Istituto Nazionale Previdenza Sociale

giugno 2019 – numero 24

WorkINPS *Papers*

What Firms Do: Gender Inequality in Linked Employer – Employee Data*

Alessandra Casarico

Salvatore Lattanzio

ISSN 2532 -8565

Lo scopo della serie WorkINPS papers è quello di promuovere la circolazione di documenti di lavoro prodotti da INPS o presentati da esperti indipendenti nel corso di seminari INPS, con l'obiettivo di stimolare commenti e suggerimenti.
Le opinioni espresse negli articoli sono quelle degli autori e non coinvolgono la responsabilità di INPS.

The purpose of the WorkINPS papers series is to promote the circulation of working papers prepared within INPS or presented in INPS seminars by outside experts with the aim of stimulating comments and suggestions.
The views expressed in the articles are those of the authors and do not involve the responsibility of INPS.

In copertina: uno storico "Punto cliente" a Tuscania
INPS, Direzione generale, Archivio storico

I WORK*INPS* PAPER

Le basi dati amministrative dell' *INPS* rappresentano una fonte statistica unica per studiare scientificamente temi cruciali per l' economia italiana, la società e la politica economica: non solo il mercato del lavoro e i sistemi di protezione sociale, ma anche i nodi strutturali che impediscono all'Italia di crescere in modo adeguato. All' interno dell'Istituto, questi temi vengono studiati sia dai funzionari impiegati in attività di ricerca, sia dai *VisitInps Scholars*, ricercatori italiani e stranieri selezionati in base al loro curriculum vitae e al progetto di ricerca presentato.

I WORK*INPS* hanno lo scopo di diffondere i risultati delle ricerche svolte all' interno dell'Istituto a un più ampio numero possibile di ricercatori, studenti e policy makers. Questi saggi di ricerca rappresentano un prodotto di avanzamento intermedio rispetto alla pubblicazione scientifica finale, un processo che nelle scienze sociali può richiedere anche diversi anni. Il processo di pubblicazione scientifica finale sarà gestito dai singoli autori.

Pietro Garibaldi

What Firms Do:
Gender Inequality in
Linked Employer – Employee Data*

Alessandra Casarico
Bocconi University CESifo and Dondena

Salvatore Lattanzio
University of Cambridge

What Firms Do:
Gender Inequality in Linked Employer-Employee Data*

Alessandra Casarico† Salvatore Lattanzio‡

June 5, 2019

Abstract

This paper investigates the contribution of firms to the gender gap in earnings on average, at different quantiles of the earnings distribution, and over time to shed light on the role of firm pay policies in hindering or reinforcing the gender wage gap and to identify how their impact comes about. Using a linked employer-employee dataset for Italy, we show that the gap in firm pay policies explains on average 30% of the gender pay gap in the period 1995-2015. Sorting of women in low pay firms explains a larger fraction of the gender pay gap than differences in bargaining, on average and at the bottom of the distribution, whereas the latter dominates at the top. Moreover, differences in bargaining have increased in importance over the two decades. To explain sorting, we investigate whether women have a lower probability of moving towards firms with higher pay rates, and find that this is indeed the case. This differential mobility penalises, in particular, highly skilled women and can be related to the variability in wages in destination firms, with women not moving to those with high (unexplained) variance in pay. We also find some evidence that the firm environment as captured by exogenous changes in the gender balance in leadership positions influences the bargaining power of women, indicating that the latter is partly institution-driven.

Keywords: Bargaining, Sorting, Linked Employer-Employee Data, Mobility gap, Gender quotas

***JEL* codes**: J16, J31, J71

*We would like to thank Toke Aidt, Massimo Anelli, Alex Bryson, Carlo Devillanova, Kai Liu, Hamish Low, Gabriella Santangelo and audiences at University of Cambridge, Collegio Carlo Alberto, Italian Social Security Institute (INPS), University of Luxembourg, German Employment Agency (IAB), XII workshop in Public Policy Design at the University of Girona, 2017 ECINEQ Conference, 2018 AIEL Conference, 2018 SIEP Conference and 2019 RES Conference for comments and suggestions. This project was realised thanks to the VisitINPS Scholars B programme. We thank the Research Division of the Italian Social Security Institute (INPS), especially Massimo Antichi, Mariella Cozzolino, Edoardo Di Porto, and Paolo Naticchioni, for providing help with data access. The views expressed here are of the authors only and do not involve the responsibility of INPS. We thank Angela Ciavarella (CONSOB) for providing data on the board renewals of listed companies and Luca Brugnara for excellent research assistance.

†Bocconi University, CESifo and Dondena. Email: alessandra.casarico@unibocconi.it.

‡University of Cambridge. Email: sl828@cam.ac.uk.

Sintesi non tecnica

I differenziali salariali di genere sono una delle variabili cruciali in base alle quali valutare la presenza di disuguaglianze di genere nel mercato del lavoro. La letteratura economica ha messo in evidenza diversi fattori, sia dal lato della domanda di lavoro che da quello dell'offerta, che possono spiegare l'apertura e la persistenza del gender pay gap: differenze nei livelli di istruzione e di produttività, penalità legate alla nascita di un figlio e diverse caratteristiche psicologiche o norme sociali tra uomini e donne sono tra gli elementi messi in luce.

Parte del divario è riconducibile al comportamento delle imprese. Le politiche salariali o di progressione di carriera adottate, l'organizzazione dei tempi di lavoro, la presenza o assenza di servizi complementari al lavoro stesso possono infatti contribuire alla disuguaglianza di genere.

In questo studio, analizziamo l'impatto delle imprese sulla determinazione del divario salariale di genere in media, lungo la distribuzione dei salari e nel tempo. In particolare, misuriamo quanto di questo impatto dipende dal fatto che le donne lavorino in imprese che pagano mediamente salari inferiori – un effetto di sorting – oppure siano meno abili degli uomini a contrattare salari migliori – un effetto di bargaining.

Per quantificare il contributo delle imprese utilizziamo le banche dati INPS e ci focalizziamo sulle storie contributive dei lavoratori nel settore privato tra il 1995 e il 2015. Troviamo che, dato un differenziale di genere nei salari settimanali pari in media al 21%, esso è spiegato per il 30% dalle differenze nelle politiche salariali applicate a uomini e donne. Effettuando un esercizio di scomposizione del contributo delle imprese in sorting e bargaining, le nostre analisi rivelano che il fatto che le donne si concentrino in imprese che pagano mediamente di meno tutti i lavoratori spiega il 20% del divario salariale di genere, mentre il loro minor potere contrattuale è responsabile del rimanente 10%. Il contributo delle imprese non è omogeneo lungo la distribuzione dei salari. In particolare, nella parte alta della distribuzione l'effetto di bargaining domina quello di sorting, indicando come, anche quando le donne raggiungono posizioni di vertice, esse pagano comunque una penalità rispetto ai colleghi maschi per via della loro minore abilità a contrattare. Nel corso del tempo, il contributo delle imprese al divario salariale di genere è aumentato come percentuale del differenziale salariale di genere ed in particolare è cresciuta l'importanza dell'effetto di bargaining, probabilmente a causa della maggiore diffusione di forme di contrattazione decentrata.

Una spiegazione in termini di mobilità verso imprese "migliori"

Perché il sorting concorre in misura determinante al contributo delle imprese al divario salariale di genere? Le donne sono impiegate in imprese con politiche salariali meno generose sin dal primo ingresso nel mercato del lavoro, ma l'impatto del sorting tende a peggiorare lungo la carriera lavorativa. Questa evidenza è compatibile con uno scenario in cui la probabilità di spostarsi in imprese con politiche salariali "migliori" sia minore per le donne rispetto agli uomini durante la carriera. Abbiamo testato formalmente questa ipotesi, stimando la probabilità che uomini e donne hanno di spostarsi verso imprese migliori e troviamo che le donne hanno una probabilità inferiore di 3 punti percentuali rispetto ai colleghi maschi. Questo differenziale di mobilità è eterogeneo a seconda dell'abilità degli individui. In particolare, per le donne più "abili" (più istruite o produttive) il gap rispetto ai colleghi è maggiore rispetto a quello misurato per le donne meno produttive. Indaghiamo il ruolo della avversione al rischio, attitudine a competere e costi di commuting come fattori che influenzano le traiettorie di mobilità di lavoratori e lavoratrici e quindi i divari salariali.

Un possibile impatto delle quote di genere?

Nella parte finale del nostro studio ci chiediamo se un cambiamento dell'equilibrio di genere al vertice delle imprese possa favorire un ribilanciamento delle differenze nelle politiche salariali. A tal fine, sfruttiamo l'introduzione delle quote di genere nei consigli di amministrazione delle società quotate, realizzata con la Legge 120/2011, e analizziamo l'impatto che esse hanno avuto sul potere contrattuale delle lavoratrici. I risultati dicono che l'impatto c'è stato ed è concentrato sulle nuove assunte. Inoltre, l'effetto è maggiore nelle imprese che avevano una minore rappresentanza femminile nei consigli di amministrazione nel periodo antecedente all'introduzione delle quote.

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,[1] with large cross-country differences and with smaller reductions at the top of the distribution.

In this paper we study the contribution of firms to the gender gap in earnings on average, at different percentiles of the earnings distribution and over time, to shed light on the role of firm pay policies in hindering or reinforcing the gender wage gap. To identify how firms' impact comes about, we distinguish between sorting across firms and bargaining within firms, and investigate gender differences in mobility towards firms with more generous pay policy as a driver of sorting. In addition, we exploit exogenous variation in the gender composition of board of directors to study the impact of firm environment on gender differences in bargaining power. 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 and it contains more than 22 million workers employed by approximately 1.6 million firms.

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 depend not only on the characteristics and behaviour of workers, but also on those of the firms which employ them.[3] The choice of some countries to impose disclosure policies that require firms to report, among other things, the pay level of men and women (e.g. Equality Act in the UK or the Decree Law 254/2016 on Non-financial reporting in Italy) speaks to the emerging role of firms as key actors in generating or taming gender inequality. Firm-related gender 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; Card et al., 2016; Cardoso

[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., 2018, for evidence on the US, Card et al., 2013, for Germany).

et al., 2016).[4] In addition, women may show lower bargaining power compared to men working at the same firm: women may negotiate less aggressively (Bowles et al., 2007, 2005; Babcock et al., 2006; Rozada and Yeyati, 2018) and this can result in gender pay gaps and different standards of promotion, even when wages tend to be equal within the same occupations (Petersen and Morgan, 1995; Blau, 2012).

In our paper, we first measure the firm contribution to the gender earnings gap. To develop our analysis, 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 key to measure the link between firms and the gender earnings gap, since in our model firm effects are proportional to firm rents, and represent a premium paid by firms to their workforce.[5] Then, following Card et al. (2016), we explain the impact of firm pay policy on the gender earnings gap, by separating the role of sorting across firms and bargaining within firms, and study their impact not only at the mean, but also along the distribution of earnings and over time. Once we have established the evidence on the firm contribution, we explore what drives sorting and bargaining. We hypothesise that sorting is the outcome of gender gaps in mobility across firms. More precisely, we propose a novel definition of gender mobility gap, which takes into account where the origin and destination firms are located in the firm fixed effect distribution, and we study whether high (unexplained) earnings dispersion in high firm effect firms or their geographical distribution can contribute to explain the male-female gap in the probability of moving towards more generous firms. As to bargaining, we ask ourselves whether the lower bargaining power of women captures a psychological trait or whether it is, at least partly, driven by the firm environment. The dimension of firm environment we focus on is the gender balance at the top of the hierarchy. To obtain exogenous variation in the gender composition at the top of the corporate ladder, 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. The research focusing on the impact of gender quotas on worker outcomes, and in particular on female wages and employment, finds little to no effect. Examining the introduction of gender quotas in boards in Norway, Bertrand et al. (2019) find a positive impact on the qualification level of appointed female board members, but no robust evidence of trickle-down effects on female employees. Similarly, Maida and Weber (2019) find no significant impact of the introduction of gender quotas in Italy on female wages or on women's progression towards the top echelons of the firms' hierarchy in Italy. Our focus

[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]See Card et al. (2018) for a survey of the literature on the elasticity of wages to firms' rents.

on bargaining is new and can help us unpacking the effects of gender quotas on worker outcomes by providing a potential mediating factor.

We find that differences in firm-specific premia account for approximately 30% of the Italian gender pay gap at the mean. We show that sorting 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, and this determines one third of the firm contribution and slightly less than one tenth of the mean gender pay 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 differences in bargaining dominate 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 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. We argue that this may reflect the spreading of more decentralised wage setting in the Italian labour market, highlighting that it has differentially affected men and women, to the detriment of the latter. When investigating 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 ability. We provide evidence that women tend not to move towards firms with high (unexplained) earnings dispersion, indicating that gender differences in risk aversion, attitudes towards competition or cost of effort may be at play. As to bargaining power, under some specifications we find that women hired after the implementation of gender quotas see a more marked increase in their bargaining power compared to men. This result is associated with a more positive selection of women into firms in the post-reform period, a relevant trickle down effect which, to the best of our knowledge, is documented for the first time. Overall, this evidence provides some support to the view that bargaining power results from a combination of employer and employee characteristics and behaviour, rather than depending (exclusively) on innate traits.

The contribution of this paper to the literature is four-fold. First, we show that the impact of firms on the gender pay gap is non-negligible and remains fairly constant over time, with differences in bargaining power increasing in importance in recent years. The almost unchanged gender gap in firm pay policy, coupled with a declining gender pay gap, suggests that the firm contribution is gaining importance and the policy focus on firms is appropriate.[6] Second, we show that there is considerable heterogeneity in the impact of firms along the earnings distribution, with sorting playing a major role in the bottom and middle part of the distribution, and bargaining dominating at the top. This evidence suggests that the relative absence of women from the top of the earnings distribution documented by the literature[7] can partly be explained by gender differences in bargaining power. The increasing importance of bargaining power over time can also provide an explanation for the smaller decrease over time of the male-female earnings gap at the top.[8] Third, we propose a novel definition of gender mobility gap, pointing out the lower probability for women of moving towards firms adopting more generous pay policies. This penalty – which hits especially high skilled women – can be related to gender differences in preferences and cost of effort. While the former are hardly directly malleable, the latter can be an indirect target of policy through, for example, the promotion of a more flexible job organisation or incentives to stronger fathers' involvement, with the goal of overcoming obstacles to upper mobility of women. Last, we are the first to investigate the impact of an exogenous increase in the share of female members in the board of directors of listed companies on the relative bargaining power of female employees. The evidence that gender quotas affect the bargaining power of newly hired women – who are also more skilled – reveals that some trickle down effects are present; it also points to the importance of strengthening gender balance at the top of the firm hierarchy and pinpoints a specific channel – that of bargaining – through which legislation can address gender gaps in earnings.

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.

---

[6] Note that this evidence, which also includes the aftermath of the Great Recession, differs from the one on West Germany provided by Bruns (2019) for the period 1995-2008: there, the gender gap in firm pay policies has increased, rather than having been constant, providing an explanation for the stall in the decline of the gender wage gap.

[7] A rich literature investigates the gender pay gap across the wage distribution and shows the presence of larger gaps at the top – providing evidence of a *glass ceiling* (Albrecht et al., 2003, 2015; Arulampalam et al., 2007) 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).

[8] Blau and Kahn (2017) and Goldin (2014). 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

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 1995-2015.[9] 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,[10] the number of days and weeks worked in a given year, the type of contract (whether full-time or part-time), the province 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,[11] 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. Specifically, we use information on sales and value added.

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.[12] We further

[9] 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.

[10] 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.

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

[12] 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.

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 4.2. 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. *Experience* is the labour market experience of workers, computed as the difference between the current year and the first year of work. *Adjusted weeks* are the number of weeks worked in a year, standardised to account for part-time work (see text for details). *Weekly earnings* are expressed in 2010 real prices. The number of workers per firm is computed as the average of the yearly male and female workforce at each firm.

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.[13]

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.[14] The average real weekly earnings – the measure of pay we focus on – are larger for men, with a 22% gender gap. The average number of male workers in a firm is 8 and of female workers is 5, both

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

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

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

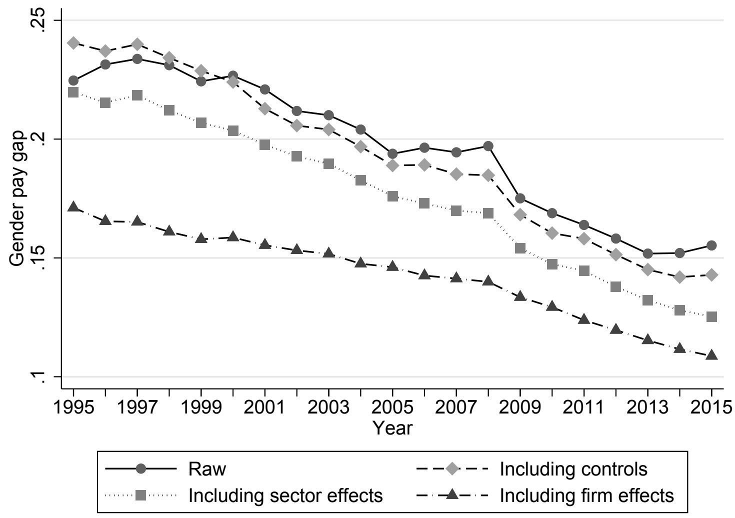

Notes. The figure plots coefficients of a dummy for male workers from log wage regressions, run for each year in four different specifications: without controls ("Raw"); controlling for observable characteristics of workers, i.e. cubic polynomials in age, experience and tenure, a dummy for full-time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

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.[15]

2.2 Evidence on the Gender Earnings Gap in Italy

Figure 1 reports the evolution of the gap in log average 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. The raw average gender pay gap was approximately 22.5 log points in 1995 and 15.5 log points in 2015.

[15]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.

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 average 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, experience and tenure, a dummy for full-time contract, the number of weeks worked) and a full set of year, occupation and province dummies. The inclusion of these controls leaves the main coefficient of interest on the male dummy practically unchanged. Column (3) includes 2-digit sector fixed effects. Their inclusion reduces the coefficient on the male dummy by 1.6 log points with respect to column (2), indicating that women tend to sort into low-pay sectors. Even within sectors, women tend to work for low-pay firms, as shown by the specification in Column (4), which includes firm fixed effects. The coefficient on the male dummy decreases by 3.2 log points relative to column (3) and by 5 log points with respect to the unconditional estimate. 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 analysis: in Figure 1, besides the raw gender gap in earnings, we plot the coefficients of the male dummy from regressions that control for individual observable characteristics (as in column (2) of Table 2), for sector fixed effects (as in column (3) of Table 2) and for firm fixed effects, in addition to individual observables (as in column (4) 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 no controls (solid line), a set of observable individual characteristics (dashed line), sector fixed effects (dotted line), and that additionally controls for firm fixed effects (dashed-dotted line).[16] The figure shows the presence of a strong *glass ceiling* effect,

[16]Following Canay (2011), we estimate fixed effects quantile regressions in two steps. In the first

Table 2: Regression results

	(1)	(2)	(3)	(4)
Male	0.192***	0.190***	0.174***	0.142***
	(0.003)	(0.002)	(0.001)	(0.001)
Covariates	No	Yes	Yes	Yes
Year effects	No	Yes	Yes	Yes
Province effects	No	Yes	Yes	Yes
Sector effects	No	No	Yes	Yes
Firm effects	No	No	No	Yes
R-squared	0.040	0.514	0.546	0.708
Observations	208,669,170	208,669,170	207,788,391	207,788,391

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). Sectors are taken from 2-digit NACE Rev2. Robust standard error, clustered at firm level, in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01.

meaning that the gap increases at the top of the distribution of earnings. At the 99th percentile, the raw gap between male and female weekly earnings is approximately 47 log points against a value slightly above 13 at the median. When firm effects are included, the gender gap in earnings decreases, especially in the middle and top portions of the distribution. The impact of firms at the very top (above top 1 per cent) is smaller, though, 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.

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.

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.

3.1 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):

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

where w_{ijt} is the natural logarithm of real weekly earnings, for worker i in firm j at time t, with $i \in \{1, ..., N\}$, $j \in \{1, ..., J\}$, 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 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_{ijt} 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.[17]

To estimate (1), we construct connected sets of firms and workers separately by gender and focus on the largest connected set for female and male workers.[18]

3.2 Normalisation of Firm Effects

Since male and female fixed effects are estimated separately, to compare their levels we need to normalise them with respect to a common criterion. For this purpose, we consider 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.

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

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

[18] 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.

in the accommodation and food industry, which is usually identified in the literature as a low-surplus sector (Card et al., 2016; 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 effect in the accommodation and food sector.[19]

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:

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

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

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.[20]

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

[19] We report an alternative normalisation procedure in Appendix C, where we empirically identify the set of firms that pay zero rents to their workers. Results do not change and we leave to Appendix C a more thorough discussion of this alternative normalisation procedure.

[20] The two quantities are computed taking the average of the normalised firm effects across men and women. 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.

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.[21]

We often choose to report the results as averages of sorting and bargaining computed from equations (4) and (5). Hence, unless otherwise specified, we refer to:

$$\begin{aligned} \text{Sorting} &= \frac{1}{2} \sum_{x \in \{F,M\}} \left\{ \mathbb{E}\left[\psi_j^x \mid g = M\right] - \mathbb{E}\left[\psi_j^x \mid g = F\right] \right\}, \\ \text{Bargaining} &= \frac{1}{2} \sum_{x \in \{F,M\}} \mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = x\right]. \end{aligned} \tag{6}$$

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:

$$\begin{aligned} &\mathbb{E}\left[\psi_j^M \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in p_k^F\right] \\ =&\mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = M, i \in p_k^M\right] + \mathbb{E}\left[\psi_j^F \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^F \mid g = F, i \in p_k^F\right] \quad (7) \\ =&\mathbb{E}\left[\psi_j^M - \psi_j^F \mid g = F, i \in p_k^F\right] + \mathbb{E}\left[\psi_j^M \mid g = M, i \in p_k^M\right] - \mathbb{E}\left[\psi_j^M \mid g = F, i \in p_k^F\right]. \quad (8) \end{aligned}$$

In both equations (7) and (8), 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. When reporting the results, we average sorting and bargaining as resulting from the two alternative decompositions of equations (7) and (8), akin to what we do in (6).

[21]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.

4 Results

4.1 Estimation of two-way models

We estimate (1) separately for the largest connected groups of female and male workers. We include as controls cubic polynomials in age,[22] tenure and experience, occupation dummies (blue-collar, white-collar, executive, middle manager and apprentice) and a full set of year dummies. Panel A 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.

Panel B 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.[23] Worker and firm effects display negative or no correlation (-0.04 and 0 in the male and female sample, respectively). 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.

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

4.2.1 Average Decomposition

Overall sample We focus on the double connected set of workers and firms. 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.

[22] 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).

[23] 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.

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

Panel A. 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%
Panel B. 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(worker effects, firm effects)	-0.043	0.000

Notes. The Table reports summary statistics for the largest connected sets used for the estimation of the AKM two-way models. Panel A reports sample sizes for the largest connected sets of male and female workers. Panel B reports summary statistics from the estimation of equation (1), separately for men and women.

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, Jewell et al. (2018) for UK and by Coudin et al. (2018) for France. Thus, sorting is the main factor behind the different premia men and women receive on average.

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)

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

	(1) Total	(2) Appr.	(3) Blue collar	(4) White collar	(5) Middle man.	(6) 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. The table reports results of the Oaxaca-Blinder decomposition of equations (4) and (5). Firm effects are normalised with respect to the average gender-specific firm effects in food and accommodation. Column (1) shows results for all workers. 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.

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), 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: Gender pay gap, firm effects gap, sorting and bargaining by age and cohort

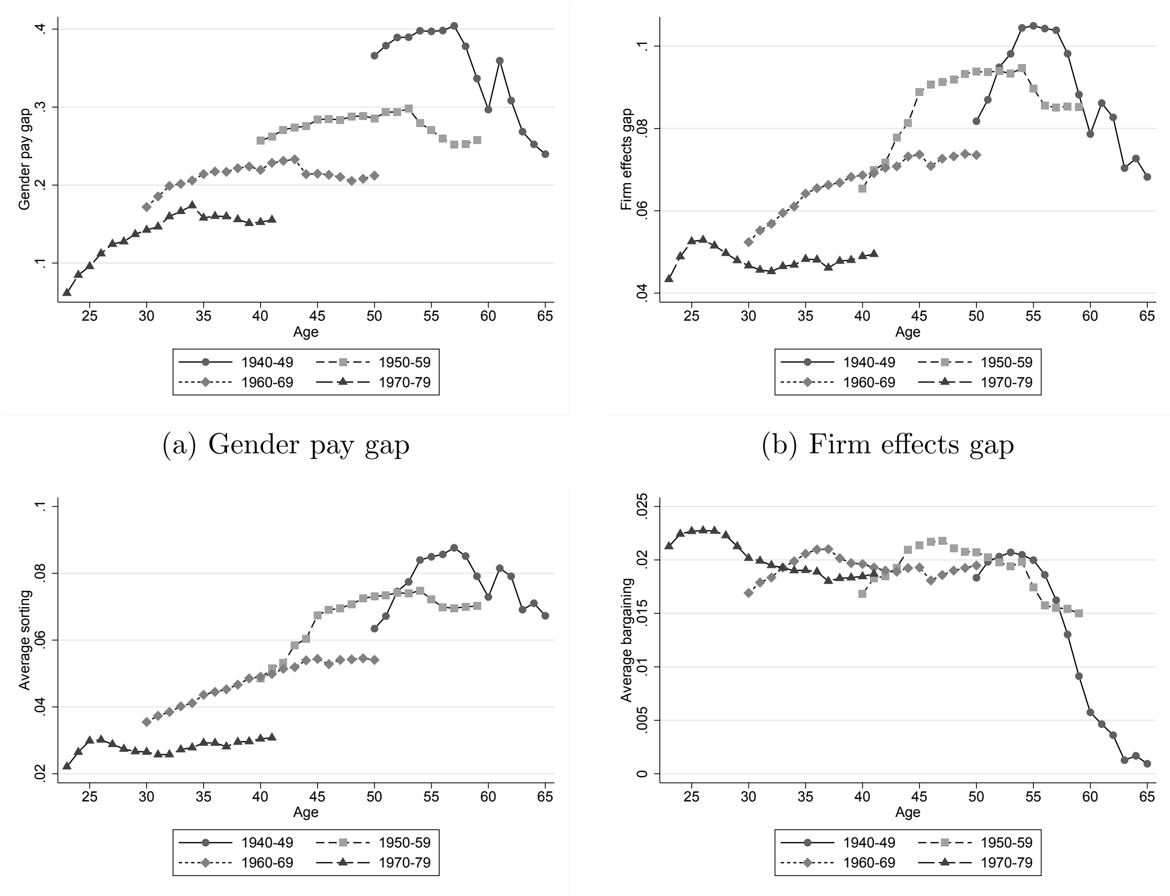

(a) Gender pay gap

(b) Firm effects gap

(c) Sorting

(d) Bargaining

Notes. The horizontal axis reports the mean age in each year in 1995-2015 for each cohort.

Age and cohorts Figure 2 shows the evolution of the gender pay gap (panel a), firm effects gap (panel b), sorting (panel c) and bargaining (panel d), by age and cohorts. We identify four cohorts: 1940-49, 1950-59, 1960-69, 1970-79. We construct this Figure by first defining the age by cohort cells. To do so, we compute the mean age of each cohort in each year in our data. We then compute the gap in firm effects and its decomposition into sorting and bargaining in each age by cohort cell. The Figure shows the presence of sizeable cohort effects in the evolution of the gender pay gap over the life cycle. Older cohorts display higher gender gaps in earnings than younger cohorts, even at the same age. The same holds for the gap in firm effects and sorting, whereas the bargaining power effect remains fairly stable across cohorts and over the life cycle until age 60, when it suddenly drops to values close to 0. Moreover, the gap in firm effects remains stable over time for the youngest cohort (1970-79), but tends to increase within cohort for each of the other cohorts (except for ages close to retirement). The same pattern characterises the evolution of sorting, which is flat for the youngest cohort, but increasing in age for the other cohorts. This result shows that the rising importance of firm components over the life cycle, highlighted also by Card et al. (2016) and Bruns (2019), is not only an age

Figure 3: Decomposition by sector

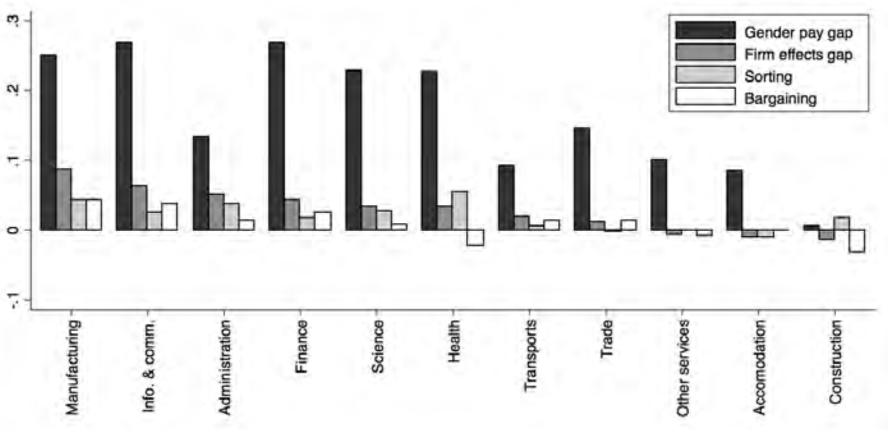

Notes. Sectors are defined according to Ateco 2007 sectoral codes and ordered according to the gap in firm effects (highest to smallest). We exclude sectors that employ less than 1% of the total workforce in all years considered. The sectors reported in the figure represent 95% of the total person-year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

effect, as they argue, but it is the outcome of both a cohort effect and an age effect *within* cohort.

In Figure D.2 in the Appendix, we show the fraction of the gender pay gap explained by firm components. As a share of the gender pay gap, the firm effects gap (panel a) is higher for the youngest cohort – which is the only one we can observe before 30 – and it flattens out for all cohorts after that age. For the oldest cohort, however, the gap in firm effects explains a lower fraction of the gender pay gap, indicating that individual unobservable and observable characteristics are more important in determining gender differences in pay for this specific cohort. When we look at the evolution of sorting over the life cycle (panel b), again we find that the behaviour of the oldest cohort is different from that of the others. In addition, within cohort sorting tends to increase in importance. On the other hand, there are little between-cohort differences in the impact of bargaining on the gender pay gap, which stays constant or tends to decline monotonically with age.

Sectors We investigate the role of sectors in Figure 3, which shows the sectorial decomposition of the gender pay gap and the gap in firm effects along with the estimated bargaining and sorting effects. We recall that 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).[24] Overall, the gender gap in earnings is the highest in ICT and finance. 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, administra-

[24]We exclude sectors that comprise less than 1% of the total person-year observations.

Figure 4: Firm effects along the wage distribution

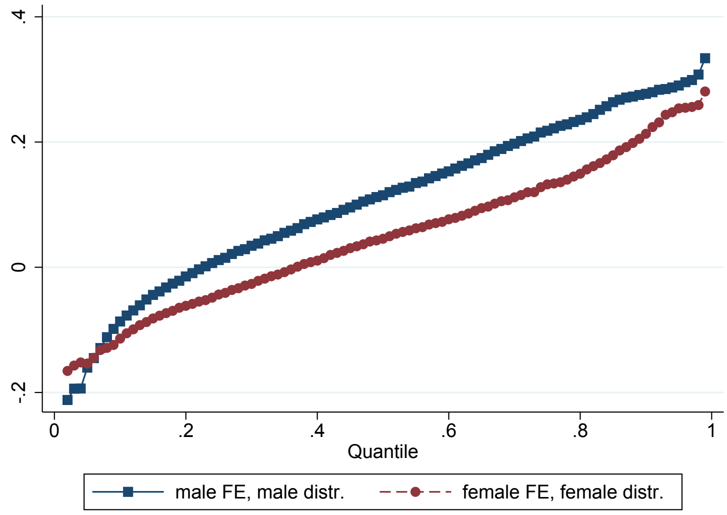

Notes. The figure shows the average male firm effects across percentile bins of the male distribution of earnings and average female firm effects across percentile bins of the female distribution of earnings.

tion and health.[25] In finance and ICT, on the contrary, bargaining power explains a larger share of the firm effects gap relative to sorting.

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 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 smoothed 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 earnings distribution, and smaller at lower and higher quantiles. As to the determinants of this contribution, sorting is

[25] 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 5: Impact of firm components on the gender pay gap along the earnings distribution

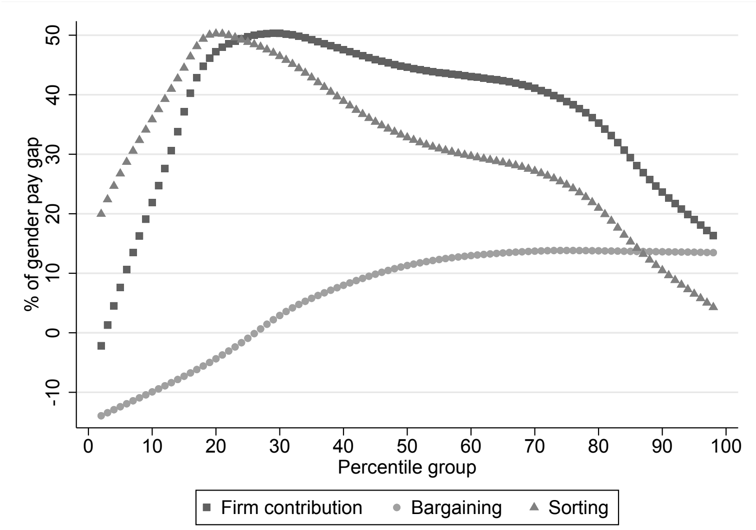

Notes. The Figure shows smoothed differences between male and female firm effects and their decompositions into sorting and bargaining across percentile bins of the distribution of earnings.

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.[26] 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.[27] High-pay workers are also likely to work in sectors with higher earnings on average. Consistently, we have shown that bargaining explains a larger portion of the gender pay gap in Finance and ICT. In Appendix Figure D.3 we report the same decomposition for 1995, 2000, 2005 and 2010. The results are very similar.

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

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

[27]For apprentices, a category of workers which is likely to be placed at 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 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, who are also more likely to occupy the bottom part of the earnings distribution.

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

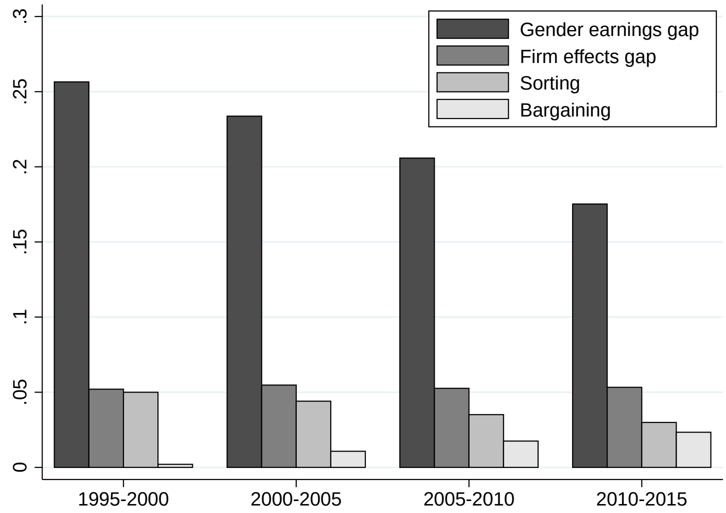

Notes. The Figure shows the evolution of the gender pay gap, firm effects gap, sorting and bargaining over time. We estimate firm effects in each of the four overlapping time intervals. See text for details.

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., 2018; Alvarez et al., 2018) and could impact differently men and women (Bruns, 2019). 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 are slowly revealed over time or triggered by changes in the composition of peers (Mas and Moretti, 2009).

We allow here for additional flexibility in the evolution of individual and firm unobservable heterogeneity over time. The availability of a long panel enables us to recover individual fixed effects and gender-specific estimates of the firm fixed effects in subintervals 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.[28] For each subinterval we build a double 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

[28] We have checked that the conditional random mobility assumption holds in each of the sub-intervals. Results are available upon request.

firm effect gap, sorting and bargaining in each of the four sub-intervals. As explained in section 3.3, we present results averaging sorting and bargaining as in equation (6). The gap in firm effects remains unchanged, but since the gender pay gap declines over time, as a share of the latter the firm effects gap increases in importance.[29] 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.[30] 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 gender gap in firm policies 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 have been binding for employers and workers: they have been signed by unions and employers' associations at the industry level and have provided wage floors for each job title. Firms could not 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.

Increased female labour force participation can be another explanation. Female employment was 41.1% in 1995 against a value of 50.6% in 2015.[31] 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.

A possible concern with the results that we find in this section is that the rising importance of bargaining over time is the outcome of a composition effect, due to the

[29] Table D.1 in the Appendix reports the values used to produce Figure 6. It shows that both male and female firm effects increase especially after 2005, but they grow at the same pace, leaving the difference unaltered.

[30] 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.1.

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

fact that the youngest cohort is more represented in the last subinterval with respect to the previous periods and for this cohort bargaining explains a larger fraction of the gap in firm effects – and thus in earnings – compared to older cohorts. If younger cohorts are more represented in recent sub-intervals, the results that we find may be driven by the different composition of our samples. This is, however, not the case. Indeed, Figure D.4 shows sorting and bargaining as a percentage of the gap in firm effects in the first and last subinterval. Even though there are clear differences between cohorts, the relative importance of bargaining and sorting in determining the gap in firm effects has changed for all cohorts. As an example, sorting in 1995-2000 (panel a) accounts for approximately 95% of the gap in firm effects for the 1940 cohort, against a value around 80% in 2010-2015 (panel b). At the same time, bargaining importance has increased for this cohort from values around 4% (panel c) to 20% (panel d). The same holds for other cohorts (for which differences over time are more marked). Hence, our results do not seem to reflect only an age/cohort composition effect.

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". In estimating it, we condition on overall mobility and investigate differences in mobility rates of men and women towards firms offering higher pay premia. Gender gaps in mobility are shown to be an important driver of the gender gap in wage growth (Del Bono and Vuri, 2011; Loprest, 1992), especially early in the career (Manning and Swaffield, 2008). However, the literature lacks evidence on the gender difference in mobility across firms ranked according to the generosity of their pay policy. Since we have shown that firm components account for a sizeable fraction of the gender pay gap, it is worth investigating 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 *conditional* on workers' time-varying observable and unobservable characteristics, which we control for by estimating firm effects conditional on age, experience, tenure, occupation, time trends and individual 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 procedure. What may threaten the estimates are firm or worker transitory and permanent shocks that determine a change in earnings before the move and trigger mobility. We show in 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.[32] 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 components observable to the researcher, it can be correlated with workers' characteristics.

Empirical strategy Our estimation strategy relies on the following probit model:

$$\Pr\left\{\mathbf{1}\left[Q^g_{f_1} > Q^g_{f_0}\right]\right\} = \Phi(\alpha + \gamma F_i + \delta Z_{it} + \delta_s + \lambda_t) \tag{9}$$

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), δ_s are sector fixed effects 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.

Figure D.5 in the Appendix shows mobility rates for the full sample of movers (panel a) – i.e. the sample of all workers who change job between two *consecutive* years.[33] The figure shows that the mobility rate is slightly higher for men, with large differences by age classes: young workers tend to move more often than old ones, for whom male-female differences are close to zero. In order to abstract from seasonal jobs and fixed-term

[32] 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.

[33] We thus do not consider gaps in the work histories of individuals as mobility: these can be periods out of the labour force, in self-employment, or in the public sector, which we are not able to identify separately.

contracts that are not converted into open-ended contracts, we restrict our sample and retain workers that move to a new firm between two consecutive years and, in addition, 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 "individual". 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. In this restricted sample – Figure D.5 (panel b)– the mobility rate of men and women is lower, but the male-female gap is of comparable size to that in the full sample. In particular, in the restricted sample we do not consider many moves that happen early in the career, when workers are likely to change jobs more frequently and end up in seasonal jobs or fixed-term contracts. The issue of selection of workers into mobility seems not to be of particular concern: even if men move more often than women, in the restricted sample the difference in mobility rate is only 0.2 percentage points and this is the sample we focus on for the mobility analysis.

Results on the gender mobility gap Table 5 shows average marginal effects from the estimation of equation (9). The first column shows the results for all types of job moves. Female workers are 3 percentage points less likely to move to a better firm within sectors. This differential probability is lower for "individual" moves (column 2) than "firm" moves (column 3): while the former difference is 1.6 percentage points, the latter is 4.4 percentage points. 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 within a given sector. Overall, a gender gap in the likelihood of upward firm mobility is present, and it is 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. We also add age at the moment of job move, which has a negative impact on the probability of upward mobility.

In Appendix Figure D.6 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

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

	(1) All	(2) Firm	(3) Individual
Female	-0.030***	-0.044***	-0.016***
	(0.005)	(0.007)	(0.005)
Age	-0.002***	-0.000*	-0.003***
	(0.000)	(0.000)	(0.000)
Change province	0.018***	0.015**	0.011***
	(0.004)	(0.006)	(0.004)
Change occupation	0.027***	0.024***	0.017***
	(0.004)	(0.005)	(0.004)
Change to full-time	0.043***	0.016***	0.042***
	(0.007)	(0.005)	(0.008)
Observations	5,216,076	2,259,559	2,956,517
Sector FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Baseline Probability	0.286	0.268	0.298

Notes. The table reports average marginal effects from probit regressions where the dependent variable is the probability of moving to a firm in a higher firm effect quartile. Column (1) shows results for all moves in the restricted sample defined in the main text. Column (2) shows results for moves not determined by firm closure. Column (3) shows results for moves happening because of firm closure. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: * p<0.1, ** p<0.05, *** p<0.01

moves. Moreover, the gender mobility gap tends to be lower or insignificant for older workers. The gap is approximately 6 percentage points for workers aged 19-25 years old and it becomes approximately zero for workers aged 56-65.

Unobserved worker ability 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" (left panel) versus "high" (right panel) 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 26.7%, whereas for a woman is 24.0%, a 2.7 percentage points gap. In the high individual fixed effect group, these probabilities are 30.3% and 27.0%,

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

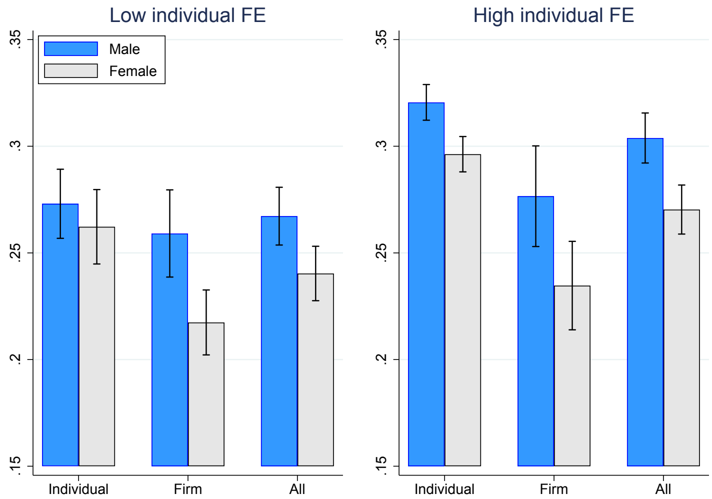

Notes. Each bar represents the average probability for men and women of moving to a firm in a higher firm effect quartile for different types of moves and workers. Low (high) individual fixed effect workers are defined as those having an individual fixed effect below (above) the median of the individual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

respectively, a 3.3 percentage points gap. As to "firm" moves, the gap remains fairly constant at around 4.2 percentage points, but it increases from 1.1 percentage points to 2.4 percentage points for high fixed effect workers in the case of "individual" 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. 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 82% of the corresponding rise for men. One reason why the gender mobility gap increases for high ability women might be the time cost of child-rearing, which disproportionately affects women (Kleven et al., 2019). 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. Alternatively, women may value non-pecuniary benefits (such as flexible time arrangements or firm provision of welfare services) more. As a consequence, they may be willing to stay in or move to firms that have lower firm effects, because they are compensated for the loss of part of their earnings potential by a better balance between family and work.[34]

[34] 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.

Geographical concentration and earnings dispersion We here explore two further reasons for why women pay a penalty in their probability of moving up in the firm effect distribution. First, we investigate whether firms offering high pay policies to female workers are geographically concentrated in some provinces or cities, whereas those offering high pay policies to men are more geographically dispersed, providing the opportunity to improve earnings to a potentially wider set of (male) workers. We graphically test this hypothesis by studying the geographical distribution of average male and female firm effects. Figure D.7 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 average 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, indicating that the gender mobility gap we identify in the data cannot be explained by a different distribution of firms with high generous pay policy across provinces. We then test whether the difference in mobility rates to better firms is a within or between province phenomenon. Appendix Figure D.8 plots the marginal effects from probit models estimated as in equation (9) within each province (thus, excluding the dummy for change of province). Different colours indicate the sign and significance (at 95% confidence level) of coefficients. They 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 effect of the female dummy is "less" negative in Southern provinces. This is probably due to the fact that most of the low fixed effect firms are located in the South, where most movements are between low quartiles of the fixed effect distribution (e.g. from 1st to 2nd quartile) and, thus, easier to achieve. Note that this evidence could be consistent with women having a higher cost of commuting relative to men (even within the smaller province scale).

A second channel that may explain why women tend to move less frequently to high-pay firms is that these firms may display higher earnings dispersion, for instance because they use incentive pay more,[35] or because a larger share of earnings paid comes from overtime. If women are more risk averse or less inclined to compete, or if they have higher cost of effort,[36] they may be less willing than men to move to higher quartile firms if these firms have a higher dispersion of earnings. In Figure 8 we divide our sample of moves in four groups, defined by the level of earnings dispersion of the origin-destination

[35] Albanesi et al. (2015) show that 93% of the gender gap in executive compensation in the United States is due to differences in incentive pay.

[36] For instance, women may be less willing to work overtime or unconventional hours (Goldin, 2014), because of household responsibilities.

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

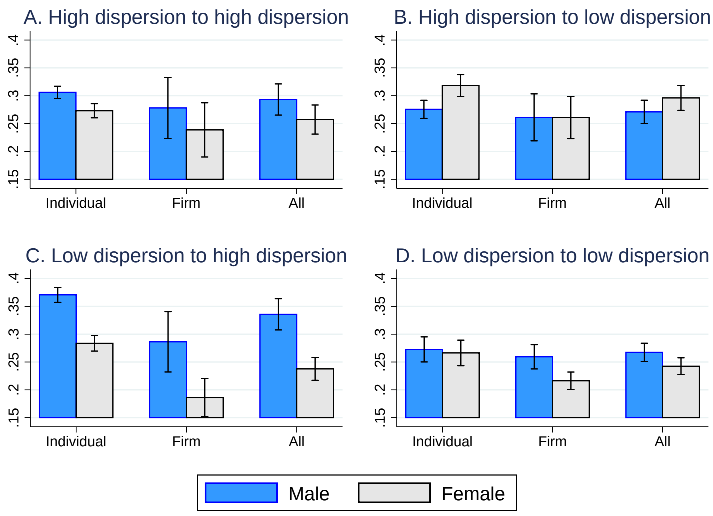

Notes. Each bar represents the average probability for men and women of moving to a firm in a higher gender-specific firm effect quartile for different types of moves. We define high (low) dispersion firms those having a standard deviation of residual weekly earnings higher (lower) than the 75th percentile of the distribution of standard deviations of residual earnings of the firms in our data (see text for details on the derivation of residual earnings). Panel A displays the probability of moving to a firm in a higher quartile of the firm fixed effect distribution for movements that happen from high dispersion firms to high dispersion firms. Panel B displays the probabilities of moving from high dispersion firms to low dispersion firms. Panel C displays the probabilities of moving from low dispersion firms to high dispersion firms. Panel D displays the probabilities of moving from low dispersion firms to low dispersion firms.

firm. We define high earnings dispersion firms as follows. We first compute residuals from regressions of weekly earnings on a full set of sectoral, occupation and full-time dummies. We then compute the firms' standard deviation of residual earnings across all workers and periods. We define high earnings dispersion firms those having a standard deviation of residual weekly earnings higher than the 75th percentile of the distribution of standard deviations of residuals of the firms in our data.[37] We finally estimate the probit model (9) for each of the four groups of workers that move between firms with different levels of earnings dispersion.

We find evidence that the gap in the probability of moving to higher quartile firms widens when movements happen from low to high-dispersion firms (panel C), whereas women are more likely to move to a "better" firm when moving from high- to low-dispersion firm (panel B), at least when the movement is not determined by firm closure. In other words, women tend to move proportionally less to firms with high earnings dispersion compared to men. The magnitude of the difference is not negligible: the gap in the probability of moving from low- to high-dispersion firms is -9.8 percentage points,

[37] We use residual earnings to capture differences in pay across firms that happen within sector and occupational composition of the workforce, thus reflecting differences in dispersion between firms rather than between sectors or groups of firms.

whereas the move in the opposite direction (from high- to low-dispersion firms) implies a gap in favour of women equal to 2.5 percentage points. This evidence is consistent with women having a lower preference for competing, or higher risk aversion and cost of effort compared to men, which make them less prone to work in firms with high earnings dispersion.

Overall, this evidence highlights that sorting comes from the lower probability of women to move to better firms. 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 – remains an open question. We have explored some potential explanations, finding evidence consistent with higher female risk aversion and lower attitude to compete, or higher cost of effort. Clearly, the mobility gap can also affect the difference in bargaining power between men and women. If women are less likely to quit a firm and move towards one with better pay policy, firms find it easier to extract rents from them, rather than from men. Thus, mobility gaps can drive both sorting and differences in bargaining.

6 Bargaining and Gender Balance in Corporate Governance

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, where opportunities to bargain are potentially present or more widespread. 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. Do women have an innate lower ability to bargain or does the firm environment influence bargaining power? The measure of firm environment we focus on is the extent of gender balance at the top of the firm hierarchy. For instance, 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. To assess if and to what extent bargaining power within firms can be influenced by the degree of gender equality at the top of firm governance, we exploit a recent Italian law which prescribes gender quotas in boards of listed firms. This law provides exogenous variation in the gender composition of boards, allowing us to identify the causal impact of a change in our measure of firm environment on the relative bargaining power of female and male workers. From this angle, we also contribute to the literature which evaluates the impact of gender quotas on worker performance (Bertrand et al., 2019; Maida and Weber, 2019) by focusing on mediating factors through which effects on workers' outcomes can show up.[38]

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.[39] 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 7.4%, only 1.4 percentage points higher than the share in 2008. The first year of implementation of the law, 2012, the share jumped to 11.3%, and it kept rising until 33.3% in 2017.

6.2 Empirical Analysis

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

[38] Other studies examine the impact of female-led firms on labour market outcomes of female employees and on the extent of gender gaps in earnings within the firm, without relying on exogenous variation in gender composition of directors, e.g. Bell (2005), Cardoso and Winter-Ebmer (2010), Flabbi et al. (2016), Gagliarducci and Paserman (2015)

[39] 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 2017.

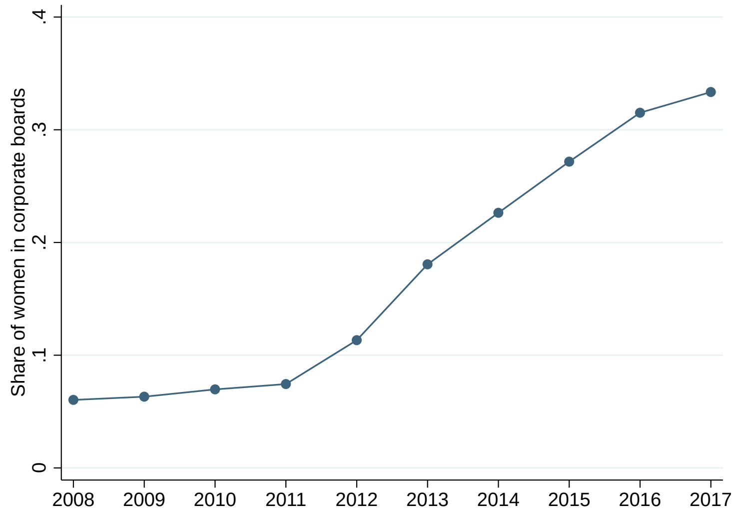

effect can be expressed as in equation (2), that is $\psi_j^g = \gamma^g \bar{S}_j$. The ratio or the difference between female and male gender-specific shares, γ^F and γ^M, captures female bargaining power relative to men directly.

Figure D.9 in the Appendix shows that women have indeed a lower bargaining power compared to men according to this definition. It plots female firm effects against male firm effects, both averaged across percentile bins of log value added per worker (our proxy for firm surplus $\bar{S}_j$). The slope of the linear fit of the relationship in Figure D.9 is an estimate of the relative bargaining power of women, γ^F/γ^M, and equals 0.85, meaning that firms share a lower fraction of increases in value added with female employees relative to males.[40]

In order to estimate the causal impact of a change in the gender composition of board of directors on bargaining power, we use data for the period 2008-2017,[41] and we adopt two alternative empirical strategies. First, we use a static/canonical specification (Borusyak and Jaravel, 2017), focusing on listed firms only, and exploit the staggered timing of the first board renewal after the reform becomes binding. We depart from our model with two fixed effects: we collapse data at the firm level, and estimate a "reduced-form" model in which we regress log average male and female weekly earnings on log average value added per worker to recover a rent-sharing coefficient and, with that, measure the

[40] A well established literature, surveyed in Manning (2011) and Card et al. (2018), has investigated the magnitude of the elasticities of wages to firms' financial conditions, as a way to depart from the hypothesis of competitive labour markets.

[41] Data for 2016 and 2017 have very recently become available and we use them for this part of the analysis in order to evaluate the medium-run effects of the gender quota policy.

bargaining power of workers. The firm-level equation is:

$$w_{jt}^{g} = \kappa + \gamma^{g} Treat_{jt} \times \overline{S}_{j}^{pre} + \eta_{t}^{g} + \phi_{j}^{g} + \varepsilon_{jt}^{g} \tag{10}$$

where w_{jt}^{g} represents average weekly earnings, in firm j at time t, with $g = \{M, F\}$. κ is a constant; $Treat_{jt}$ is a dummy equal to 1 for firm j in the year of board renewal and all subsequent years, and 0 before. $\overline{S}_{j}^{pre}$ is the surplus of firm j, as measured by the log average value added per worker at each firm over the period 2008-2011 – the source is AIDA-Bureau Van Dijk. We average value added in the period before the implementation of gender quotas, since there is evidence that the reform affected firm productivity (Bruno et al., 2018). η_{t}^{g} and ϕ_{j}^{g} are, respectively, gender-specific year and firm fixed effects. The parameter of our interest is γ^{g}, which measures the change in bargaining power by gender after the first board renewal. We also provide results of a dynamic specification where we interact average log value added per worker in 2008-2011, $\overline{S}_{j}^{pre}$, with distance dummies relative to the year of first board renewal, and plot the related coefficients. In this way, we ensure that pre-trends are absent and we check whether the reform had an impact in a specific period.

The second strategy we adopt compares the outcome of listed and non-listed companies. The literature on the effects of gender quota highlights the importance of selecting an appropriate control group for the firms targeted by the reform (Comi et al., 2019; Bertrand et al., 2019; Ferrari et al., 2018; Maida and Weber, 2019). We identify a control group by matching listed companies, treated by the reform, with a subset of non-listed corporations, selected according to a Mahalanobis metric on the following set of firm characteristics averaged over the period 2008-2011: log weekly earnings, female log weekly earnings, value added per worker, sales per worker, male and female worker effects, share of part-time workers and female part-time workers, share of permanent workers, share of executives and female executives, share of women above the 90th percentile of the firm distribution of weekly earnings, female hiring rate, log of firm size and log of firm size squared, share of workers aged 35-54 and over 55, Ateco sector dummies and region dummies. We then estimate firm-level regressions of the form:

$$w_{jt}^{g} = \kappa + \gamma^{g} Treat_{j} \times Post_{t} \times \overline{S}_{j}^{pre} + \delta^{g} Post_{t} \times \overline{S}_{j}^{pre} + \eta_{t}^{g} + \phi_{j}^{g} + \varepsilon_{jt}^{g} \tag{11}$$

where $Treat_{j}$ is a dummy equal to 1 for listed companies and equal to 0 for matched non-listed companies. $Post_{t}$ is a dummy equal to 1 starting from 2012. The other variables are defined as above. Again, the parameter of interest is γ^{g}, which measures the change after the reform in bargaining power by gender in treated versus control firms. The interaction $Post_{t} \times \overline{S}_{j}^{pre}$ controls for differences in pre-period value added per worker between treated and control firms in the post reform years. Similarly to the first empirical strategy, we also report results of an event study specification, in which we interact the treatment status $Treat_j$ and average surplus $\overline{S}_j^{pre}$ with year dummies.

In both strategies, in order to reduce the risk of selection of firms into listing or delisting due to the implementation of the reform, we focus our analysis on firms that are continuously listed between 2011 and 2014.[42,43]

We test the balance of covariates between listed and non listed firms, before and after the matching, by regressing the treatment status (being listed and, thus, subject to the quota) on the whole set of covariates, averaged over the period 2008-2011. Table D.2 shows the results. Column (1) reports regression results for the unmatched sample, whereas Column (2) provides results for the matched sample. There is no evidence that the two groups are different in the pre-reform period after matching.

6.3 Results

Panel A of Table 6, in columns (1) to (3), reports results from the canonical specification (10), which exploits the staggered introduction of the reform. The outcome variable in column (1) is log average weekly earnings of male and female workers. Columns (2) and (3) split between joiners and stayers. We do not find any significant effect for men, but rent-sharing significantly increases for female joiners, that is, female workers who join the firm in the period 2008-2017. In particular, a 10 percent increase in value added determines a 0.2 percent increase in average wages of female joiners after the first board renewal. This result provides evidence that an increase in the share of women in corporate boards causes a rise in bargaining power of women who join the firm, although wages for existing contracts are not affected.

In Panel B of Table 6, columns (1) to (3), we perform heterogeneity analysis according to the intensity of the treatment. Intensity is high when the pre-reform share of women in board of directors is below 10%, medium when between 10 and 20%, low when above 20%.[44] Again, we do not find any significant effect for males, whereas we highlight that the increase in female bargaining power is concentrated in high intensity firms, where the gender imbalance at the top of the hierarchy was more marked before the reform. Furthermore, in firms with a low intensity of the treatment, women suffer a loss in rent-sharing. Hence, if we consider low intensity firms as terms of comparison, the increase in female bargaining power in high intensity firms is even more pronounced.

The results are confirmed by the dynamic specification, as shown in Panel (a) of Figure D.10 in the Appendix. In particular, for joiners, we can exclude the presence of

[42]Nothing changes if we consider the set of firms that are continuously listed over the period 2008-2017.

[43]The number of listed companies in Italy over the period 2008-2017 ranges between a minimum of 323 in 2012 to a maximum of 421 in 2017 (source: *Borsa Italiana*). We use information on 212 continuously listed firms in the period 2011-2014, of which 167 have no missing balance-sheet information.

[44]We have information on the share of women in the pre-reform period for 158, out of 167 firms. 109 firms are classified as high intensity, 34 as medium and 15 as low.

Table 6: Impact of gender quotas on bargaining power of workers

	(1)	(2)	(3)	(4)	(5)	(6)
	Static/Canonical DD			Matched DDD		
	Total	Joiner	Stayer	Total	Joiner	Stayer
			Panel A: Aggregate effect			
			Males			
γ^M	-0.003	-0.004	-0.002	0.005	-0.006	0.005
	(0.007)	(0.011)	(0.007)	(0.003)	(0.006)	(0.004)
Obs.	1665	1461	1660	3933	3256	3911
R^2	0.033	0.025	0.024	0.028	0.018	0.024
N. firms	167	166	167	394	390	394
			Females			
γ^F	-0.008	0.020***	-0.006	0.000	-0.007	0.002
	(0.006)	(0.006)	(0.005)	(0.003)	(0.004)	(0.002)
Obs.	1654	1320	1645	3902	2919	3874
R^2	0.052	0.038	0.066	0.073	0.028	0.077
N. firms	167	161	167	394	384	393
			Panel B: Treatment intensity			
			Males			
γ^M_{High}	-0.003	0.001	-0.003	0.003	-0.000	0.002
	(0.007)	(0.012)	(0.006)	(0.004)	(0.007)	(0.004)
γ^M_{Medium}	-0.000	-0.016	0.004	0.007	-0.012	0.009*
	(0.007)	(0.012)	(0.008)	(0.005)	(0.009)	(0.005)
γ^M_{Low}	0.012	-0.020	0.016	0.026	-0.010	0.021
	(0.021)	(0.017)	(0.021)	(0.021)	(0.014)	(0.016)
Obs.	1580	1396	1578	3638	3023	3623
R^2	0.038	0.027	0.032	0.036	0.018	0.030
N. firms	158	157	158	364	360	364
			Females			
γ^F_{High}	-0.004	0.023***	-0.003	0.002	-0.002	0.004
	(0.006)	(0.007)	(0.004)	(0.003)	(0.005)	(0.003)
γ^F_{Medium}	-0.010*	0.014	-0.007	-0.000	-0.014*	0.001
	(0.005)	(0.009)	(0.005)	(0.005)	(0.008)	(0.004)
γ^F_{Low}	-0.037**	-0.020**	-0.031**	-0.016**	-0.033***	-0.021*
	(0.018)	(0.010)	(0.015)	(0.007)	(0.011)	(0.012)
Obs.	1564	1271	1555	3609	2709	3588
R^2	0.080	0.052	0.090	0.078	0.0035	0.088
N. firms	158	152	158	364	354	364

Notes. The table shows coefficients from the estimation of equation (10) in columns (1) to (3) and (11) in columns (4) to (6). Panel A reports aggregate effects and Panel B reports results for treatment intensity, according to the pre-reform share of female board directors. The dependent variables are average weekly earnings for all workers (*Total*), for workers joining the firm (*Joiner*) and for incumbents (*Stayer*). The number of observations is lower in Panel B because we do not have information on pre-reform board composition for all firms. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

different pre-trends for men and women, whereas women experience a significant jump in their bargaining power right after the board renewal.

We now turn to the second specification based on ex-ante matching. Panel A of Table 6 shows, in columns (4) to (6), the estimates recovered from equation (11), none of which is significant, for both males and females. The heterogeneity analysis, in Panel B, columns (4) to (6), confirms a negative effect on bargaining in low intensity firms and a F-test for the equality of coefficients for high and low intensity rejects the null of equality. This suggests that an increase of the share of women in corporate boards at least prevents the decline in female rent-sharing in firms most affected by the reform.

Panel (b) of Figure D.10 in the Appendix shows the results of the event-study specification, confirming the absence of significant effects.

Overall, there is some evidence that the reform affects the rent-sharing of female workers, especially in high-intensity firms, suggesting that the lower bargaining power of women compared to men is partly institution-driven.

Bargaining and the skill composition of the workforce How does the increase in female bargaining power ensue? Does it come with a change in the skill composition of the workforce or does it happen independently of it? To address these questions, similarly to the analysis on wages, we adopt two empirical strategies. First, we estimate the following canonical specification:

$$\alpha^g_{jt} = \kappa + \zeta^g Treat_{jt} + \eta^g_t + \phi^g_j + \varepsilon^g_{jt} \tag{12}$$

Second, we estimate a difference-in-differences model in the ex-ante matched sample:

$$\alpha^g_{jt} = \kappa + \zeta^g Treat_j \times Post_t + \eta^g_t + \phi^g_j + \varepsilon^g_{jt} \tag{13}$$

where α^g_{jt} are average gender-specific AKM worker effects estimated in equation (1) which measure the skill composition of the female and male workforce. ζ^g is, in both equations, the coefficient of our interest. All the other variables have the same definition as before.

Panel A of Table 7 reports the estimates of equation (12) in columns (1) to (3). We find that the overall skill composition of the workforce has improved thanks to a rise in the average skill level of female workers. The effect is driven by changes within high and medium intensity firms, as Panel B shows. In columns (4) to (6), we display results for the matched difference-in-differences. The results are qualitatively similar, although the coefficients for female skill composition, in both Panel A and Panel B, are no longer significant.[45]

[45] If we restrict the sample to replicate the one used for the analysis in Table 6, i.e. we exclude firms for which we have no balance sheet information, we find positive and significant coefficients for female skill composition, overall and in the high intensity group.

Table 7: Impact of gender quotas on skill composition of workers

	(1)	(2)	(3)	(4)	(5)	(6)
	Static/Canonical DD			Matched DDD		
	All	Male	Female	All	Male	Female
	Panel A: Aggregate effect					
ζ^g	0.024*	0.017	0.030***	0.008	0.006	0.019
	(0.015)	(0.019)	(0.011)	(0.012)	(0.015)	(0.012)
Obs.	2120	2106	2091	4490	4488	4477
R^2	0.169	0.128	0.17	0.208	0.161	0.180
N. firms	212	212	212	449	449	449
	Panel B: Treatment intensity					
ζ^g_{High}	0.017	-0.001	0.032**	0.004	-0.012	0.025
	(0.018)	(0.023)	(0.013)	(0.015)	(0.019)	(0.015)
ζ^g_{Medium}	0.032	0.045*	0.039**	0.013	0.032	0.016
	(0.020)	(0.024)	(0.020)	(0.016)	(0.020)	(0.017)
ζ^g_{Low}	0.059	0.092	-0.000	0.004	0.042	-0.020
	(0.040)	(0.061)	(0.030)	(0.044)	(0.055)	(0.028)
Obs.	2020	2011	1991	4180	4179	4170
R^2	0.172	0.133	0.184	0.211	0.160	0.189
N. firms	202	202	202	418	418	418

Notes. The table shows coefficients from the estimation of equation (12) in columns (1) to (3) and (13) in columns (4) to (6). Panel A reports aggregate effects and Panel B reports results for treatment intensity, according to the pre-reform share of female board directors. Dependent variables are average AKM worker effects for *All*, *Male* and *Female* workers. The number of observations is lower in Panel B because we do not have information on pre-reform board composition for all firms. Robust standard errors, clustered at the firm level, in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

Therefore, part of the impact on female bargaining power is mediated by a change in the skill composition of women. The reform has attracted more skilled women towards firms with the largest increase in female board members, a relevant trickle down effect which, to the best of our knowledge, is documented for the first time.

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 and find that firm effects play a significant role. Firm characteristics account for approximately 30% of the total average gender pay gap, with sorting accounting for roughly 20-22% of the total gender pay gap and bargaining playing the dominant role at the top of the earnings distribution.

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. To explain this result, we provide evidence consistent with the presence of gender differences in preferences and cost of effort. Finally, exploiting an exogenous change in the firm environment as measured by the gender composition at the top of the firm hierarchy, we show that the relative bargaining power of women can be enhanced.

Our analysis contributes to the understanding of the role of firms in influencing the level and dynamics of the gender wage gap. The importance of gender differences in firm pay policy has increased over time as a share of the overall gender earnings gap, making the behaviour of firms critical to any attempt of tackling the gender pay gap. Differences in bargaining, in particular, play an important role in explaining what happens at the top of the pay distribution, where women advancement has been more limited. We have also highlighted avenues for policy to affect the gender earnings gap, identifying gender differences in upward mobility and gender balance in the corporate structure as important factors behind sorting and bargaining.

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 allow the identification and exploration of different channels, providing a solid ground on which to build policy recommendations to reduce obstacles to further women's advancements in the labour market.

References

Abowd, J., R. H. Creecy, and F. Kramarz (2002). Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. Longitudinal employer-household dynamics technical papers, Center for Economic Studies, U.S. Census Bureau.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–333.

Albanesi, S., C. Olivetti, and M. J. Prados (2015). Gender and Dynamic Agency: Theory and Evidence on the Compensation of Top Executives. In *Gender in the Labor Market*, Chapter 1, pp. 1–59.

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. ZEW Discussion Papers 17-023, ZEW - Zentrum fur Europaische Wirtschaftsforschung / Center for European Economic Research.

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*(2), 163–186.

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). 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 (2019). Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway. *The Review of Economic Studies 86*(1), 191–239.

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

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

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.

Blau, F. D. and L. M. Kahn (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature 55*(3), 789–865.

Blinder, A. (1973). Wage Discrimination: Reduced Form and Structural Estimates. *Journal of Human Resources 8*(4), 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 111*, 122–133.

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.

Bruno, G. S. F., A. Ciavarella, and N. Linciano (2018). Boardroom Gender Diversity and Performance of Listed Companies in Italy. Working Papers 87, CONSOB.

Bruns, B. (2019). Changes in Workplace Heterogeneity and How They Widen the Gender Wage Gap. *American Economic Journal: Applied Economics 11*(2), 74–113.

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*(3), 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 (2019). Where Women Make a Difference: Gender Quotas and Firms' Performance in Three European Countries. *ILR Review* (forthcoming).

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: Evidence from Europe. 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. D. Pronzato (2018). Do Board Gender Quotas Matter? Selection, Performance and Stock Market Effects. IZA Discussion Papers 11462, Bonn.

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.

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.

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.

Jewell, S., G. Razzu, and C. Singleton (2018). Who Works for Whom and the UK Gender Pay Gap? Ese discussion papers, Edinburgh School of Economics, University of Edinburgh.

Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and Gender Inequality: Evidence from Denmark. *American Economic Journal: Applied Economics* (forthcoming).

Loprest, P. J. (1992). Gender Differences in Wage Growth and Job Mobility. *American Economic Review 82*(2), 526–32.

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.

Maida, A. and A. Weber (2019). Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform. CEPR Discussion Papers 13476, C.E.P.R. Discussion Papers.

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). 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 (2018). Firming Up Inequality. *The Quarterly Journal of Economics 134*(1), 1–50.

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:

$$w_{it} = a_{it} + \gamma^g S_{ijt}, \tag{A.1}$$

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 $g = M, F$.[1] We assume that S_{ijt} can be written as follows:

$$S_{ijt} = \overline{S}_j + \phi_{jt} + m_{ij}, \tag{A.2}$$

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:

$$a_{it} = \theta_i + X'_{it}\beta^g + u_{it}, \tag{A.3}$$

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 (A.2) and (A.3) into (A.1), we get:

$$w_{it} = \theta_i + \psi^g_j + X'_{it}\beta^g + \varepsilon_{it} \tag{A.4}$$

where

$$\psi^g_j = \gamma^g \overline{S}_j \tag{A.5}$$
$$\varepsilon_{it} = \gamma^g \left(\phi_{jt} + m_{ij}\right) + u_{it} \tag{A.6}$$

Equation (A.4) is consistent with the two-way fixed effects model *á la* Abowd et al. (1999) presented in equation (1) in the main text.

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

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:

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

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. $\boldsymbol{\varepsilon}$ 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 (A.6). 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.[2] 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 (A.6) 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

[2]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. In 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 do 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.1: Mean weekly earnings of movers across firm effects quartiles

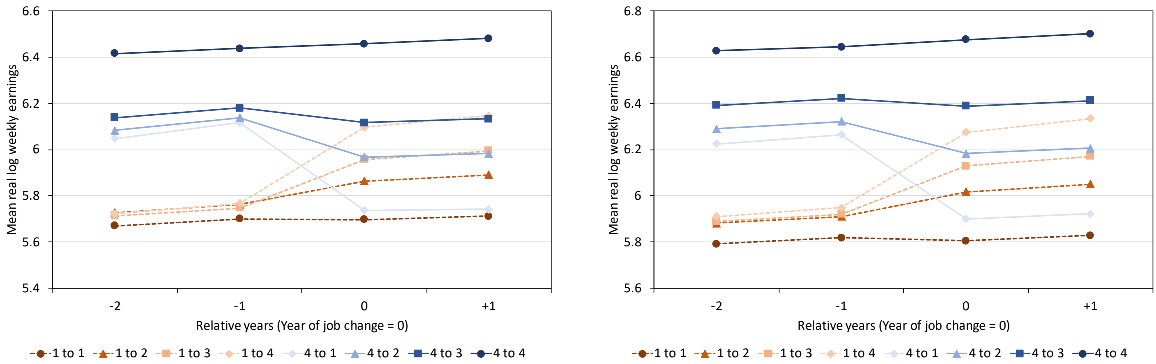

(a) Female sample (b) Male sample

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

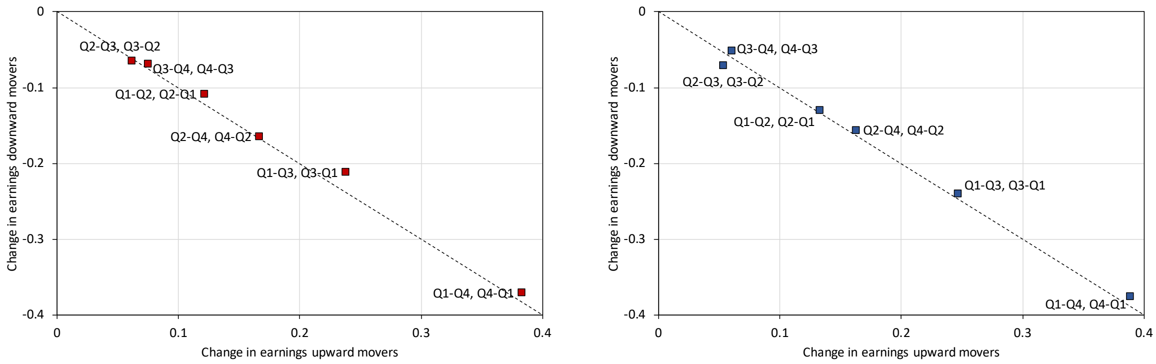

(a) Female sample (b) Male sample

gains and losses with each opposite move is 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.[3] The 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[4] 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 a reasonable assumption.

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

[4] Adjusted earnings changes equal raw earnings changes minus the earnings change for within-quartile movers: that is, we subtract the change for movers from quartile q to quartile q from the raw change for movers from quartile q to quartile q', with $q \neq q'$.

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

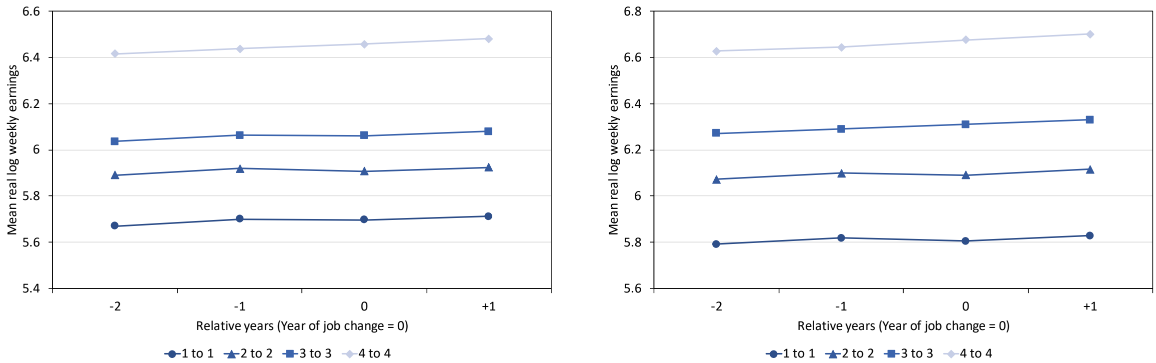

(a) Female sample (b) Male sample

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 characterised by almost no earnings gains. This is indeed the case: both panels show that the 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 (A.6) 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 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 by misspecification the most. 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

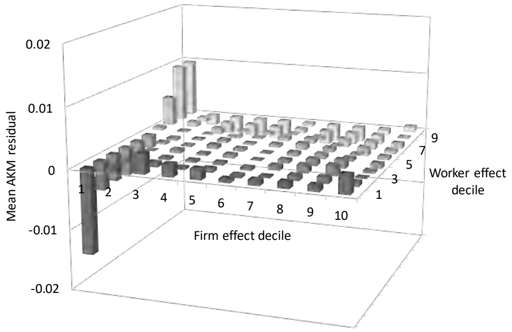

(a) Female sample

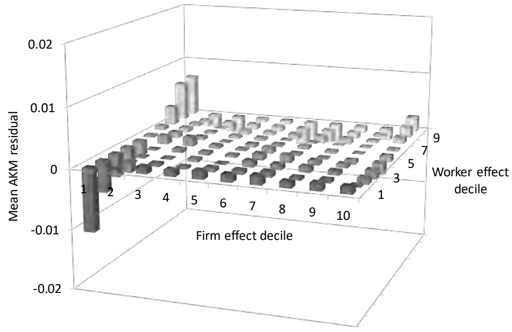

(b) Male sample

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 – that is, the fraction of firm's surplus shared with employees – 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 the relationship between firm effects and firm's average surplus. We measure the latter with average log value added per worker over the longest period available for each firm.[5] Figure C.1 plots the

[5]The coverage of balance sheet data in AIDA-Bureau Van Dijk is limited in the 1990s and early 2000s. The use of average value added allows us to impute average quantities to missing values. 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 balance-sheet information for 39,986,670 person-year observations.

relationship between male and female firm effects against average log value added per worker.[6] The relationship is clearly positive and, as value added increases, female firm effects increase less than male firm effects. Moreover, the relationship is rather flat in the first 10 percentiles of value added and only after this threshold it starts to be increasing.[7] 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.

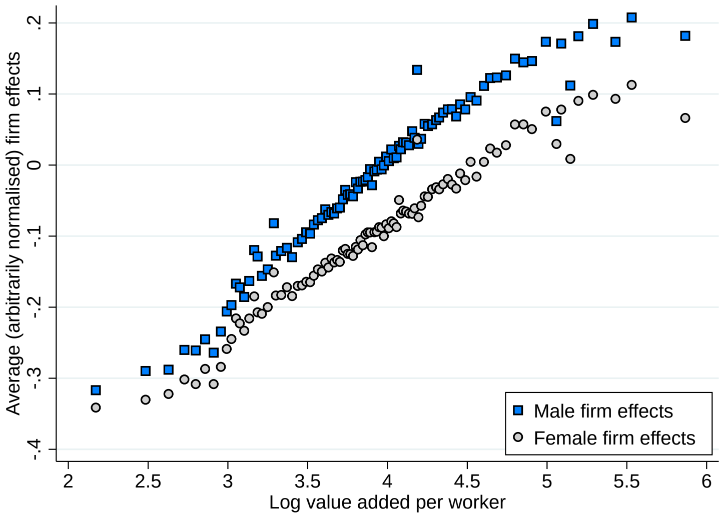

We decompose firm effects as in equations (4) and (5). Results are reported in Table C.1. With this alternative normalisation, the impact of firm components on the gender pay gap increases. The difference in firm effects accounts for 38% of the gap in weekly earnings, a 7.3 percentage points rise with respect to our preferred normalisation in the main text. Though sorting still dominates, the increase in bargaining explains the larger impact of firm effects,[8] 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 72% and 80% of the gender pay gap. For blue and white collar workers the impact is lower, around 15% and 11-14%, respectively. For middle managers, bargaining is the main driving force behind the firm contribution to the gender pay gap, as already highlighted in the main text. 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

[6]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 bins of log value added per worker.

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

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

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 qualitatively unchanged, we choose to keep as many observations as possible in the analysis.

Table C.1: Gender pay gap, firm effects gap, 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	0.246	0.168	0.207	0.300	0.408	0.355
Female firm effects across females	0.166	0.131	0.102	0.214	0.368	0.282
Firm effects gap	0.081	0.036	0.105	0.086	0.040	0.074
% of gender pay gap	*38.0%*	*88.1%*	*46.5%*	*31.8%*	*32.5%*	*31.5%*
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	0.037	0.033	0.036	0.037	0.049	0.048
% of gender pay gap	*17.4%*	*80.2%*	*15.8%*	*13.6%*	*39.7%*	*20.3%*
Using female distribution	0.032	0.029	0.035	0.030	0.044	0.026
% of gender pay gap	*15.2%*	*71.6%*	*15.4%*	*10.9%*	*35.6%*	*11.2%*

Notes. The table reports results of the Oaxaca-Blinder decomposition of equations (4) and (5). Firm effects are normalised with respect to the average firm effects in the group of firms in the first decile of the distribution of average log value added per worker. Column (1) shows results for all workers. Columns (2) to (6) report results for subsamples defined by occupation categories: apprentice, blue-collar, white-collar, middle manager and executive.

D Additional Figures and Tables

Table D.1: 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. The table reports results of the Oaxaca-Blinder decomposition of equations (4) and (5) in four overlapping time intervals, indicated in the column headers. Firm effects are estimated separately in each time interval and normalised with respect to the average firm effects in food and accommodation in each period.

Table D.2: Covariate balance, before and after matching

	(1) Unmatched	(2) Matched
Value added per worker	0.000***	-0.000
	(0.000)	(0.000)
Sales per worker	-0.000***	0.000
	(0.000)	(0.000)
Male worker effects	0.016***	0.104
	(0.004)	(0.109)
Female worker effects	0.005*	-0.087
	(0.003)	(0.122)
Share women above 90th perc.	0.082***	1.833
	(0.031)	(1.697)
Share permanent workers	0.022***	0.138
	(0.005)	(0.225)
Share part-time workers	-0.009	0.096
	(0.008)	(0.483)
Share female part-time workers	-0.005*	0.086
	(0.003)	(0.283)
Female hiring rate	0.011***	0.096
	(0.004)	(0.098)
Share workers 35-54 years old	0.003	-0.187
	(0.007)	(0.190)
Share workers older than 55	-0.008	0.120
	(0.012)	(0.395)
Log weekly earnings	0.031***	0.105
	(0.007)	(0.197)
Log female weekly earnings	0.007	0.001
	(0.006)	(0.182)
Share executives	0.244***	0.620
	(0.040)	(0.407)
Share female executives	-0.099***	-0.337
	(0.035)	(0.373)
Log firm size	-0.044***	-0.039
	(0.007)	(0.054)
Log firm size squared	0.008***	0.007
	(0.001)	(0.005)
Observations	57,117	1,780
R-squared	0.097	0.053

Notes. The Table reports estimates from regressions where the dependent variable is a dummy for treated firms, i.e. continuously listed firms over the period 2011-2014. All regressors are average values over 2008-2011. Column (1) and (2) shows results for unmatched and matched samples, respectively. Robust standard errors in parentheses. Significance levels: *** p<0.01, ** p<0.05, * p<0.10.

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

Moves	Frequency	Mean Log Real Weekly Earnings -2	-1	0	+1	4 Year Change 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' moves 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 the change for movers from quartile q to quartile q from the raw change for movers from quartile q to quartile q', with $q \neq q'$.

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

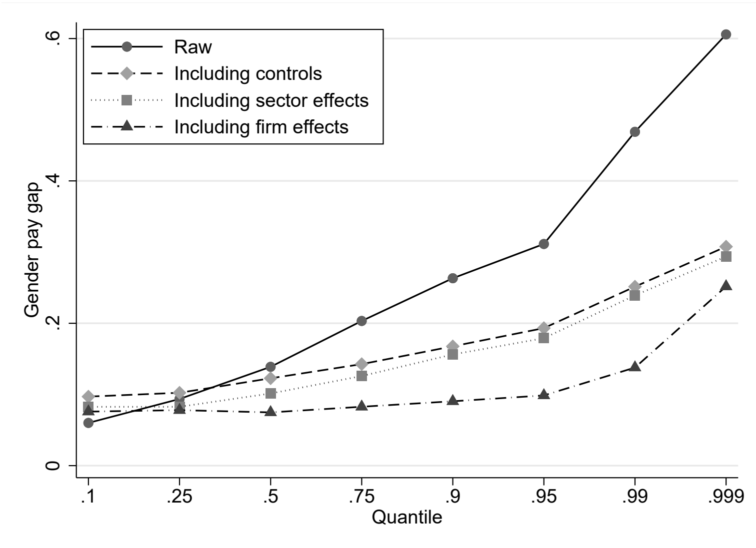

Notes. The graph plots the coefficients on the male dummy in a quantile regression in four different specifications: without controls ("Raw"); controlling for observable characteristics of workers, i.e. cubic polynomials in age, experience and tenure, a dummy for full-time contract, the number of weeks worked, occupation and province of work fixed effects ("Including controls"); controlling for observable characteristics and, additionally, for sector fixed effects ("Including sector effects"); controlling for observable characteristics and, additionally, for firm fixed effects ("Including firm effects"). Fixed effect quantile regressions are estimated in two steps, following Canay (2011). The first step consists in running an OLS regression of weekly earnings on observables and fixed effects. The second step consists in running a canonical conditional quantile regression, where the dependent variable is the residual of earnings from fixed effects computed in the first step.

Figure D.2: Firm effects gap, sorting and bargaining as a percentage of the gender pay gap by age and cohort

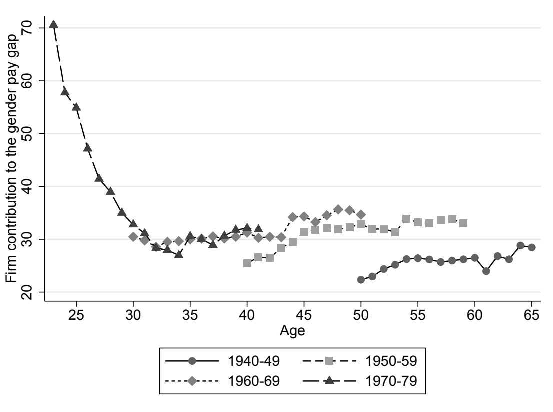

(a) Firm effects gap

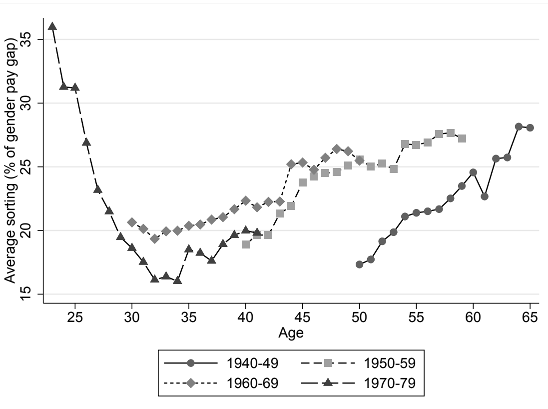

(b) Sorting

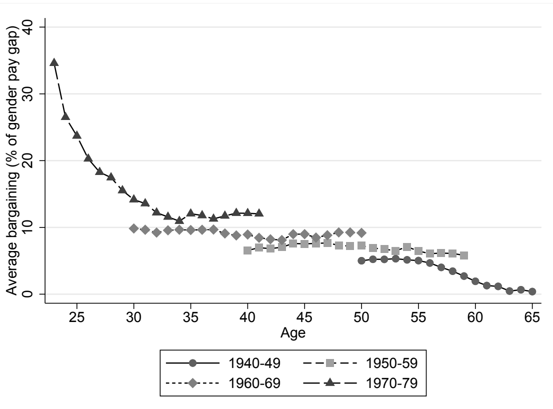

(c) Bargaining

Figure D.3: Impact of firms on the gender pay gap along the earnings distribution in 1995, 2000, 2005, 2010

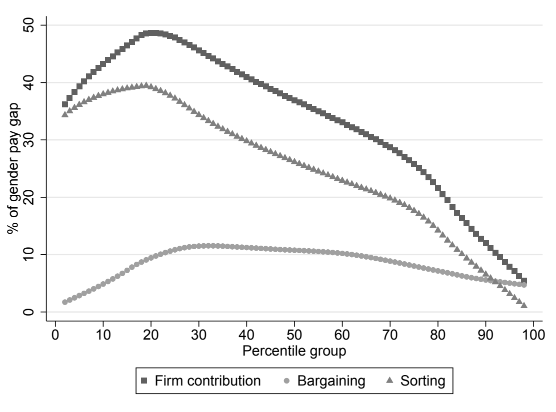

(a) 1995

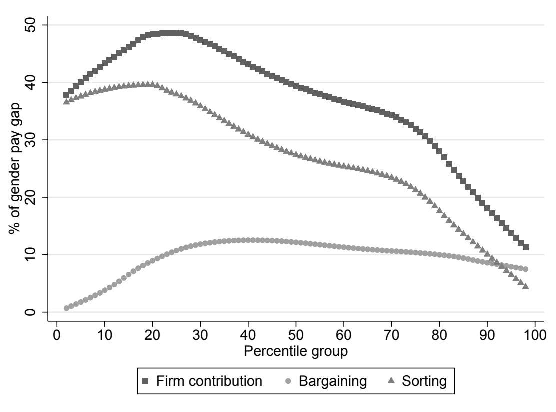

(b) 2000

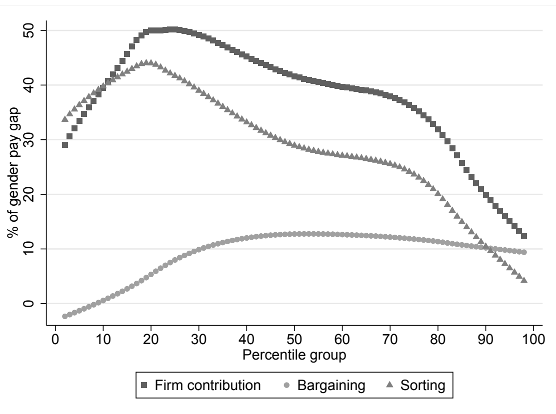

(c) 2005

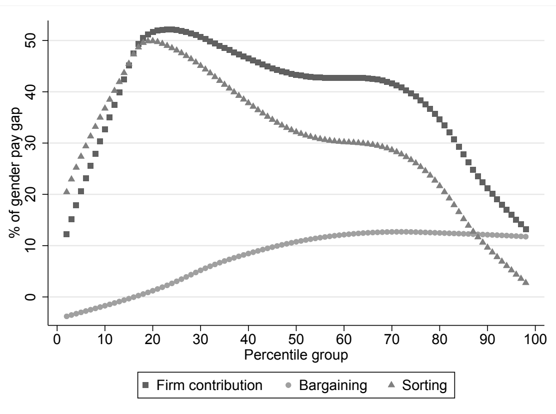

(d) 2010

Figure D.4: Sorting and bargaining as a percentage of the gender pay gap in 1995-2000 and 2010-2015

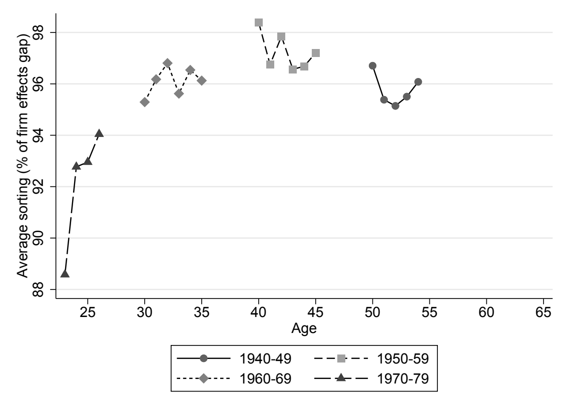

(a) Sorting 1995-2000

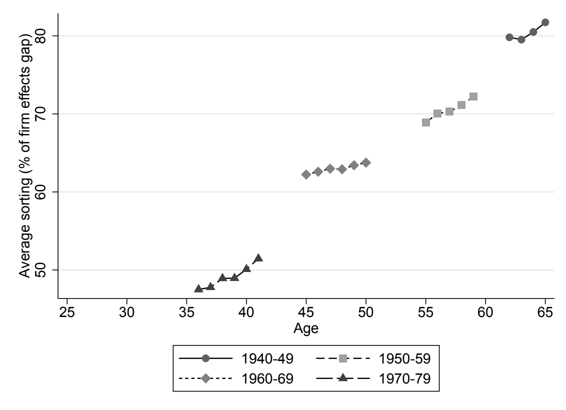

(b) Sorting 2010-2015

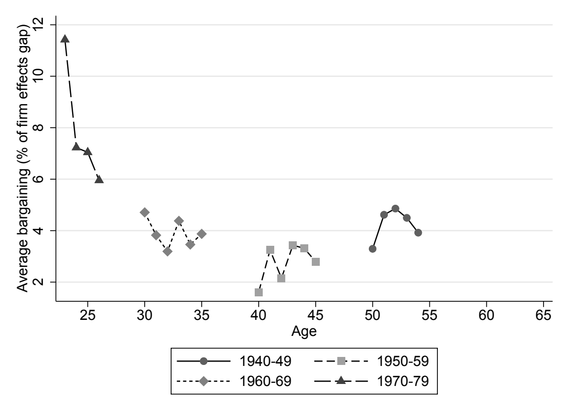

(c) Bargaining 1995-2000

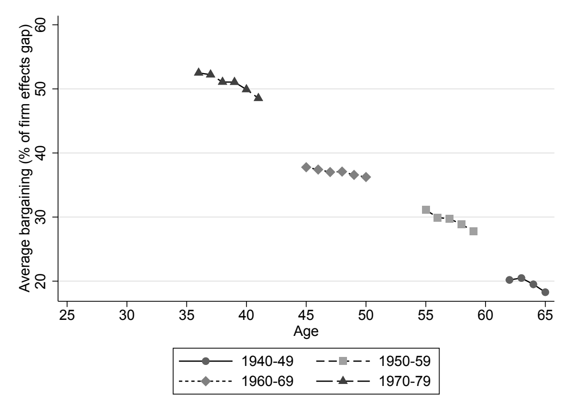

(d) Bargaining 2010-2015

Figure D.5: Mobility rates

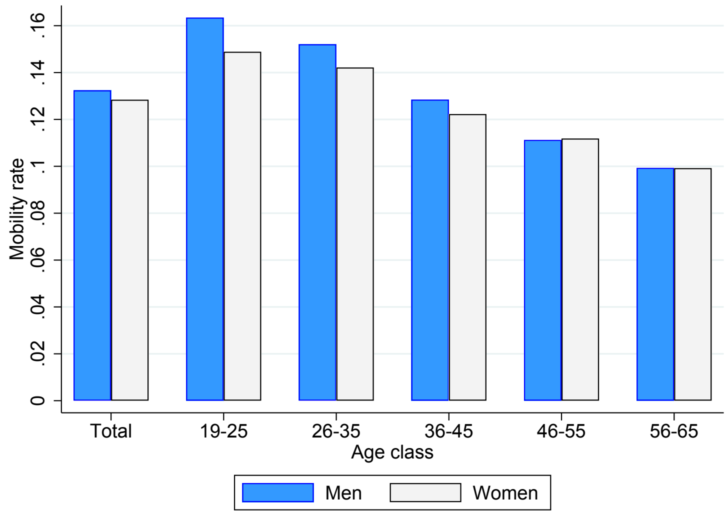

(a) Full sample

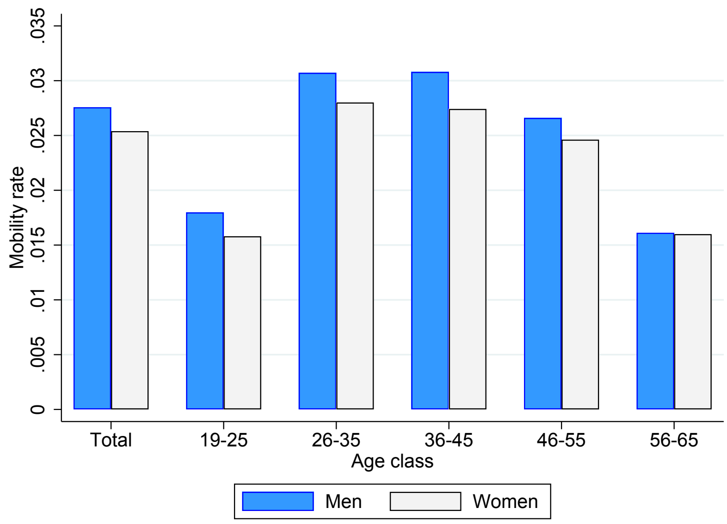

(b) Restricted sample

Notes. The mobility rate is defined as the share of workers changing employer between two *consecutive* years. The full sample (panel a) considers all moves. The restricted sample retains only moves such that the worker stays in the destination firm for at least two years after the move: this is the sample used for the analysis in section 5. All differences are statistically significant at 1% level.

Figure D.6: Gender-specific probabilities of moving to higher-quartile firm by age

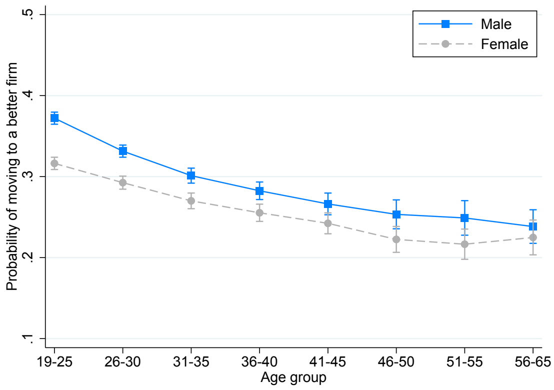

(a) All

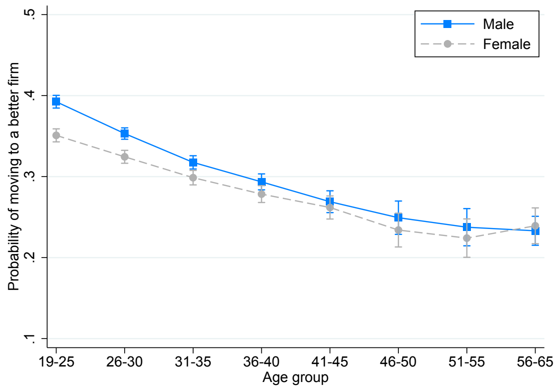

(b) Individual

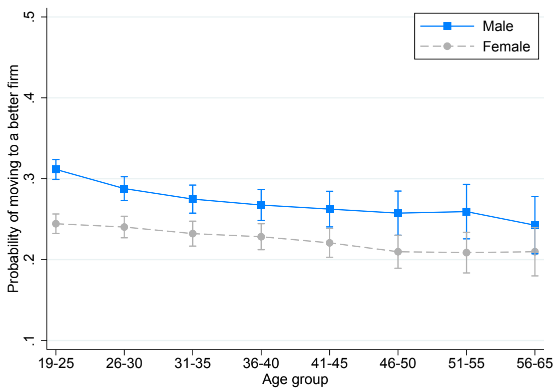

(c) Firm

Figure D.7: Average male and female firm effects in Italian provinces

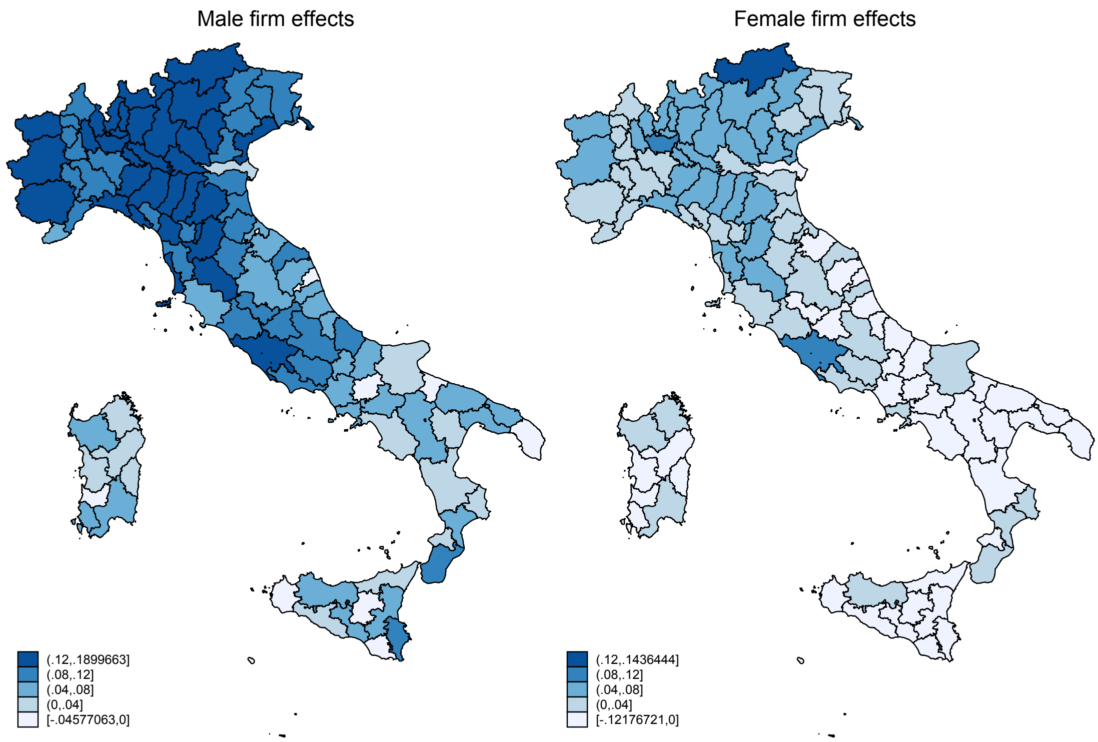

Figure D.8: Gender mobility gap within province

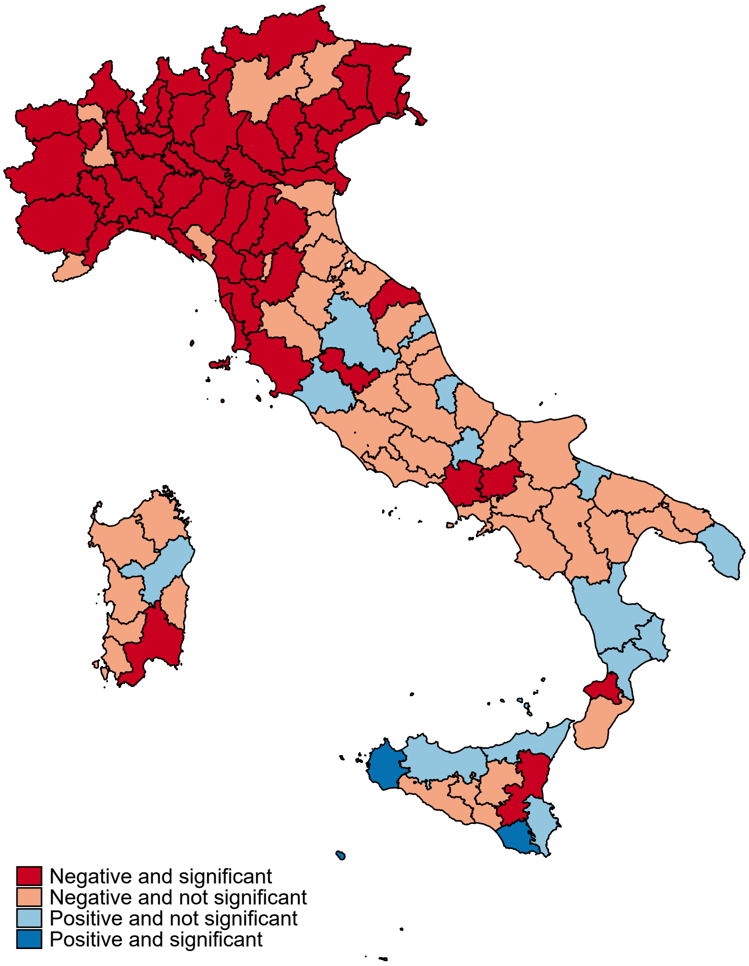

Notes. The Figure plots with different colours the marginal effect of the female dummy in a probit regression as in equation (9), estimated for each Italian province. Red (blue) areas denote provinces where the coefficient on the female dummy is negative (positive). Dark (light) areas indicate significant (not significant) coefficients at 95% confidence level.

Figure D.9: Female vs male firm effects

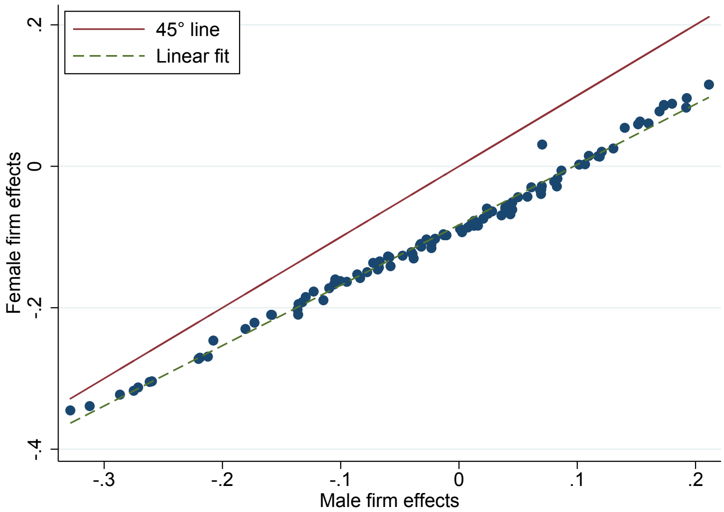

Notes. Both female and male firm effects are averaged across percentile bins of log value added. Slope of the linear fit is 0.85.

Figure D.10: Dynamic and event-study results

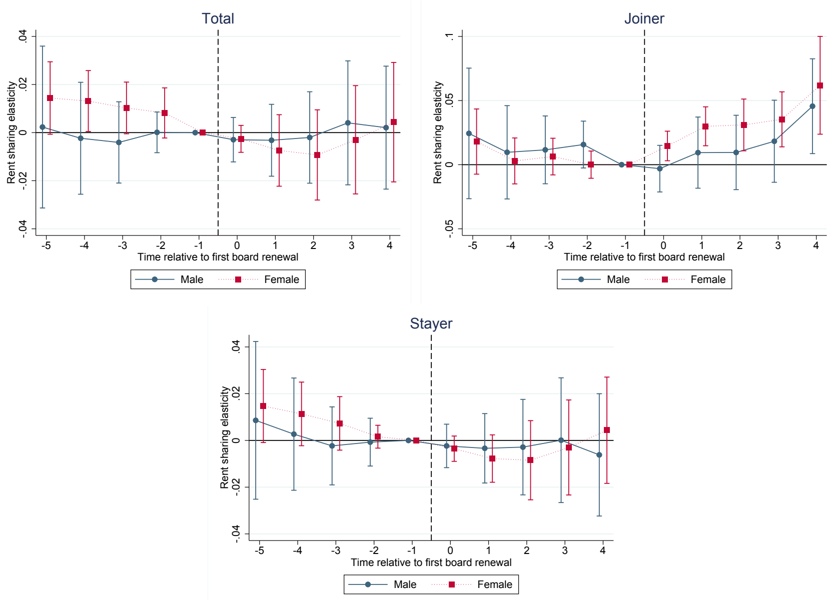

(a) Dynamic specification, listed firms

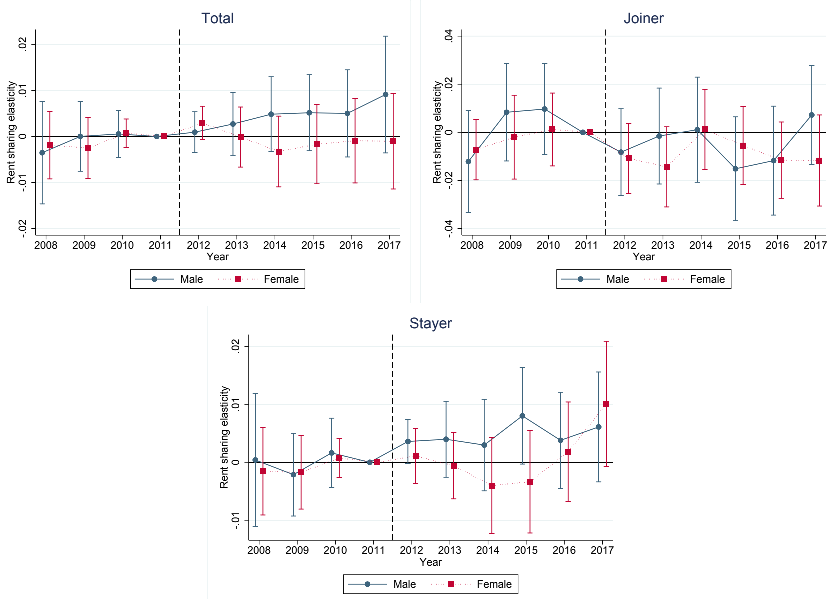

(b) Event study specification, listed vs matched non-listed firms

Notes. The Figure plots the coefficients from the estimation of dynamic (panel a) and event-study (panel b) specifications akin to equations (10) and (11). Specifically, panel (a) plots γ_k^g coefficients estimated from the following dynamic specification: $w_{jt}^g = \kappa + \sum_{k \neq -1} \gamma_k^g \times \mathbf{1}(t = k) \times \overline{S}_j^{pre} + \eta_t^g + \phi_j^g + \varepsilon_{jt}^g$. Panel (b) plots γ_k^g coefficients estimated from the following event-study specification: $w_{jt}^g = \kappa + \sum_{k \neq 2011} \gamma_k^g \times \mathbf{1}(t = k) \times Treat_j \times \overline{S}_j^{pre} + \sum_{k \neq 2011} \delta_k^g \times \mathbf{1}(t = k) \times \overline{S}_j^{pre} + \eta_t^g + \phi_j^g + \varepsilon_{jt}^g$. Vertical lines are 95% confidence intervals.