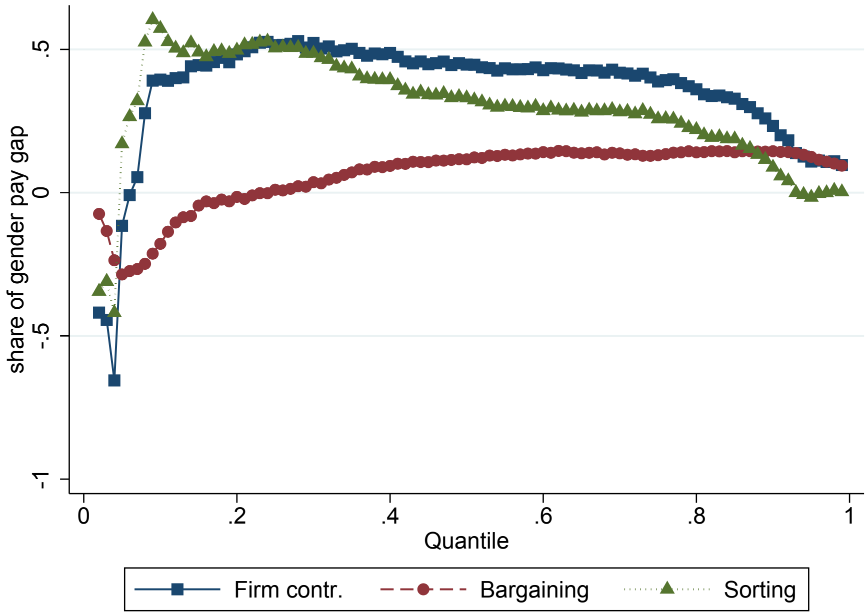

Summarising, for low income women a substantial portion of the gender pay gap is explained by where they work, whereas for high income women a larger share of the gender pay gap is due to their lower bargaining power within the firm.

4.2.3 Evolution of Firm Components over Time

Up to now, we have assumed that firm effects, individual ability and the returns to observable worker characteristics are fixed over time. However, these wage components may evolve over time and contribute to rising or declining wage inequality (Card et al., 2013, Barth et al., 2016, Song et al., 2017, Alvarez et al., 2018) and could impact differently men and women (Bruns, 2018). For example, firm effects may evolve over time due to changes in the productivity of firms or more productive firms increasingly sharing a higher portion of their rents with workers. On the other hand, individual unobserved ability may decrease over time, due to ageing (Grund and Westergaard-Nielsen, 2008), or increase thanks to components of individual productivity that slowly reveal over time or that are triggered by changes in the composition of peers (Mas and Moretti, 2009).

The availability of a long panel enables us to recover gender-specific estimates of the firm fixed effects in sub-intervals between 1995 and 2015. Specifically, we run separate AKM regressions in four overlapping intervals of six years each: 1995-2000, 2000-2005, 2005-2010, 2010-2015. For each sub-interval we build a dual connected sample as we do for the main analysis. We normalise firm effects with respect to the food and accommodation sector and analyse the evolution of the gap in firm effects and its decomposition into sorting and bargaining for each subinterval.

Results are summarised in Figure 6, where we plot the average gender pay gap, the firm effect gap, sorting and bargaining in each of the four subinterval. As explained in section 4.2.2, we take the average of sorting and bargaining estimated using equations (4) and (5).[27] Over time, as already high-

[27]Table D.2 in the Appendix reports the values used to produce Figure 6.

Figure 6: Evolution of gender pay gap, firm effects gap, sorting and bargaining over time.

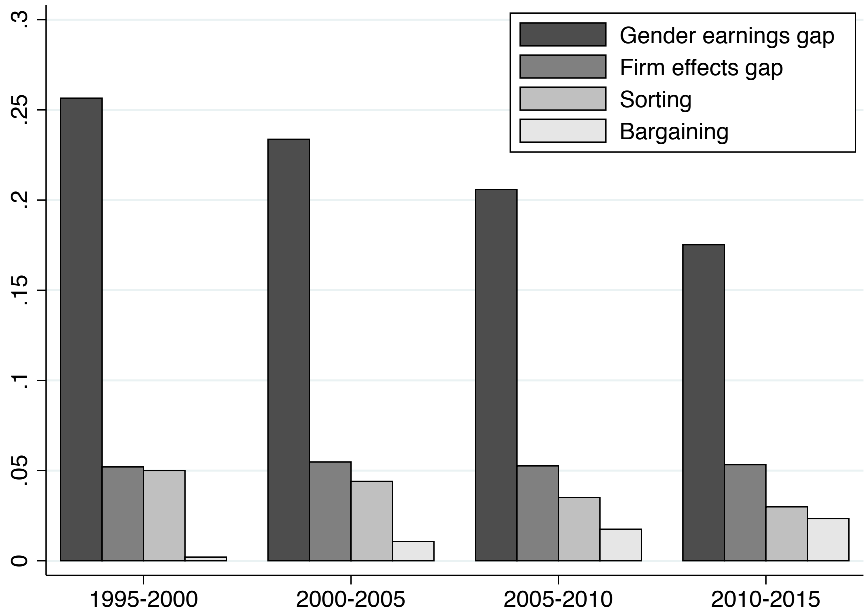

lighted, the gender gap in earnings decreases. On the contrary, the gap in firm effects is practically unchanged. Table D.2 in the Appendix shows that both male and female firm effects increase especially after 2005, but they grow at the same pace, leaving the difference unaltered. Interestingly, the impact of sorting declines over time. In the first sub-interval, sorting explains almost entirely the firm contribution to the gender pay gap (which amounts to approximately 20%), whereas very little is due to within firm differences in firm pay policies. During the period 2010-2015, the two channels have approximately equal weights in explaining the differences between male and female firm effects.[28] This result sheds light on important features of the Italian labour market. Women tend to be employed in "better" firms in more recent years, i.e. in firms with more generous pay policies towards all employees. However, the overall contribution of firm policies to the gender gap has remained unaltered because women now pay a higher penalty with respect to their male colleagues within the same firms, given the increased role of bargaining.

A possible explanation for this phenomenon is the increased role of decentralised wage setting in the Italian labour market. Historically, Italy has been characterised by a quite strongly centralised wage setting. Collective contracts are binding for employers and workers: they are signed by unions and employers' associations at the industry level and provide wage floors for each job title. Firms cannot opt-out. In 1993 a reform allowed for "top-up" agreements that can be negotiated at the regional or firm-level, usually depending on firm performance or productivity. The impact of the reform on the flexibility of bargaining agreements has been positive, although limited (Devicienti et al., 2008). Yet, additional room for firm-level bargaining can differentially impact men and women, if women have on average a lower bargaining power than men, as we have extensively shown in previous sections.

[28] In this period, whether sorting or bargaining is the main driving force behind the firm contribution to the gender pay gap in the fourth interval depends on the decomposition method adopted. See column (4) of Table D.2.

Increased female labour force participation can be another explanation. Female employment was 41.1% in 1995 against a value of 50.6% in 2015.[29] This increase may be associated with the entry of less skilled women in the labour market, whom firms may be less willing to share their rents with. If so, the estimated average bargaining power of female employees at firm level would decrease over time. At the same time, the entry of less skilled women may have favoured a reallocation of women across firms, with more skilled women moving to firms and/or jobs that better suited their competences, which would explain the reduced importance of sorting.

5 Firm-to-firm Mobility and Sorting

In this section, we further investigate the sorting channel by estimating a gender gap in the probability of moving to a "better" firm, i.e. to a firm belonging to a higher quartile of the gender-specific firm effects distribution. This is a novel definition of mobility, which takes into account the features of the origin/destination firms and we label it "gender mobility gap". The importance of gender gaps in mobility is shown to be a main driver of the gender gap in wage growth (Del Bono and Vuri, 2011), especially early in the career (Manning and Swaffield, 2008). However, the literature lacks evidence on the difference in gender mobility across firms ranked according to their degree of rent-sharing. Since we have shown that firm components account for a sizeable fraction of the gender pay gap, it is worth investigating how the difference in firm premia comes about, by focusing on mobility across different quartiles of the firm effects.

It is important to stress a point about the identification of firm effects in the AKM model. We analyse here how mobility across firms with different firm effects varies by gender. This may seem in contrast with the random mobility assumption required for the identification of firm effects in the AKM model, discussed in Appendix B. Note, however, that mobility in AKM has to be random *conditionally* on workers' time-varying observable and unobservable characteristics, which we control for by estimating firm effects conditionally on age, experience, tenure, occupation, time trends and worker fixed effects. It is therefore consistent with the random mobility assumption required by AKM that more skilled or more experienced workers move to higher fixed effect firms, because these wage components are controlled for in the estimation process. What may threaten the estimates are firm or worker transitory and permanent shocks that determine a change in earnings before the move and that determine mobility. We show in the Appendix B that these shocks are not a threat to identification in our context. Furthermore, mobility based on non-wage characteristics of firms is not problematic.[30] Mobility may also be determined by different risk preferences of workers (Argaw et al., 2017), different networks of family, friends and coworkers or different effort in on-the job search (Card et al., 2016). Hence, as long as mobility is related to non-wage components or wage components that do not change over time and are thus absorbed by the individual fixed effect or time-varying wage com-

[29]Source: Istat, Labour force survey. Employment rate for age group 20-64.

[30]Card et al. (2013), discuss, for example, mobility determined by firm amenities, proximity to home or better recruiting effort; Van Der Berg (1992) discusses the role of a number of non-wage amenities related to job changes, such as fringe benefits, moving costs and adjustment costs to a new work environment. Sorkin (2018) uses job-to-job flows to estimate the value of non-pay characteristics in earnings dispersion and find that they explain up to 15% of the variance of earnings in the United States.

ponents observable to the researcher, it can be correlated with workers' characteristics.

Our estimation strategy relies on the following probit model:

$$\Pr\left\{\mathbf{1}\left[Q^g_{f_1} > Q^g_{f_0}\right]\right\} = \Phi(\alpha + \gamma F_i + \delta Z_{it} + \lambda_t) \tag{8}$$

where Q^g_j indicates the gender-specific quartile of the distribution of firm effects to which firm $j = \{f_1, f_0\}$ belongs. $\mathbf{1}[\cdot]$ is an indicator function, equal to 1 if the destination firm f_1 belongs to a higher quartile than the origin firm f_0. F_i is a dummy for females, Z_{it} includes additional covariates (age and dummies for changing province, occupation and type of contract) and λ_t are year fixed effects.

One important aspect to take into account when analysing mobility patterns is to distinguish mobility determined by firms' closures from mobility determined by other reasons. The INPS data record for each firm the date of opening and closure. Following Del Bono and Vuri (2011), we define "firm" moves those happening in the year of firm closure or in the year before it. These are certainly constrained job moves. The other moves are classified as "individual". This does not necessarily capture a voluntary choice of the worker, since they can comprise also moves related to, say, occupations disappearing due to technological change or job downgrading following childbirth.

In order to abstract from seasonal jobs, we retain workers that move to a new firm and are observed in that firm for at least 2 consecutive years. After these restrictions, we are left with a set of 5.2 million job moves, 2.9 million of which are classified as due to individual choice. Workers can move more than once over their work career. Overall, 68% of moves in our sample refer to workers who changed job once, 28% twice. Only 4% of moves refer to workers who move three times or more (at most five) between 1995 and 2015.

Table 5 shows average marginal effects from the estimation of equation (8). The first column shows the results for the entire sample of job moves. Female workers are 4.6 percentage points less likely to move to a better firm. This differential probability is lower for "individual" moves (column 2) than "firm" moves (column 3): while the former difference is 3.4 percentage points, the latter is 5.8 p.p. In other terms, the gender gap in the probability of moving to a better firm is smaller when moves are "individual". When women are constrained to move by their firm closure, they are much less likely than men to end up in a firm with a more generous pay policy. Overall, a gender mobility gap in the likelihood of upward firm mobility is present, and economically and statistically significant. In each specification we include as additional covariates a dummy for changing province, occupation and type of contract (specifically from part- to full-time). Each of these covariates raises the probability of moving to a "better" firm. The effect of switching to full-time is the strongest across all moves, especially "individual" moves, as it raises the probability of moving to a better firm by approximately 3 percentage points.

In Appendix Figure D.3 we plot the probabilities for male and female workers of moving to higher-quartile firms by age groups. Females are less likely to move to a better firm at each age and the mobility gap is always higher for "firm" moves. It is interesting to note that the probability for workers of both genders to move to a better firm is much higher when the decision comes from individual choice rather than firm closure when the worker is young, but the decline with age in such probability is faster for individual moves. Moreover, the gender mobility gap tends to be lower for

Table 5: Probit model for job moves to a firm in a higher fixed effect quartile

	(1)	(2)	(3)
	All	Individual	Firm
Female	-0.046***	-0.034***	-0.058***
	(0.005)	(0.006)	(0.008)
Change province	0.017***	0.014***	0.004
	(0.005)	(0.005)	(0.007)
Change occupation	0.024***	0.015***	0.022***
	(0.004)	(0.004)	(0.005)
Change to full-time	0.032***	0.030***	0.011*
	(0.008)	(0.009)	(0.006)
Baseline probability	0.292	0.305	0.274
Year dummies	Yes	Yes	Yes
Observations	5,219,098	2,958,423	2,260,675

Notes. Average marginal effects from probit regressions where the dependent variable is the probability of moving to a firm in a higher firm effect quartile. Each specification includes age as control and a full set of year dummies. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: * $p<0.1$, ** $p<0.05$, *** $p<0.01$

older workers, as panel (b) shows. The gap is approximately 6 percentage points for workers aged 19-25 years old and it becomes approximately zero for workers aged 56-65.

The likelihood of moving to a higher-quartile firm might be influenced by the unobserved ability of the worker. In Figure 7 we show the probabilities for male and female workers of moving to a better firm, distinguishing between workers with "low" versus "high" individual fixed effects. We define low individual fixed effect workers those below the median of the distribution of fixed effects and high individual fixed effect workers those above the median. The figure shows that high individual fixed effect workers are more likely to move to a better firm than low fixed effect workers. Furthermore, it shows that the gender mobility gap is present and persistent across all types of move and workers, and that it is more pronounced for high individual fixed effect workers. The probability of moving to a better firm for a man in the low fixed effect group is 27.1%, whereas for a woman is 23.4%, a 3.7 p.p. gap. In the high individual fixed effect group, these probabilities are 31.2% and 25.6%, respectively, a 5.6 p.p. gap. As to "individual" moves, the gap increases from 5.1 p.p. to 6.6 p.p. for high fixed effect workers, whereas it increases from 2.2 p.p. to 4.7 p.p. for the same group of workers in the case of "firm" moves. In other words, women with higher unobserved ability are more likely to end up in a better firm than women with lower ability, but relative to men they pay a higher penalty. This is because there is a large gain in the probability of moving for more able men, whereas the gain for more able women is smaller. One reason why the gender mobility gap increases for high ability women might be related to the cost of child-rearing, which disproportionately affects women. In fact, the rise in the probability of moving up in the firm distribution of fixed effects for high relative to low ability women is approximately 52% of the corresponding rise for men. If women have to dedicate a higher fraction of their time to child rearing relative to men, irrespective of their ability level, they will search less effectively for a better job, decreasing the likelihood of ending up in a firm with more generous pay policies, which can especially benefit high-ability women.

Figure 7: Gender-specific probabilities of moving to a higher-quartile firm by worker effect and type of move.

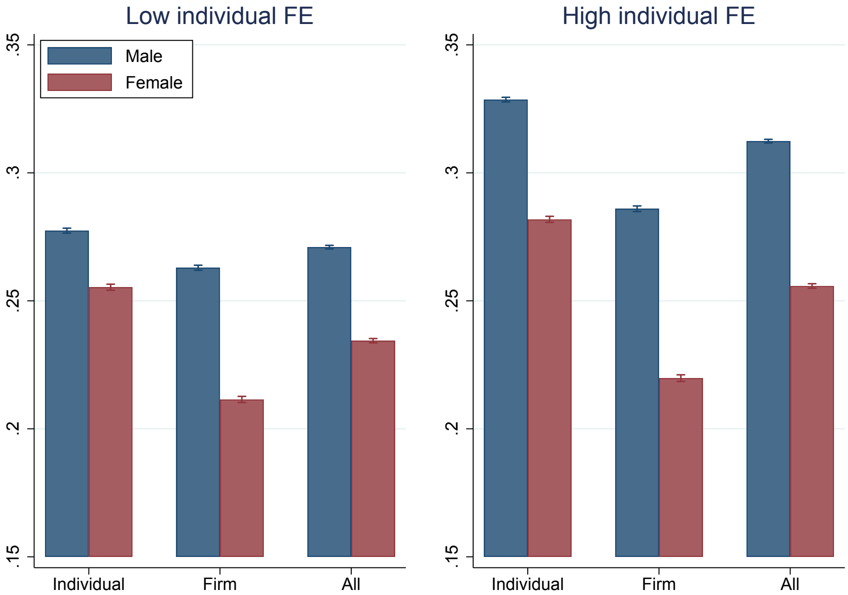

Notes. Each bar represents the average probability for men and women of moving to a firm in a higher firm effect quartile for different types of moves and workers. Low (high) individual fixed effect workers are defined as those having an individual fixed effect below (above) the median of the individual fixed effect distribution.

We explore two further reasons for why women pay a penalty in their probability of moving up in the firm effect distribution. First, women may have higher commuting costs than men. In addition, firms offering high pay policies to female workers may be geographically concentrated in some provinces or cities, where the best opportunities for women are located. On the contrary, firms offering high pay policies to men may be more geographically dispersed, providing the opportunity to improve their earnings to a wider set of (male) workers. We graphically test this hypothesis by studying the geographical distribution of average male and female firm effects. Figure D.4 in the Appendix plots the map of average male and female firm effects across Italian provinces, the lowest geographical unit our data allow to explore. A darker colour indicates a higher firm fixed effect. The maps for men and women are drawn according to the same scale and female firm effects are lower on average than male firm effects in all provinces. However, the distribution for men and women across provinces is fairly similar, meaning that provinces that pay high firm effects to men tend to do the same with women. At first, at least visually, there seems to be no difference in the geographical distribution of firm effects for men and women. More formally, we test whether the difference in mobility rates to better firms is a within or between province phenomenon. Appendix Figure D.5 plots the marginal effects from probit models estimated as in equation (8) within each province. Each dot represents the coefficient on the female dummy, together with 95% confidence intervals. The coefficients are negative for 94 of the 110 Italian provinces, and 52 of them are significantly different from 0 at a 95% confidence level. Hence, also within province men tend to move to employers that offer more generous pay policies. Interestingly, the gap is more negative in Northern provinces. This is probably due to the fact that most of the high-wage firms are located in the North, and thus movements towards

Figure 8: Gender-specific probabilities of moving to a higher-quartile firm by firm's variance of wages.

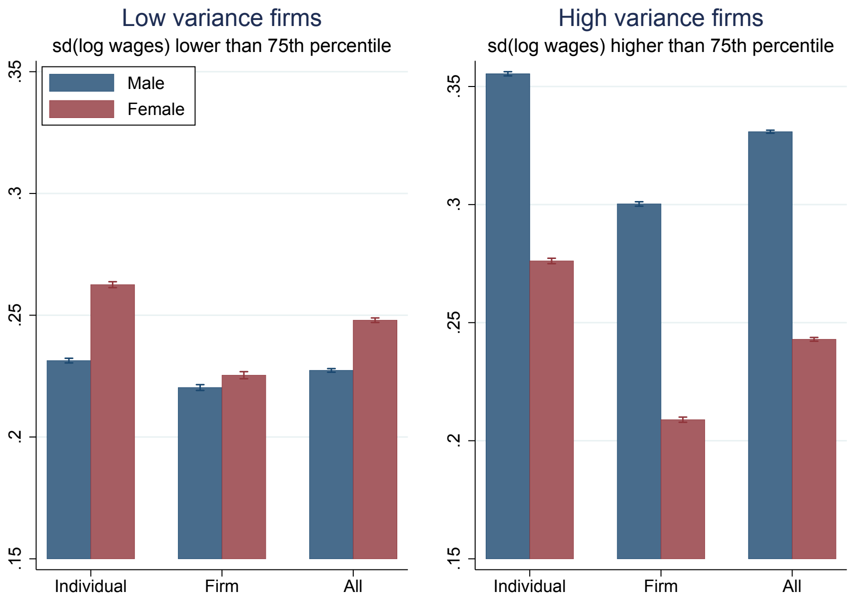

Notes. Each bar represents the average probability for men and women of moving to a firm in a higher firm effect quartile for different types of moves and firms. Low (high) variance firms are defined as those below (above) the 75th percentile of the distribution of wages' standard deviations over the period 1995-2015.

the top of the firm fixed effects distribution are more frequent there than in the South.

A second channel that may explain why women tend to move less frequently to high-wage firms is their relative risk aversion. Women may be less willing than men to move to a higher quartile firm if firms have a high dispersion of earnings. In Figure 8 we divide our sample of moves in two groups, defined by the standard deviation of firms' wages. We define as "high variance" those firms that are above the 75th percentile of the distribution of the standard deviations of earnings of all firms in the dual connected sample. Interestingly, the probability of moving to a better firm is higher for women in the low variance groups, whereas it is higher for men in the high variance group. In other terms, women tend to refrain from moving to a high-pay firm when the firm has a high dispersion of earnings, which we take as a proxy of the firm's riskiness for the worker. This result may reflect different underlying features of the labour market. First, high-variance firms may be concentrated in sectors where the proportion of women is traditionally low (e.g. Finance) and, thus, female moves in this set of firms may be quantitatively fewer than male moves. Second, high-variance firms may have a high dispersion of earnings because they pay high salaries to executives, an occupational category with a strong under-representation of women (see Table 1). Women may refrain from moving to these firms, even if they have generous pay policies, because they would not have a higher probability of accessing the top of the job ladder. Third, high-variance firms are likely those with higher competition among employees. If women suffer competition more than men (Bertrand, 2011, Heyman et al., 2013) they may be less likely to move to this type of firms.

Overall, this evidence highlights that sorting comes from the lower probability of women to move to better firms. In addition, the gender mobility gap increases as women age, irrespective of their (unobserved) skills. These results contribute to explaining the pattern in Figure 2: the gender pay gap is related to the type of firms women and men end up working in, but along the career the impact of the sorting channel tends to increase because of a gender mobility gap in the probability of changing the current job for another one in a better firm.

Understanding the reasons for the lower probability of women to move to firms with better pay policy – the more so for women with high individual fixed effects – and the age pattern behind it remains an open question. We have explored some potential explanation, finding some support for risk aversion. A change in preferences following childbirth can be another driver. Women may value more non-pecuniary benefits (such as, part-time contracts or proximity to the workplace) or may be more time-constrained in the job search. As a consequence, they may be willing to stay in or move to firms that have lower firm effects, as long as they are compensated for the loss of part of their earnings potential by a better balance between family and work.[31]

6 Gender Quotas and Rent-Sharing

We have shown that bargaining is the most important factor explaining the impact of firm pay policy on the gender pay gap at the top of the pay distribution. In addition, we have established that the role of bargaining in driving firm fixed effects has grown in importance over the two decades considered. The contribution of bargaining to gender differences in pay signals that for women it may be harder to contract not just on pay rises for a given job, but also for promotions within firms. The fact that corporate boards are male-dominated may be behind the adoption of more generous pay policy towards male employees or the fact that men are at the top of the managerial pipeline. A change in the gender composition of corporate boards may therefore modify the relative bargaining power of men and women, to the advantage of the latter, if a stronger presence of women on corporate boards increases the firm's attention towards female workers, or if female workers are more inclined to ask for increases in pay or for promotions, when the top of the corporate hierarchy is more gender balanced. Our goal is thus to analyse whether a more gender balanced board of directors has any impact on the relative bargaining power of female and male workers, as measured by a rent-sharing coefficient. Whilst some studies have focused on the impact of female-led firms on labour market outcomes of female employees and on the extent of gender gaps within the firm,[32] very little is known on the impact on their relative bargaining power.

In our modelling framework, outlined in section A in the Appendix, we show that the firm fixed effect can be rationalised in a wage equation as a rent-sharing coefficient, i.e. as the share of surplus that a firm pays to its employees. As a consequence, the firm fixed effect can be expressed as in equation (2), that is $\psi_j^g = \gamma^g \bar{S}_j$. Following Card et al. (2016), we assume that rent-sharing differs by gender, meaning that the share of surplus paid to male employees is different from the share of surplus paid to female employees.[33] We have already highlighted in Figure **??** that female firm

[31] Fanfani (2018) shows that there exists a correlation between the availability of flexible work arrangements (proxied by the share of part-time contracts) and the gap in firm pay policies in a sample of Italian manufacturing firms.

[32] For example, Bell (2005), Cardoso and Winter-Ebmer (2010), Flabbi et al. (2016), Gagliarducci and Paserman (2015).

[33] A well established literature, surveyed in Manning (2011) and Card et al. (2018), has investigated the magnitude

effects are lower than male firm effects within percentiles bins of log value added per worker. The ratio or the difference between female and male gender-specific shares, γ^F and γ^M, captures directly female bargaining power relative to men. To obtain exogenous variation in the gender composition of boards, we exploit a recent Italian law which prescribes gender quotas in boards of listed firms and which we describe in more detail in the next section. The question we ask is whether the gender quota reform had any impact on the relative bargaining power of female employees, raising their rent-sharing coefficient relative to men.

6.1 The Italian Gender Quota Reform

In 2011 the Italian parliament passed the law 120/2011 (*Golfo-Mosca reform*) with the goal of increasing the number of women present on board of directors and supervisory bodies of listed companies and state-owned not listed companies. In particular, the law requires that the Boards of Directors and the Board of Statutory Auditors must ensure "gender balance". The law is temporary, since it applies only for three consecutive board renewals (approximately 9 years) and gradual: for the first of the three board mandates, the law requires that a fifth of the seats in the board must be reserved for the least represented gender, whereas for the second and third mandates, the quota goes up to a third. Firms have to comply with the law requirements starting from the first renewal of the board after August 2012. The reform had a *phase in* period between August 2011 and August 2012, i.e. from when the law entered into force to when the requirements it prescribed became mandatory. During this period firms could comply with the law but were not required to. After August 2012, if a firm does not comply with the law, it first incurs in a warning from CONSOB, the National Commission for Companies and the Stock Exchange. After four months since the first warning, there is a fine of up to 200,000 Euro. If after three additional months the firm has not changed its board to make it compliant with the law, the elected board members lose their office.[34] The policy had a clear impact on the share of women in the boards of listed companies, as Figure 9 shows. Until 2011, the share of women in the board of directors was approximately 8%, only 1% higher than the share in 2008. The first year of implementation of the law, 2012, the share jumped to approximately 12%, and it kept rising until 32% in 2016.

6.2 Empirical Analysis

To test whether the reform had any impact on the relative rent-sharing of female employees, we depart from our model with two fixed effects and adopt instead a "reduced-form" approach which allows us to recover the rent-sharing coefficient directly from the wage equation. In order to capture the impact of the reform, the literature on the effects of gender quota highlights the importance of selecting an appropriate control group for the firms targeted by the reform (Ferrari et al., 2016; Comi et al., 2017; Bertrand et al., 2018). As in the majority of European countries that promoted legislation favouring

of the elasticities of wages to firms' financial conditions, as a way to depart from the hypothesis of competitive labour markets.

[34]For a comprehensive description of the Law, see Profeta et al. (2014).

Figure 9: Share of women in the board of directors of listed companies between 2008 and 2016.

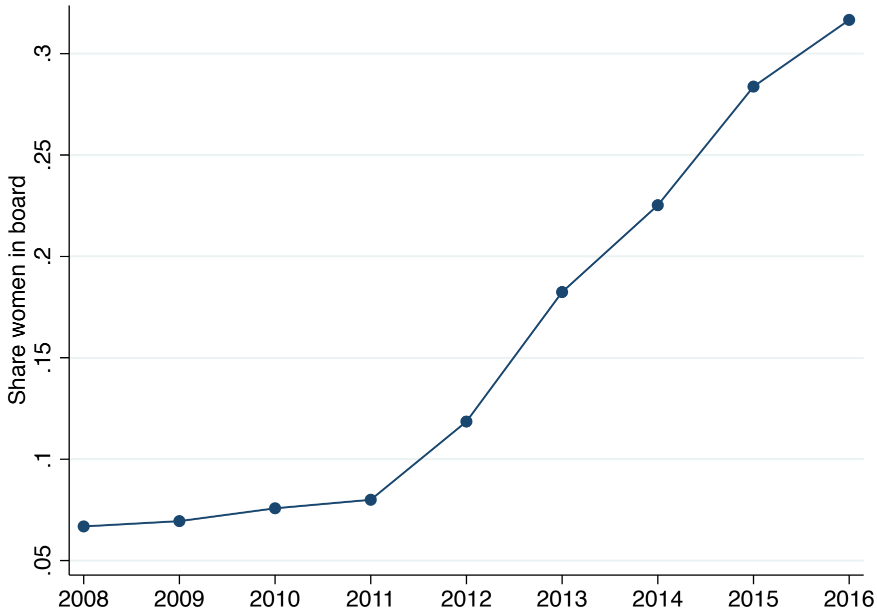

a higher presence of women at the top of the firms' hierarchy,[35] the Italian reform targeted publicly listed companies. Comparing the outcomes of listed and non-listed companies would be misleading: listed firms are different in several aspects from non-listed firms and controlling for observable characteristics may not be enough to avoid the problem. Ex-ante matching techniques can make firms in the treated and control group as comparable as possible. However, even ex-ante matching can present problems if listing is correlated with unobserved firms' characteristics that the researcher cannot control for. In this paper, we adopt a triple difference approach and build treatment and control groups by exploiting the different timing according to which the law applies to listed companies, which depends on the timing of the first board renewal after the introduction of the law. Since here we focus on employees' outcomes rather than firm-wide outcomes and we have information on the universe of workers employed by each listed company, we can restrict the sample to listed companies and still have enough power to detect sizeable effects of the reform. More specifically, we have information on boards' renewals taking place between 2012 and 2014.[36] According to the Italian legislation, firms have to renew their boards every three years, hence during the time interval 2012-2014, all firms changed their boards. We know which firms changed their boards in each of the three years, the gender composition of each board, firms' characteristics from AIDA-Bureau Van Dijk dataset and workers' characteristics from the INPS dataset. We consider all firms renewing their boards in 2012 and 2013 as treated, and firms not renewing their board in these two years, but in 2014, as control. We look at outcomes in leads. For firms renewing in 2012, we consider earnings in 2013 and 2014

[35] See Comi et al. (2017) and Profeta et al. (2014) for a detailed description of the reforms' characteristics in European countries.

[36] Since the law came into force in August 2011, but its requirements became mandatory in August 2012, firms renewing their boards in the period between August 2011 and August 2012 were not required to comply with the law, but in principle they could. We consider all firms renewing their boards in 2012 as treated. Hence, the parameter we estimate captures an intention to treat (ITT) effect. We perform robustness checks on the definition of treatment and control groups in the next section.

to estimate the rent-sharing coefficient we are interested in. For firms renewing in 2013, we look at earnings in 2014. We consider unrealistic that a board renewal in a given year has an impact on workers' outcomes in that same year. We do not consider earnings in 2015 because they may change also for firms renewing their boards in 2014, confounding the analysis.

In order to avoid potential selection effects of firms into listing or delisting due to the implementation of the reform, we focus only on firms that are continuously listed in the time window between 2009 and 2014.[37] On the workers' side, to start with, we want to abstract from potential selection effects due to the hiring of more skilled workers. Thus, we select a balanced panel of *stayers* in these firms over the period 2009-2014, in order to avoid potential biases coming from a change in the skill composition of the pool of workers. We end up with a sample of 146 firms[38] and a total number of approximately 900,000 observations.

It is important to stress that there may be potential selection of firms into treatment. In particular, firms renewing their boards before August 2012 may decide to comply with the law even if they are not required to do so.[39] If firms renewing before August 2012 and complying with the law are those with better attitudes towards women, our estimates of the effects of the reform on the rent-sharing coefficient may be upward biased. On the other hand, firms that do not anticipate compliance with the law may renew their board before August 2012 when the requirements of the law are not binding. Their inclusion in the treatment group may cause a downward bias in our estimates. As a robustness check, we define alternative treatment groups that (*i*) exclude firms renewing their board in 2012 *tout-court* or (*ii*) exclude firms with less than 20% of women in their board of directors in 2012.

We test the balance of covariates between treated and control firms, by computing the difference in means of individual and firm characteristics in the pre-reform period, i.e. between 2009 and 2011. Table 6 shows the results. The first two columns of the Table report means for each variable in treated and control units. Column (3) reports the difference between the mean for the treated and the mean for the control group. Column (4) reports the p-value of the difference. Most of the covariates show differences that are not statistically significant at conventional levels. Women tend to be more often employed with part-time contracts in control firms, whereas the share of male and female workers employed as executives is higher in treated firms. As to firm wide characteristics, average wages tend to be higher in treated firms. As to this last point, we show in Appendix Figure D.6 that there are no differential trends in the evolution of earnings in the pre-reform period, for both men and women in treated and control firms. Overall, this balance test reassures us as to the comparability of treated and control units in the pre-reform years.

We recover the DDD estimate by running the following regression:

$$\ln w_{ijt+1} = \alpha + \gamma \cdot Post_t \cdot Treat_j \cdot F_i \cdot \ln \overline{VA}_j + \zeta X_{it} + \delta Z_{jt} + \varepsilon_{ijt}. \tag{9}$$

[37] We select 2009 as a starting year, in order to have the same number of renewals in the pre and post period: firms renewing during 2012-2014 had the previous board renewal between 2009 and 2011. Hence, our choice of the time interval.

[38] The number of listed companies in Italy over the period 2009-2014 ranges between a minimum of 323 in 2012 to a maximum of 342 in 2014 (source: *Borsa Italiana*). We have information on balance sheets, board composition and renewals for 163 publicly listed companies. Of these, 17 firms are not continuously listed between 2009 and 2014.

[39] Board renewals happen during the board meeting that approves the budget. The meeting takes place within 120 days of the closing date of the fiscal year, which coincides with the calendar year in most cases.

Table 6: Balance test of covariates in the pre reform period.

	(1)	(2)	(3)	(4)
	Treated	**Control**	**Diff.**	**P-value**
Women				
Age	41.42	41.16	0.26	0.77
Experience	19.95	19.44	0.51	0.49
Part-time	0.15	0.29	-0.14	0.04**
Blue-collar	0.08	0.21	-0.12	0.49
White-collar	0.67	0.69	-0.01	0.94
Middle Manager	0.01	0.01	0.00	0.78
Executive	0.23	0.10	0.14	0.02**
Men				
Age	44.20	43.87	0.33	0.71
Experience	22.84	22.30	0.54	0.51
Part-time	0.01	0.01	-0.01	0.30
Blue-collar	0.15	0.12	0.02	0.81
White-collar	0.47	0.69	-0.22	0.12
Middle Manager	0.05	0.03	0.01	0.38
Executive	0.34	0.15	0.18	0.04**
Firms				
Board female share	0.07	0.09	-0.02	0.42
Firm size	1195.44	2498.53	-1303.08	0.34
Log value added per worker	5.05	4.86	0.18	0.44
Log weekly earnings	6.75	6.60	0.15	0.05*
North	0.13	0.16	-0.03	0.71
Centre	0.84	0.81	0.02	0.75
South	0.03	0.03	0.00	0.95

Notes. P-values computed from standard errors clustered at the firm level. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

In the above equation, $\ln w_{ijt+1}$ are log weekly earnings of worker i employed in firm j at time $t+1$ (i.e. we look at earnings in leads), $Post_t$ indicates the after reform period (i.e. after 2011, excluded), $Treat_j$ indexes listed companies that renew their board in 2012 or 2013, F_i is a dummy for female employees and $\ln \overline{VA}_j$ is the log average value added per worker at each firm over the entire period 2009-2014. We use average value added, since we think of rent-sharing as independent from firm transitory shocks that affect the companies in specific years. Taking the average over the period considered enables us to recover a clean estimate of the rents paid to the firm's workforce. We also control for individual and firm characteristics in X_{it} and Z_{jt}, respectively. We include observable worker characteristics that enter in a standard Mincerian wage equation. Namely, a cubic polynomial in age and experience, interacted with three occupation dummies,[40] and a dummy for part-time workers. As to firms, we control for region of location fixed effects, the share of women in the board pre and post reform, firm size, measured by the yearly number of employees, and sector fixed effects. We also include all the possible interactions between the treatment, time and female dummies with $\ln \overline{VA}_j$,

[40] We exclude apprentices from the analysis (approximately person-year 34,000 observations).

which are not shown in the equation, apart for the triple interaction of our interest.

Because of the inclusion of the gender dummy in the regression, we cannot control for individual unobserved heterogeneity, because gender would be collinear with an individual fixed effect. However, the sample restrictions that we perform should ensure that the firm skill composition is quite homogeneous in both treated and control group.

The coefficient of interest is γ, which measures the change in rent-sharing of female employees relative to men in treated firms after the reform was implemented. If controls are excluded, this parameter can be recovered by simply computing rent-sharing coefficients for treated and control firms in the period before and after the reform and then taking the difference between the coefficients estimated for men and those for women. More precisely:

$$\left[\left(\gamma^{F,T}_{Post}-\gamma^{M,T}_{Post}\right)-\left(\gamma^{F,T}_{Pre}-\gamma^{M,T}_{Pre}\right)\right]-\left[\left(\gamma^{F,C}_{Post}-\gamma^{M,C}_{Post}\right)-\left(\gamma^{F,C}_{Pre}-\gamma^{M,C}_{Pre}\right)\right],$$

where T and C are short-hands for treated and control groups, and M and F stand for male and female.

Results are reported in Table 7. Column (1) includes no controls. It thus represents the average change of rent-sharing for female employees in treated firms after the reform. It amounts to -0.02 and it is not significant at conventional levels. We interpret this coefficient in the following way. A 10% increase in firm log average value added translates in a 0.2 percentage points decrease in female wages with respect to male wages for treated firms in the post reform period. Column (2) adds controls for firm and individual characteristics and region fixed effects. The coefficient of interest increases slightly in magnitude and is significant at the 10% level. When we add sector fixed effects in column (3), the coefficient decreases in absolute values to 0.016 but it is not statistically significant. Overall, there is no clear evidence that the reform has an impact on female employees' rent-sharing (and, if any, the impact is negative). In column (4) we exclude executives from the analysis, the more unbalanced occupational category between treated and control firms (see Table 6). The exclusion of executives does not change the significance of the estimate, which only slightly decreases in magnitude with respect to the specification in column (3).

We perform some robustness checks by modifying the definition of treatment and control groups. In column (5) we exclude from the treatment group all firms changing the composition of their boards in 2012. Thus, we measure the average effect on the treated for firms renewing the board in 2013. The sample size reduces considerably, as expected, and the coefficient remains small and insignificant. As an alternative, in column (6) we include in the treatment group only those firms renewing their boards in 2012 whose share of women in the board of directors is above or equal to 20 per cent, and include among the controls those not satisfying the quota. Also with this alternative treatment group the coefficient is not different from zero at any conventional significance level.

Our results indicate that the reform which brought about a larger share of women in corporate boards does not improve rent-sharing for women compared to men. There is mild evidence in the opposite direction, which is however not significant when sector fixed effects are included.

In order to study whether the reform had any impact in one specific year and to exclude the presence of pre-reform trends, we report the results of an event study analysis, which shows that our outcome variable does not differ across the treated and control firms in the pre-reform period, whereas

Table 7: DDD estimates of relative rent-sharing after the gender quota reform.

	(1) No controls	(2) Controls	(3) Sector FE	(4) No exec.	(5) No 2012	(6) Alt. treat.
$F_i \cdot Post_t \cdot Treat_j \cdot \ln \overline{VA}_j$	-0.020	-0.023*	-0.016	-0.012	-0.003	-0.008
	(0.013)	(0.013)	(0.012)	(0.011)	(0.012)	(0.013)
Observations	750,480	750,480	750,395	726,330	600,316	750,395
R-squared	0.159	0.694	0.721	0.648	0.718	0.721
Controls	No	Yes	Yes	Yes	Yes	Yes
Region fixed effects	No	Yes	Yes	Yes	Yes	Yes
Sector fixed effects	No	No	Yes	Yes	Yes	Yes

Notes. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

differences in rent-sharing – if any – appear after the reform. More precisely, in Figure 10 we plot the coefficients of an interaction between a dummy for female workers, a dummy for treated firms, the average log value added per worker of the firm and year dummies. We normalise coefficients with respect to the year before the reform came into force, i.e. 2011. The left panel reports results from a specification similar to column (2) in Table 7, thus controlling for observable firm and individual characteristics and region fixed effects. The pre-reform coefficients are positive but not statistically significant using 90% confidence intervals. This provides evidence of the absence of pre-reform differences in the trends of our coefficient of interest between treated and control firms. In the first year of the reform the coefficient drops and is statistically significant and remains negative, but not significant, in the following year. The right panel additionally controls for sector fixed effects. As in Table 7, adding sector fixed effects reduces the coefficients in the post reform period, which become not significantly different than zero at any conventional level. However, we can reassuringly reject the hypothesis of differential trends in the pre reform period.

Overall, the impact of the reform on the relative female-male rent-sharing is null. This may be explained by the limited time we observe the effects of the reform and by the fact that a change in the remuneration policies in a firm does not immediately react to a change in the gender composition of the board, especially in a rigid labour market as Italy. It could also signal that male and female workers do not modify significantly their bargaining behaviour towards the firm when the board composition changes. Though we are the first to investigate the impact of gender quotas on bargaining at the worker level, null results on workers' outcomes are also found by Bertrand et al. (2018) in the case of Norway. This evidence does not provide a conclusive answer on the effectiveness of gender quotas on the relative bargaining power of female employees.

7 Concluding Remarks

Thanks to a large matched employer-employee dataset on the universe of Italian workers and private sector firms for the period 1995-2015, we investigate the contribution of firms to the gender pay gap

Figure 10: Event study analysis.

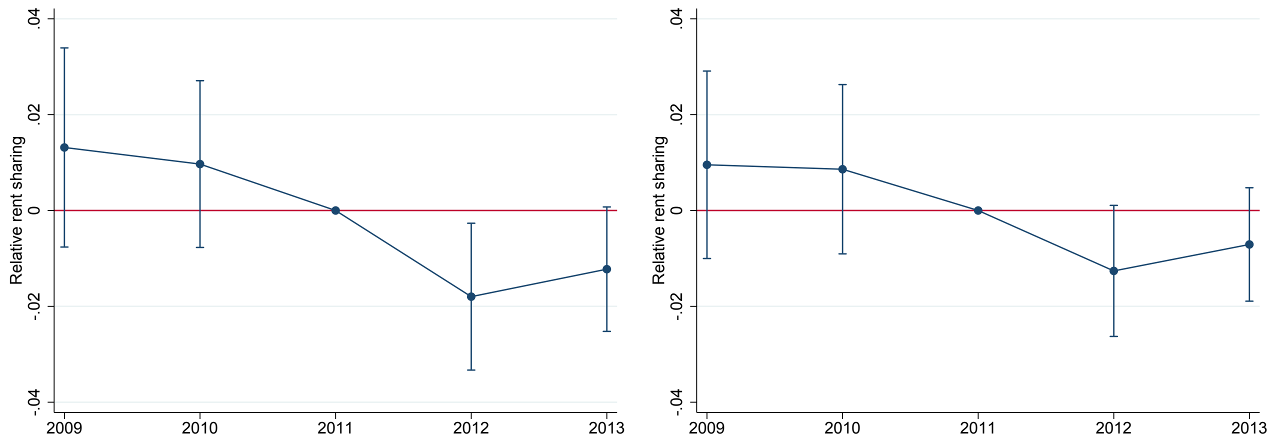

(a) Sector fixed effects not included

(b) Sector fixed effects included

and we find that firm effects play a significant role. Firm characteristics account for approximately 30% of the total average gender pay gap. We single out the primary determinants of this contribution into sorting and bargaining, by applying a standard Oaxaca-Blinder decomposition. Results show that firm effects operate mainly via sorting, which accounts for roughly 20-22% of the total gender pay gap. The contribution of bargaining is around 8-9%. When we look at different age groups and cohorts, we show that firm components are more important when workers are older. Sorting is the main driver of the firm contribution to the gender pay gap also across sectors, with important exceptions in high-pay industries, such as finance.

As to occupations, blue-collar workers are the most affected by sorting. Managers are, instead, a category of workers for which bargaining explains a larger share of the gender pay gap. This is consistent with our findings across the distribution of earnings: the effect of sorting is uniformly larger than that of bargaining at lower and middle quantiles, whereas the opposite is true at the top of the distribution, where it is more likely that managers are placed.

The importance of sorting and bargaining in explaining the firm contribution to the gender gap in earnings evolves over time. We show that, in the first years in our data, sorting explains entirely the gender gap in firm pay policies. More recently, bargaining has gained importance. Changes in sorting and bargaining are such that the overall gap in firm pay policies is unaltered during the period under analysis.

When we analyse the drivers of sorting, we find that a gender mobility gap is present and persistent, with women – especially high ability ones – displaying a lower likelihood of moving to better paying firms, compared to men with similar characteristics. We investigate gender differences in risk aversion and in access to more generous firms as possible channels and find some evidence that the former is at work. As to bargaining, we show that the law 120/2011, which requires gender balance in the composition of corporate boards of listed companies, did not modify the bargaining power of female versus male workers, as measured by the change in the rent-sharing coefficient in firms affected by the reform.

Other mechanisms may drive differences in workplace-related inequality beyond those analysed in this paper. For example, better peers can boost the productivity of workers and, thus, their wages. If better peers impact men and women differently, we can observe gender differences in earnings.

The increased availability of linked employer-employee data will help identifying and exploring the different channels, providing a solid ground on which to build policy recommendations to reduce obstacles to female career advancements.

References

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002). Computing person and firm effects using linked longitudinal employer-employee data. *Longitudinal Employer-Household Dynamics Technical Papers 2002-06*.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 2*, 251–333.

Albrecht, J., A. Bjorklund, and S. Vroman (2003). Is There a Glass Ceiling in Sweden? *Journal of Labor Economics 21*(1), 145–177.

Albrecht, J., P. S. Thoursie, and S. Vroman (2015). Parental Leave and the Glass Ceiling in Sweden. In *Gender Convergence in the Labor Market*, Chapter 3, pp. 89–114.

Altonji, J. G. and R. M. Blank (1999). Race and Gender in the Labor Market. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3 of *Handbook of Labor Economics*, Chapter 48, pp. 3143–3259. Elsevier.

Alvarez, J., F. Benguria, N. Engbom, and C. Moser (2018). Firms and the Decline in Earnings Inequality in Brazil. *American Economic Journal: Macroeconomics 10*(1), 149–89.

Argaw, B. A., M. F. Maier, and O. J. Skriabikova (2017). Risk attitudes, job mobility and subsequent wage growth during the early career. Technical report.

Arulampalam, W., A. L. Booth, and M. L. Bryan (2007). Is There a Glass Ceiling Over Europe? Exploring the Gender Wage Gap Across the Wage Distribution. *Industrial and Labor Relations Review 60*.

Atkinson, A. B., A. Casarico, and S. Voitchovsky (2018). Top incomes and the gender divide. *The Journal of Economic Inequality 16*(2), 225–256.

Azmat, G. and B. Petrongolo (2014). Gender and the Labor Market: What Have We Learned from Field and Lab Experiments? *Labour Economics 30*, 32–40.

Babcock, L., M. Gelfand, D. Small, and H. Stayn (2006). Gender differences in the propensity to initiate negotiations. In M. Z. D. De Cremer and J. Murnighan (Eds.), *Social psychology and economics*, Chapter 13, pp. 238–259. Mahwah, NJ: Lawrance Erlbaum.

Barth, E., A. Bryson, J. C. Davis, and R. Freeman (2016). It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States. *Journal of Labor Economics 34*(S2), S67–S97.

Barth, E., S. P. Kerr, and C. Olivetti (2017). The Dynamics of Gender Earnings Differentials: Evidence from Establishment Data. Working Paper 23381, National Bureau of Economic Research.

Bayard, K., J. Hellerstein, D. Neumark, and K. Troske (2003). New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data. *Journal of Labor Economics 21*(4), 887–922.

Bell, L. A. (2005, July). Women-Led Firms and the Gender Gap in Top Executive Jobs. IZA Discussion Papers 1689, Institute for the Study of Labor (IZA).

Bertrand, M. (2011). New Perspectives on Gender. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4, Chapter 17, pp. 1543–1590. Elsevier.

Bertrand, M., S. E. Black, S. Jensen, and A. Lleras-Muney (2018). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway. *The Review of Economic Studies*, forthcoming.

Black, S. E. and E. Brainerd (2004). Importing Equality? The Impact of Globalization on Gender Discrimination. *Industrial and Labor Relations Review 57*(4), 540–559.

Blau, F. and L. Kahn (2003). Understanding International Differences in the Gender Pay Gap. *Journal of Labor Economics 21*(1), 106–144.

Blau, F. D. (2012). *Gender, Inequality, and Wages*. Edited by Anne C. Gielen, and Klaus F. Zimmermann, Oxford University Press.

Blau, F. D. and L. Kahn (2016). The Gender Wage Gap: Extent, Trends, and Explanations. *IZA Discussion Papers 9656*.

Blau, F. D. and L. M. Kahn (1997). Swimming Upstream: Trends in the Gender Wage Differential in the 1980s. *Journal of Labor Economics 15*(1), 1–42.

Blau, F. D. and L. M. Kahn (2000). Gender differences in pay. *Journal of Economic Perspectives 14*(4), 75–99.

Blau, F. D. and L. M. Kahn (2006). The U.S. Gender Pay Gap in the 1990s: Slowing Convergence. *Industrial and Labor Relations Review 60*(1), 45–66.

Blinder, A. (1973). Wage Discrimination: Reduced Form and Structural Estimates. *Journal of Human Resources 8*, 436–455.

Bøler, E. A., B. Javorcik, and K. H. Ulltveit-Moe (2018). Working across Time Zones: Exporters and the Gender Wage Gap. *Journal of International Economics, forthcoming*.

Bowles, H. R., L. Babcock, and L. Lai (2007). Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask. *Organizational Behavior and Human Decision Processes 103*(1), 84–103.

Bowles, H. R., L. Babcock, and K. L. McGinn (2005). Constraints and triggers: Situational mechanics of gender in negotiation. *Journal of Personality and Social Psychology 89*(6), 951–965.

Bruns, B. (2018). Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap. *American Economic Journal: Applied Economics*, forthcoming.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *The Econometrics Journal 14*(3), 368–386.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70.

Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *Quarterly Journal of Economics 131*(2), 633–686.

Card, D., J. Heining, and P. Kline (2013). Workplace Heterogeneity and the Rise of West German Wage Inequality. *Quarterly Journal of Economics 128*, 967–1015.

Cardoso, A. R., P. Guimaraes, and P. Portugal (2016). What drives the gender wage gap? A look at the role of firm and job-title heterogeneity. *Oxford Economic Papers 68*(2), 506–524.

Cardoso, A. R. and R. Winter-Ebmer (2010). Female-led firms and gender wage policies. *Industrial and Labor Relations Review 64*(1), 143–163.

Comi, S., M. Grasseni, F. Origo, and L. Pagani (2017). Where Women Make The Difference. The Effects of Corporate Board Gender Quotas on Firms' Performance across Europe. Working Papers 367, University of Milano-Bicocca, Department of Economics.

Coudin, E., S. Maillard, and M. To (2018). Family, firms and the gender wage gap in France. IFS Working Papers W18/01, Institute for Fiscal Studies.

Del Bono, E. and D. Vuri (2011). Job Mobility and the Gender Wage Gap in Italy. *Labour Economics 18*(1), 130–142.

Denk, O. (2015). Financial Sector Pay and Labour Income Inequality. OECD Economics Department Working Papers 1225, OECD Publishing, Paris.

Devicienti, F., A. Maida, and L. Pacelli (2008). The resurrection of the Italian wage curve. *Economics Letters 98*(3), 335 – 341.

Fanfani, B. (2018). Tastes for Discrimination in Monopsonistic Labour Markets. Department of Economics and Statistic Working Paper Series 54, University of Torino.

Ferrari, G., V. Ferraro, P. Profeta, and C. Pronzato (2016). Gender Quotas: Challenging the Boards, Performance, and the Stock Market. CESifo Working Paper 6084, Munich.

Flabbi, L., M. Macis, A. Moro, and F. Schivardi (2016). Do Female Executives Make a Difference? The Impact of Female Leadership on Gender Gaps and Firm Performance. Working Paper 22877, National Bureau of Economic Research.

Fortin, N. M., B. Bell, and M. Bohm (2017). Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom. *Labour Economics 47*, 107–123.

Gagliarducci, S. and D. Paserman (2015). The Effect of Female Leadership on Establishment and Employee Outcomes: Evidence from Linked Employer-Employee Data. *Research in Labor Economics 41*, 341–372.

Gobillon, L., D. Meurs, and S. Roux (2015). Estimating Gender Differences in Access to Jobs. *Journal of Labor Economics 33*(2), 317–363.

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

Gregory, M. B. (2009). Gender and Economic Inequality. In W. Salverda, B. Nolan, and T. M. Smeeding (Eds.), *The Oxford Handbook of Economic Inequality*. Oxford University Press.

Groshen, E. (1991). The Structure of the Female/Male Wage Differential: Is It Who You Are, What You Do, or Where You Work? *Journal of Human Resources 26*(3), 457–472.

Grund, C. and N. Westergaard-Nielsen (2008). Age structure of the workforce and firm performance. *International Journal of Manpower 29*(5), 410–422.

Guvenen, F., G. Kaplan, and J. Song (2014). The Glass Ceiling and The Paper Floor: Gender Differences among Top Earners, 1981-2012. Working Paper 20560, National Bureau of Economic Research.

Heyman, F., H. Svaleryd, and J. Vlachos (2013). Competition, Takeovers, and Gender Discrimination. *Industrial and Labor Relations Review 66*(2), 409–432.

Ludsteck, J. (2014). The Impact of Segregation and Sorting on the Gender Wage Gap: Evidence from German Linked Longitudinal Employer-Employee Data. *Industrial and Labor Relations Review 67*(2), 362–394.

Manning, A. (2011). Imperfect Competition in the Labor Market. Volume 4 of *Handbook of Labor Economics*, Chapter 11, pp. 973–1041. Elsevier.

Manning, A. and J. Swaffield (2008). The Gender Gap in Early-Career Wage Growth. *Economic Journal 118*(530), 983–1024.

Mas, A. and E. Moretti (2009, March). Peers at work. *American Economic Review 99*(1), 112–45.

Oaxaca, R. L. (1973). Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review 14*, 693–709.

Olivetti, C. and B. Petrongolo (2016). The Evolution of Gender Gaps in Industrialized Countries. *Annual Review of Economics 8*(1), 405–434.

Petersen, T. and L. A. Morgan (1995). Separate and Unequal: Occupation-Establishment Sex Segregation and the Gender Wage Gap. *American Journal of Sociology 101*(2), 329–365.

Piketty, T., E. Saez, and G. Zucman (2018). Distributional National Accounts: Methods and Estimates for the United States. *Quarterly Journal of Economics 133*(2), 553–609.

Ponthieux, S. and D. Meurs (2015). Gender Inequality. In A. B. Atkinson and F. Bourguignon (Eds.), *Handbook of Income Distribution*, Volume 2 of *Handbook of Income Distribution*, Chapter 12, pp. 981 – 1146. Elsevier.

Profeta, P., L. Aliberti, A. Casarico, M. D'Amico, and A. Puccio (2014). *Women Directors: The Italian Way and Beyond*. Palgrave Macmillan UK.

Rozada, M. G. and E. L. Yeyati (2018). Do Women Ask for Lower Salaries? The Supply Side of the Gender Pay Gap. Working Paper 4/2018, School of Government, Universidad Torcuato Di Tella.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2017). Firming Up Inequality. Working Paper 21199, National Bureau of Economic Research.

Sorkin, I. (2018). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics 133*(3), 1331–1393.

Van Der Berg, G. J. (1992). A Structural Dynamic Analysis of Job Turnover and the Costs Associated with Moving to Another Job. *The Economic Journal 102*(414), 1116–1133.

Appendix

A Modelling Framework

The model follows Card et al. (2016). We assume that log earnings of workers can be written as:

$$w_{it} = a_{it} + \gamma^g S_{ijt}, \tag{10}$$

where a_{it} is an outside option for worker i at time t, S_{ijt} is the match surplus between worker i and firm j at time t and γ^g is the share of this surplus paid to worker i of gender $g = M,F$.[41] We assume that S_{ijt} can be written as follows:

$$S_{ijt} = \overline{S}_j + \phi_{jt} + m_{ij}, \tag{11}$$

i.e. as the sum of average surplus $\overline{S}_j$ for all employees at firm j (due to, say, market power or brand recognition), time-varying factors ϕ_{jt} that raise or lower average surplus for all employees, and a match specific component m_{ij}.

We also assume that the outside option a_{it} can be written as:

$$a_{it} = \theta_i + X'_{it}\beta^g + u_{it}, \tag{12}$$

where θ_i is individual ability (and, in our specific case, returns to education as well), X'_{it} are time-varying observable characteristics and u_{it} is a transitory component.

Replacing (11) and (12) into (10), we get:

$$w_{it} = \theta_i + \psi_j^g + X'_{it}\beta^g + \varepsilon_{it} \tag{13}$$

where

$$\psi_j^g = \gamma^g \overline{S}_j \tag{14}$$

$$\varepsilon_{it} = \gamma^g \left(\phi_{jt} + m_{ij}\right) + u_{it} \tag{15}$$

Equation (13) is consistent with the two-way fixed effects model *à la* Abowd et al. (1999) presented in equation (1) in the main text.

B Non Parametric Tests of Conditional Random Mobility

One important feature of Abowd et al. (1999) two-way fixed effects model is the assumption of *conditional random mobility*. This is a requirement for the validity of the OLS estimation of model (1), which provides consistent estimates if and only if:

$$\mathbb{E}(\mathbf{D}\varepsilon) = \mathbb{E}(\mathbf{F}\varepsilon) = \mathbb{E}(\mathbf{X}\varepsilon) = 0$$

[41] We use j as a shorthand for $J(i,t)$, i.e. the firm that employs worker i at time t.

Figure B.1: Mean weekly earnings of movers across firm effects quartiles

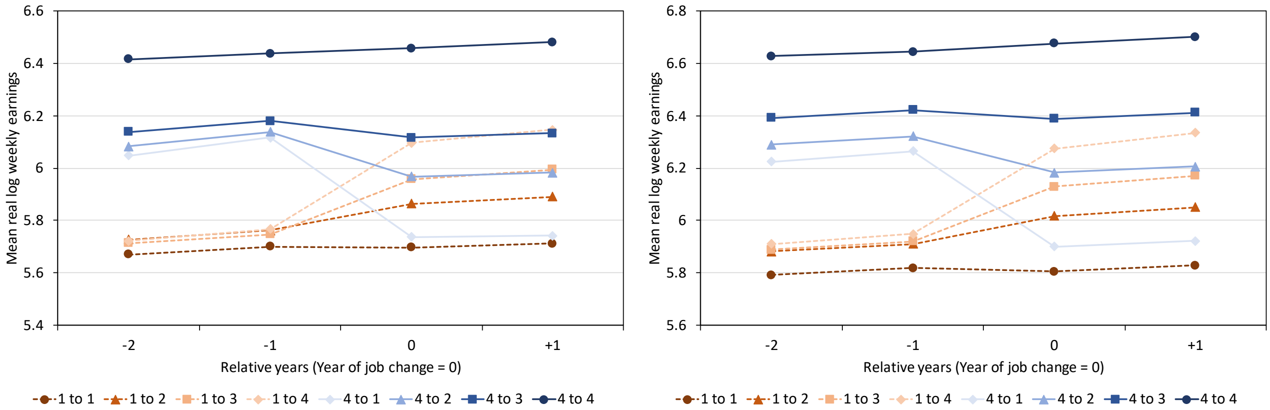

(a) Female sample

(b) Male sample

where $\mathbf{D}$ is a $(N^* \times N)$ matrix of dummies for the N individuals in the sample (N^* is the total number of person-year observations), $\mathbf{F}$ is a $(N^* \times J)$ matrix of dummies for the J firms constituting the sample, $\mathbf{X}$ is the $(N^* \times K)$ matrix of regressors. ε is the matrix of errors, where observations are stacked across individuals and time.

We focus here on the restriction imposed on the matrix of firms' dummies. Following Card et al. (2013), there are three main channels through which conditional random mobility may be violated. First, workers employed at firms that are experiencing negative shocks may decide to move to firms that are experiencing positive shocks: this generates correlation between ϕ_{jt} and the probability that worker i is employed at firm j at time t in equation (15). If this is the case, workers would experience a drop in earnings before the move, and a sudden rise in pay after. We show in Figure B.1 that this is not the case. Specifically, we build a sample of moves and compute mean weekly earnings associated with changes from the first and the last quartile of firm effects.[42] We see that for both women and men, shown in panels (a) and (b), respectively, there are no changes in the evolution of mean earnings before or after the move.

A second threat to identification comes from the presence of match effects, if workers decide to move because they think that joining a new firm would deliver a better match between their personal characteristics and the firm characteristics compared to the firm of origin. This violation implies that the match component m_{ij} in equation (15) is correlated with the probability that worker i is employed at firm j at time t. In the presence of correlation, movers would experience in any case a wage gain, irrespective of whether they move from a high-wage to a low-wage firm, or the opposite. On the other hand, if match effects are unimportant in determining mobility, then the earnings gain associated with moves from low- to high-earnings firms should be roughly comparable in magnitude to the earnings loss for moves in the opposite direction. This symmetry in gains and losses with each opposite move is

[42] We identify low-wage and high-wage firms on the basis of the quartiles of the estimated firm effects. We then assign each job mover to the corresponding quartile of the origin and destination firm. This way we identify sixteen cells of movers, each one corresponding to the pair origin-destination quartile (4×4 cells). Within each cell, we compute the mean log real weekly earnings of movers. We just retain movers that are continuously observed in the two years prior to the move and in the two years after, similarly to what we have done in section 5. Means are computed within each year. Data on the mean earnings for all the moves are reported in Table D.3.

Figure B.2: Adjusted change in earnings of symmetric job moves across firm effects quartiles

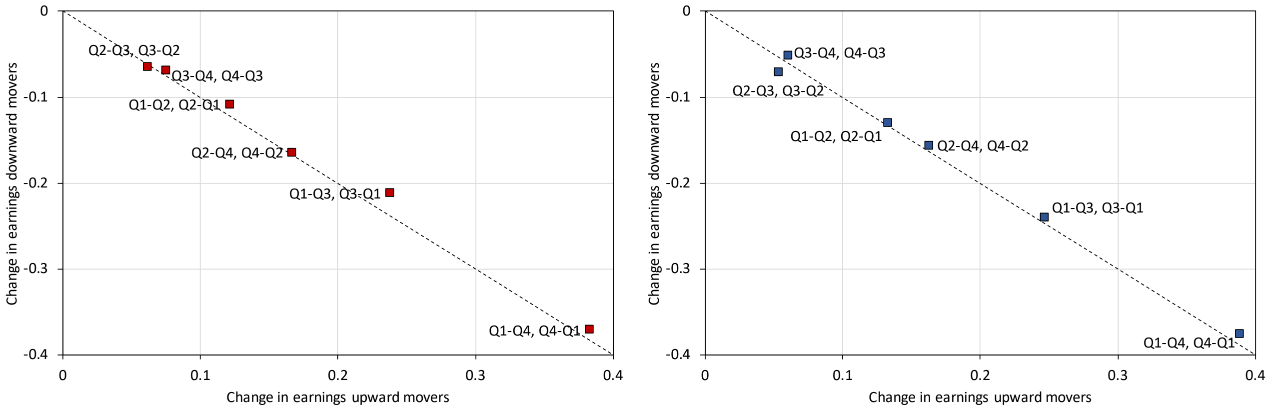

(a) Female sample

(b) Male sample

better assessed examining the magnitudes of such changes over the entire 4 year period under analysis and for all possible moves, looking at the difference in earnings from the first period considered (2 years prior to the move) to the last period (one year after). This boils down to comparing the overall earnings change (earnings one year after *minus* earnings two years before) for opposite moves.[43] Such comparisons are displayed, for the female and male sample respectively, in panels (a) and (b) of Figure B.2, where we plot the adjusted earnings changes[44] for downward movers against the adjusted earnings changes for upward movers. In both panels opposite moves display the expected degree of symmetry, that is, they are in all cases of opposite sign. Moreover, all scatter points cluster very close to the 45 degrees line, meaning that each symmetric move, both upward and downward, generates an earnings change of a similar magnitude. Therefore, we deem symmetry as a reasonable assumption.

As an additional check, panels (a) and (b) of Figure B.3 report the earnings evolution for the movers within the same quartile in the origin and destination firms. If it is true that there are no match effects in mobility, then these movements should be characterized by almost no earnings gains. This is indeed the case: both panels show that earnings evolution is basically flat for within-quartile movements. This is clearly inconsistent with specific worker-firm match gains related to job changes.

A last threat to the identification of firm effects comes from individual transitory shocks, that generate correlation between the transitory component u_{it} in equation (15) and the probability that worker i is employed at firm j at time t. If workers are experiencing an increase in their earnings before the move because of some productivity premium associated to a transitory change in their characteristics or to some of their skills showing up after an accumulation period, then they might move to other firms that reward these characteristics more, with a larger gain from the move compared to that obtained in the origin firm. On the other hand, if the transitory shock is negative, workers might experience an earnings decline in their origin firm and therefore move to firms that would limit such decline, because better suited to reward their characteristics. We can refer again to Figure B.1, where, if mobility is driven by individuals recognising their higher (lower) productivity we should

[43] Opposite moves are those from a quartile k to a quartile j and the other way around.

[44] Adjusted earnings changes are obtained by subtracting the earnings change for movers within the same quartile to each of the other raw earnings changes, as in Flabbi et al. (2016): that is, we subtract the change for movers from quartile q to quartile q to the change for movers from the q-*th* quartile.

Figure B.3: Mean weekly earnings of movers within same firm effects quartiles

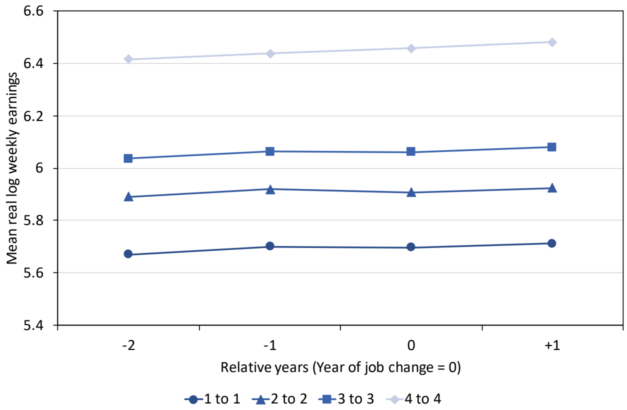

(a) Female sample

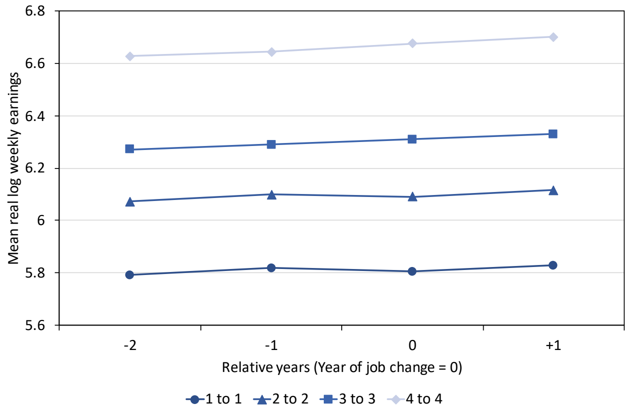

(b) Male sample

see unusual earnings growth before the move for people moving towards the top and unusual earnings decrease for people moving in the opposite direction. Nothing like that happens in the data. Both pictures show no trend before the movements.

As a final check, we follow again Card et al. (2013) and examine residuals from model (1) for different groups of individual effects in different groups of firm effects. Namely, we define deciles of both person and firm effects and compute the mean estimated AKM residuals in each of the 100 cells defined by the combination of worker and firm deciles. If our model is incorrectly specified, because, for instance, it is missing some important match component between specific individuals and firms, we would expect to find high mean residuals in those cells that are threatened the most by misspecification. Figure B.4 plot the mean residuals for each of the person-firm cells for females and males in panels (a) and (b), respectively. For both samples the deviations are really small in magnitude and exceed 1 log point only in one case (the cell defined by the first decile of both person and firm effects). Overall, we find no evidence against the conditional random mobility assumption in both the male and female samples.

Figure B.4: Mean AKM residuals across deciles of person and firm effects

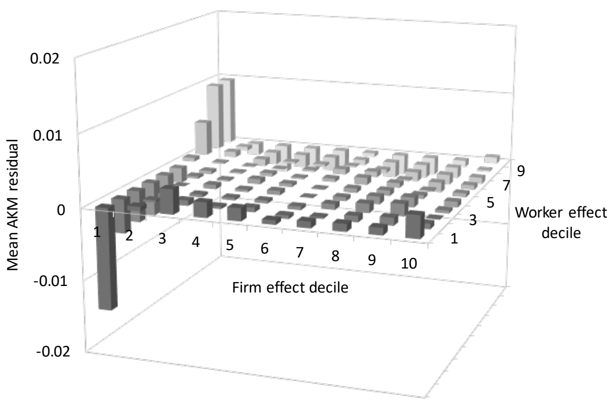

(a) Female sample

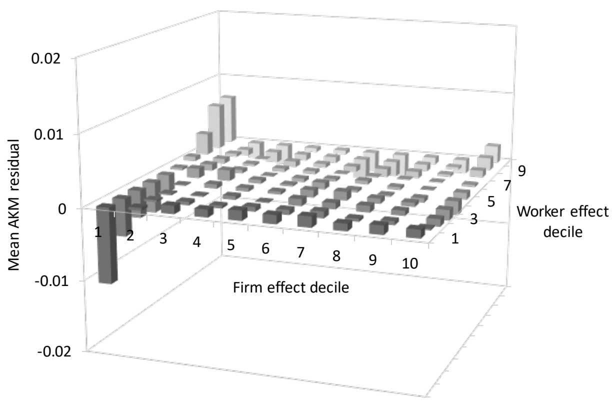

(b) Male sample

C Alternative Normalisation of Firm Effects

The magnitude of the bargaining channel depends on the specific constant chosen to normalise male and female firm effects. In the main text, coherently with our assumption that low surplus firms pay zero rents to their workforce, we identify a low surplus sector and subtract from the estimated firm effects the average firm effect in this sector. We follow the literature (Card et al., 2016, Coudin et al., 2018) and set to zero the average firm effect in the food and accommodation sector. However, the results we find on the relative contribution of firms to the gender pay gap – and its decomposition into bargaining and sorting – may depend on this normalisation choice.

We adopt here a different normalisation approach and check that our results do not change. Specifically, we assume that firm effects represent a rent-sharing component embedded in the determination of earnings (as in equation 2). Thus, we merge INPS data with balance sheet information from AIDA-Bureau Van Dijk and visually inspect this relationship. We measure firm's average surplus with average log value added per worker over the longest period available for each firm.[45] Figure C.1 plots the relationship between male and female firm effects against average log value added per worker.[46] The relationship is clearly positive and, as value added increases, female firm effects increase less than male firm effects (as highlighted also in Figure **??** in the main text). Moreover, the relationship is rather flat in the first 10 percentiles of value added and only after this threshold it starts to be increasing.[47] Hence, we choose to normalise firm effects with respect to the average firm effect of firms in the first decile of the distribution of log value added per worker.

Figure C.1: Firm effects against log value added per worker.

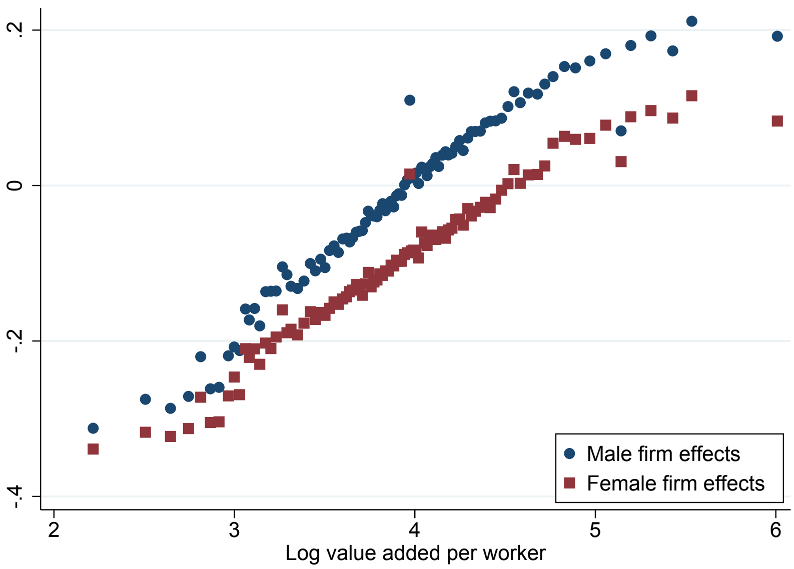

[45]The coverage of balance sheet data in AIDA is limited in the 1990s and early 2000s. The use of average value added allows us to fill missing entries with average quantities. For some firms we have no information on value added. Overall, out of the 183,062,088 person-year observations in the dual connected sample, we have missing information for 39,986,670 person-year observations.

[46]We arbitrarily normalise firm effects with respect to the largest firm in the dual connected sample in terms of number of employees in a year. To improve readability we average firm effects into percentile groups of log value added per worker.

[47]The threshold equals approximately a log value added per worker of 3.

We decompose firm effects as in equations (4) and (5). Results are reported in Table C.1. With this alternative normalisation, the impact of firm components on the gender pay gap increases. The difference in firm effects accounts for 38% of the gap in weekly earnings, a 7.3 percentage points rise with respect to our preferred normalisation in the main text. Though sorting still dominates, the increase in bargaining explains the larger impact of firm effects,[48] which accounts for as much as 17% of the gender pay gap. Also when we decompose firm components by occupation, the bargaining channel increases in magnitude. In particular, for apprentices, differences in pay policies within firm explain between 70% and 79% of the gender pay gap. For blue and white collar workers the impact is lower, around 15% and 11-13%, respectively. For middle managers, bargaining is the main driving force behind the firm contribution to the gender pay gap, as already highlighted in the text. Here the contribution of bargaining to the gender pay gap for this category of workers is slightly higher. For executives, the driver of the firm effects gap is sorting if one uses the decomposition in equation (4) and bargaining if one uses equation (5), but the estimate of bargaining is higher.

Overall, the main conclusions do not change. This alternative normalisation shows that our estimate of the bargaining channel in the main text can be interpreted as a lower bound. However, we prefer the normalisation with respect to the food and accommodation sector because we have information on sectors for *all* firms in our sample, whereas we lose around 20% of person-year observations in the normalisation based on log value added per worker. Hence, the alternative normalisation is based on a subset of firms in our data. Since the main conclusions remain unchanged we prefer keeping as many observations as possible in the analysis.

Table C.1: Gender pay gap, firm effects, sorting and bargaining with alternative normalisation.

	(1)	(2)	(3)	(4)	(5)	(6)
	Total	Appr.	Blue collar	White collar	Middle man.	Exec.
Gender pay gap	0.213	0.041	0.227	0.271	0.123	0.234
Male firm effects across males	0.240	0.161	0.201	0.293	0.401	0.349
Female firm effects across females	0.160	0.125	0.096	0.208	0.361	0.276
Firm effects gap	0.080	0.036	0.105	0.086	0.039	0.073
% of gender pay gap	*37.7%*	*86.4%*	*46.2%*	*31.5%*	*32.0%*	*31.2%*
Decomposition:						
Sorting						
Using male coefficients	0.049	0.007	0.071	0.057	-0.004	0.047
% of gender pay gap	*22.8%*	*16.6%*	*31.1%*	*20.9%*	*-3.1%*	*20.3%*
Using female coefficients	0.044	0.003	0.070	0.049	-0.009	0.026
% of gender pay gap	*20.6%*	*7.9%*	*30.7%*	*18.2%*	*-7.2%*	*11.2%*
Bargaining						
Using male distribution	0.036	0.032	0.035	0.036	0.048	0.047
% of gender pay gap	*17.0%*	*78.5%*	*15.5%*	*13.4%*	*39.1%*	*20.0%*
Using female distribution	0.032	0.029	0.034	0.029	0.043	0.025
% of gender pay gap	*14.8%*	*69.8%*	*15.1%*	*10.6%*	*35.0%*	*10.9%*

Notes. See text for details on the decomposition method. Firm effects are normalised with respect to firms in the first decile of the distribution of average log value added per worker.

[48]Estimates of sorting are unaffected by the specific normalisation chosen.

D Additional Figures and Tables

Table D.1: Gender pay gap, firm effects, sorting and bargaining by age and cohort

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	By age				**By cohort**			
	18-29	**30-39**	**40-49**	**50-65**	**1940s**	**1950s**	**1960s**	**1970s**
Gender pay gap	0.084	0.178	0.234	0.305	0.378	0.277	0.211	0.142
Male firm effects across males	0.070	0.106	0.129	0.143	0.147	0.145	0.119	0.096
Female firm effects across females	0.031	0.054	0.054	0.052	0.053	0.059	0.052	0.048
Firm effects gap	0.039	0.052	0.075	0.091	0.094	0.087	0.067	0.048
% of gender pay gap	*46.4%*	*29.2%*	*32.1%*	*29.8%*	*24.9%*	*31.4%*	*31.8%*	*33.8%*
Decomposition:								
Sorting								
Using male coefficients	0.023	0.036	0.059	0.076	0.078	0.070	0.051	0.030
% of gender pay gap	*27.4%*	*20.2%*	*25.2%*	*24.9%*	*20.6%*	*25.3%*	*24.2%*	*21.1%*
Using female coefficients	0.020	0.031	0.052	0.070	0.075	0.064	0.045	0.026
% of gender pay gap	*23.8%*	*17.4%*	*22.2%*	*23.0%*	*19.8%*	*23.1%*	*21.3%*	*18.3%*
Bargaining								
Using male distribution	0.019	0.021	0.022	0.020	0.018	0.022	0.022	0.022
% of gender pay gap	*22.6%*	*11.8%*	*9.4%*	*6.6%*	*4.8%*	*7.9%*	*10.4%*	*15.5%*
Using female distribution	0.016	0.017	0.016	0.015	0.016	0.016	0.016	0.018
% of gender pay gap	*19.0%*	*9.6%*	*6.8%*	*4.9%*	*4.2%*	*5.8%*	*7.6%*	*12.7%*

Notes. Cohorts are composed of individuals born in the indicated decade. For example, 1940s indicates people born between 1940 and 1949 included. See text for details of the decomposition method.

Table D.2: Gender pay gap, firm effects, sorting and bargaining over time

	(1)	(2)	(3)	(4)
	1995-2000	**2000-2005**	**2005-2010**	**2010-2015**
Gender pay gap	0.257	0.234	0.206	0.175
Male firm effects across males	0.087	0.088	0.099	0.100
Female firm effects across females	0.035	0.033	0.047	0.046
Firm effects gap	0.052	0.055	0.053	0.053
% of gender pay gap	*20.3%*	*23.4%*	*25.6%*	*30.4%*
Decomposition:				
Sorting				
Using male coefficients	0.049	0.045	0.038	0.036
% of gender pay gap	*19.2%*	*19.3%*	*18.6%*	*20.4%*
Using female coefficients	0.051	0.043	0.032	0.024
% of gender pay gap	*19.8%*	*18.4%*	*15.5%*	*13.7%*
Bargaining				
Using male distribution	0.001	0.012	0.021	0.029
% of gender pay gap	*0.5%*	*5.0%*	*10.0%*	*16.7%*
Using female distribution	0.003	0.010	0.014	0.018
% of gender pay gap	*1.1%*	*4.2%*	*7.0%*	*10.0%*

Notes. See text for details on the decomposition method. Firm effects and the related decompositions are computed for each of the four overlapping intervals indicated in the column headers.

Table D.3: Mean log earnings and frequencies of movers across firm effect quartiles

		Mean Log Real Weekly Earnings				4 Year Change	
Movement	Frequency	-2	-1	0	+1	Raw	Adjusted
Females							
1 to 1	292,608	5.670	5.701	5.697	5.712	0.042	0.000
1 to 2	128,899	5.728	5.763	5.865	5.891	0.164	0.121
1 to 3	60,332	5.714	5.747	5.959	5.995	0.280	0.238
1 to 4	32,348	5.722	5.767	6.098	6.148	0.425	0.383
2 to 1	130,627	5.833	5.871	5.748	5.760	-0.074	-0.108
2 to 2	233,076	5.890	5.919	5.908	5.925	0.035	0.000
2 to 3	140,290	5.942	5.975	6.011	6.038	0.096	0.062
2 to 4	65,269	6.005	6.051	6.162	6.206	0.201	0.167
3 to 1	56,456	5.926	5.979	5.742	5.756	-0.169	-0.212
3 to 2	138,182	5.972	6.010	5.937	5.950	-0.022	-0.065
3 to 3	250,809	6.037	6.064	6.062	6.080	0.043	0.000
3 to 4	153,209	6.138	6.176	6.224	6.257	0.118	0.075
4 to 1	24,302	6.049	6.118	5.737	5.743	-0.306	-0.371
4 to 2	48,828	6.084	6.140	5.968	5.984	-0.100	-0.164
4 to 3	115,656	6.139	6.181	6.117	6.134	-0.004	-0.069
4 to 4	418,917	6.417	6.438	6.459	6.481	0.065	0.000
Males							
1 to 1	478,503	5.792	5.819	5.805	5.828	0.036	0.000
1 to 2	219,074	5.882	5.911	6.017	6.051	0.169	0.133
1 to 3	114,802	5.888	5.920	6.130	6.171	0.283	0.247
1 to 4	66,192	5.910	5.950	6.276	6.335	0.425	0.389
2 to 1	190,543	5.991	6.022	5.880	5.905	-0.086	-0.130
2 to 2	384,889	6.072	6.100	6.092	6.116	0.044	0.000
2 to 3	291,559	6.161	6.183	6.230	6.257	0.097	0.053
2 to 4	138,133	6.207	6.252	6.361	6.414	0.207	0.163
3 to 1	85,678	6.095	6.127	5.892	5.914	-0.181	-0.240
3 to 2	219,818	6.182	6.207	6.150	6.170	-0.012	-0.070
3 to 3	455,806	6.271	6.291	6.310	6.330	0.059	0.000
3 to 4	306,877	6.416	6.441	6.499	6.535	0.119	0.060
4 to 1	36,610	6.225	6.265	5.901	5.922	-0.303	-0.376
4 to 2	74,026	6.291	6.322	6.182	6.207	-0.084	-0.156
4 to 3	175,613	6.392	6.422	6.389	6.413	0.021	-0.052
4 to 4	802,088	6.629	6.645	6.676	6.702	0.073	0.000

Notes. The table reports the frequency of male and female workers' movements between firm effects quartiles and the mean weekly earnings of the movers during the period between two years prior to the move to one year after. The last two columns report the overall change in earnings between the last and first period. The column labelled *Raw* is the simple difference between period "+1" and period "-2". The column labelled *Adjusted* subtracts to the raw change for movers from the *q-th* quartile the change for within-quartile *q* movers.

Figure D.1: Gender pay gap across the earnings distribution

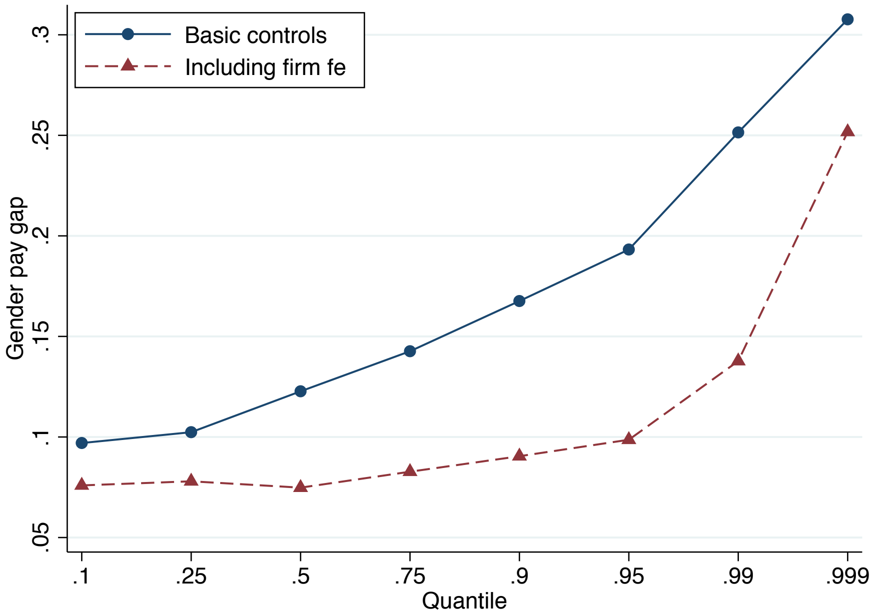

Notes. The graph plots the coefficients on a male dummy in a quantile regression with and without the inclusion of firm fixed effects (dashed and solid lines, respectively). Both quantile regressions control for observable workers' characteristics: cubic polynomials in age and experience, interacted with occupation dummies and a dummy for part-time workers.

Figure D.2: Impact of firms on the gender pay gap along the earnings distribution in 1995, 2000, 2005, 2010

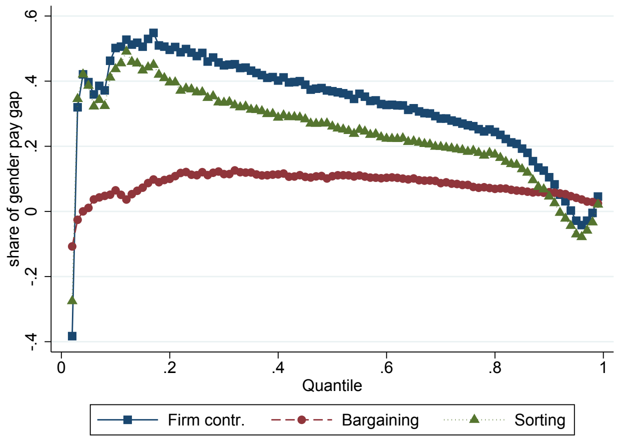

(a) 1995

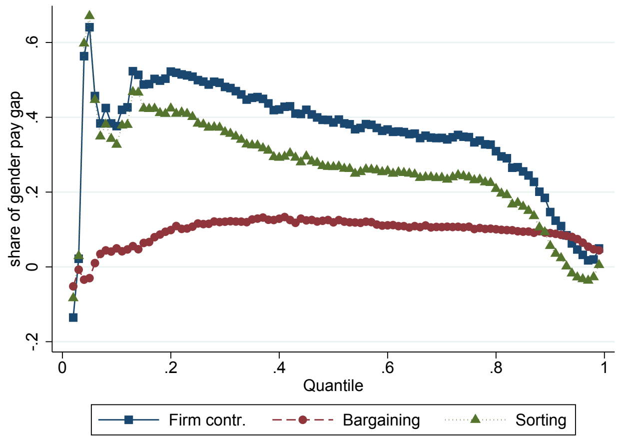

(b) 2000

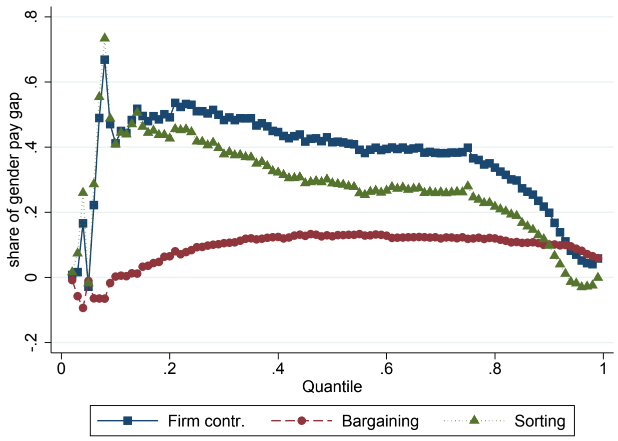

(c) 2005

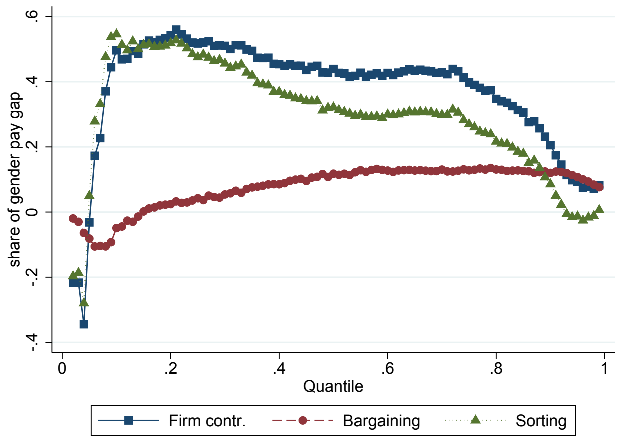

(d) 2010

Figure D.3: Gender-specific probabilities of moving to higher-quartile firm by age

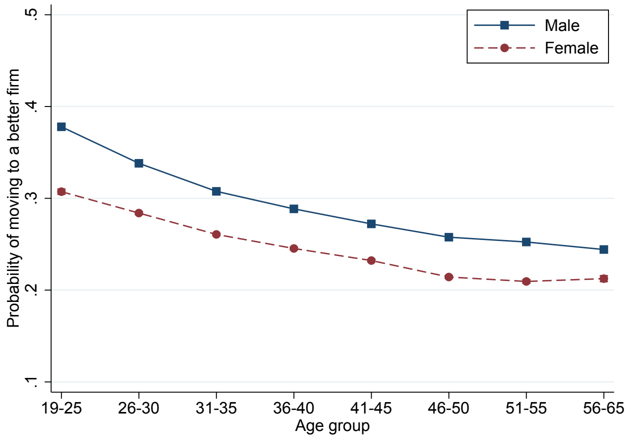

(a) All

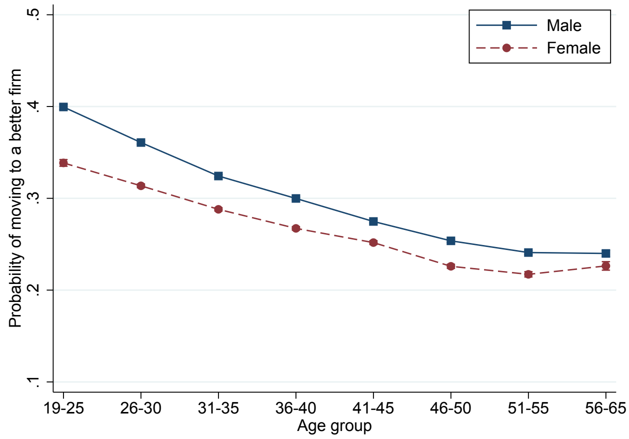

(b) Individual

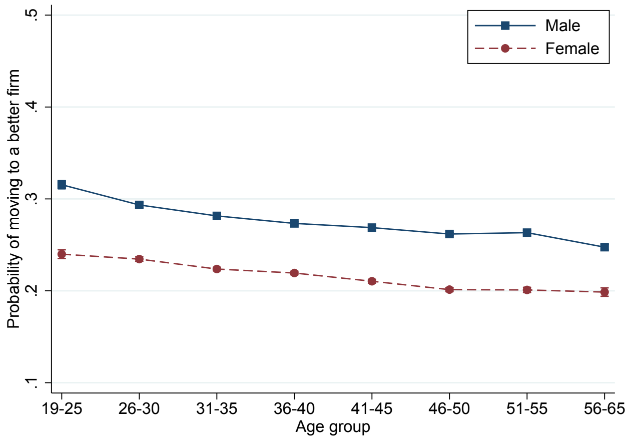

(c) Firm

Figure D.4: Average male and female firm effects in Italian provinces

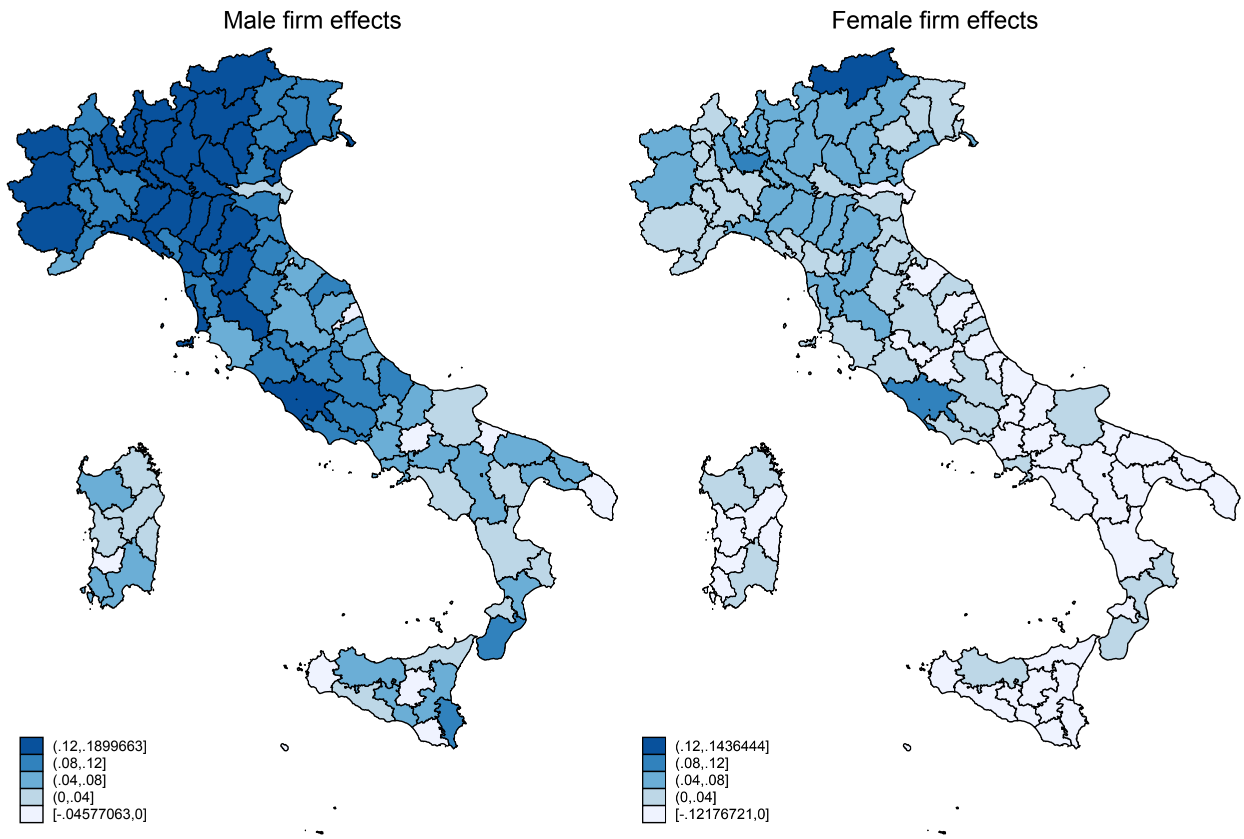

Figure D.5: Gender mobility gap within province

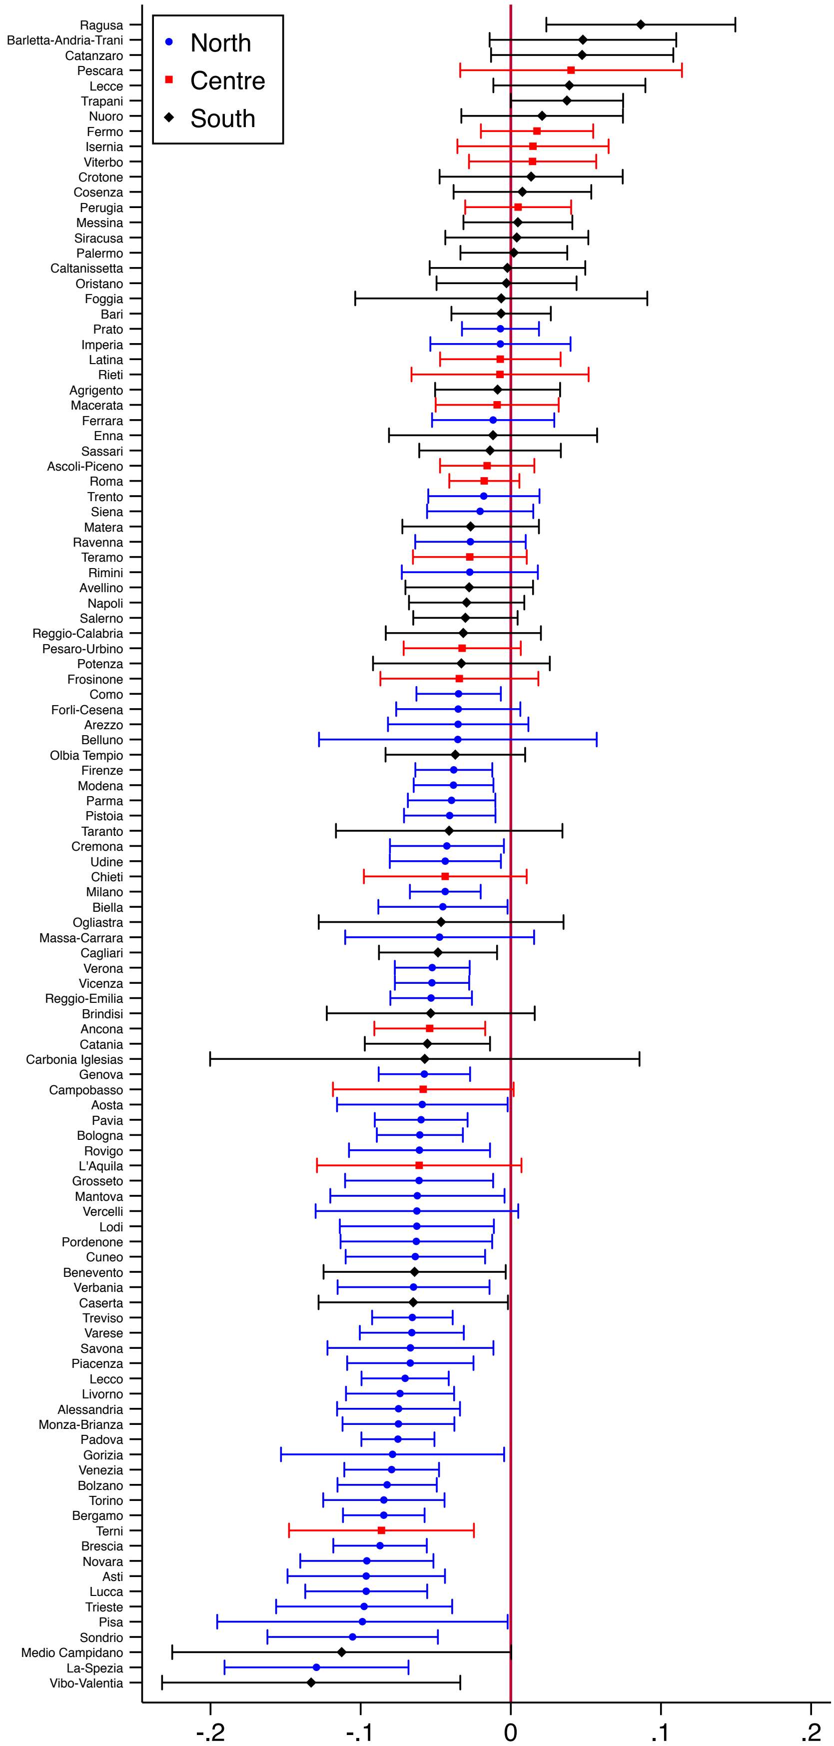

Notes. Each dot represents the marginal effect for women in a probit regression as in equation (8) estimated for each Italian province. Horizontal bars are 95% confidence intervals.

Figure D.6: Average log weekly earnings in treated and controls firms